University of Kent, 13 June 2011

Neural-Symbolic Systems for Cognitive Reasoning

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

Neural-Symbolic Computation

Cognitive Science

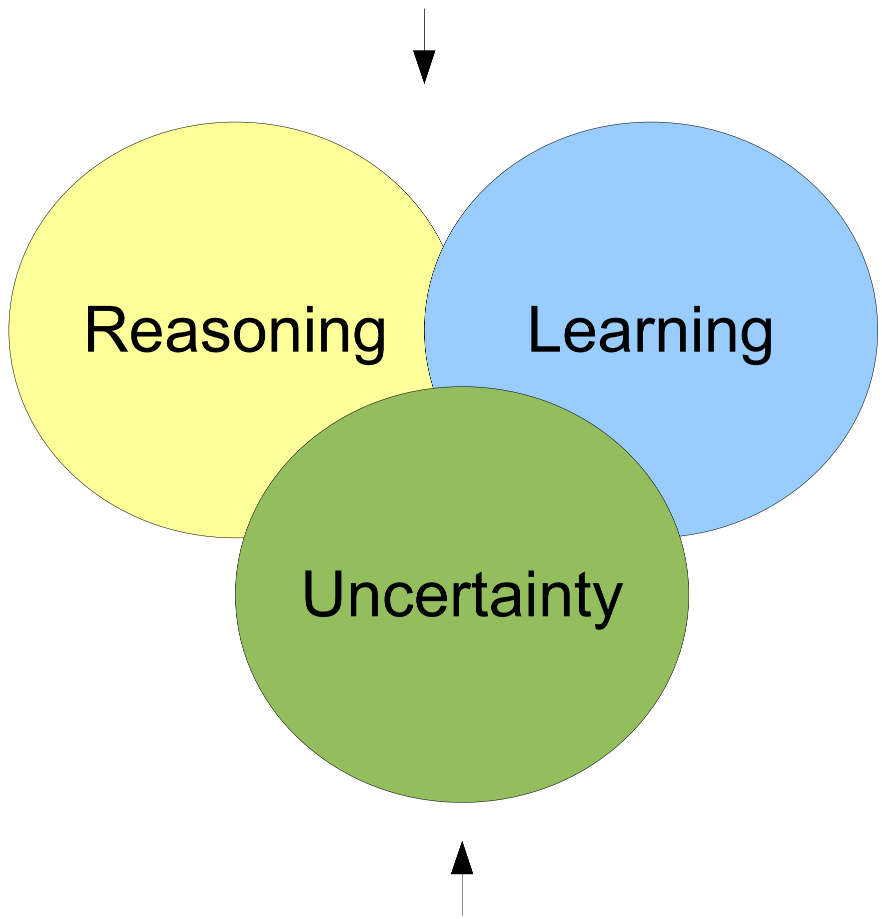

Neuroscience

Neural-Symbolic Systems

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

(New) Logic + Neural Computation

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

Artificial Intelligence

Symbolic vs. Connectionist (brain/mind) dichotomy

1960s-1980s: Expert Systems (hand-crafted rules)

1990's-present: Neural networks, Support vector machines (difficult to include domain knowledge)

New AI: Bayesian learning, probabilistic graphical models, efficient inference

IET/BCS lecture 2010, Chris Bishop

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 aspect: controlling the inevitable accumulation of errors as part of a perception-action cycle

