Nonclassical Reasoning in Neural Networks

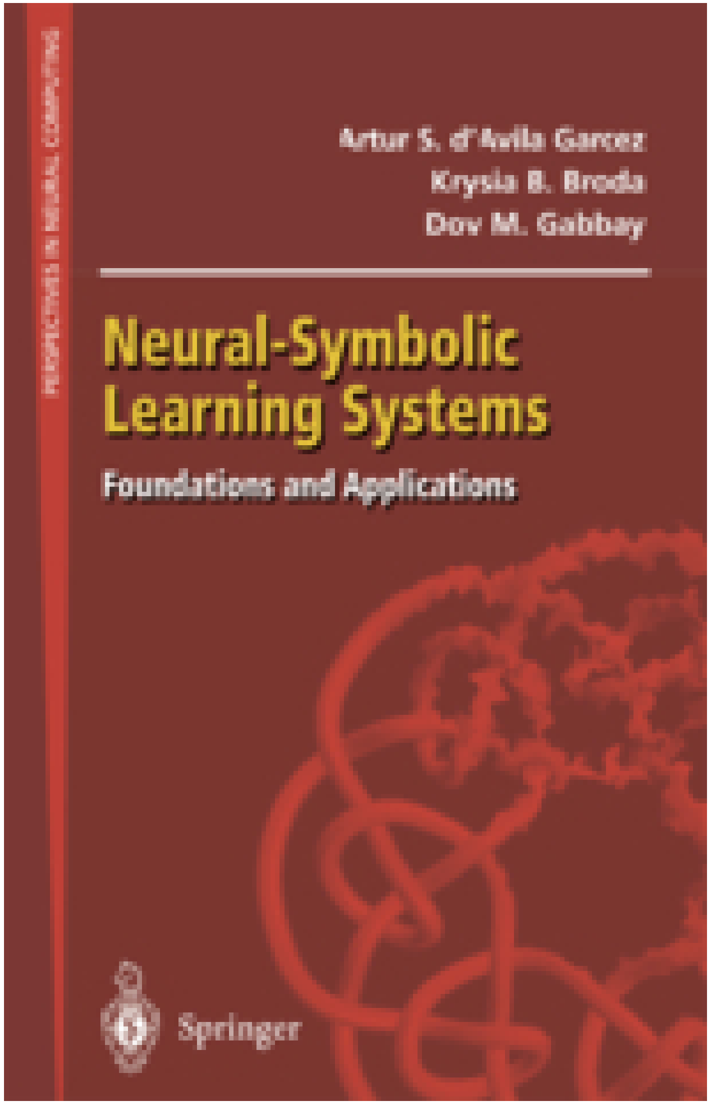

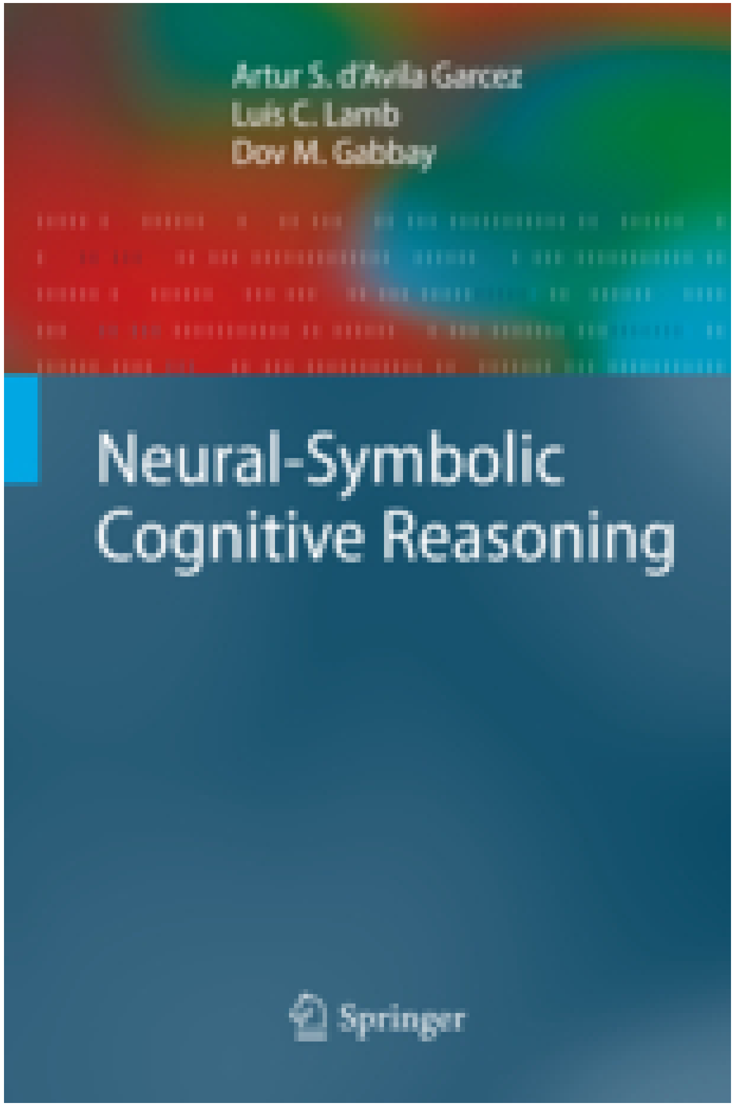

Connectionist Modal Logic

- A single CILP network can represent extended logic programs + priorities
- How about network ensembles?

