Neural-Symbolic Computing, Deep Logic Networks and Applications

Artur d'Avila Garcez
City University London
aag@soi.city.ac.uk

Schloss Dagstuhl
Leibniz-Zentrum für Informatik
17 September 2014

Neural-Symbolic Systems

Cognitive Science

Logic

Learning

Neural Computation

Neuroscience

Cognitive Neuroscience

The time is right for AI and Machine Learning informed by processes of the brain (evidence from neuropsychology, fMRI, etc.)

But computation has many desirable properties that are not present/hard to achieve in brain models:

modularity (re-use), compositionality (vs. feedback in the brain), verification (model checking), variable binding (abstraction), explanation (reasoning), usability (human-computer interaction)…

Representation precedes learning

"Logic is an attractive language of description because it has clear semantics and sound proof procedures.

However, as a basis for large systems it leads to brittleness because consistent usage of the various predicate names cannot be guaranteed.

Brittleness can be overcome by using a new kind of logic in which each statement is learnable.

But there remains the question of how such a logic can be applied when building computer systems that perform computations of a cognitive nature."

Les Valiant

Neural-Symbolic Systems

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

(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, requirements)

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 through a perception-action cycle

Applications: training in simulators, robocup, verification of software models, bioinformatics, power plant fault diagnosis, semantic web (ontology learning), general game playing, visual intelligence.

Neural-Symbolic Learning Systems

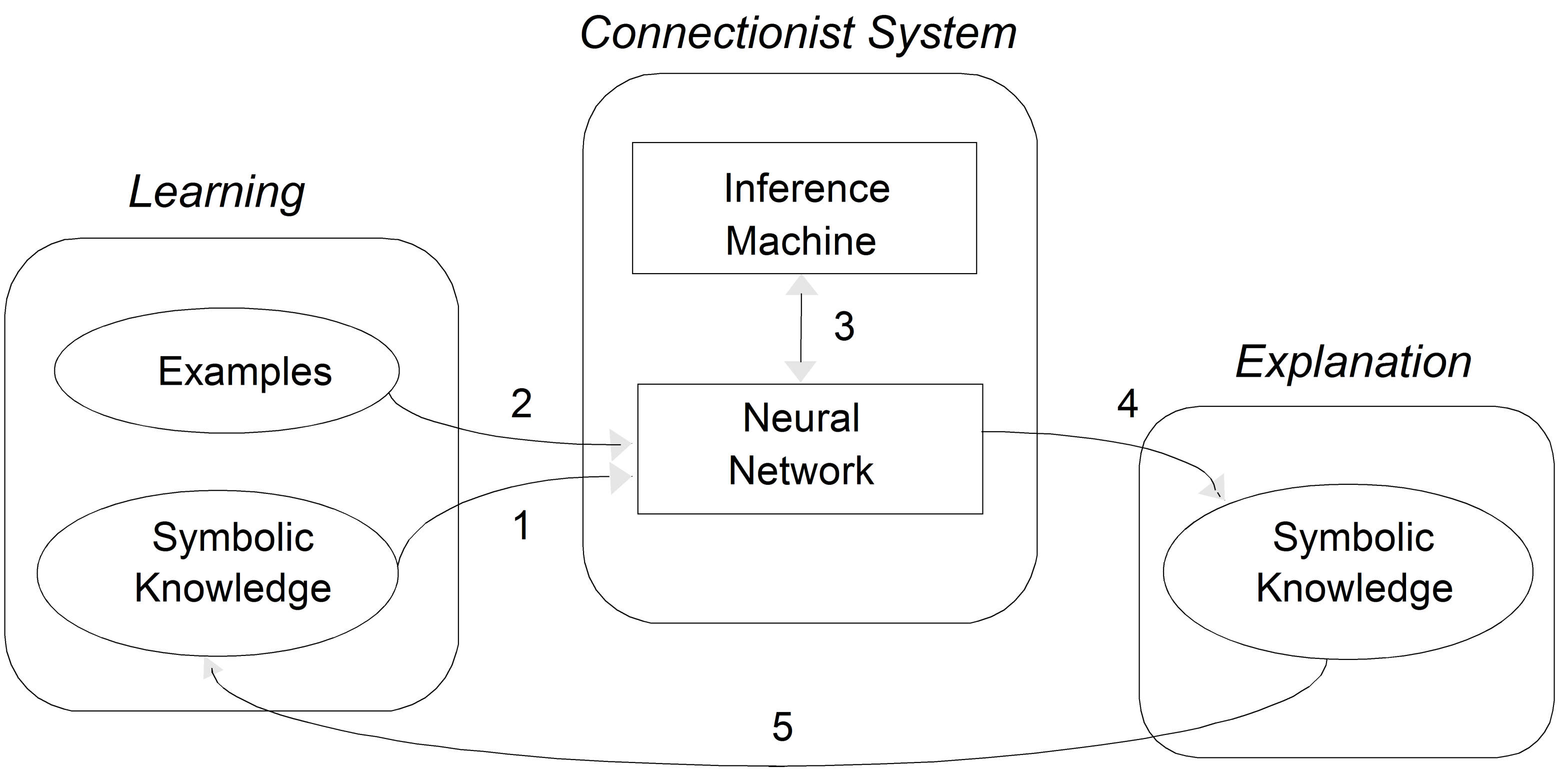

Connectionist Inductive Logic Programming (CILP) System

A Neural-Symbolic System for Integrated Reasoning and Learning

- Knowledge Insertion, Revision (Learning), Extraction

  (based on Towell and Shavik, Knowledge-Based Artificial Neural Networks. Artificial Intelligence, 70:119-165, 1994)

- Applications: DNA Sequence Analysis, Power Systems Fault Diagnosis

  (CILP using backpropagation with background knowledge:
  test set performance is comparable to backpropagation;
  test set performance on small training sets is comparable to KBANN;
  training set performance is superior than backpropagation and KBANN)

CILP Translation Algorithm

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

r_2 : $A \leftarrow E, F;$

r_3 : $B \leftarrow$

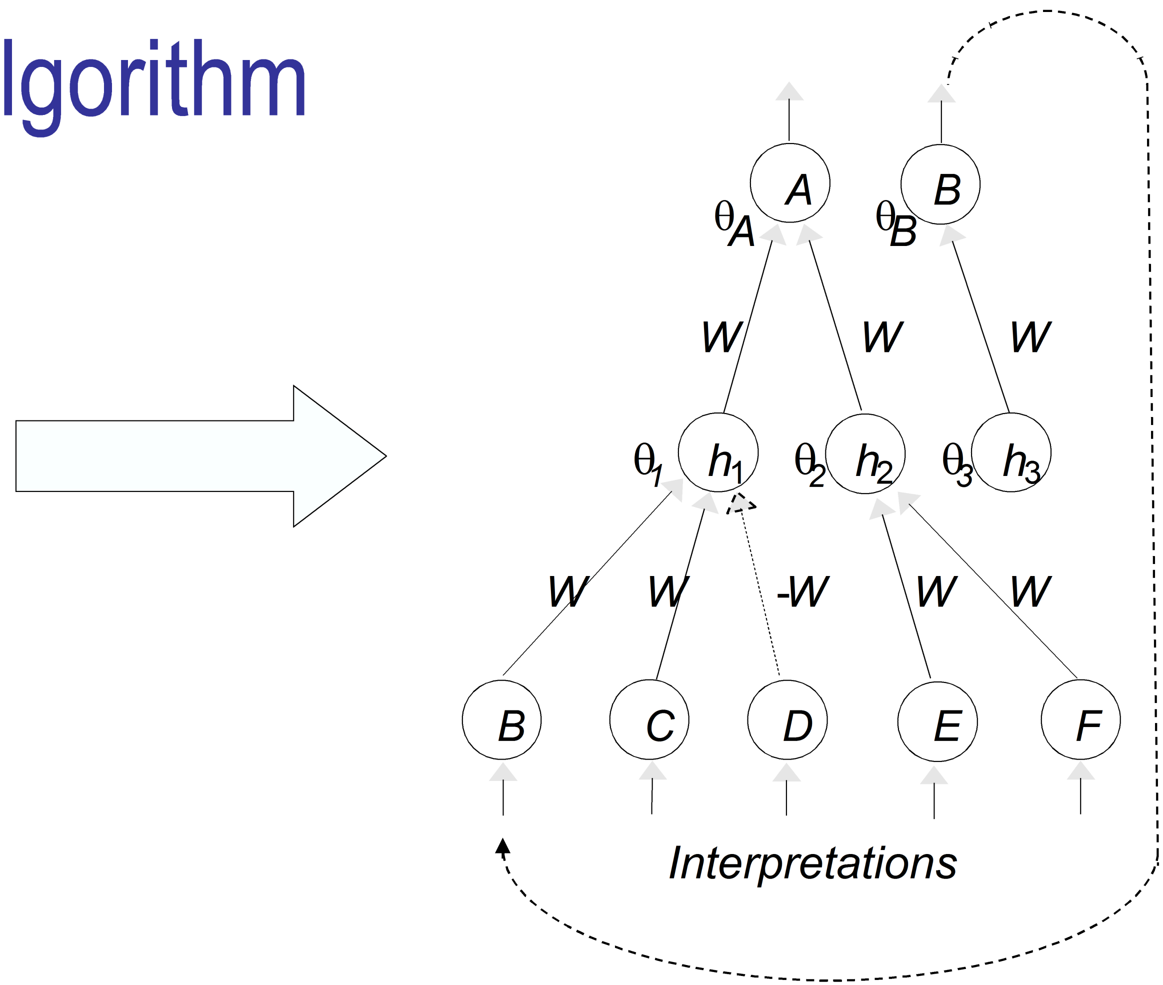

THEOREM: For any logic program P there exists a neural network N such that N computes P

based on Holldobler and Kalinke's translation, but extended to sigmoid neurons (backprop) and hetero-associative networks

Holldobler and Kalinke, Towards a Massively Parallel Computational Model for Logic Programming. ECAI Workshop Combining Symbolic and Connectionist Processing , 1994.

