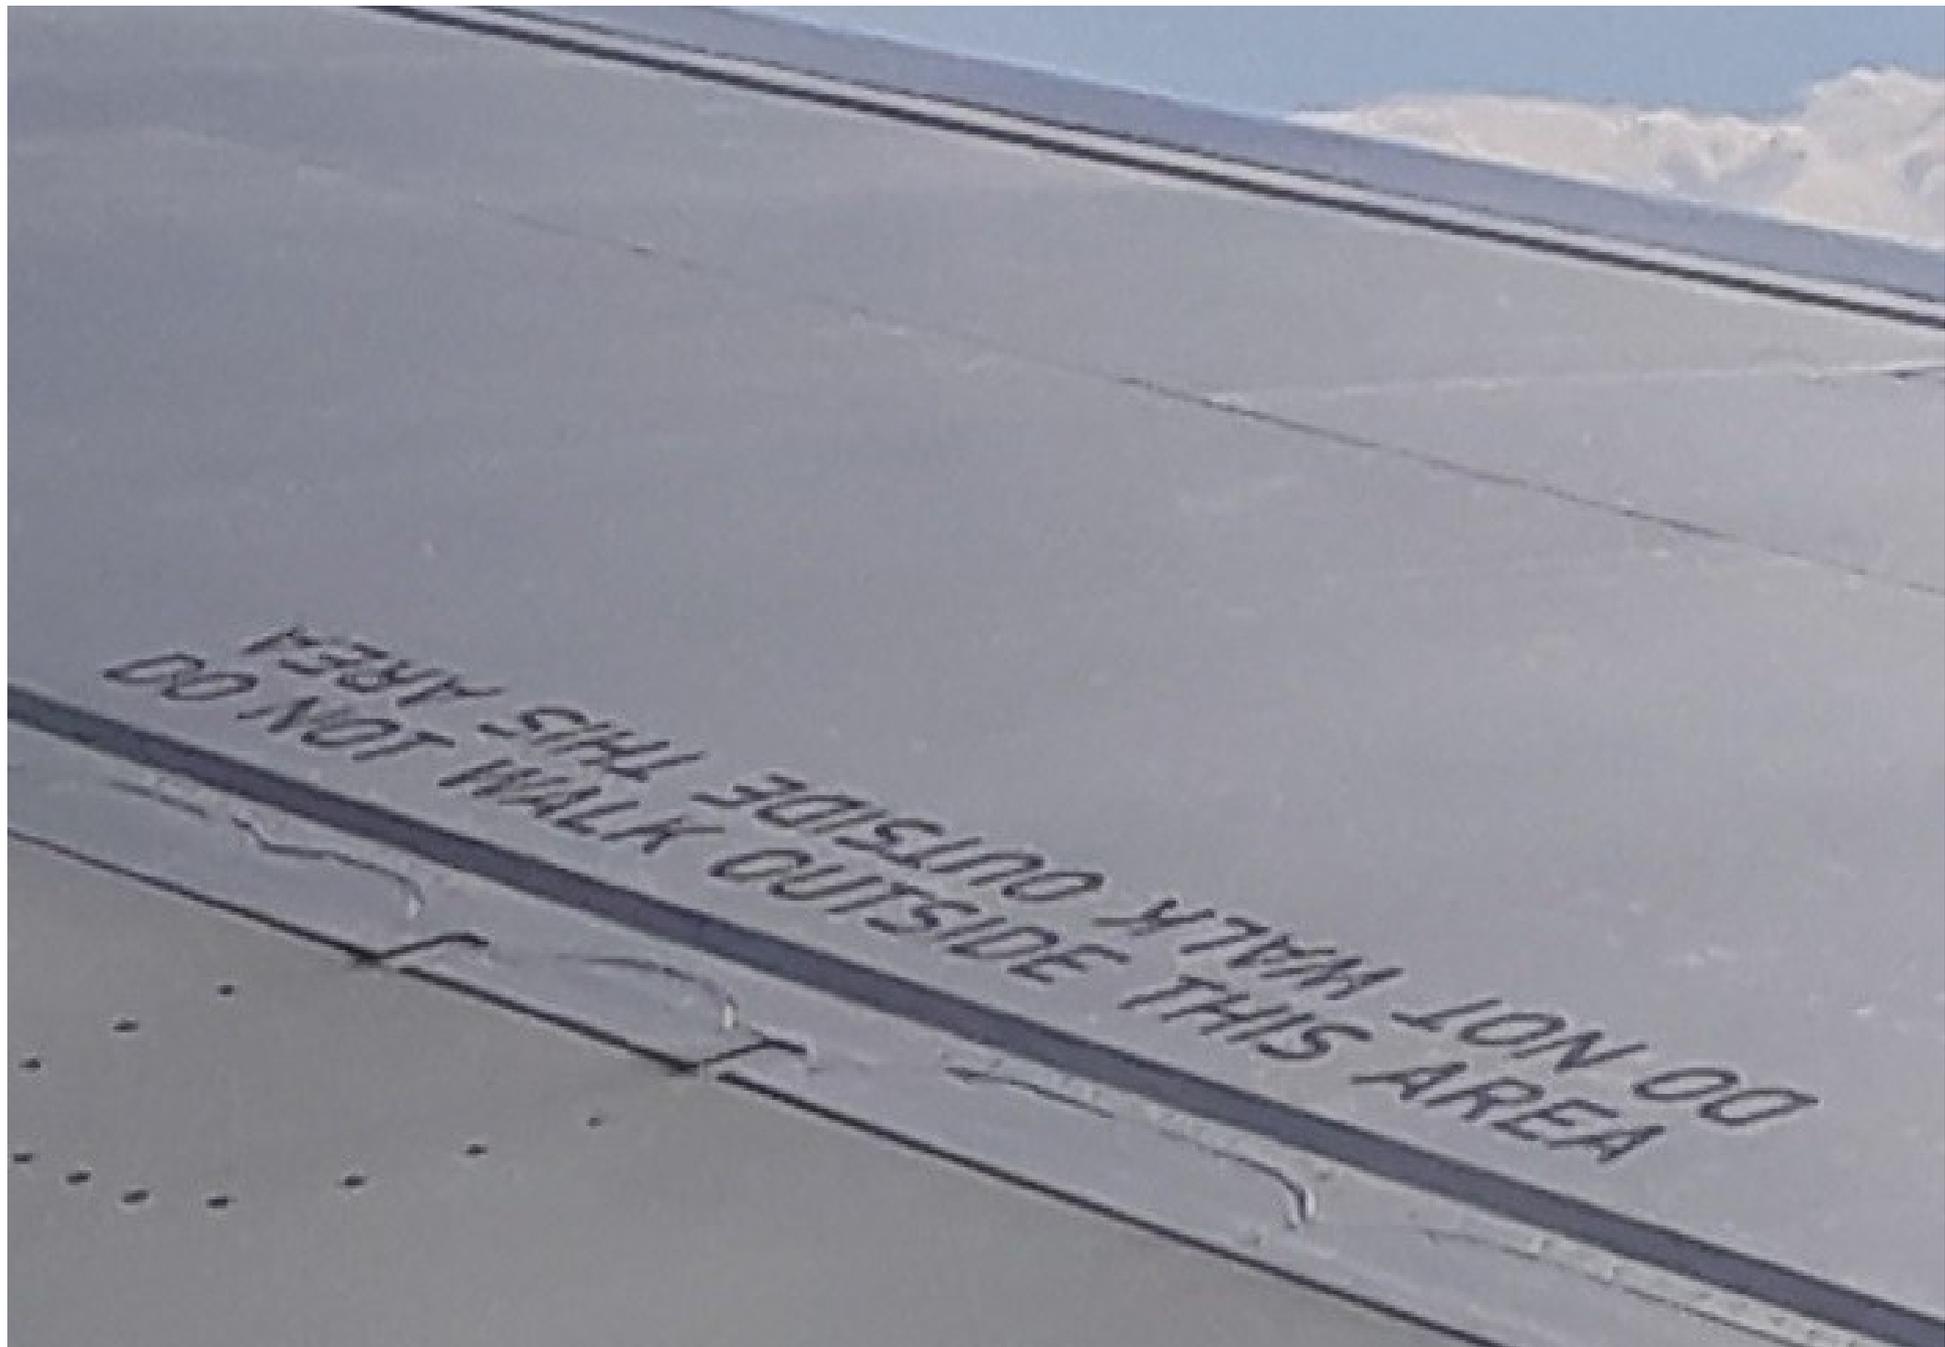


Imperial College London
27 November 2017

Neurosymbolic Computation:
Thinking beyond Deep Learning

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



हेवी सिगल जोइसिंग अटैचमेंट्स नॉन ओनली
DO NOT WALK OUTSIDE THIS AREA

Context is everything... commonsense!