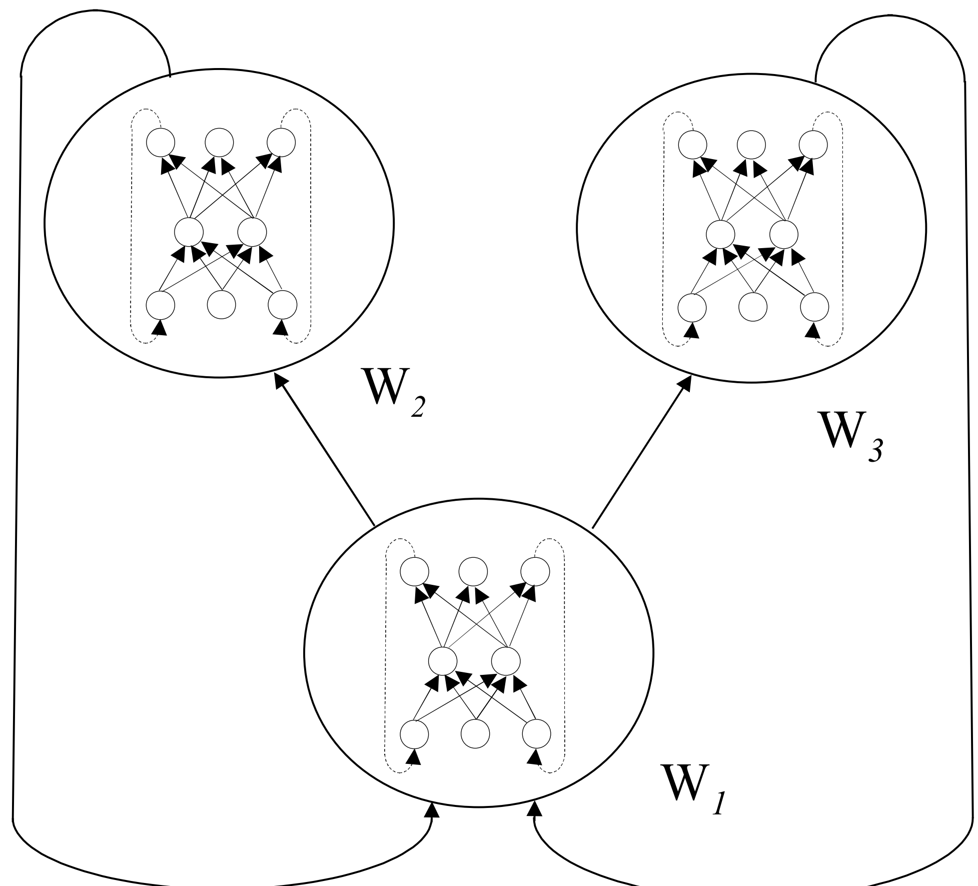

Representing Modalities and ◊

- Ensembles allow us to label concepts in possible worlds, having multiple copies, e.g. p below:

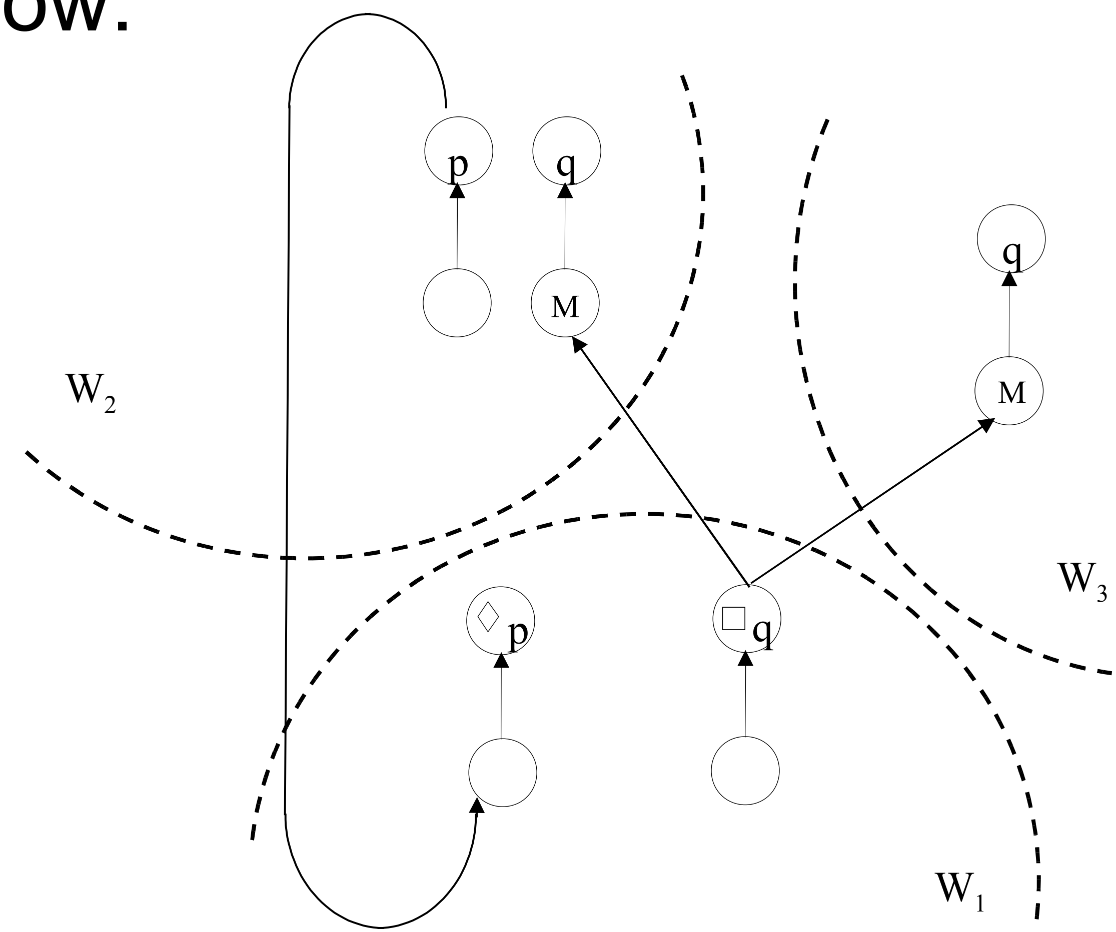

Semantics of   and ◊

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

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

Connectionist Modal Logic

$W_1 : r \rightarrow \quad q$ $\qquad$ $R(W_1, W_2)$

$W_1 : \Diamond s \rightarrow r$ $\qquad$ $R(W_1, W_3)$

$W_2 : s$

$W_3 : q \rightarrow \Diamond p$

Definition: An *extended modal program* is a finite set of clauses C of the form $W_i : ML_1,\ldots,ML_n \rightarrow MA$ and a finite set of relations $R(W_a,W_b)$ between worlds W_a and W_b in C.

$M = \{ \quad , \Diamond\}$, L_i is a literal, A is an atom.

Connectionist Modal Logic (cont.)