Power Plant Fault Diagnosis (real problem)

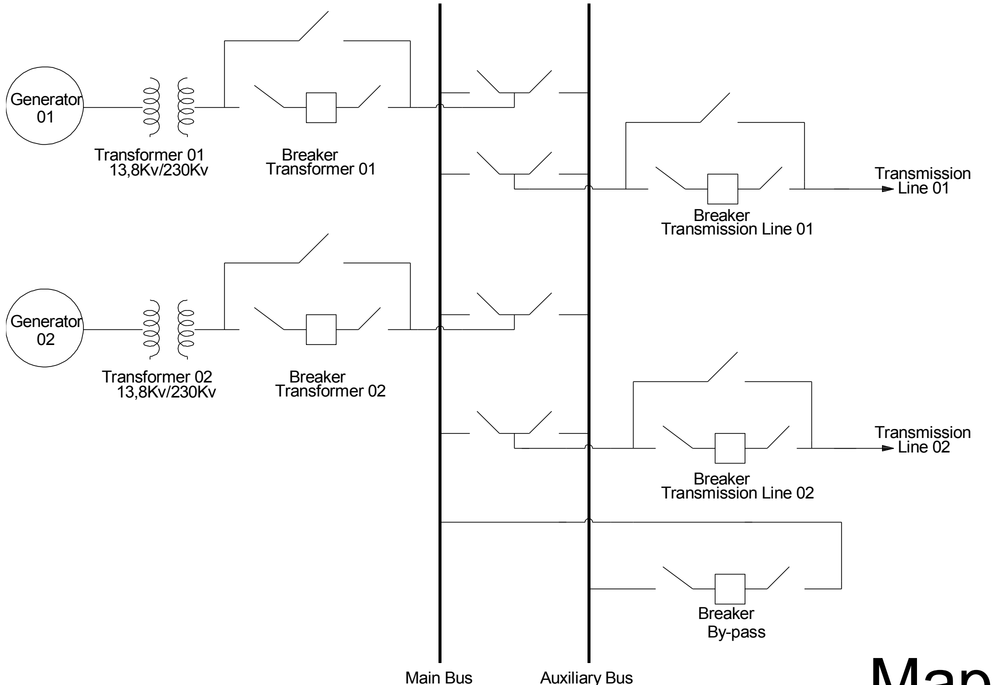

Mapping 23 alarms to 32 faults

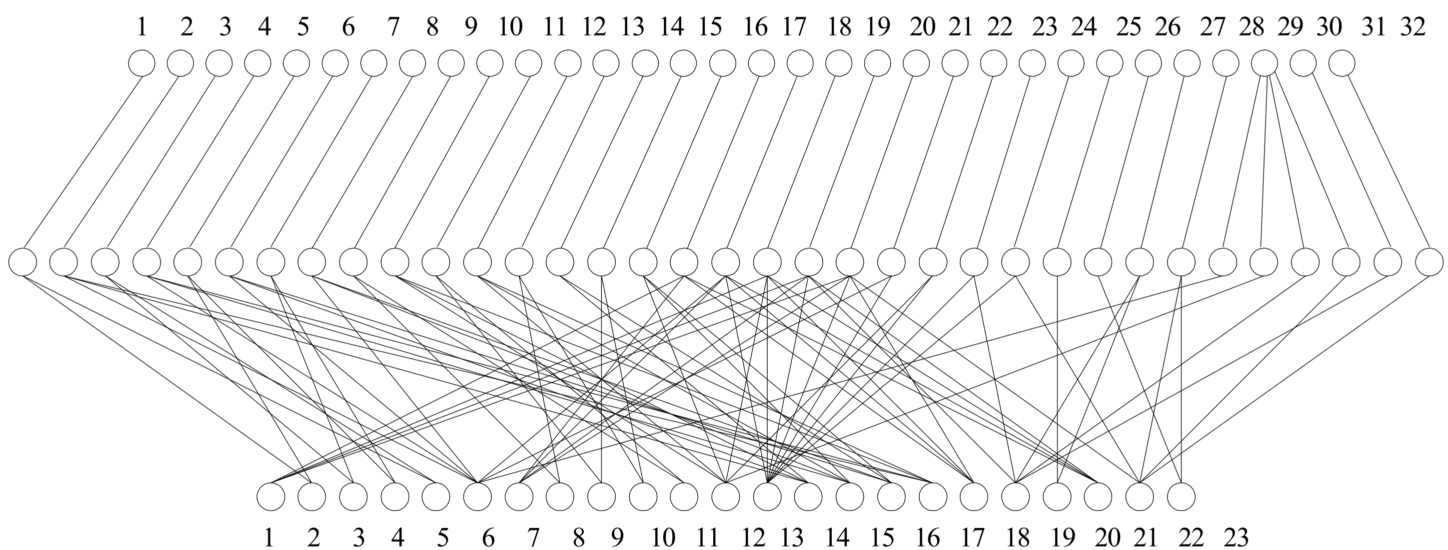

Power Plant Fault Diagnosis

Background Knowledge (35 rules with noise)

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.

Promoter Recognition (bioinformatics)

Promoter = small DNA sequence at beginning of genes

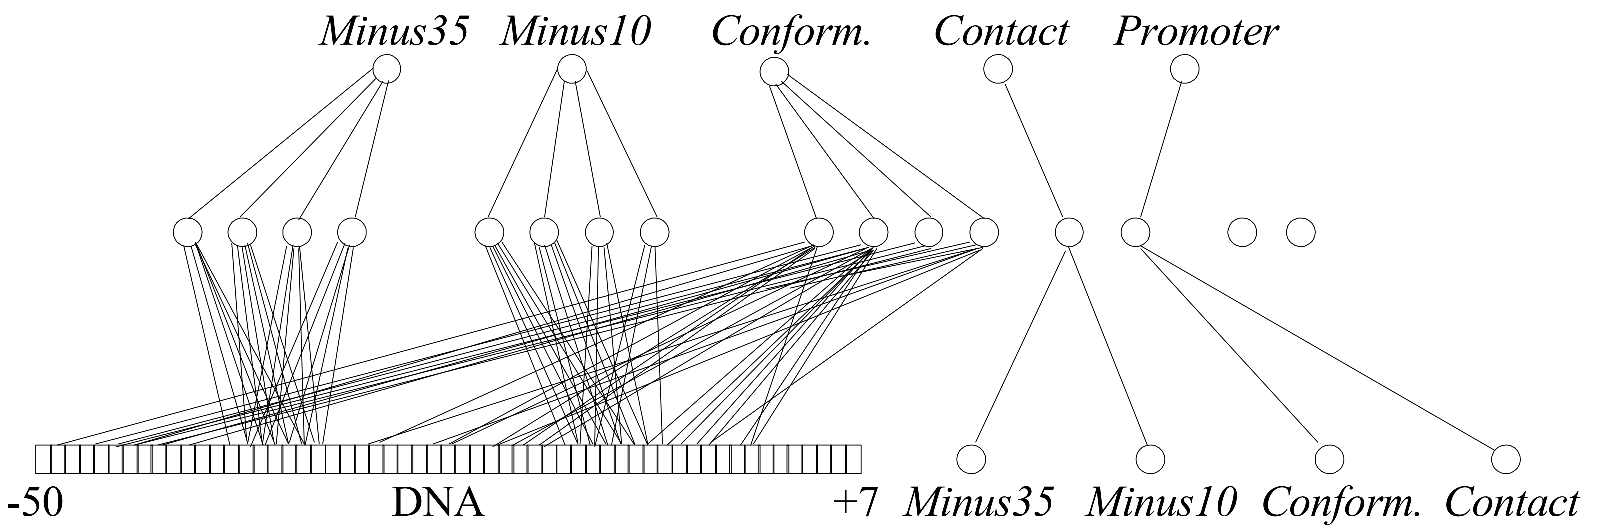

Promoter Recognition (results)

Background Knowledge (14 rules):

Promoter IF Contact AND Conformation
Contact IF Minus10 AND Minus35
Minus35 IF @-36'ttgac'
Minus35 IF @-37'cttgac'
Conformation IF @-47'caa*tt*ac' AND @-22'g***t*c' AND @-8'gcgcc*cc'

10-fold cross-validation on set of 106 examples

CILP networks learn faster than backpropagation and KBANN, and perform slightly better than backpropagation and KBANN on small training sets. We attribute this to the soundness of the CILP translation (i.e. the above theorem).