The AI revolution...

The promise of AI:

Education (active learning)

Finance (time series prediction)

Security (image and speech recognition)

Health (sensors, companions, drug design)

Telecom (infrastructure data analysis)

Games (online learning)

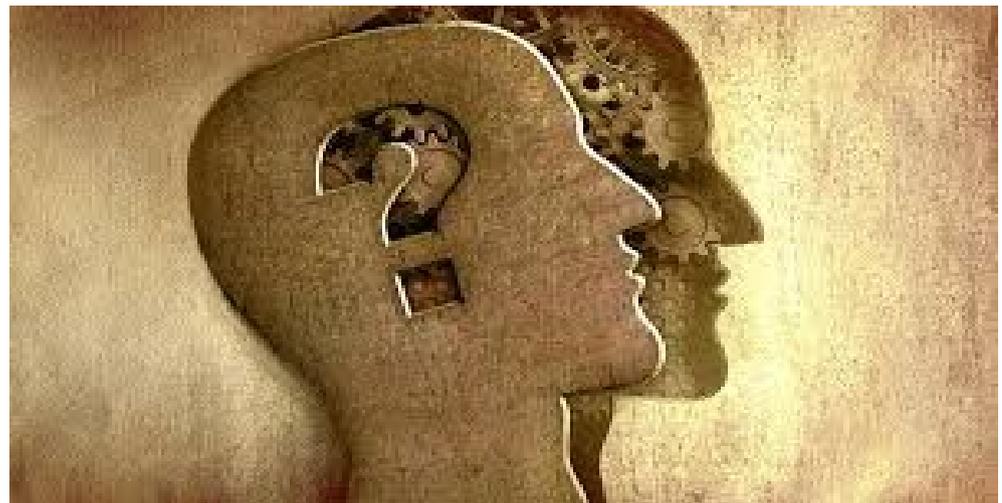
Transport (logistics optimization)

Manufacturing, Retail, Marketing, Energy, etc.

Brain/Mind dichotomy

Symbolic AI: a symbol system has all that is needed for general intelligence

Sub-symbolic AI: intelligence emerges from the brain (neural networks)



Neural-Symbolic Computing

- Neural networks for effective perception
- But perception alone is insufficient:
 - ◆ Reasoning, Explanation, Transfer
- Rich knowledge representation:
nonmonotonic, relational, variables,
recursion, time, uncertainty.
- Neural-symbolic computing: neural networks
with logical structure (in particular:
compositionality)