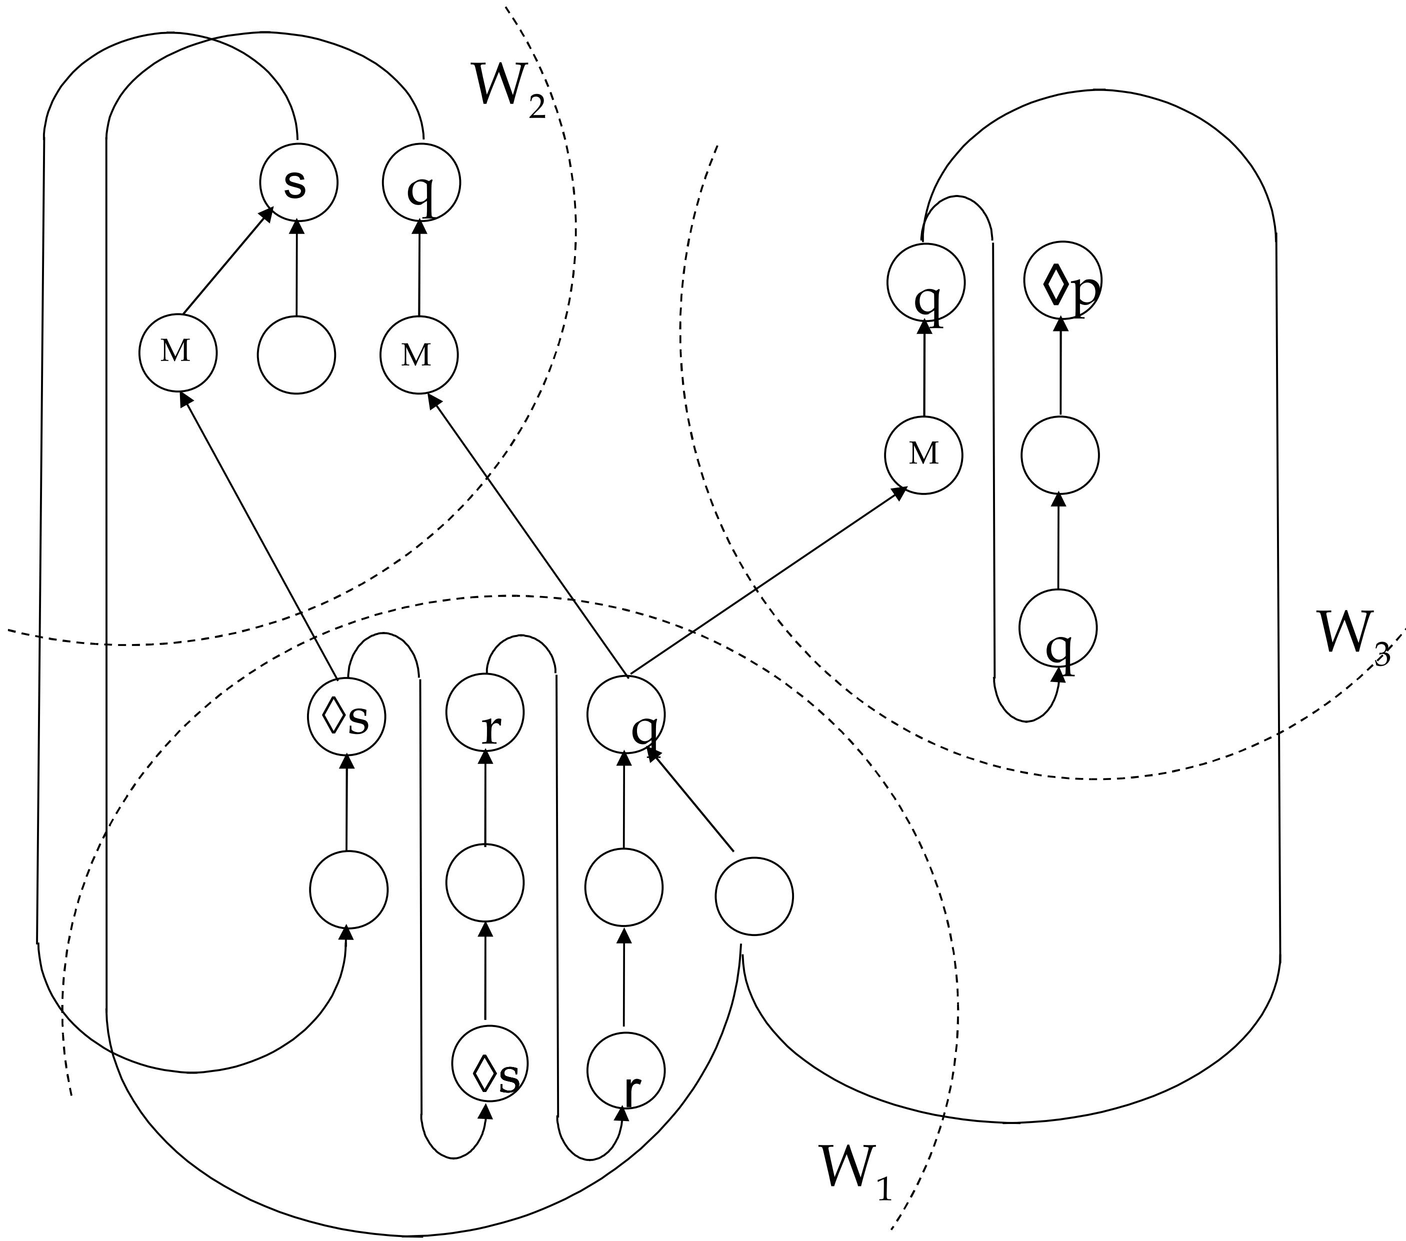

The Modalities Algorithm

Translates any Extended Modal Program into an ensemble of C-ILP networks

Theorem: For any extended modal program P there exists an ensemble of single hidden layer neural networks N such that N computes the modal fixed-point operator MT_P of P.

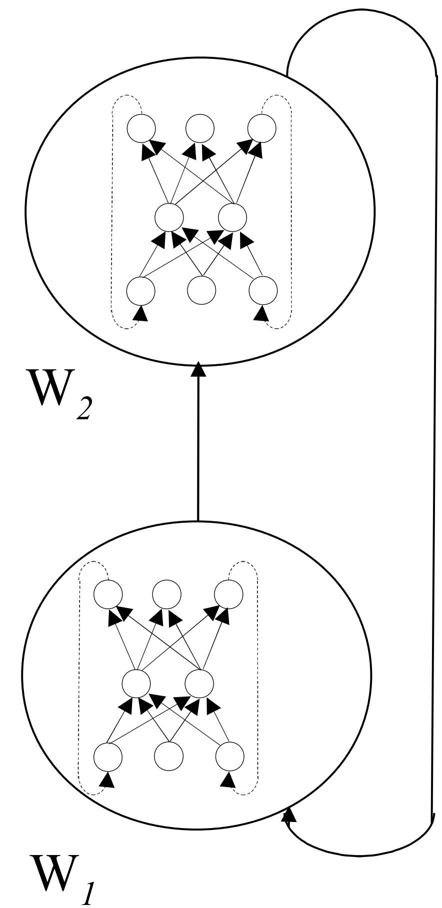

Learning with Modal BK

Each CILP network can be trained using standard Backpropagation.

One can perform learning in each possible world, but having modalities in the BK.

We have applied the Modal CILP System to the Muddy Children Puzzle.

Muddy Children Puzzle

A group of children is playing in a garden. A number of them has mud on their faces.

Each child can see if the other children are muddy, but not if she herself is muddy.

A caretaker asks: does any of you know if you have mud on your face?

$\mathbf{K}_1\, q_1 \wedge \mathbf{K}_1\, \neg p_2 \wedge \mathbf{K}_1\, \neg p_3 \rightarrow \mathbf{K}_1\, p_1$ (**K** is analogous to)