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

Neural-Symbolic Learning Systems

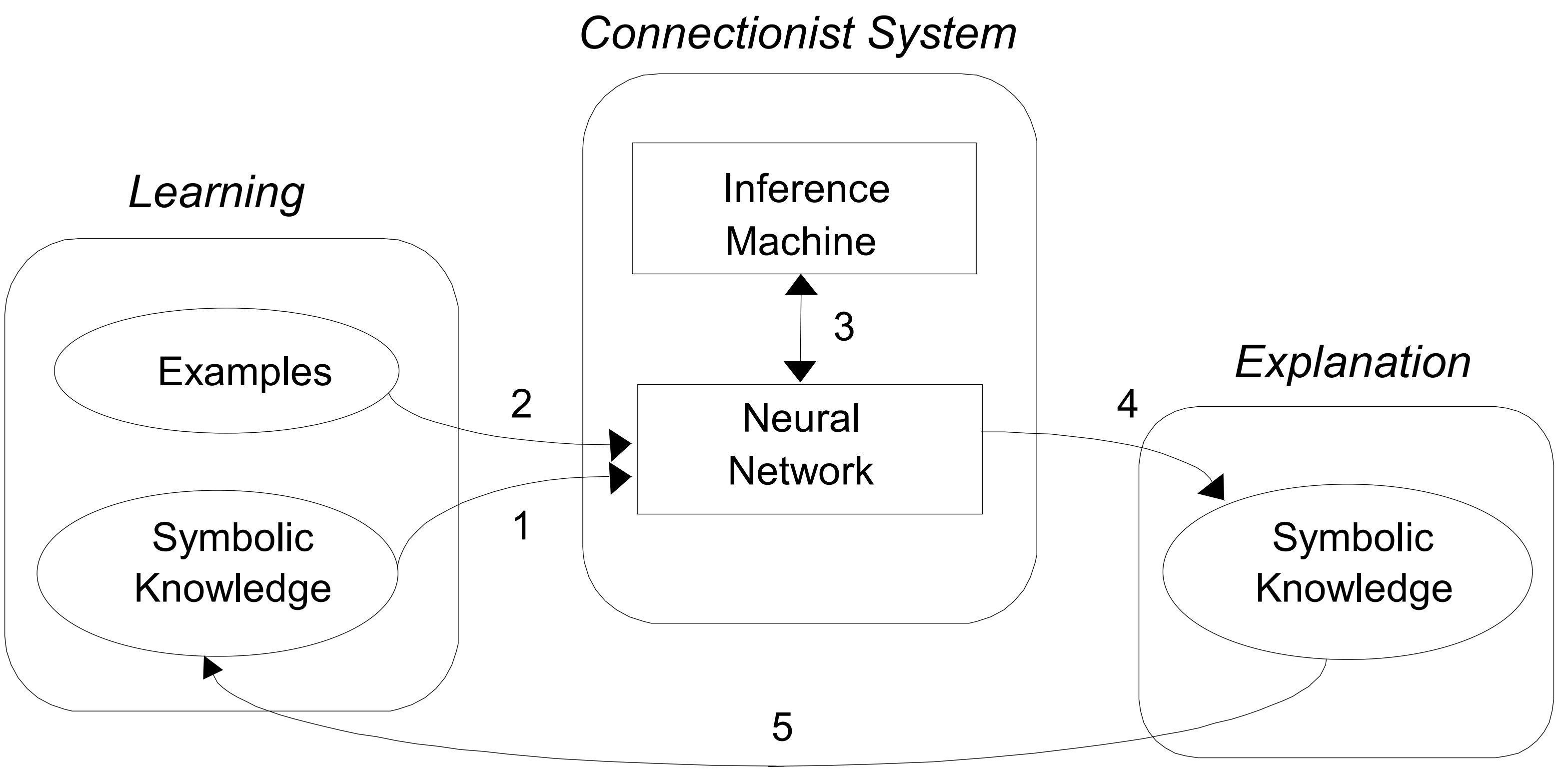

Connectionist Inductive Logic Programming (CILP) System

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 : $B \leftarrow$

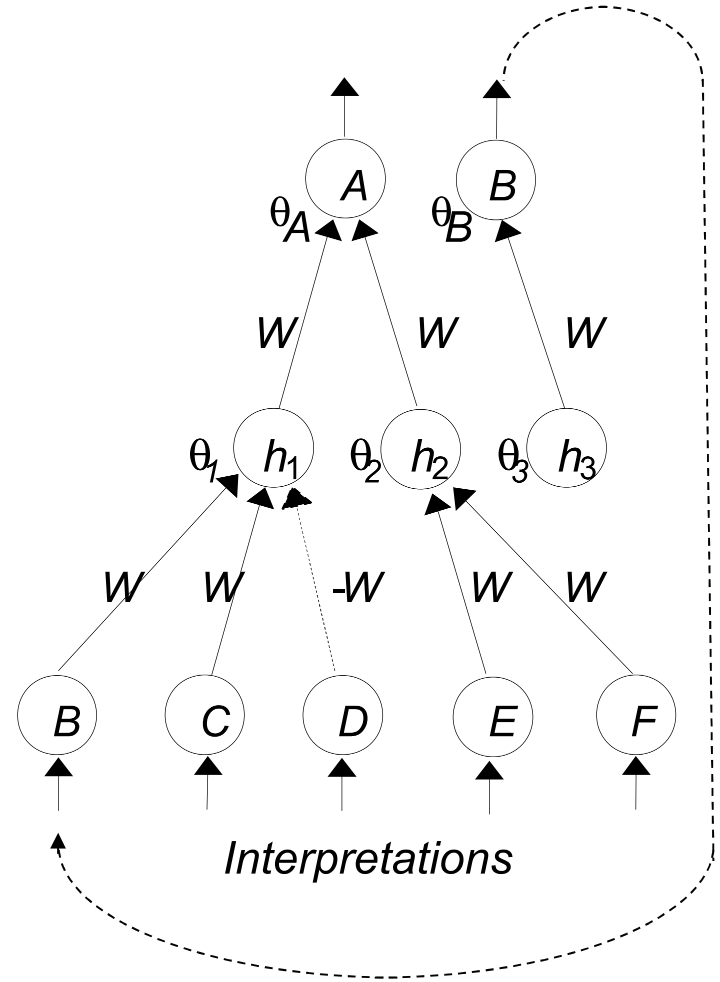

THEOREM: For any logic program P there exists a neural network N such that N computes P

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 (problem)

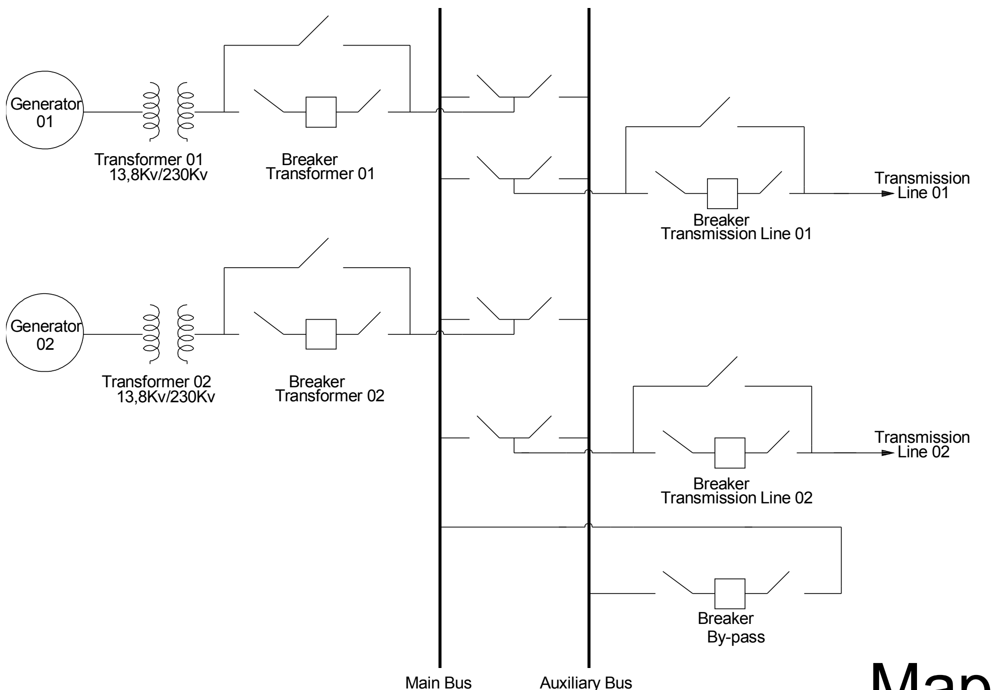

Mapping 23 alarms to 32 faults

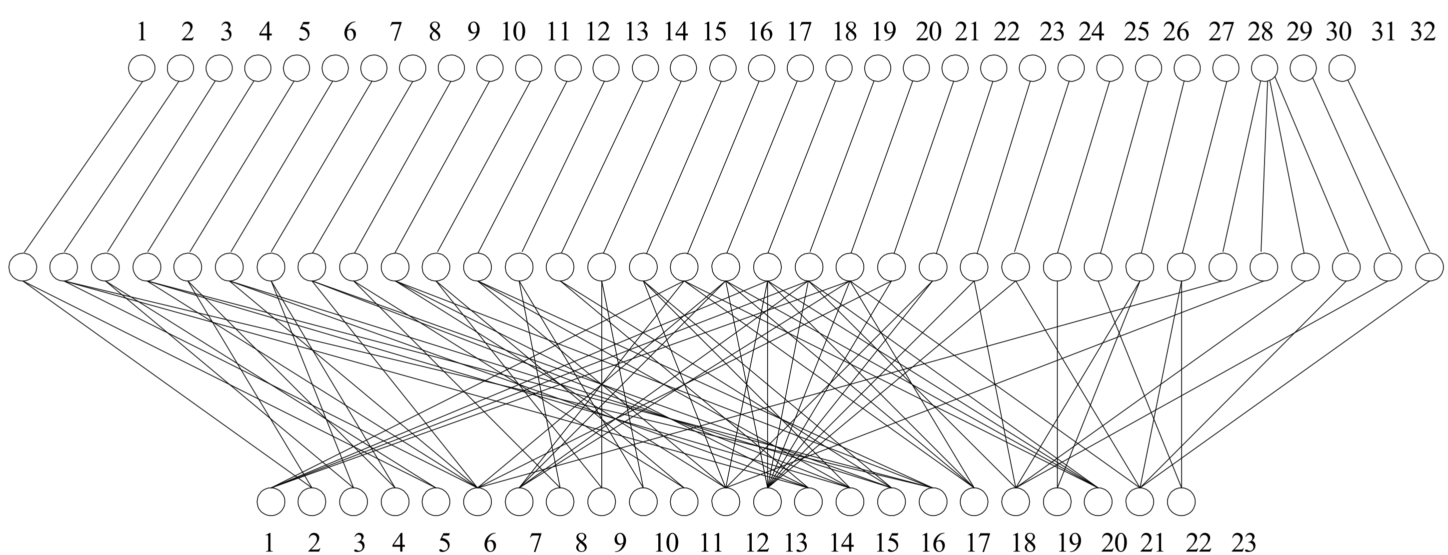

Power Plant Fault Diagnosis

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.