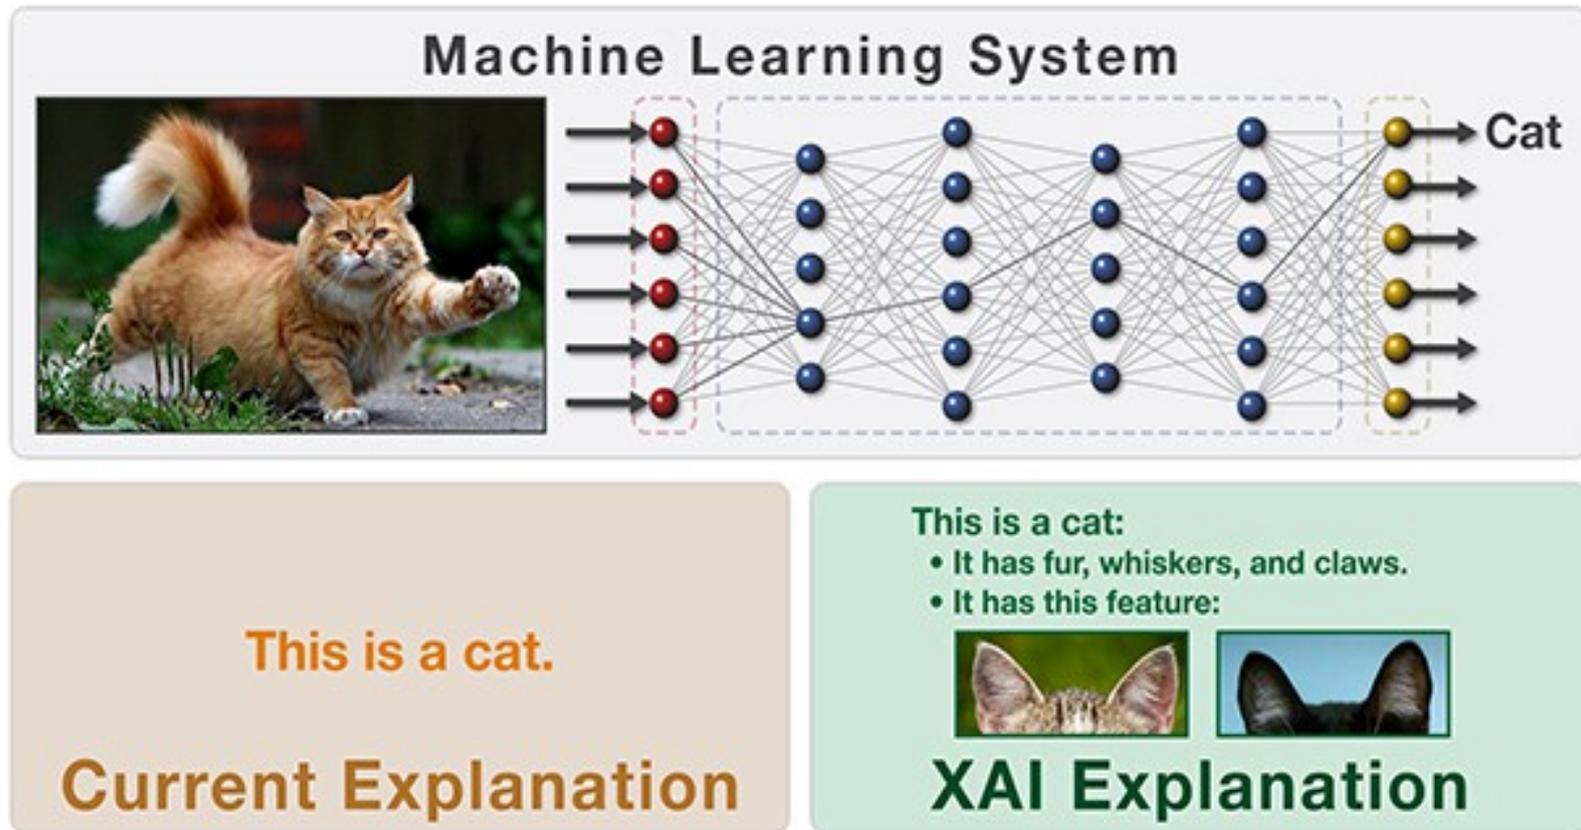
Machine Learning (ML)

Systems that improve their own performance from experience

Systems that, in addition, enable humans to improve their performance (human-machine interaction needed here)

Explainable AI: accountability, trust and transfer learning... [c.f. EU GDPR Reg. 71](#)

DARPA's Explainable AI

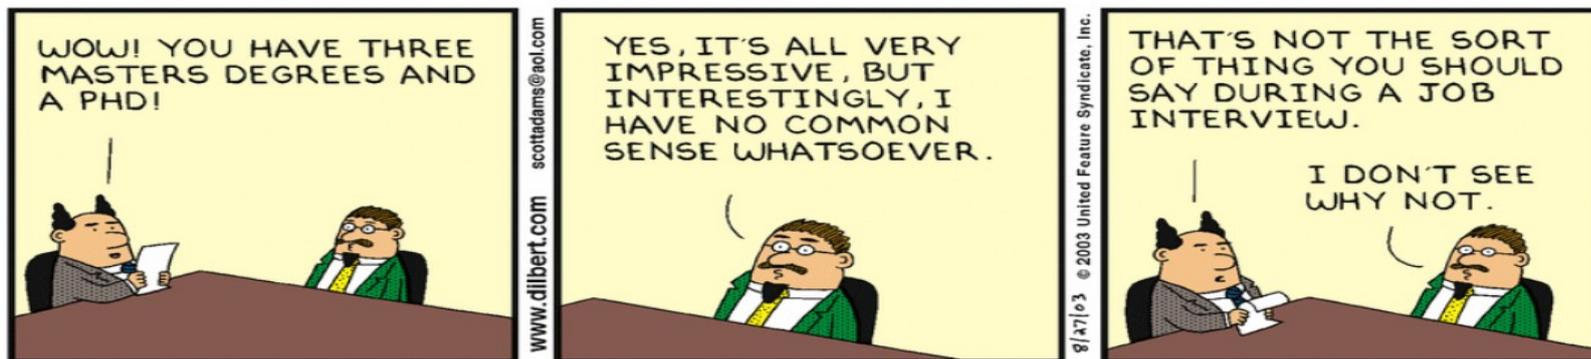


- XAI = Interpretable ML
- Explanation = proof history, not XAI

Deep Learning

Given Big Data, deep learning (based on neural nets) works better than symbolic ML!