We also ran experiments on the splice-junction determination problem obtaining similar results

CILP Rule Extraction

- Knowledge is extracted by querying/sampling the trained network;
- A partial ordering helps guide the search, reducing complexity on the average case;
- A proof of soundness guarantees that the rules approximate the behaviour of the network;
- Rule simplification and visualization techniques help experts validate the rules;
- The rules can be visualized in the form of a state transition diagram

CILP Extraction Algorithm

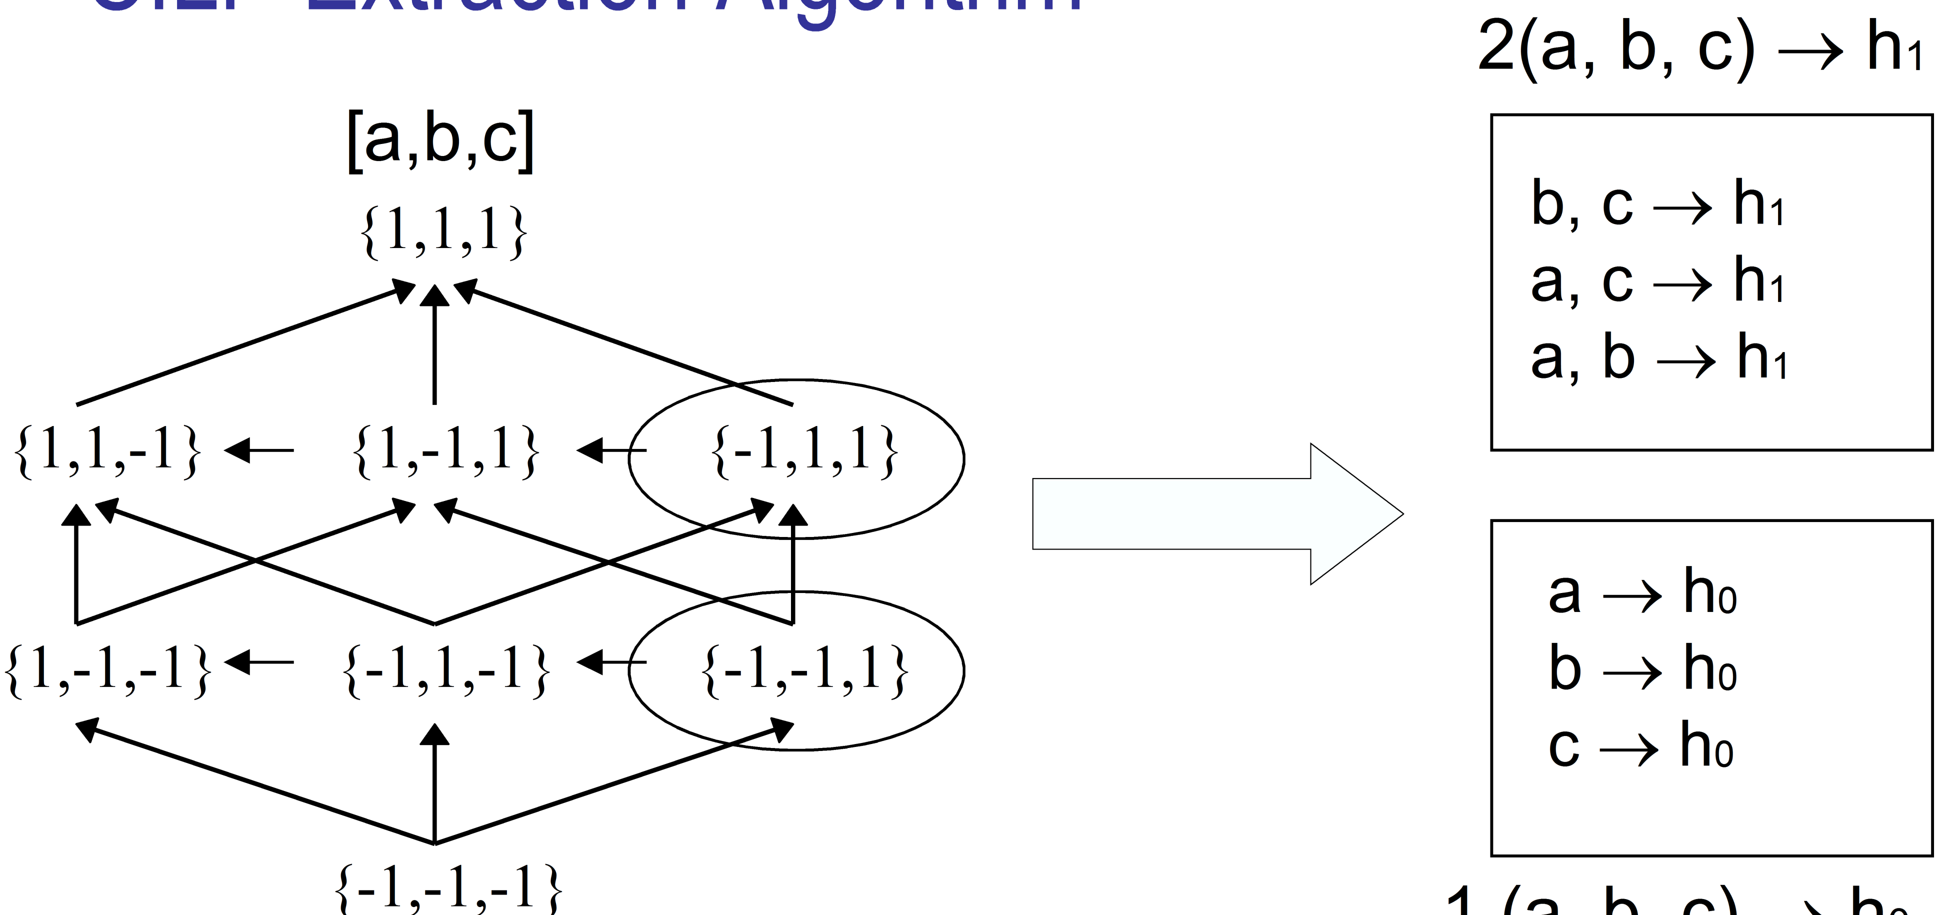

THEOREM: CILP rule extraction is sound

Challenge: efficient extraction of sound, readable knowledge from large-scale networks

Publications on CILP

d'Avila Garcez, Zaverucha. The CILP System. Applied Intelligence 11:59-77, 1999.

d'Avila Garcez, Broda, Gabbay. Knowledge Extraction from Neural Networks: A Sound Approach. Artificial Intelligence 125:153-205, 2001.

d'Avila Garcez, Broda, Gabbay. Neural-Symbolic Learning Systems. Springer-Verlag, 2002.

CILP extensions

- The importance of non-classical reasoning
- Preference, Modal, Temporal, Epistemic, Intuitionistic, Abductive Reasoning, Value-based Argumentation.
- New applications including normative reasoning (robocup), temporal logic learning (model checking), software model adaptation (business process evolution from text, e.g. email), training and assessment in simulators (driving test), visual intelligence (action classification / description in videos).

Connectionist Modal Logic (CML)

CILP network ensembles, modularity for learning, accessibility relations, disjunctive information

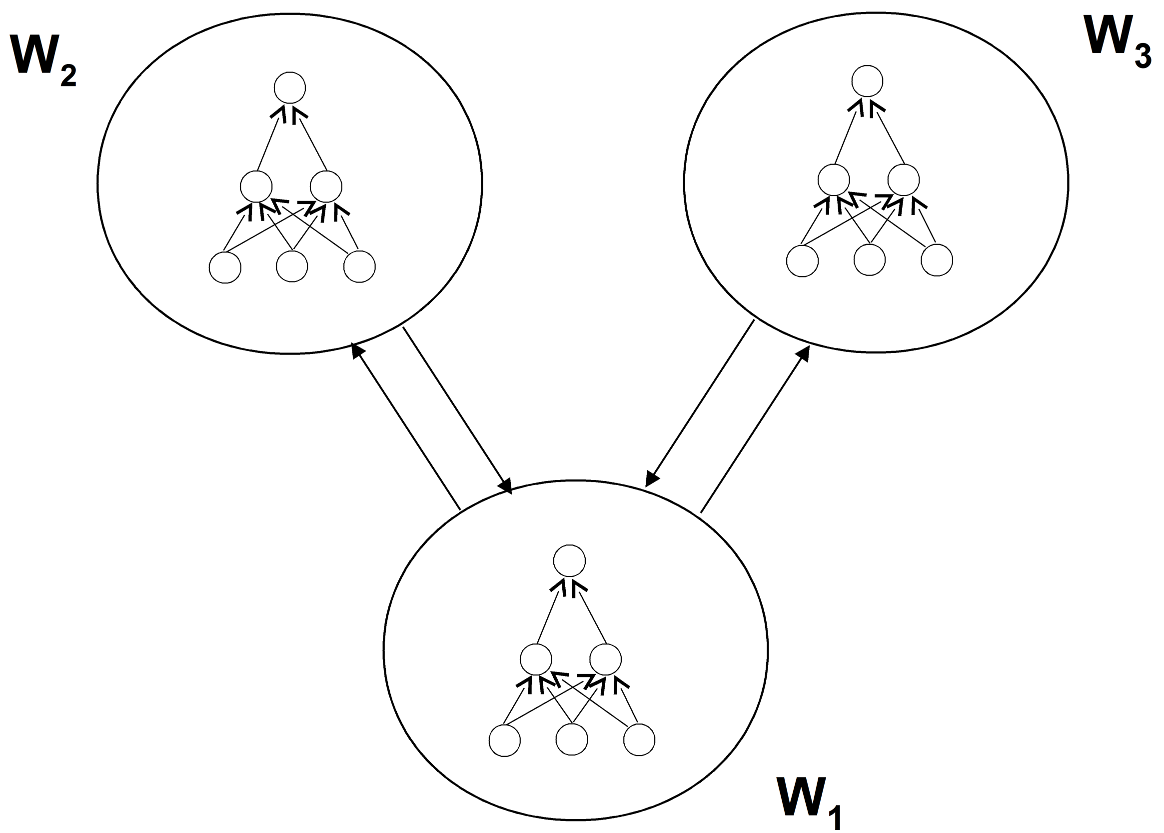

Semantics of *necessity* and *possibility*

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

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

Modalities used for reasoning about uncertainty (following J. Halpern, MIT Press).