Power Plant Fault Diagnosis (results)

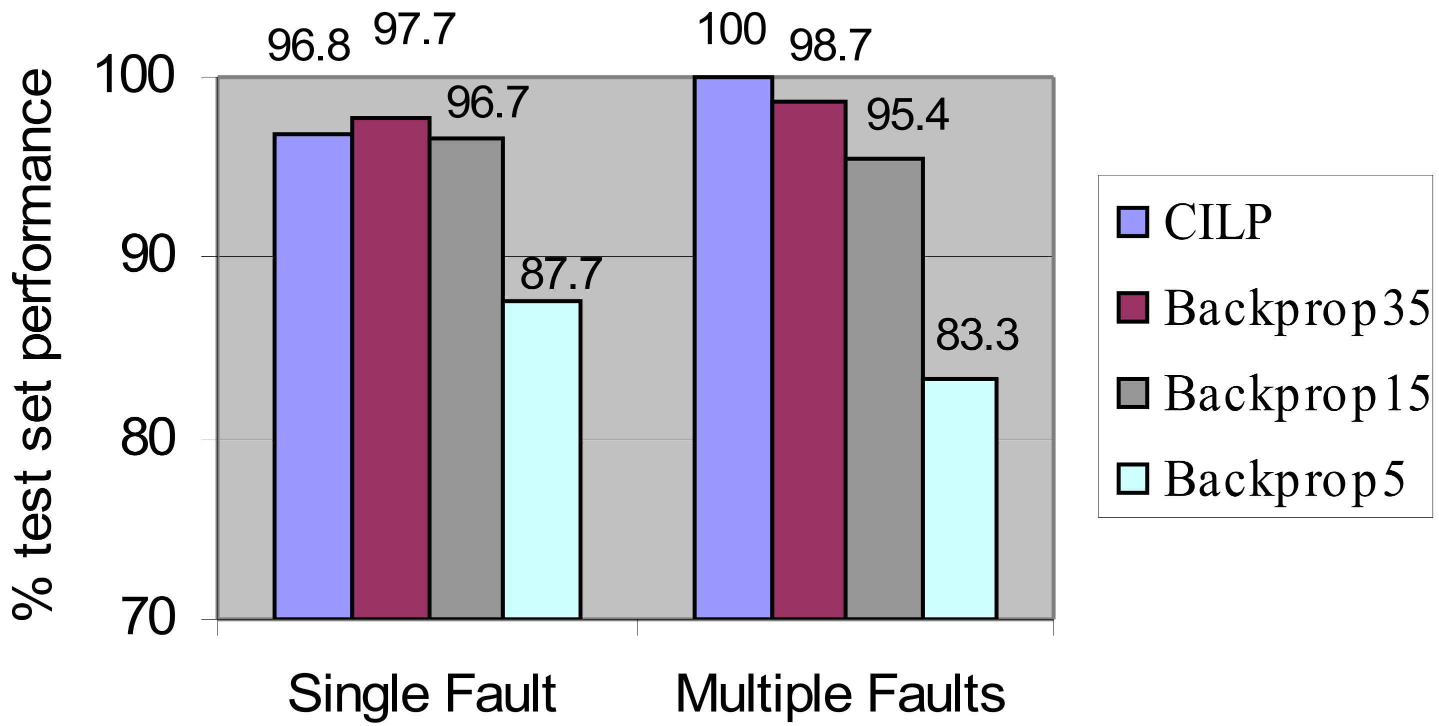

Also, as expected, CILP networks learn faster

Promoter Recognition (problem)

Promoter = small DNA sequence at beginning of genes

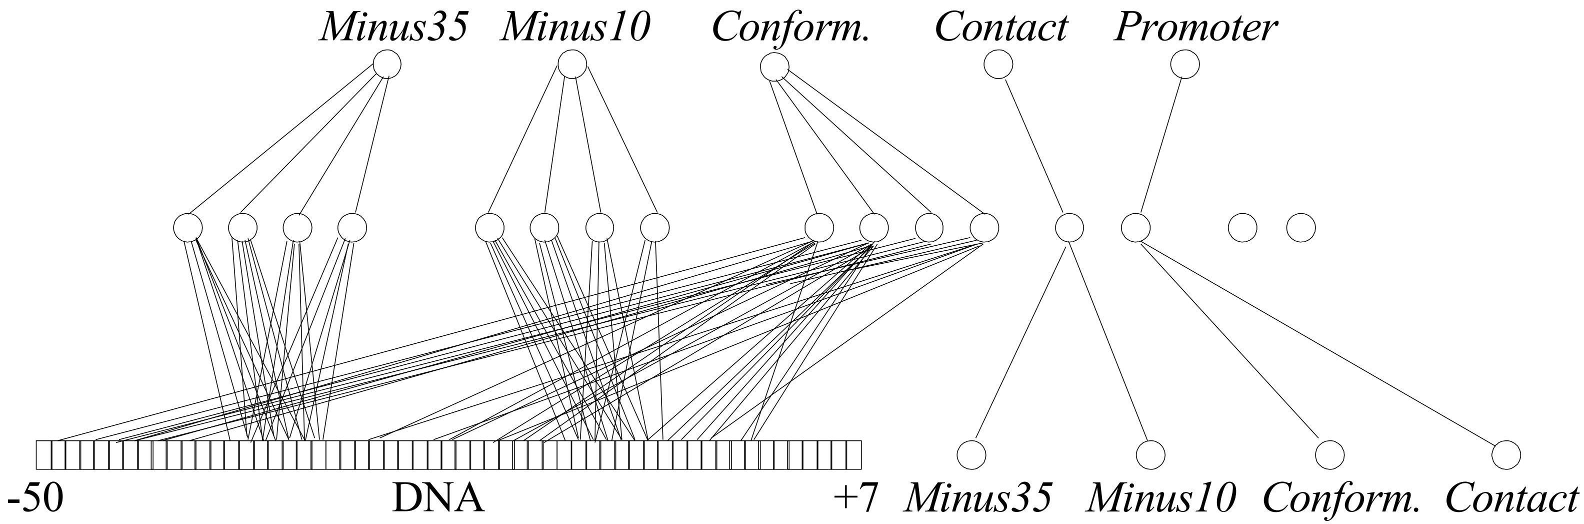

Promoter Recognition (results)