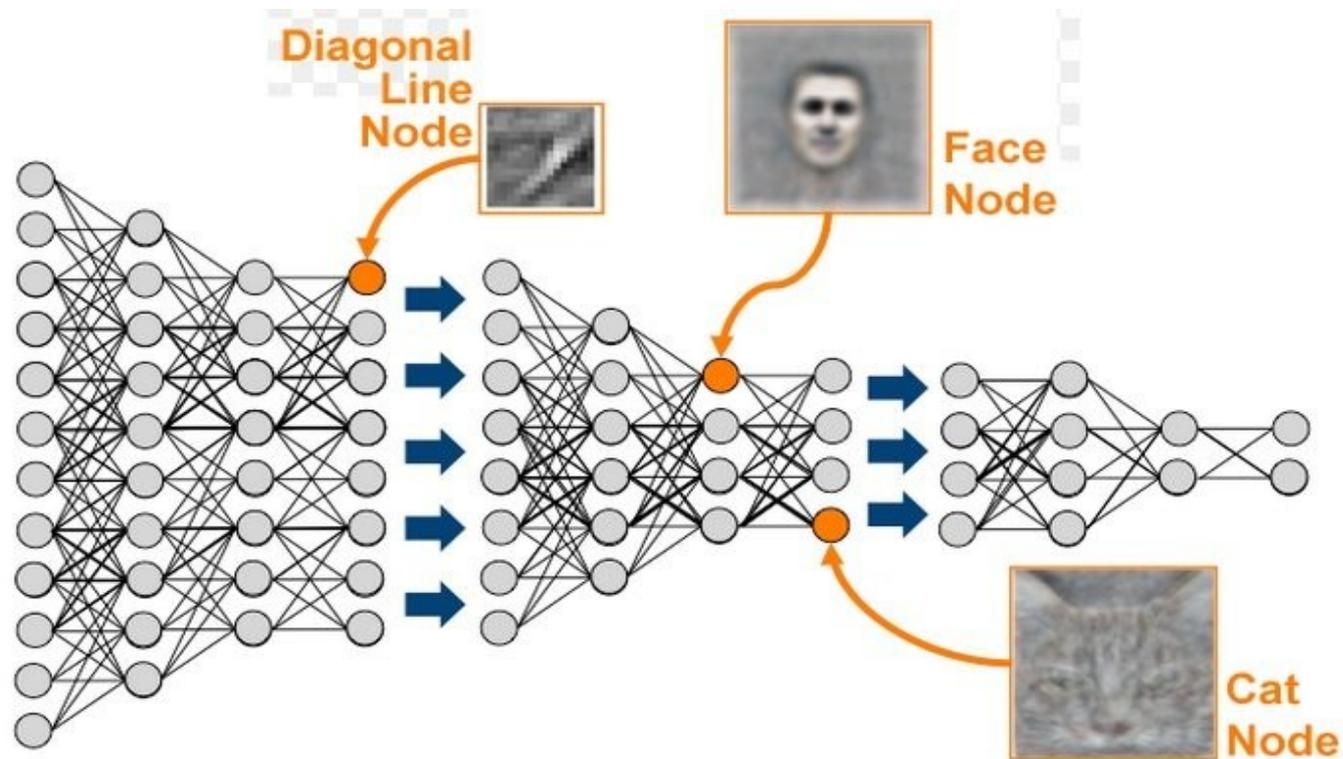
But, where is commonsense?



Deep networks

Very successful on handwritten digit classification, object recognition, speech/audio and games

How about language? Video understanding?