p_i = child i is muddy

q_j = at least j children are muddy

Muddy Children: Common Knowledge

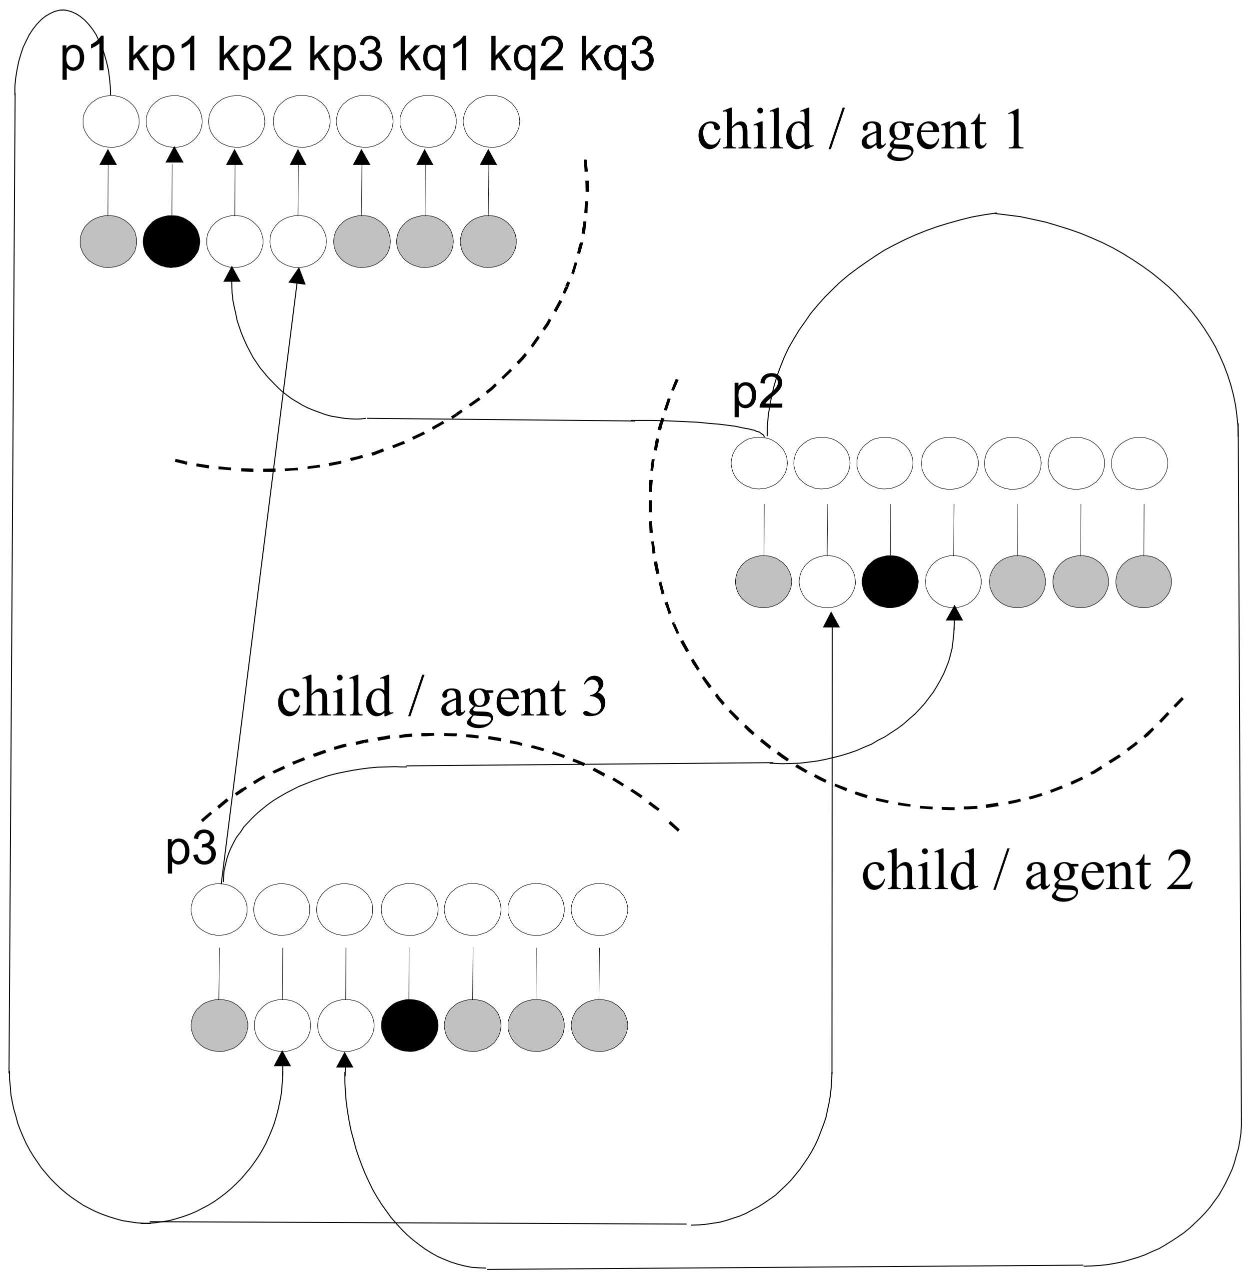

Muddy Children: Local Knowledge

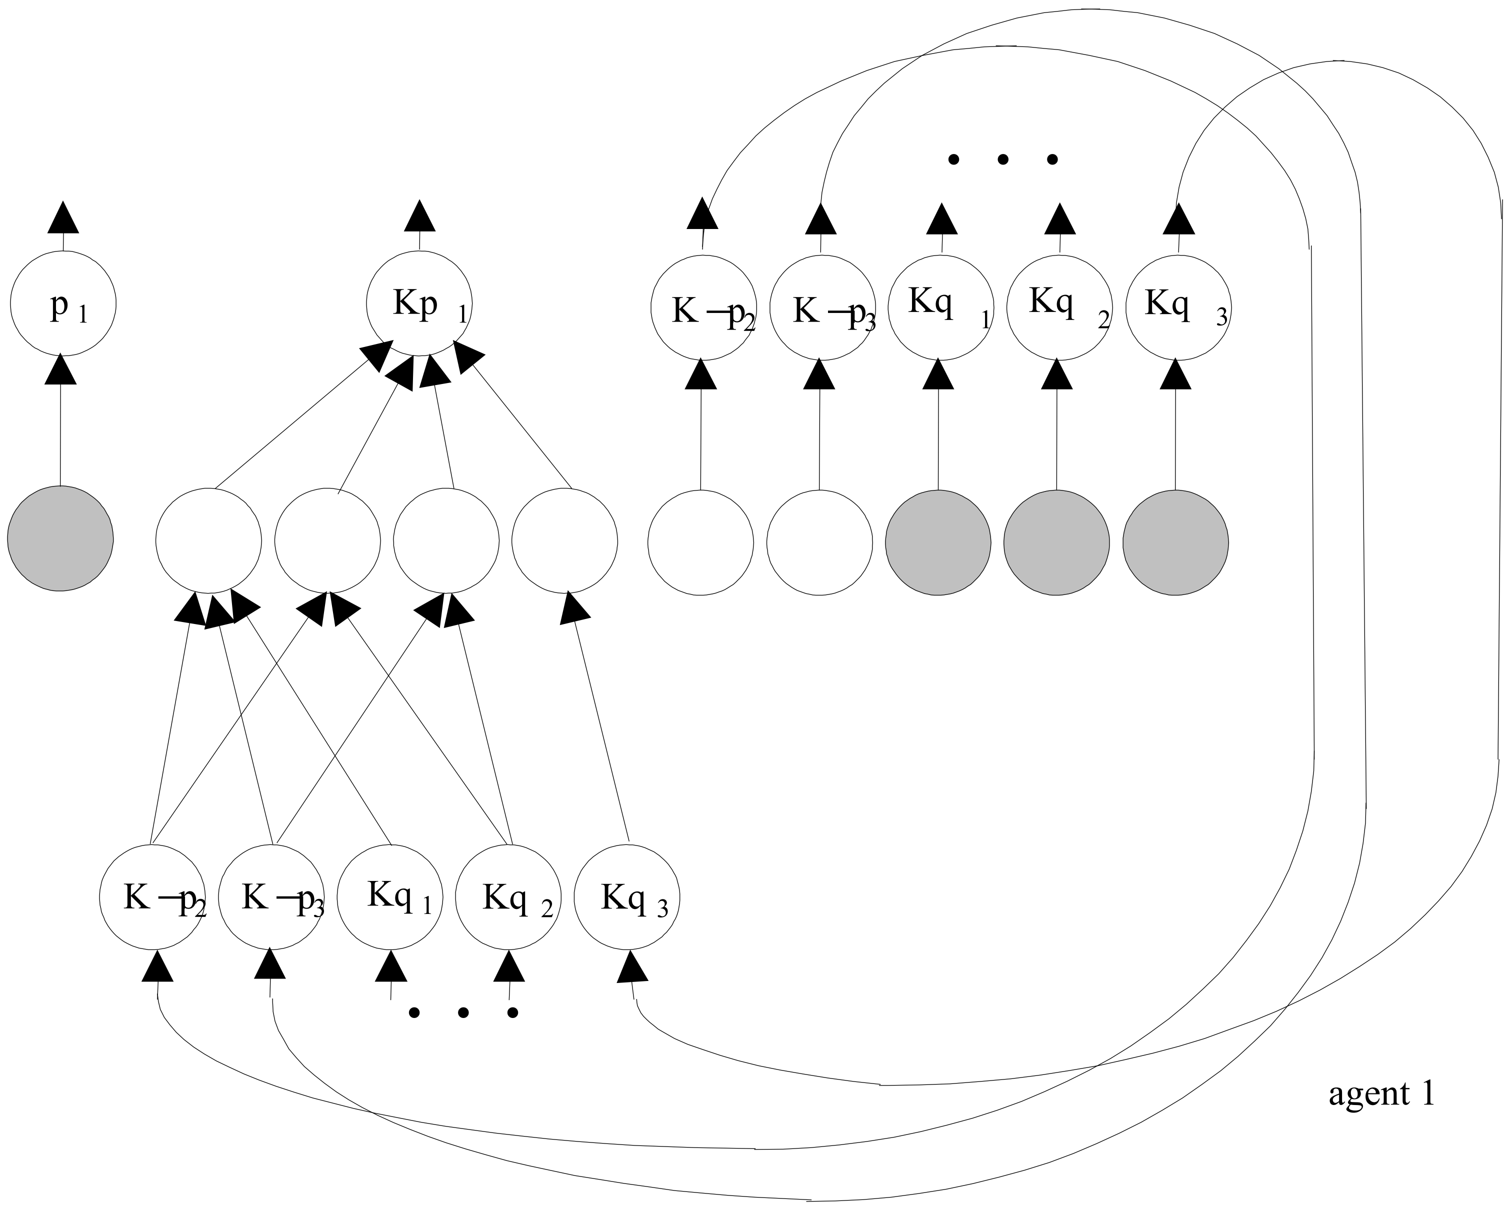

Experiments

Cross-Validation on 32 training examples.

BK with a single rule: $\mathbf{K}_1 q_1 \wedge \mathbf{K}_1 \neg p_2 \wedge \mathbf{K}_1 \neg p_3 \rightarrow \mathbf{K}_1 p_1$

Same training parameters with and without BK

- CILP ensemble without modal BK: average test set accuracy 84.37%
- CILP ensemble with modal BK: average test set accuracy 93.75%

Muddy Children: Snapshots v Full Solution

Snapshot solution: represents the knowledge held by the agents at an arbitrary time t.

Full solution: explicitly represents the fact that when it is common knowledge that there are at least n-1 muddy children (at time t); after the announcement that nobody knows if they are muddy or not, then it becomes common knowledge that there are at least n muddy children (at time t+1).