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

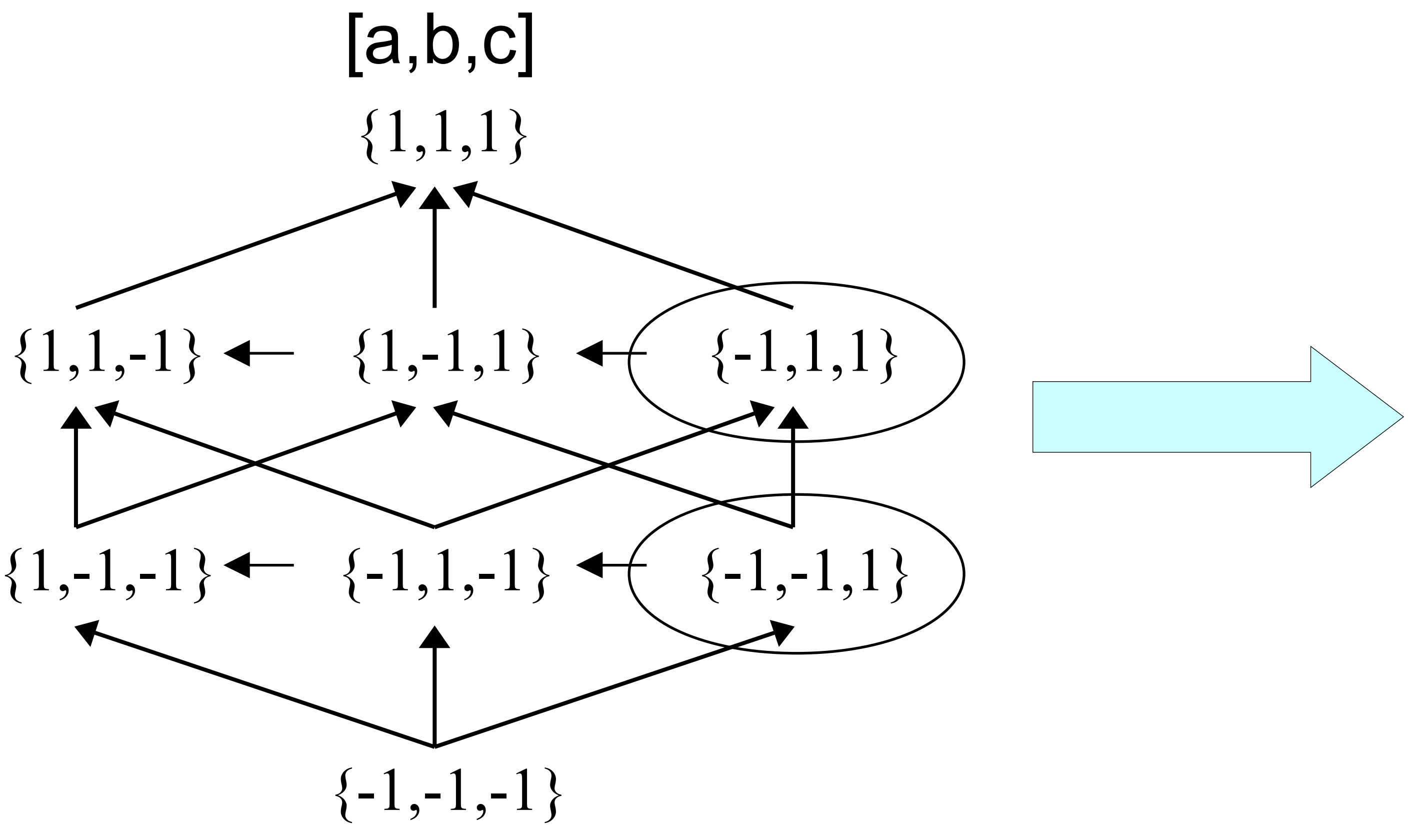

$2(a, b, c) \rightarrow h_1$

$b, c \rightarrow h_1$
$a, c \rightarrow h_1$
$a, b \rightarrow h_1$

$a \rightarrow h_0$
$b \rightarrow h_0$
$c \rightarrow h_0$

$1\ (a, b, c) \rightarrow h_0$

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 (deep networks)

- 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), training and assessment in simulators (driving test), visual intelligence (video classification), general game playing, web ontology learning.

Connectionist Modal Logic (CML)

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

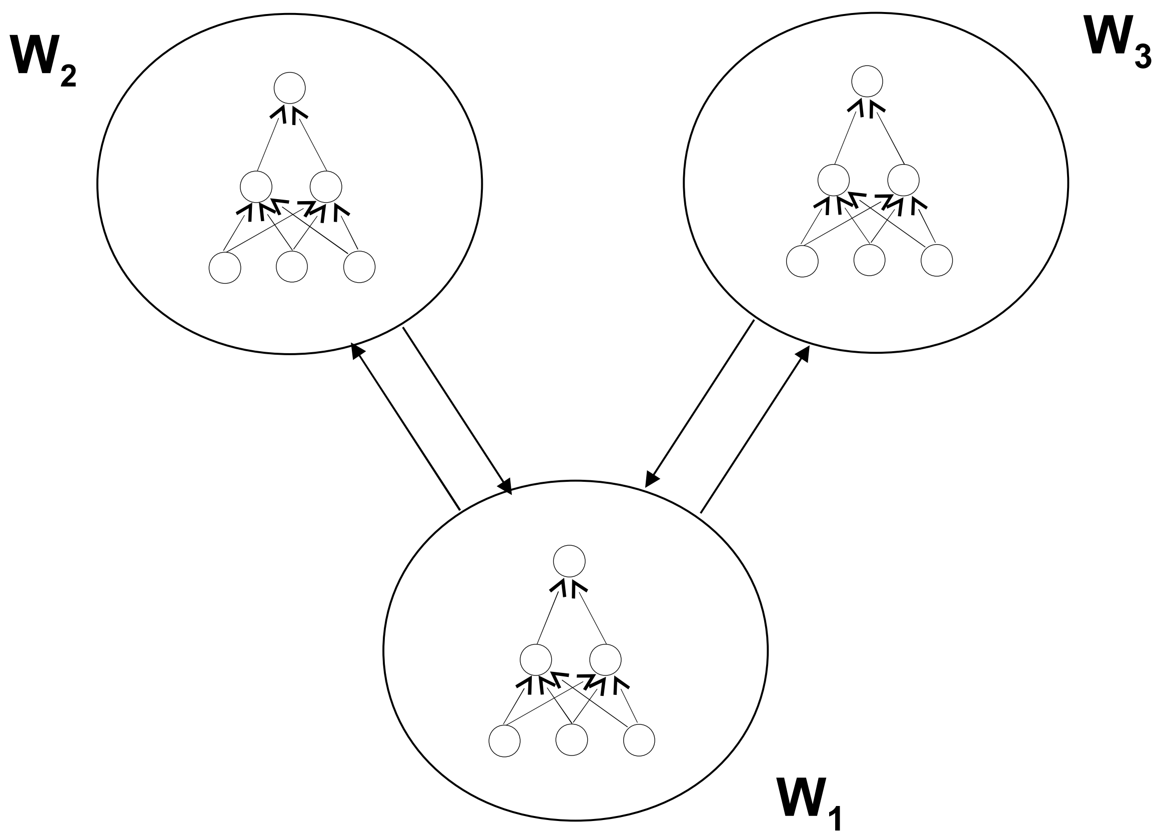

Semantics of *necessity* and *possibility*

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*

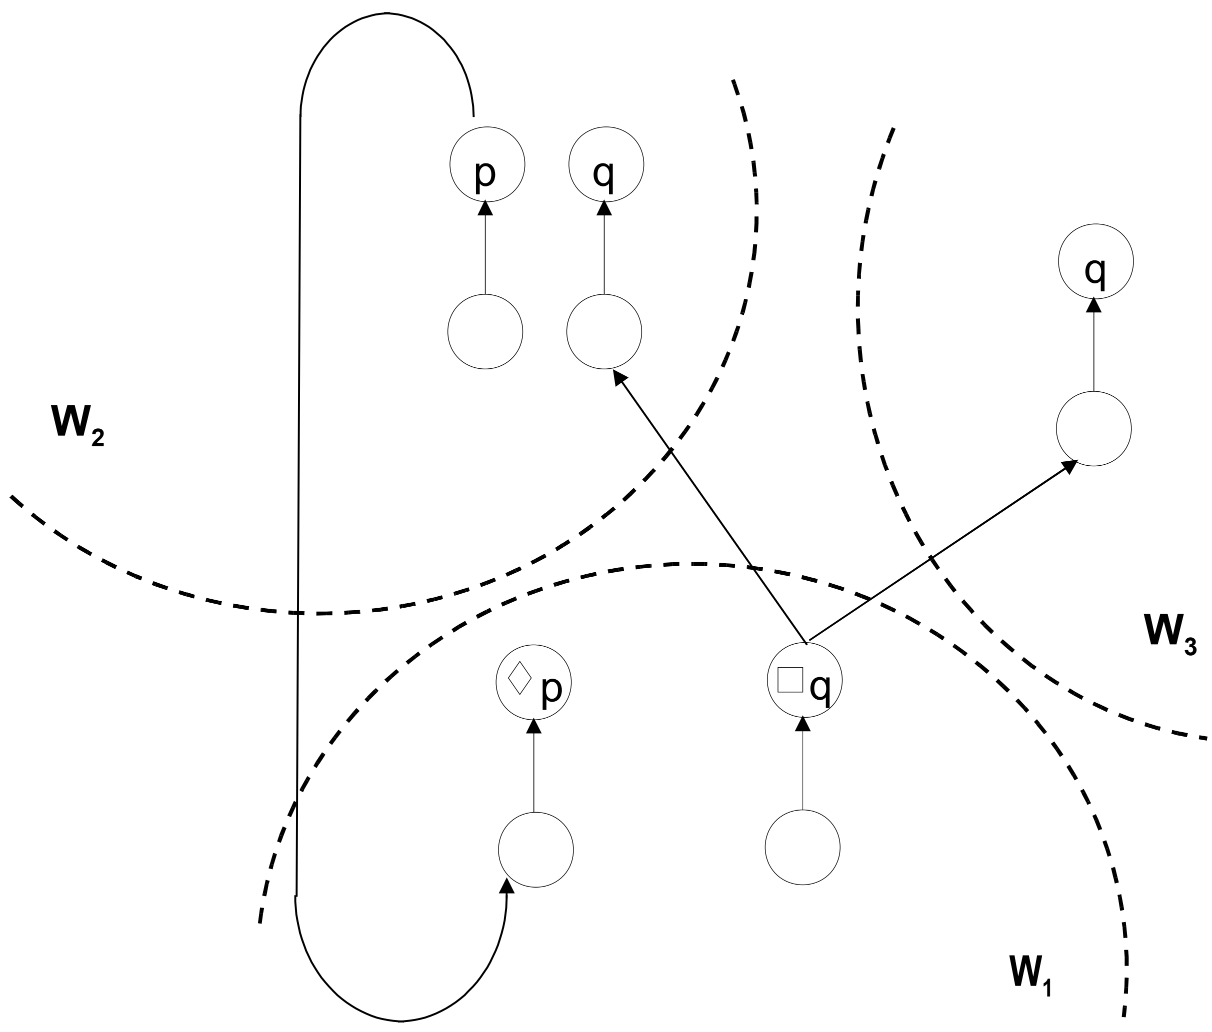

CML Translation Algorithm

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)