Representing *box* and *diamond*

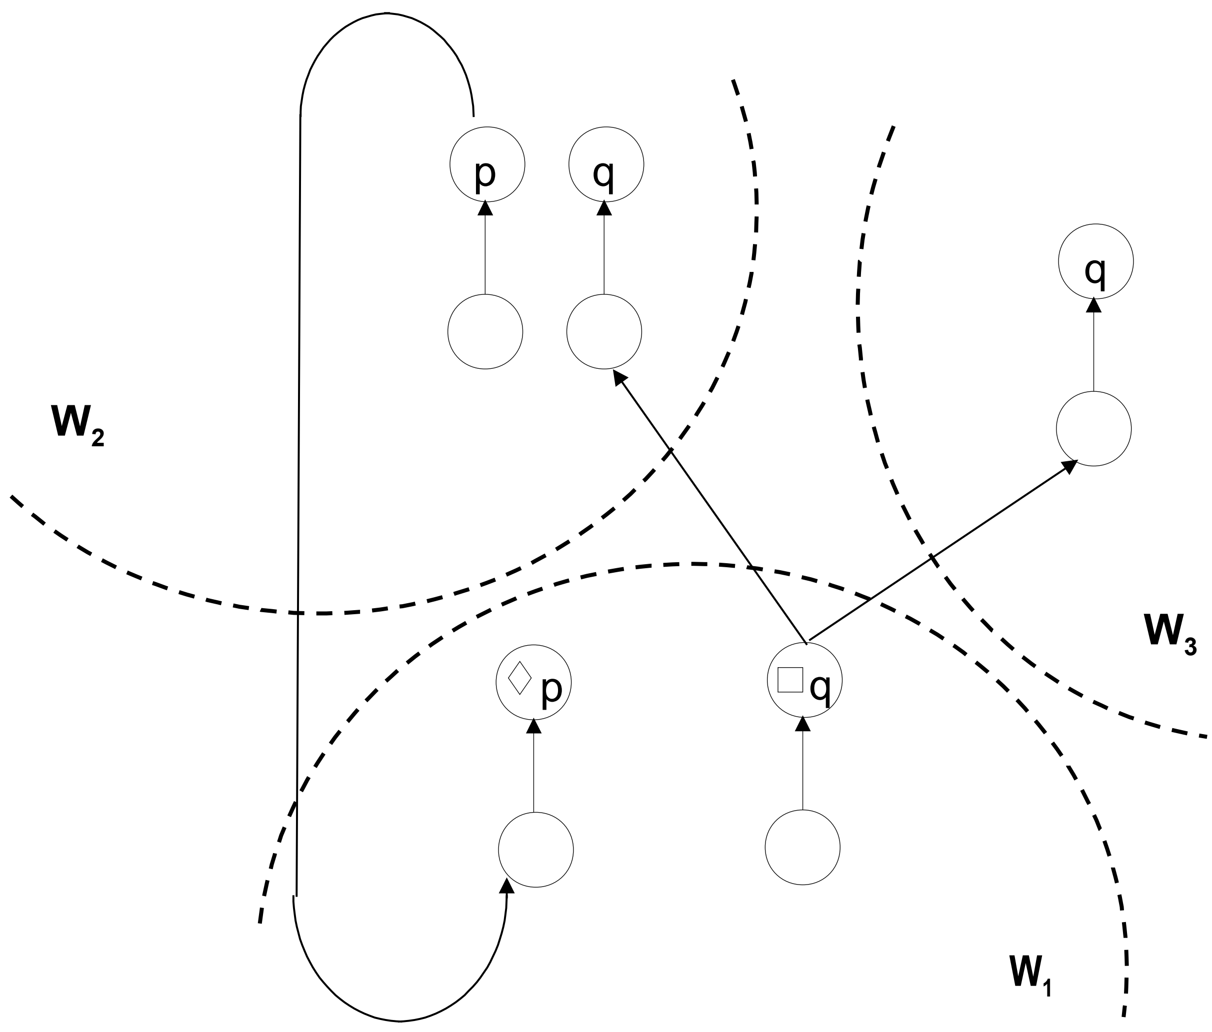

CML Translation Algorithm

Translates modal programs into ensembles of CILP networks, i.e. clauses $W_i : ML_1,...,ML_n \rightarrow MA$ and relations $R(W_a,W_b)$ between worlds W_a and W_b, with M in {*box, diamond*}.

THEOREM: For any modal program *P* there exists an ensemble of networks *N* such that *N* computes *P*.

Learning in CML

We have applied CML to a benchmark distributed knowledge representation problem: the muddy children puzzle

(children are playing in a garden; some have mud on their faces, some don’t; they can see if the others are muddy, but not themselves; a caretaker asks: do you know if you’re muddy? At least one of you is muddy)

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

Connectionist Temporal Reasoning

A full solution to the muddy children puzzle can only be given by a two-dimensional network ensemble

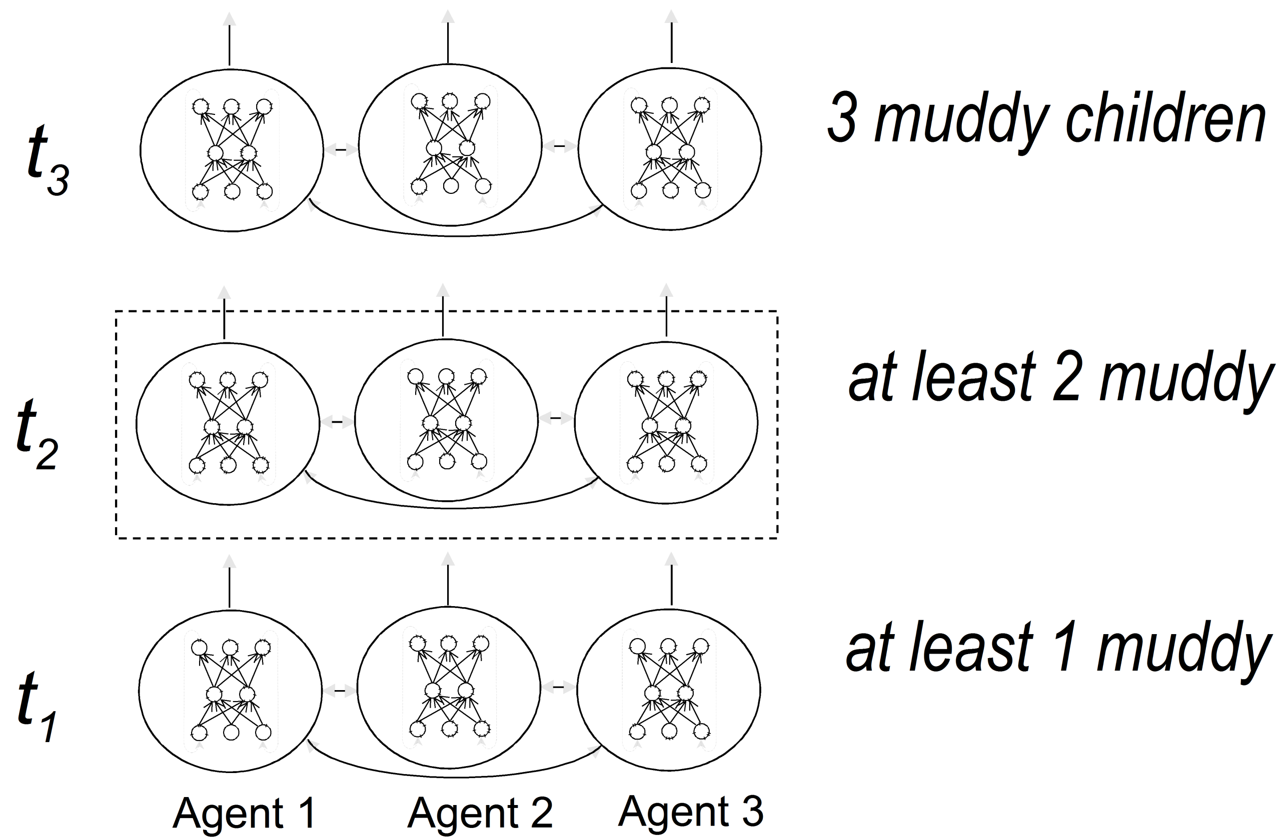

THEOREM: For any temporal program *P* there exists an ensemble of networks *N* such that *N* computes *P*.