Connectionist Temporal Logic

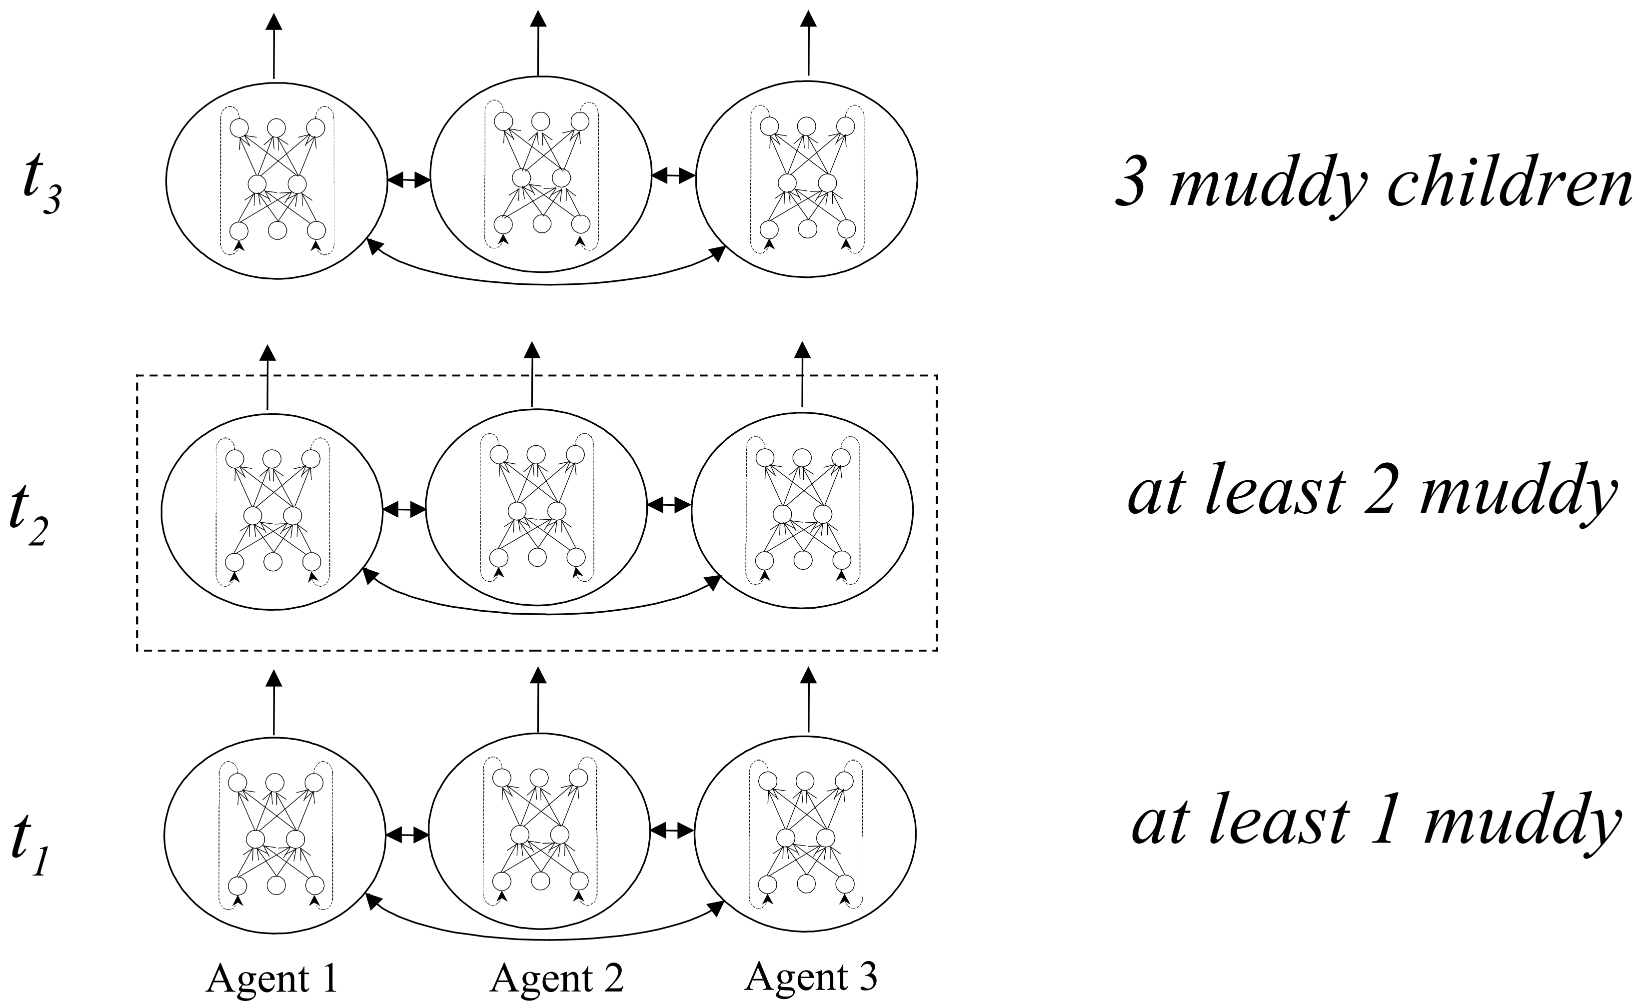

Knowledge Evolution in Time

If I know that tomorrow A will be true then tomorrow I should know that A is true

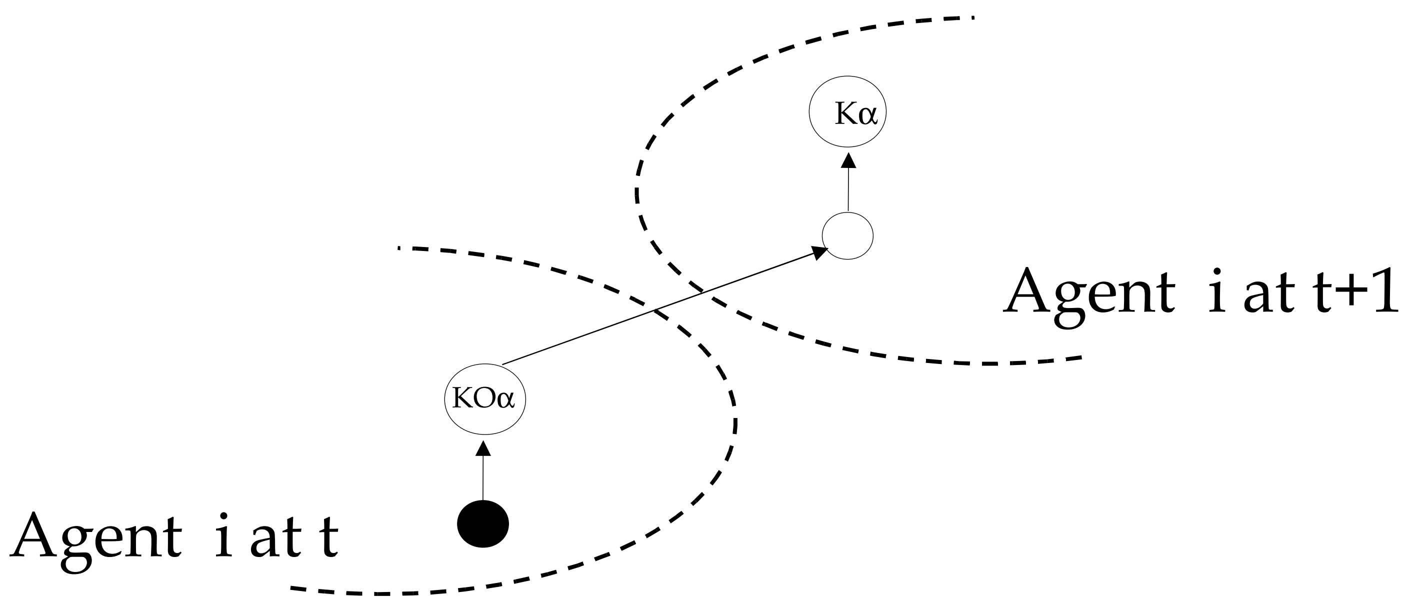

Muddy Children (full solution): Temporal Rules for Agent(child)1

t1: $\neg \mathbf{K}1\ p1 \wedge \neg \mathbf{K}2\ p2 \wedge \neg \mathbf{K}3\ p3 \rightarrow \mathrm{O}\ \mathbf{K}1\ q2$

t2: $\neg \mathbf{K}1\ p1 \wedge \neg \mathbf{K}2\ p2 \wedge \neg \mathbf{K}3\ p3 \rightarrow \mathrm{O}\ \mathbf{K}1\ q3$

Theorem: For each set of temporal rules P there exists an ensemble of single hidden layer neural networks N such that N computes the temporal fixed-point operator O*T*P of P.

Fibring Neural Networks

- A principled way of combining networks (or systems, in general)
- A principle way of integrating reasoning and learning in neural networks
- Reasoning in Network A causes learning in Network B