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

Connectionist Temporal Reasoning

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

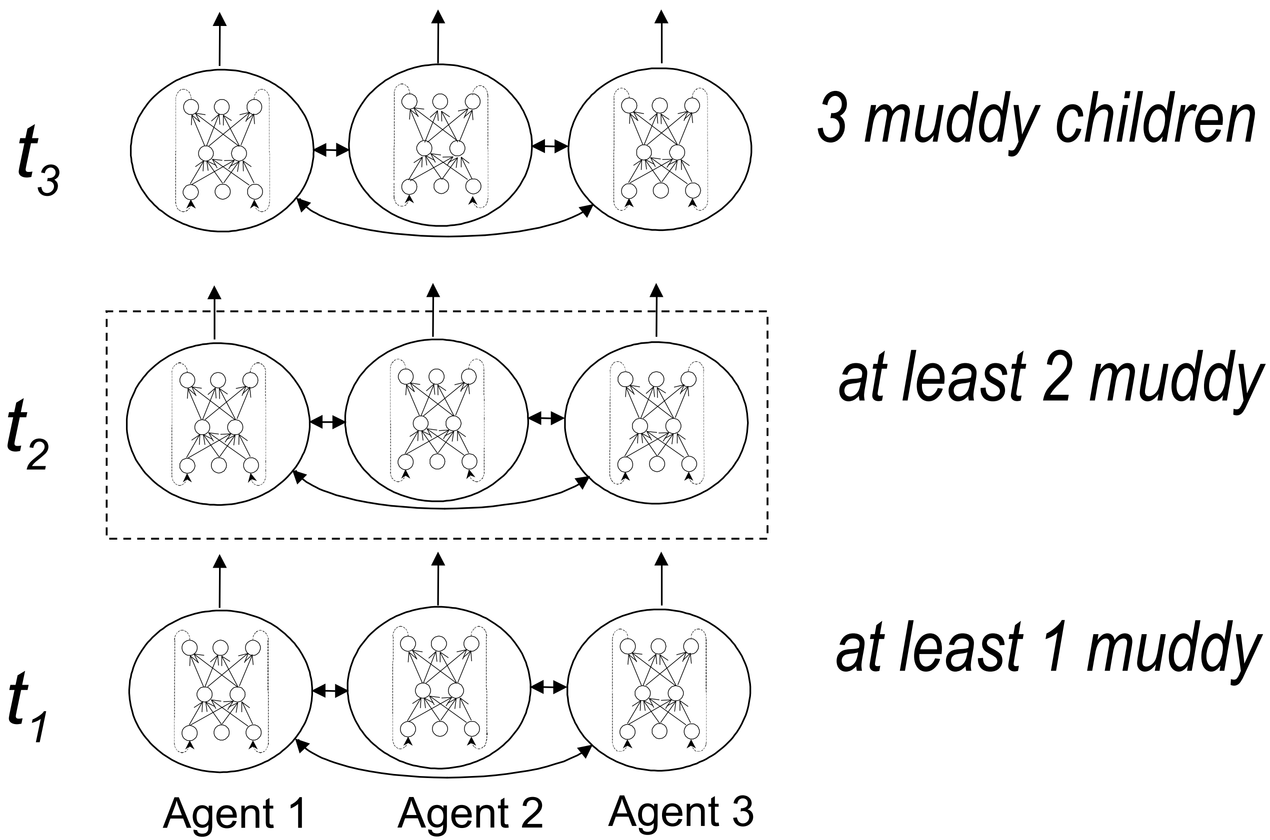

THEOREM: For any temporal program *P* there exists an ensemble of networks *N* such that *N* computes *P*.

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.

Combining (Fibring) Networks

A neuron that is a network! Modulation or recursion?

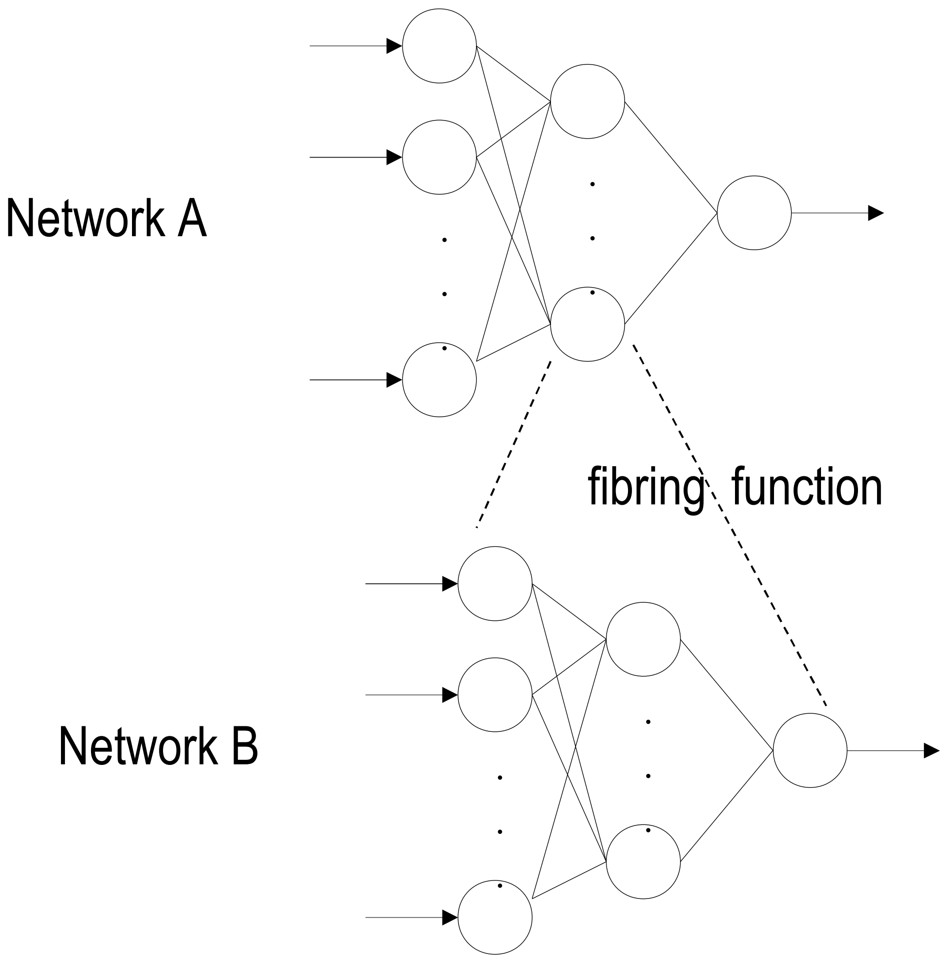

Expressiveness to represent first-order logic
Loosely-coupled integration: e.g. Network A and Legacy System B