Publications on Nonclassical Computation

Garcez, Lamb, Gabbay. Connectionist Modal Logic. TCS, 371: 34-53, 2007.

Garcez, Lamb, Gabbay. Connectionist Computations of Intuitionistic Reasoning. TCS, 358:34-55, 2006.

Garcez, Lamb. Connectionist Model for Epistemic and Temporal Reasoning. Neural Computation, 18:1711-1738, July 2006.

Lamb, Borges, Garcez. A Connectionist Cognitive Model for Temporal Synchronisation and Learning, AAAI 2007.

Recursion: Fibring Networks

A neuron that is a network! neuromodulation?

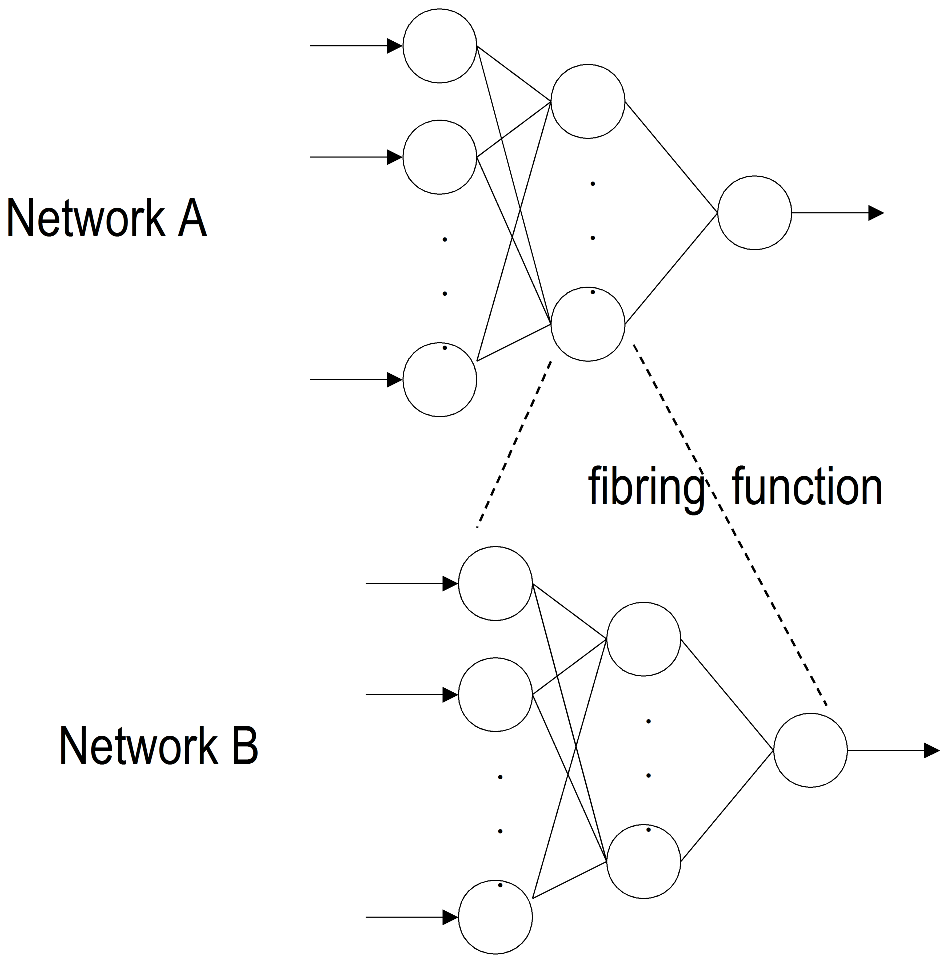

Expressiveness to represent first-order logic
Loosely-coupled integration: e.g. Network A and Legacy System B

Fibring Expressiveness

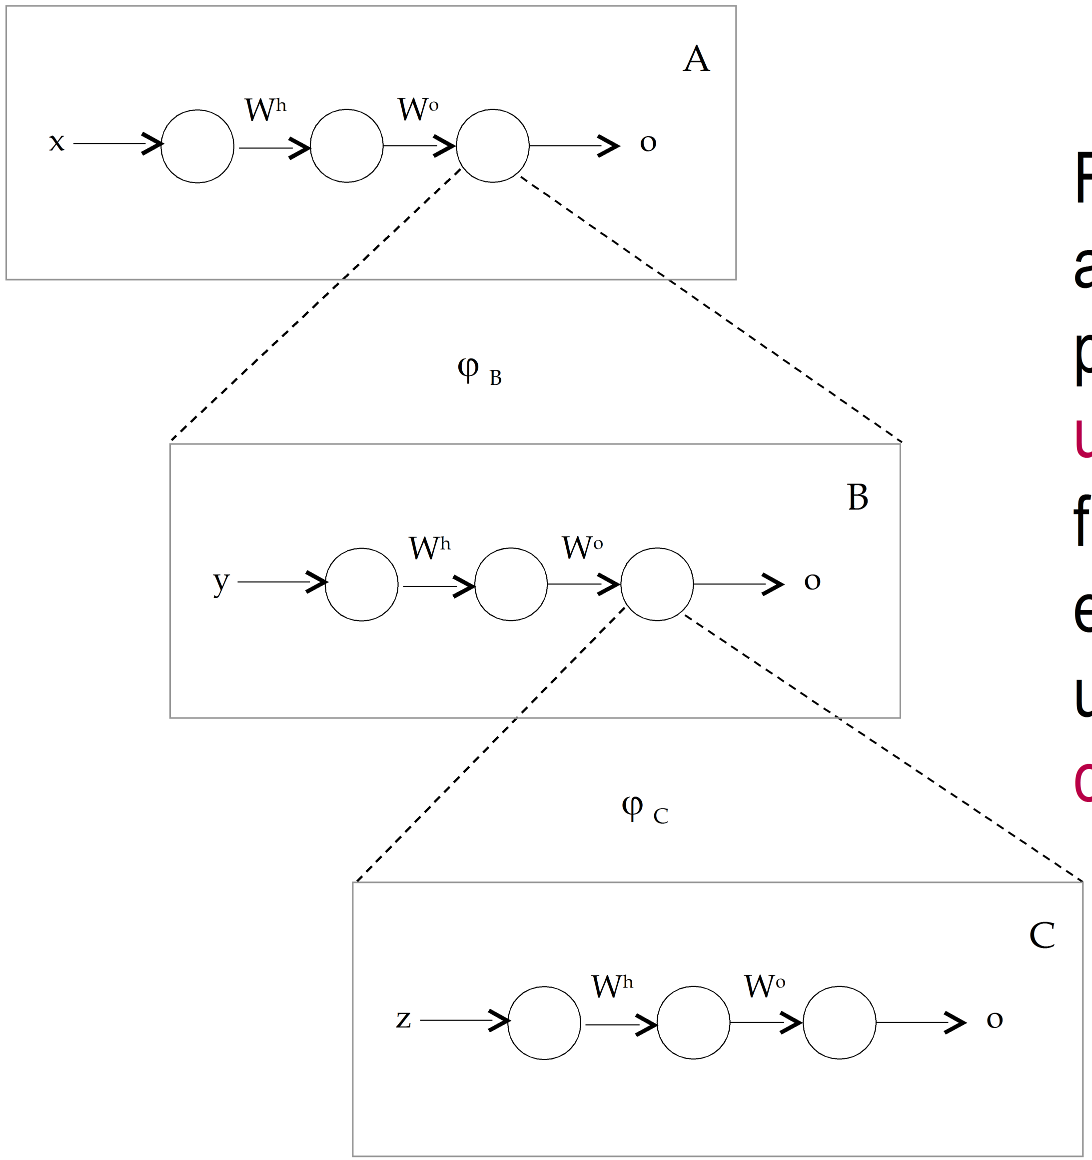

Fibred networks approximate any polynomial function in unbounded domains, e.g. $f(x)=x^2$, as opposed to each of A, B, C which are universal approximators in compact domains only.