Network A

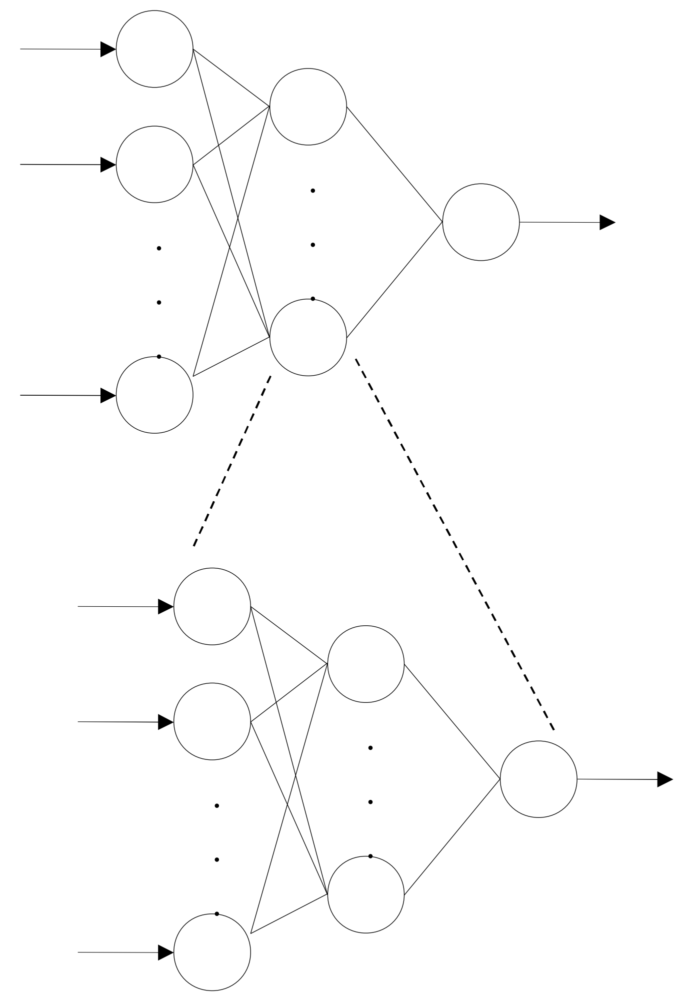

Network B

Fibring Function

- A function telling how the weights of network B will be affected by the values of network A
- Think of master/slave circuit: running A prompts learning in B

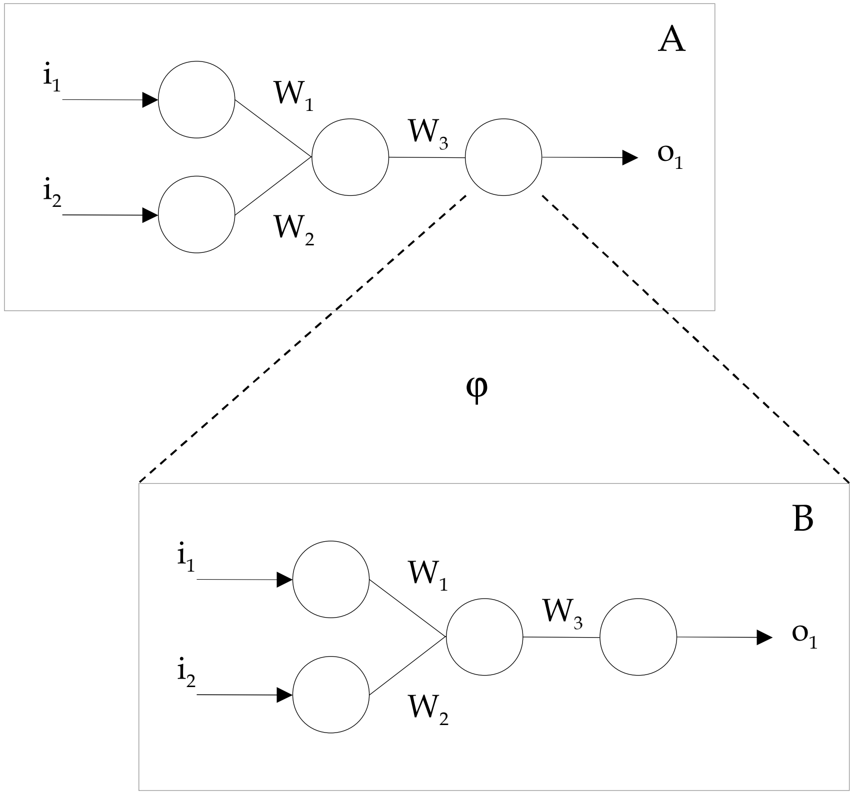

Fibring Expressiveness

Fibred Neural Networks can approximate any polynomial function (in an unbounded domain) to any desired degree of accuracy.

Standard Neural Networks can approximate functions in a compact domain only (to any desired degree of accuracy).

Fibring Expressiveness (cont.)

Any neuron can be fibred onto a network, and many levels of fibring are possible

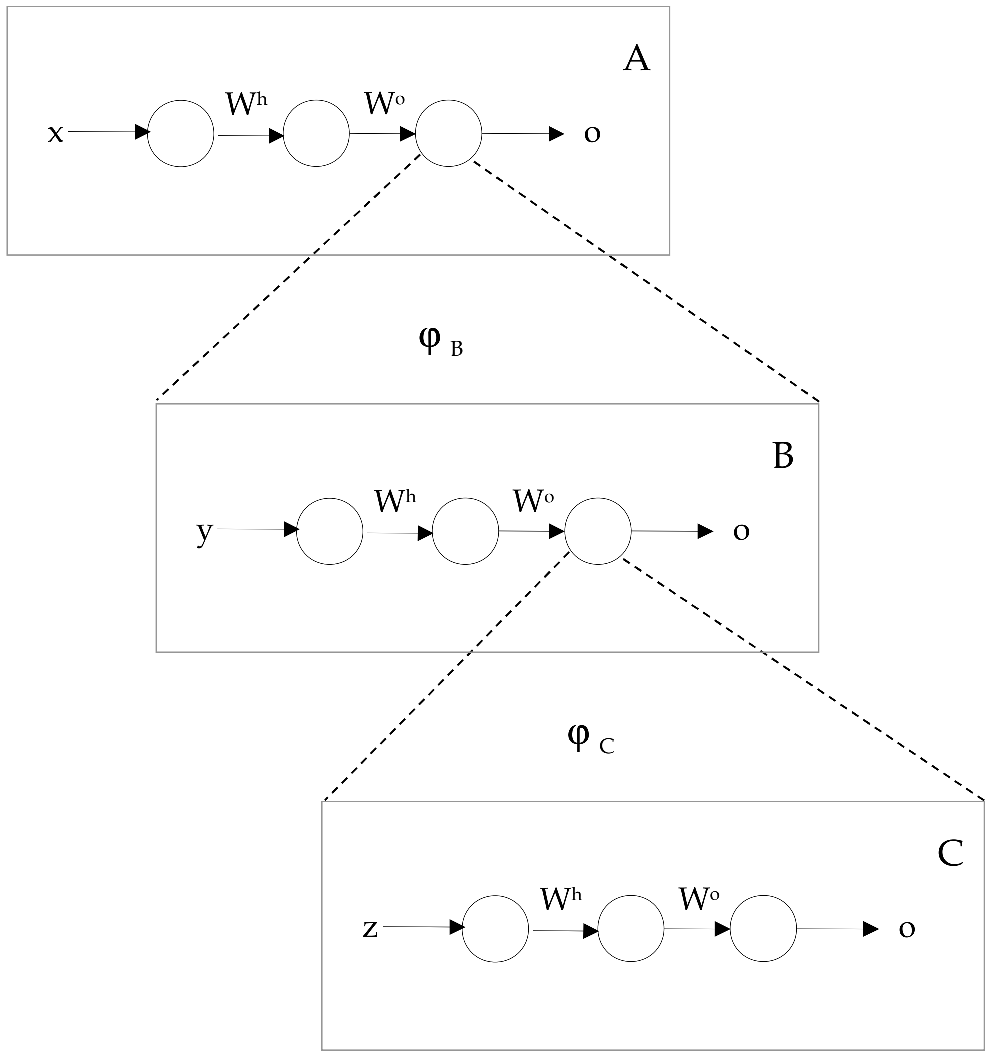

Conclusion

CILP network ensembles can represent a number of modalities in a natural way

This allows learning from examples with the use of standard backpropagation

And the use of more expressive languages, such as Modal and Temporal Logic, as BK

Also networks can be fibred so as to implement specializations and relational knowledge.