Fibring Expressiveness

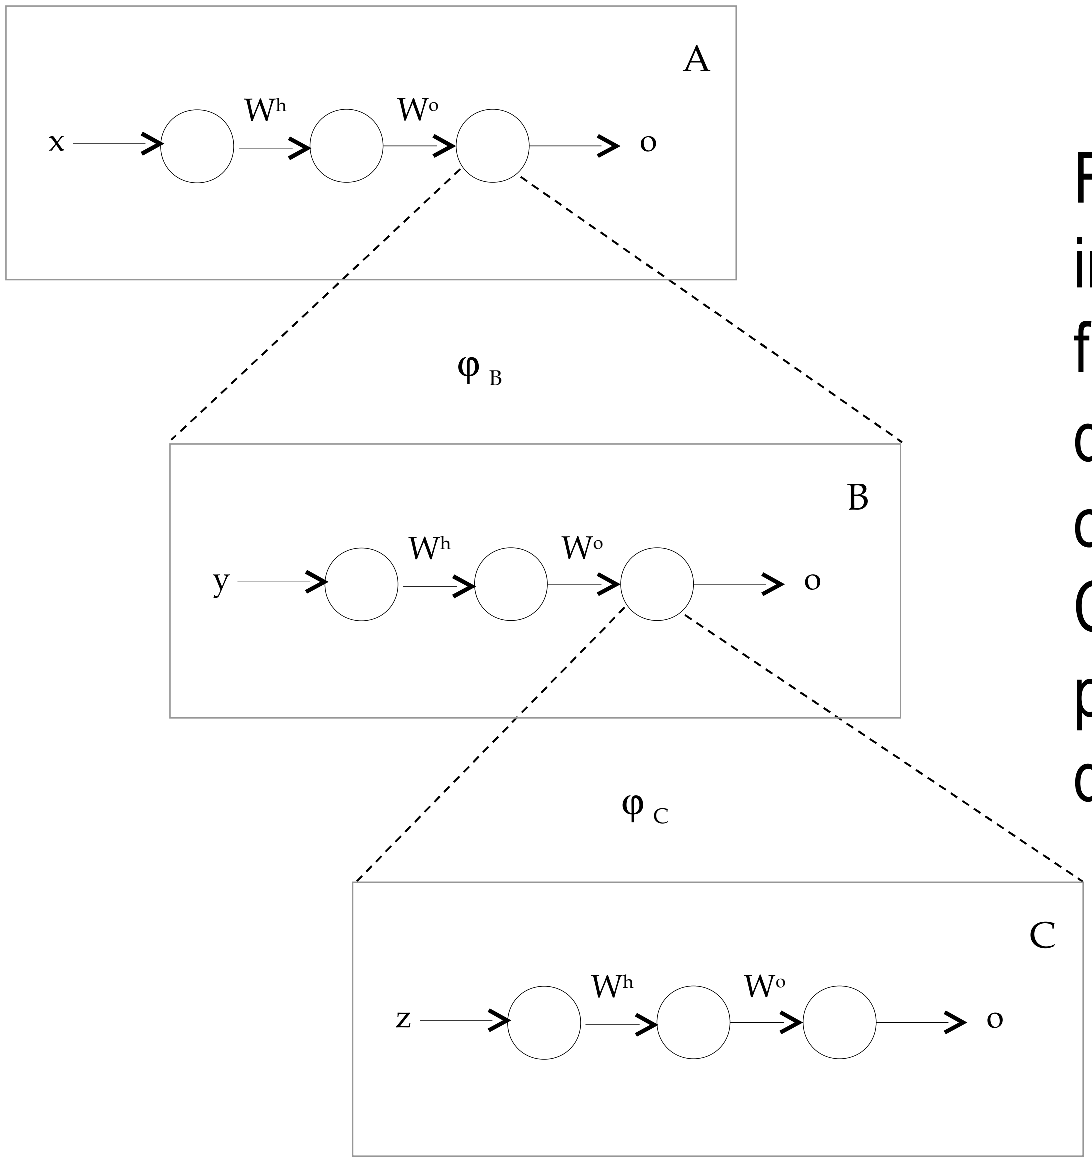

Fibred networks approximate any polynomial function in unbounded domains, e.g. $f(x)=x^2$, as opposed to each of A, B, C which are universal approximators in compact domains only.

Examples of Fibring: Relational Learning

Grandparent(X,Z) IF parent(X,Y) AND parent(Y,Z)

Inputs presented to P and Q at the same time trigger the learning process in the meta-level

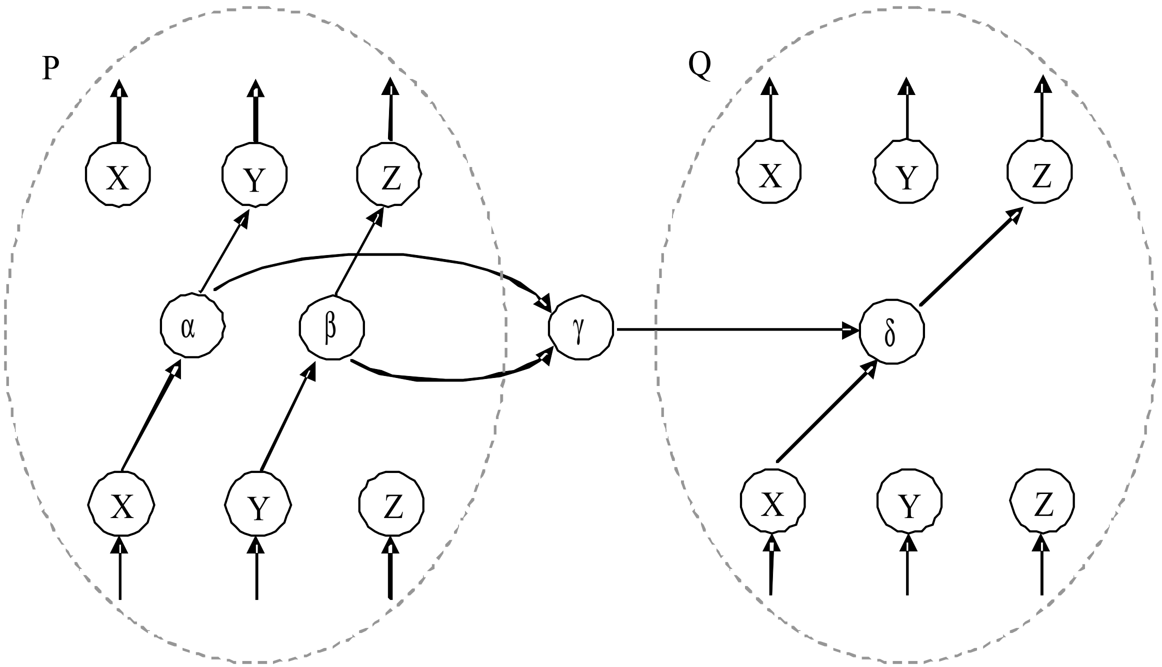

Experiments on the east-west trains dataset show an improvement from 62% (flat, propositional network) to 80% (metalevel network) on test set performance (leaving one out cross-validation)