Cognitive Model: Fibred Network Ensembles

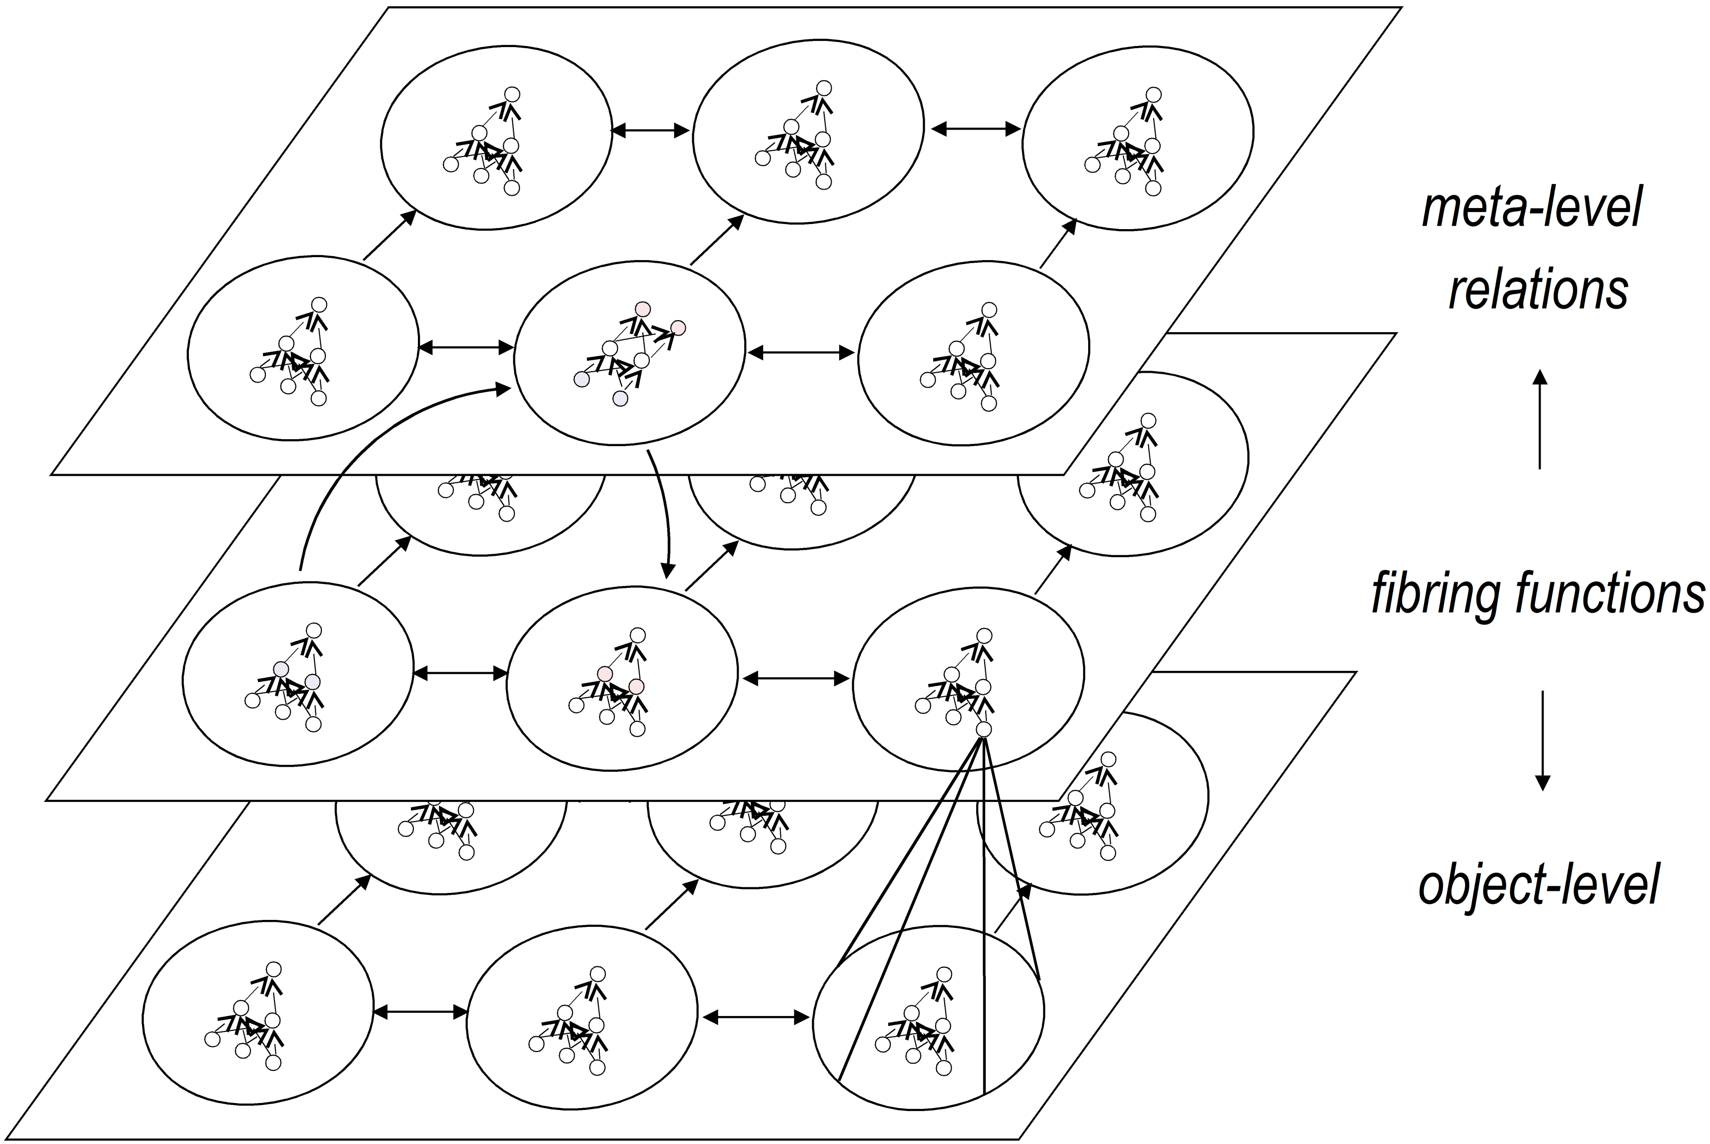

Publications on Fibring and Cognitive Model

Garcez, Lamb, Gabbay. Neural-Symbolic Cognitive Reasoning. Springer, 2009.

Garcez, Gabbay, Ray, Woods. Abductive Reasoning in Neural-Symbolic Systems. Topoi 26:37-49, 2007.

Bader, Garcez and Hitzler. Computing First Order Logic Programs by Fibring Artificial Neural Networks. FLAIRS Conference, AAAI Press, 2005.

Garcez, Gabbay. Fibring Neural Networks. AAAI 2004, July 2004.

Applications

Training and Assessment in Simulators

Learning new information from observation of experts and trainees at task execution and reasoning about this information online to provide feedback to the user

L. de Penning, A. S. d'Avila Garcez, L. C. Lamb and J. J. Meyer. A Neural-Symbolic Cognitive Agent for Online Learning and Reasoning. IJCAI'11, July 2011

Applications (cont.)

Learning Normative Rules of the RoboCup Competition

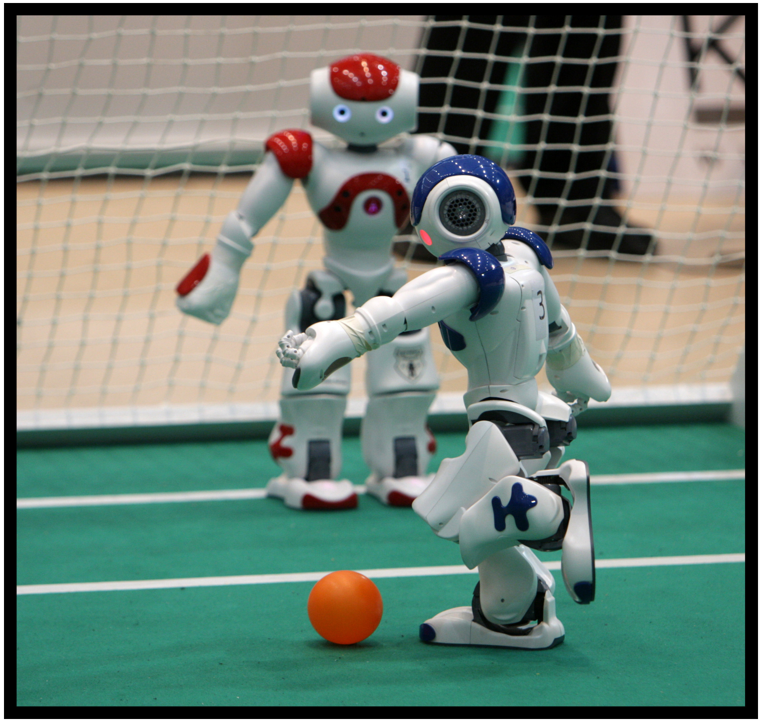

G. Boella, S. Colombo-Tosatto, A. S. d'Avila Garcez, V. Genovese, A. Perotti and L. van der Torre. Learning and Reasoning about Norms using Neural-Symbolic Systems. In Proceedings of 11th International Conference on Autonomous Agents and Multiagent Systems, AAMAS'12, Valencia, Spain, June 2012.

Applications (cont.)