Cognitive Model: Fibred Network Ensembles

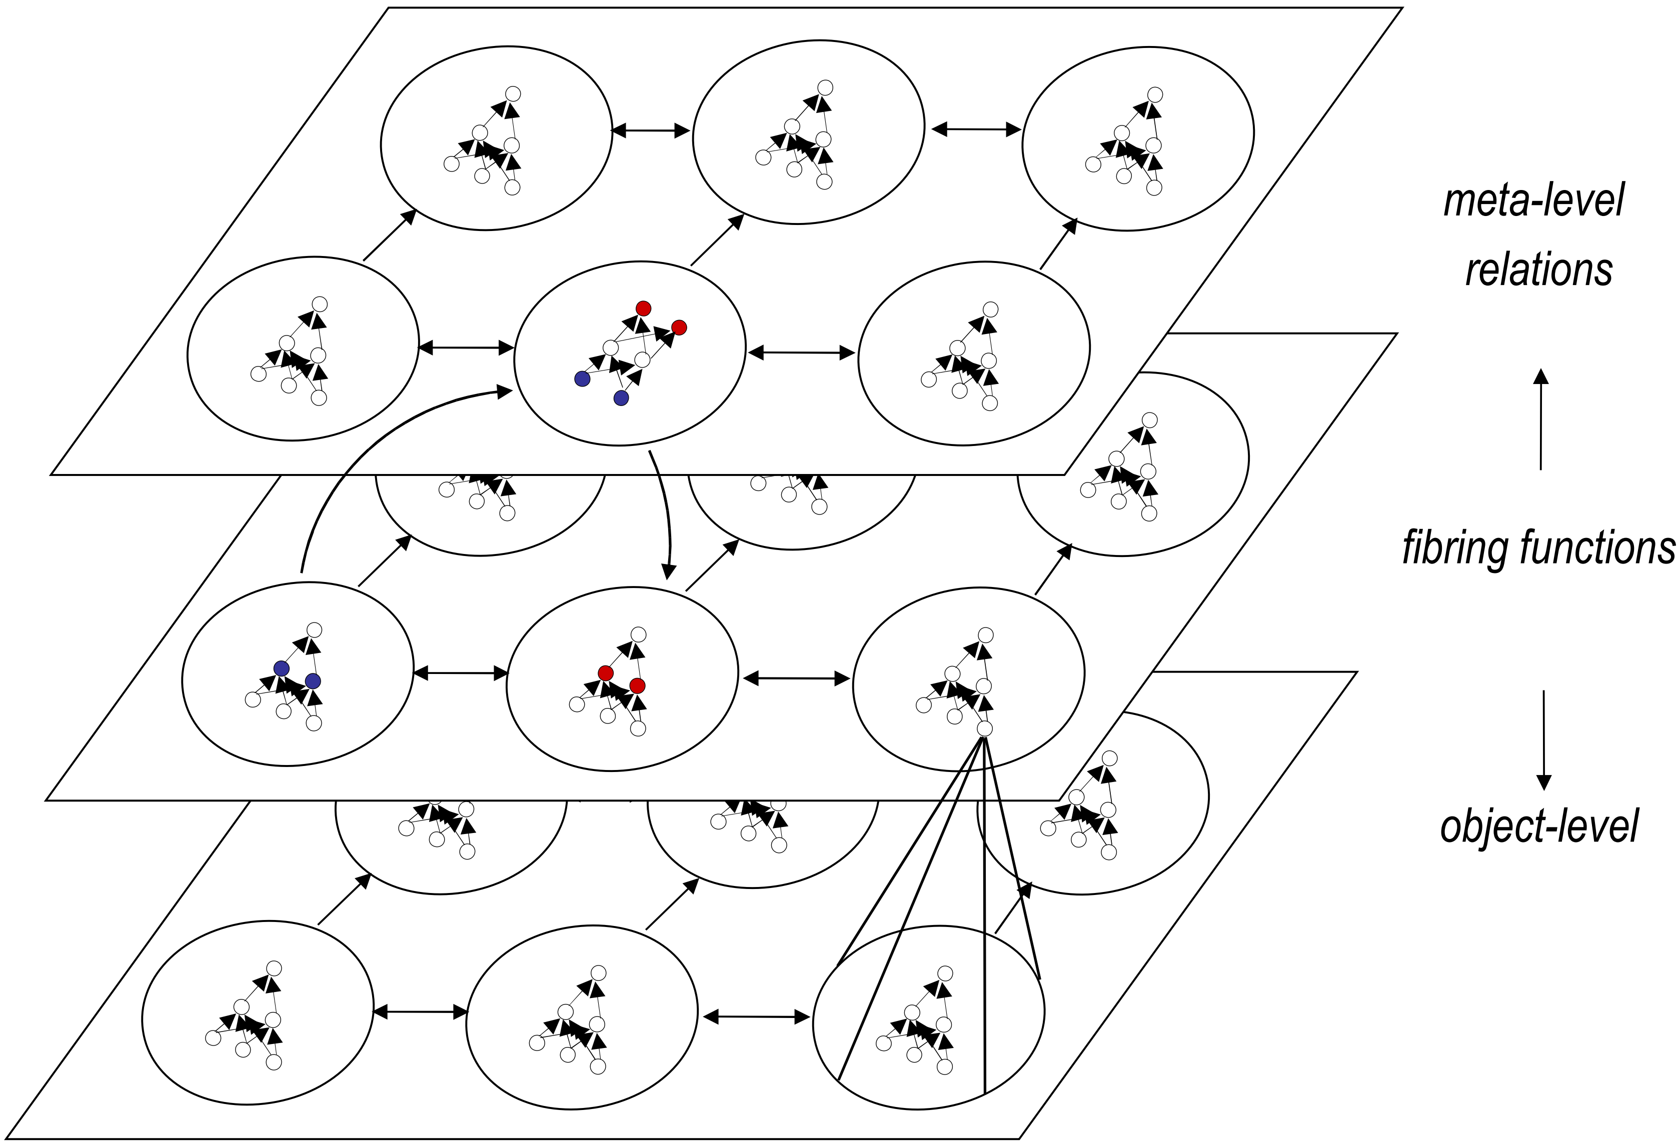

Future Work

Rule Extraction from Network Ensembles

Full First Order Logic (variables, relations and learning)

Experiments on Real-World Problems of Distributed Knowledge Representation

Reasoning about Uncertainty

Challenges

- Reconciling FOL learning and FOL reasoning
- Recurrent vs. Symmetric nets / deep nets?
- Semi-supervised / unsupervised learning
- Gains with parallelism, complexity results
- Cognitive agent: learn-reason-action cycle
- Representations for learning: nonclassical logic
- Abduction: attention, emotions, utility function
- Applications: simulation, video classification, semantic web, software systems (model checking + adaptation)