Examples of Fibring: Metaphors

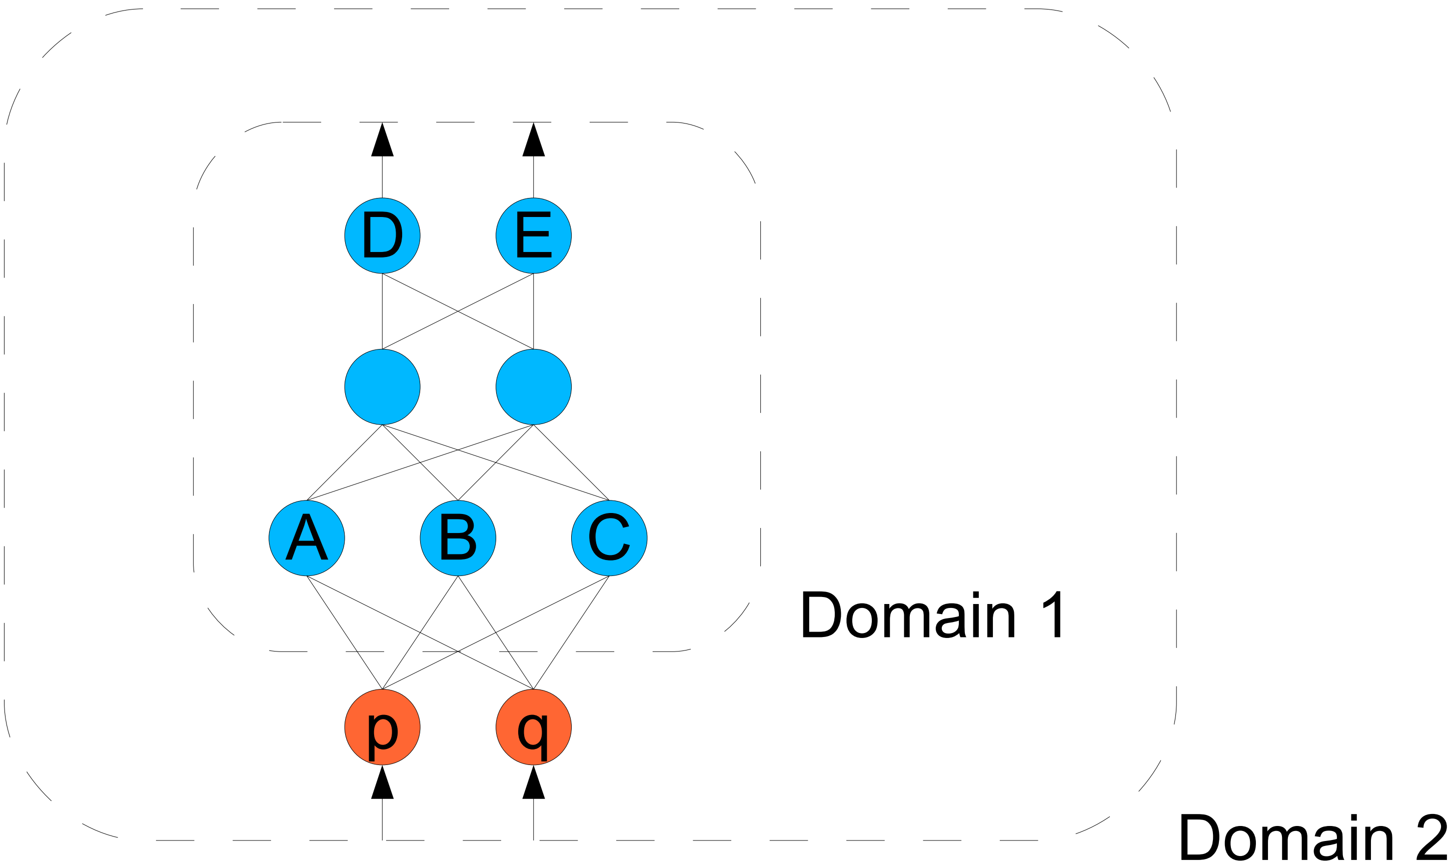

Combining a CILP network and a restricted Boltzmann machine

Cognitive Model: Fibred Network Ensembles

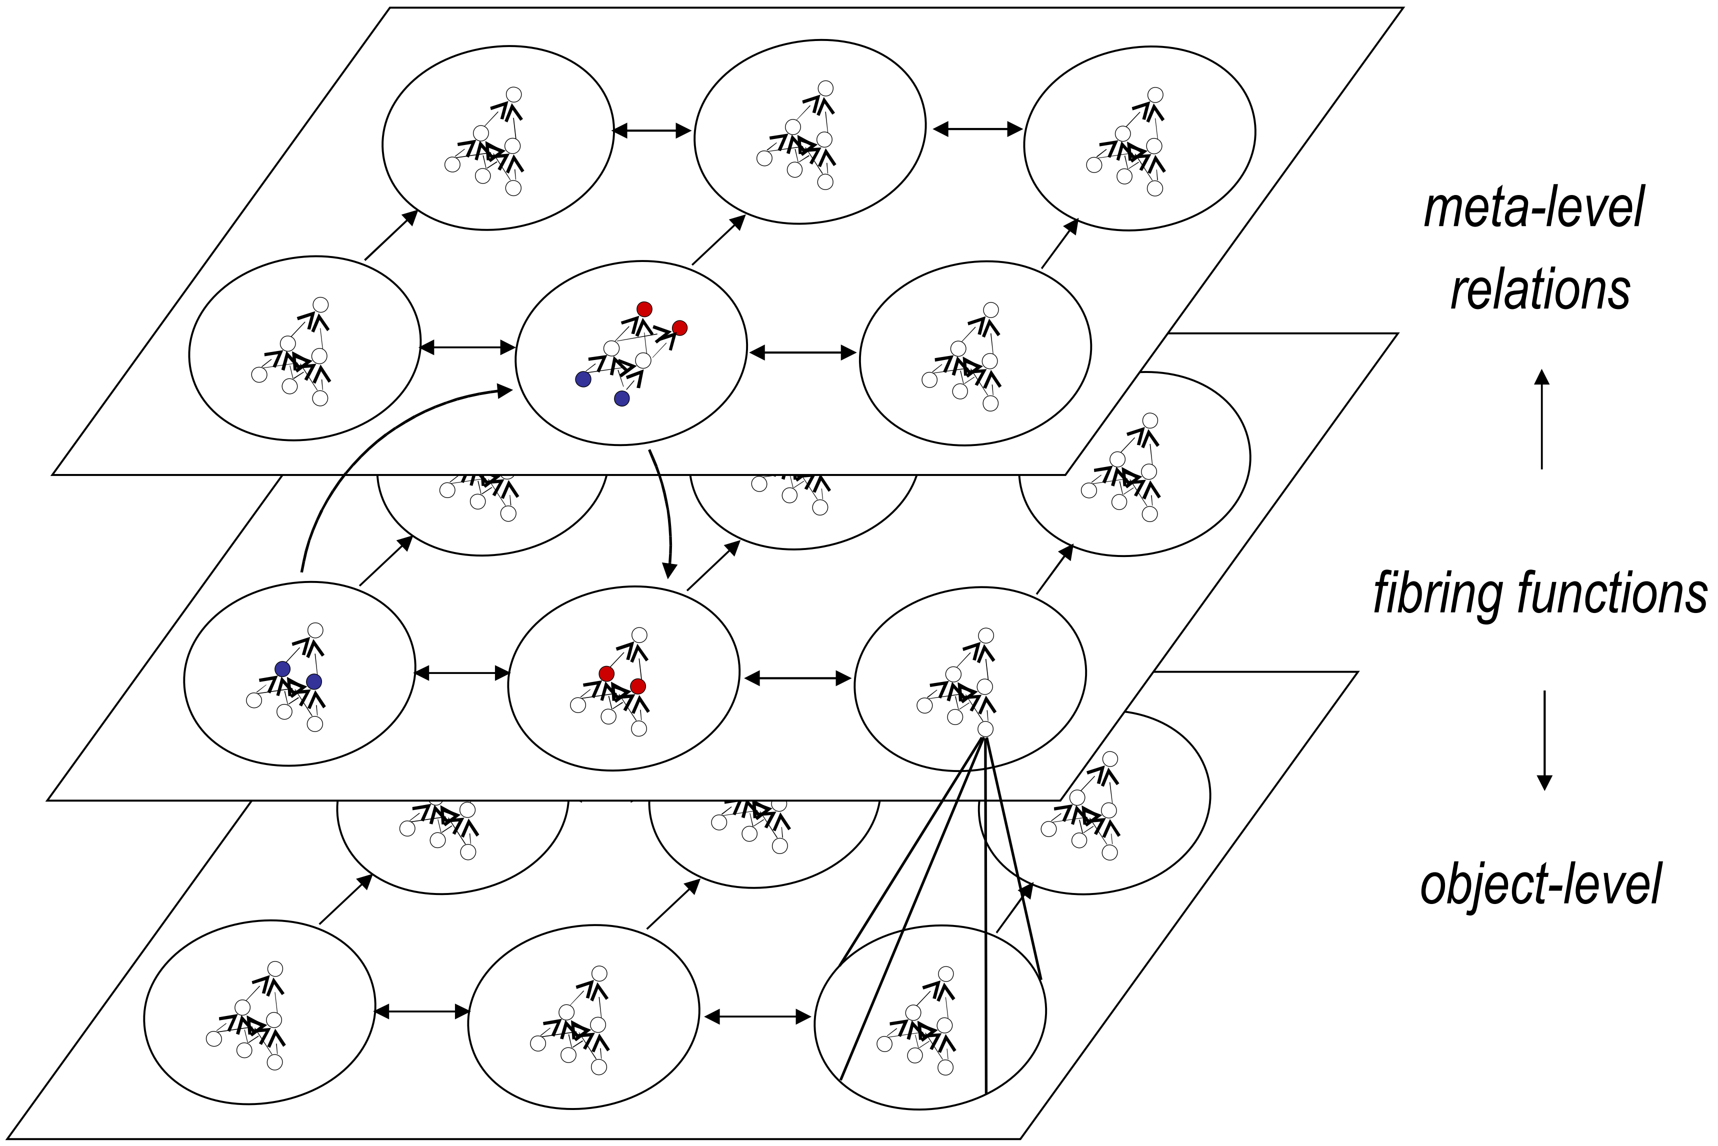

Publications on Fibring and Cognitive Model

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.

Recent 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

Recent Applications (cont.)

Learning Normative Rules of the RoboCup Competition

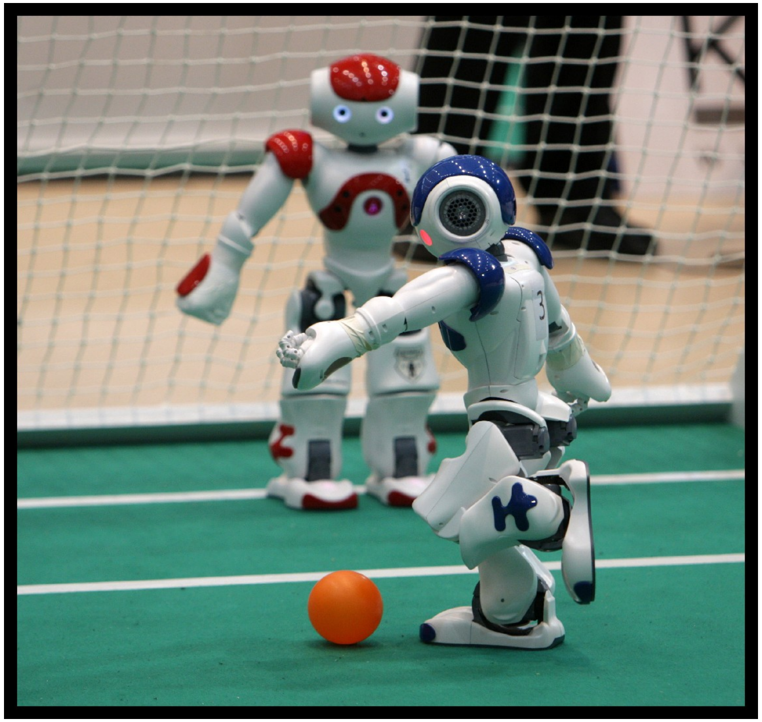

G. Boella, S. Colombo Tosatto, A. S. d'Avila Garcez, V. Genovese, D. Ienco and L. van der Torre. A Neural-Symbolic System for Normative Agents. AAMAS'11, May 2011.

Recent Applications (cont.)