Software Model Verification and Adaptation Framework

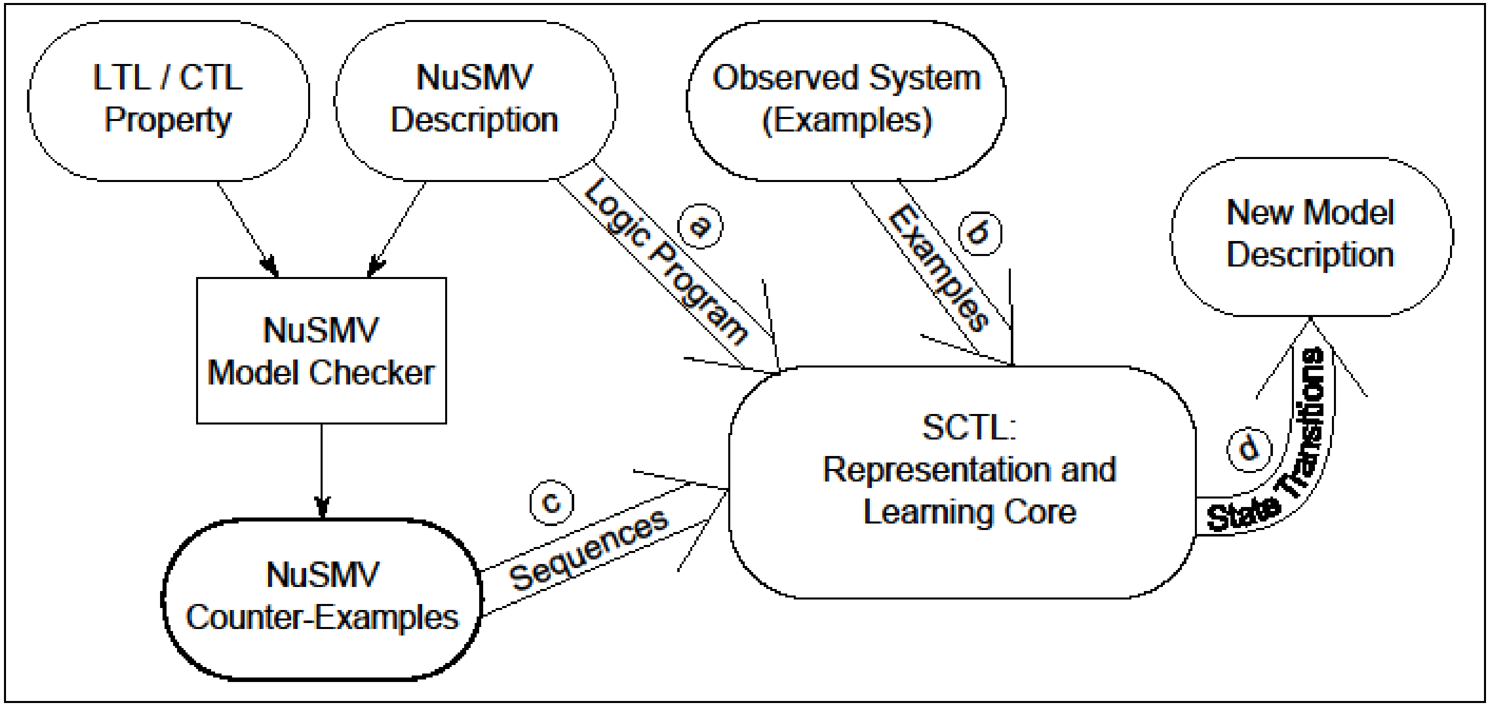

R. V. Borges, A. S. d'Avila Garcez and L. C. Lamb. Learning and Representing Temporal Knowledge in Recurrent Networks. IEEE Transactions on Neural Networks, 22(12):2409 - 2421, December 2011.

A. Perotti, A. S. d'Avila Garcez and Guido Boella, Neural Networks for Runtime Verification. The 2014 International Joint Conference on Neural Networks (IJCNN 2014), IEEE WCCI, Beijing, China, July 2014.

Applications (cont.)

Describing Actions in Video

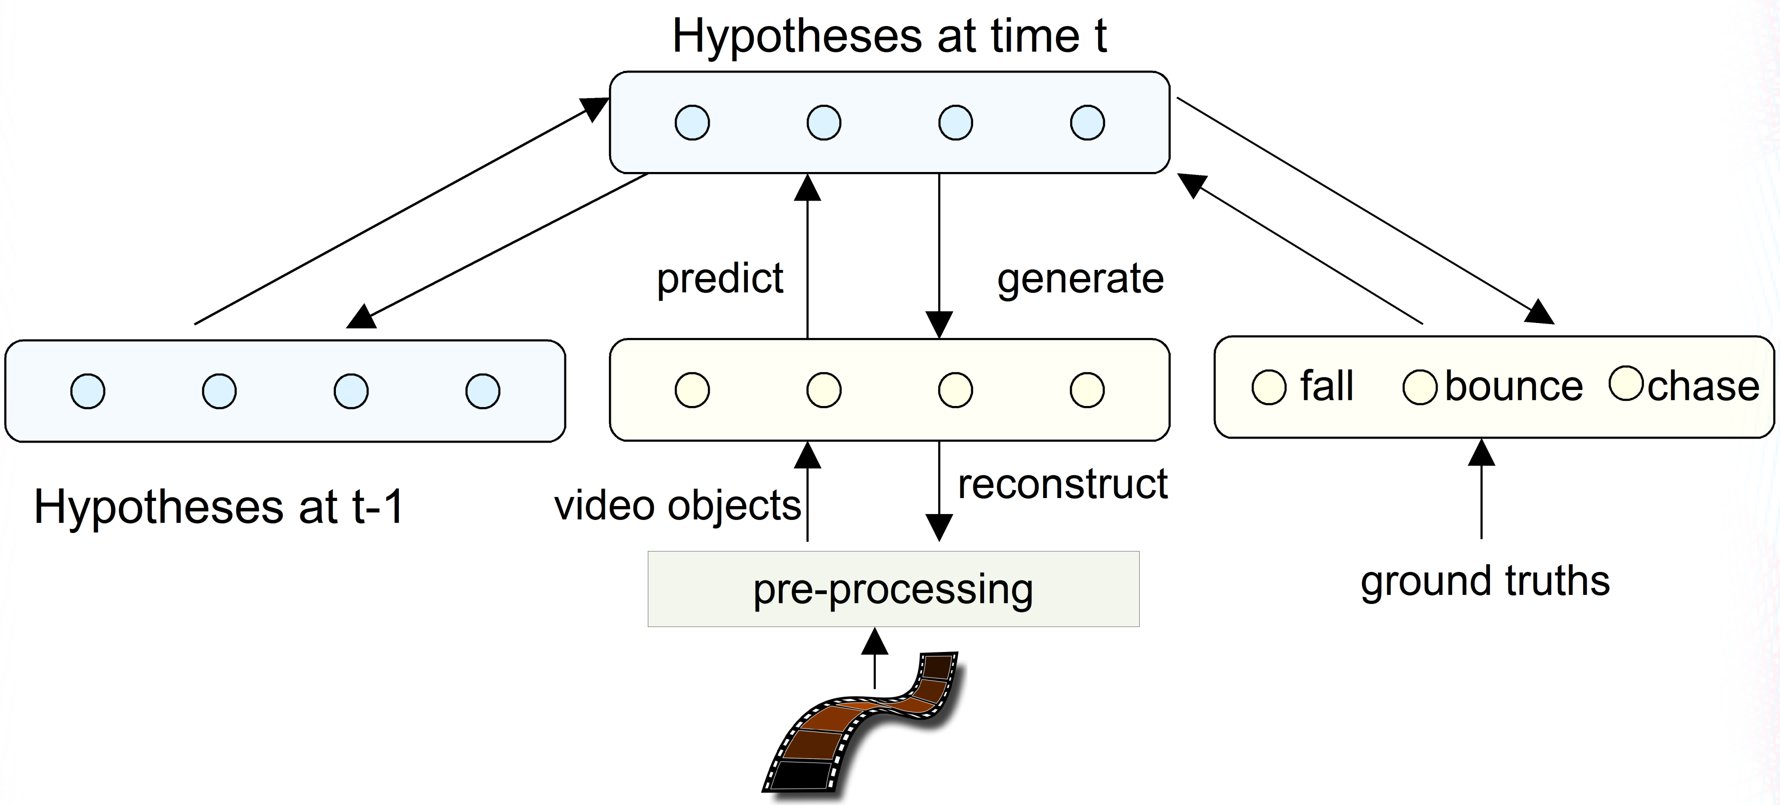

L. de Penning, A. S. d'Avila Garcez and J. J. Meyer. Dreaming Machines: On multimodal fusion and information retrieval using neural-symbolic cognitive agents. ICCSW 2013, Schloss Dagstuhl OpenAccess Series in Informatics, OASIcs, London, UK, September 2013.

Current Work (next 10 years)

Efficient First Order Logic Learning: encoding vs. propositionalisation / binding problem

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

Effective Knowledge Extraction from Very Large Networks: understanding multiple layers of abstraction, including semi-supervised learning and deep networks

Relational Learning

P. Smolensky, On the proper treatment of connectionism, BBS 11(1) 1988

J. McCarthy, Epistemological challenges for connectionism, BBS 11(1) (commentary), 1988

Andy Clark, Whatever next? Predictive brains, situated agents, and the future of cognitive science, BBS 36(13), 2013

Relevant for the analysis of complex networks: drug design in bioinformatics, link analysis in social networks, etc.

Related work: (probabilistic) ILP, deep networks, MLNs, BLOG, StarAI (lifted inference), etc.

Relational Learning in Neural Nets

A number of attempts at first-order (and higher-order) logic learning in neural networks: semantic approach (Hitzler et al), fibring, topos (Osnabruck), association (SHRUTI), unification (Pinkas, etc), etc.