Software Model Verification and Adaptation Framework

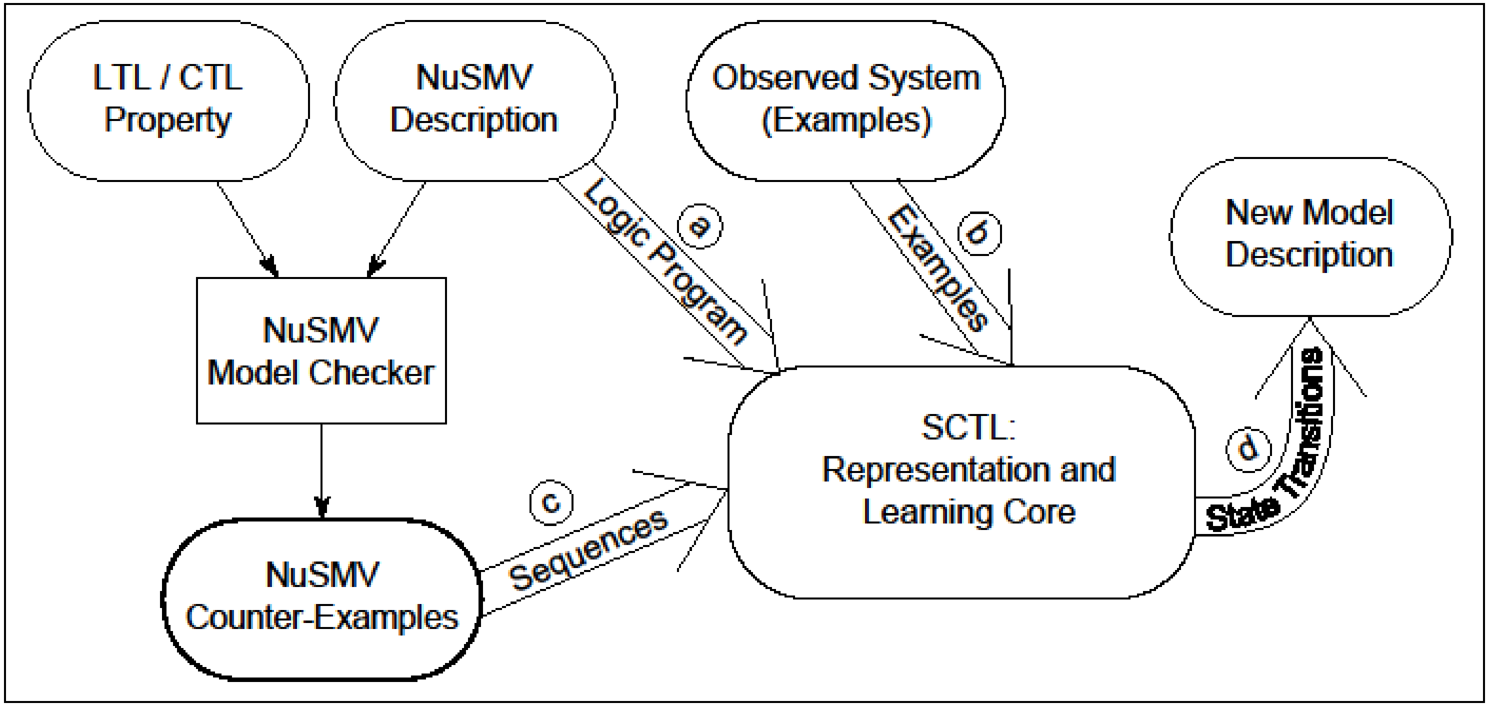

Borges, Garcez, Lamb, Nuseibeh. Learning to Adapt Requirements Specifications of Evolving Systems. ICSE (NIER Track), May 2011.

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.

Current/Future Work

- Theory: how brains make mental models / abduction / argumentation
- Practice: systems and applications (training in simulators, visual intelligence)
- First Order Logic Learning: encoding vs. propositionalisation / binding problem
- Neural Networks for Normative Systems (Robocup)
- General Game Playing (Stochastic CILP)
- Adding/Extracting domain knowledge from deep belief networks