CILP++ (available to download from sourceforge) is a neural-symbolic system that can solve ILP problems efficiently (through **propositionalization**) using a neural network trained with backpropagation (França, Zaverucha and d'Avila Garcez, Mach. Learn., 2014)

Knowledge Extraction from Deep Nets

Deep networks have shown good performance in image, audio, video and multimodal learning

We would like to know why by studying the role of symbolic reasoning in deep nets

In particular, we would like to find out:

- How knowledge is represented in deep architectures
- Relations between Deep Networks and a hierarchy of rules
- How knowledge can be transferred to analogous domains

Deep networks

The lower level layer is expected to capture low-level features (e.g. edges)

Higher level layers combine features to learn progressively more abstract concepts (e.g. shapes)

Labels can be attached to the top RBM for classification (e.g. digit classes)

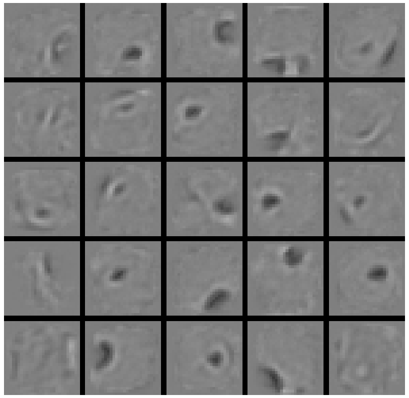

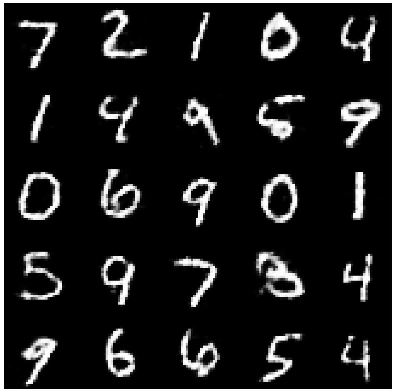

Rule extraction from (R)BMs

[Pinkas, 1995]: rule extraction from symmetric networks using penalty logic; proved equivalence between conjunctive normal form and energy functions

[Penning et al., 2011]: extraction of temporal logic rules from RTRBMs using sampling; extracts rules of the form

hypothesis(t) iff belief(1)^, ..., ^belief(n) ^ hypothesis(t-1)

[Tran and Garcez, 2012]: rule extraction from DBNs using confidence-value similar to penalty logic but maintaining implicational form; extraction without sampling

Partial model extraction

- Decompose the network into rules (pruning weights; averaging the absolute values of the remaining weights)

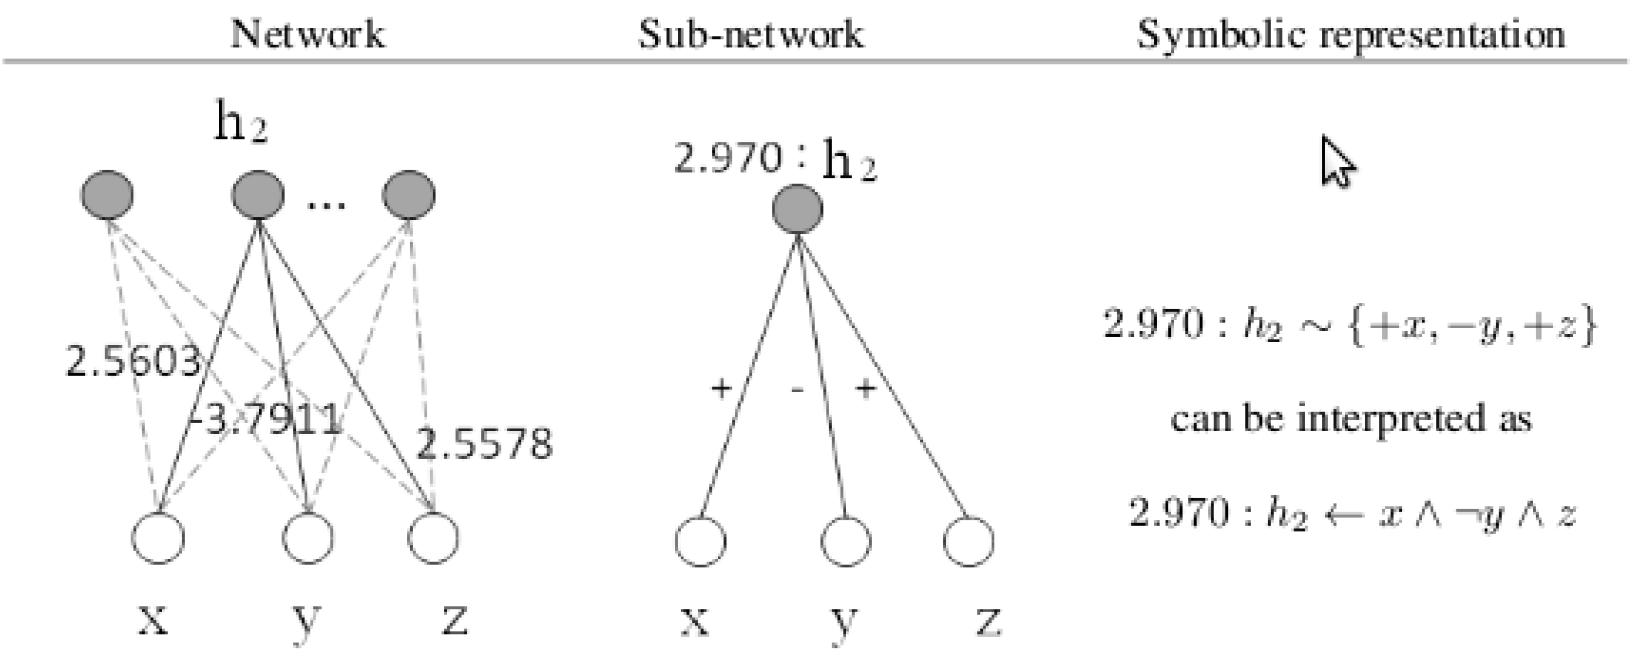

- This is guaranteed to minimize information loss [Tran and Garcez, 2013, 2014]

Inference

Given:
$c : \mathsf{h} \leftarrow \bigwedge_{\forall p \in P} \mathsf{b}_p \wedge \bigwedge_{\forall n \in N} \neg \mathsf{b}_n$
$c_{p'} : \mathsf{b}_{p'}$ where $p' \in P, c_{p'} \in [min, max]$
$c_{n'} : \mathsf{b}_{n'}$ where $n' \in N, c_{n'} \in [min, max]$
Infer:
$c_h : \mathsf{h}$ with $c_h = c \times (\sum_{p'} c_{p'} + \sum_{m \neq \forall p'}^{m \in P} c_m - \sum_{n'} c_{n'} - \sum_{k \neq \forall n'}^{k \in N} c_k)$
where $c_m = c_k = \frac{min+max}{2}$

Application: Low-cost representation for RBMs [Tran and Garcez 2014]

	TiCC	MNIST_{10K}	MNIST_{60K}	YALE face
RBM (J=500)	94.851% ± 0.033	97.198 ± 0.060	98.553% ± 0.031	95.000% ± 2.833
Rules	94.711% ± 0.072	97.240 ± 0.089	98.530% ± 0.040	94.333% ± 3.865
RBM (J=1000)	94.928% ± 0.016	97.245% ± 0.031	98.680% ± 0.024	97.000% ± 2.919
Rules	94.729% ± 0.070	97.219% ± 0.056	98.562% ± 0.035	96.667% ± 1.757

Transfer Learning

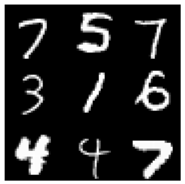

MNIST

TiCC

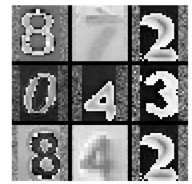

ICDAR

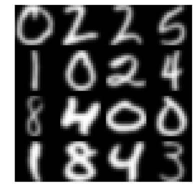

USPS

MADBASE

	MNIST_{30k} : ICDAR_d	MNIST_{30k} : TiCC_w	MNIST_{05k} : TiCC_a	MNIST_{05k} : TiCC_d	USPS : MADBASE
SVM	39.04	73.44	59.16	60.34	80.4
RBM	37.63 ± 0.505	75.20 ± 0.745	62.85 ± 0.079	63.42 ± 0.090	80.38 ± 0.120
SC STL	46.23	70.06	55.82	57.78	81.7
RBM STL	52.26 ± 0.331	72.88 ± 0.098	58.13 ± 0.205	62.08 ± 0.321	81.43 ± 0.211
RBM MIX	52.43 ± 0.132	76.49 ± 0.361	63.21 ± 0.134	65.04 ± 0.330	80.90 ± 0.253
aTPL ($\beta = 0$)	51.64 ± 0.284	77.56 ± 0.564	63.00 ± 0.160	$\mathbf{65.75 \pm 0.110}$	$\mathbf{83.607 \pm 0.151}$
aTPL ($\beta > 0$)	$\mathbf{52.68 \pm 0.104}_{\beta=5}$	$\mathbf{80.51 \pm 0.142}_{\beta=5}$	$\mathbf{63.86 \pm 0.185}_{\beta=0.01}$	$65.66 \pm 0.122_{\beta=0.01}$	$83.11 \pm 0.173_{\beta=0.01}$

Conclusion: Why Neurons and Symbols?

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

To develop effective computational systems for integrated reasoning and learning.

To apply and evaluate such systems in complex new tasks: social robotics, health informatics, big data