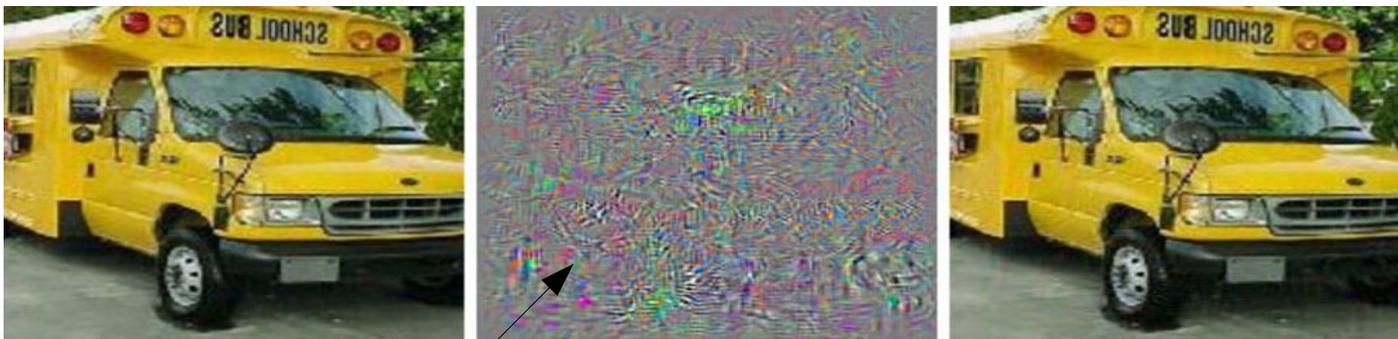


Geoff Hinton says AI needs to start over...

- Deep neural nets (CNNs) do not recognise negative images



- Adversarial networks are not doing much better

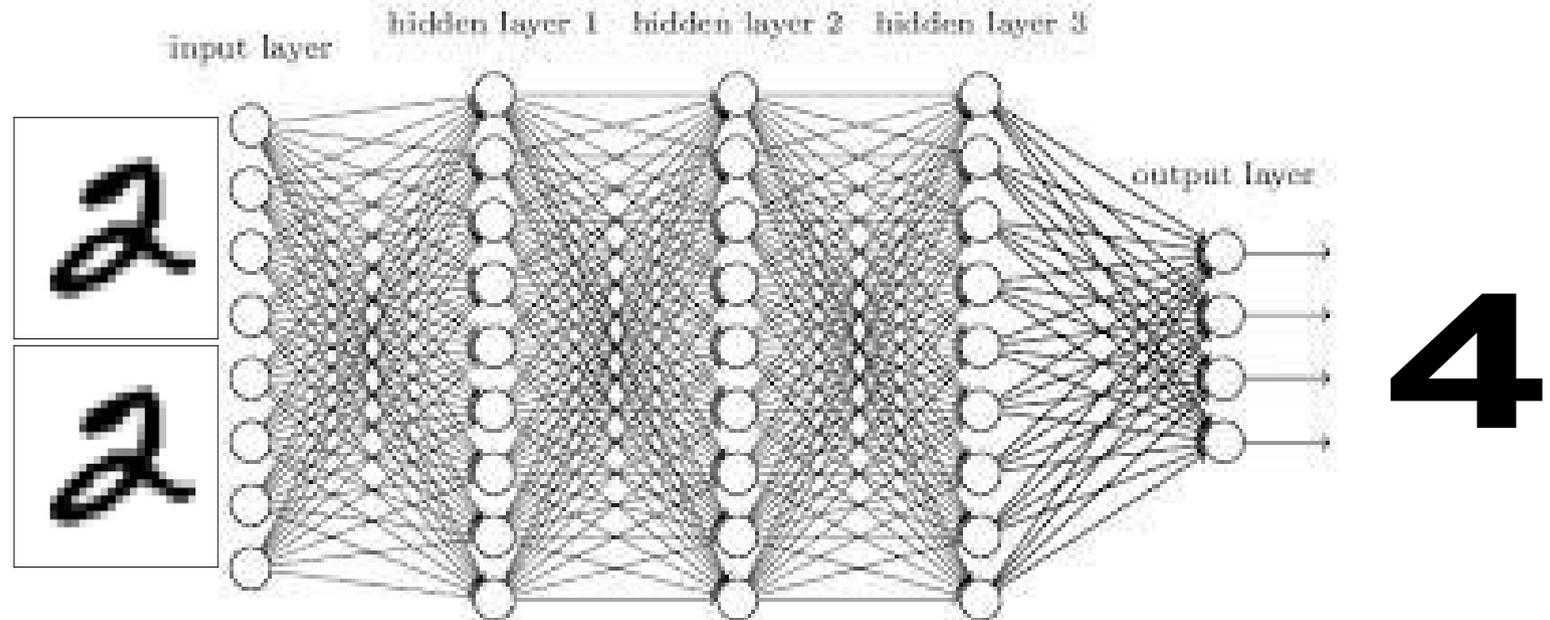


School bus

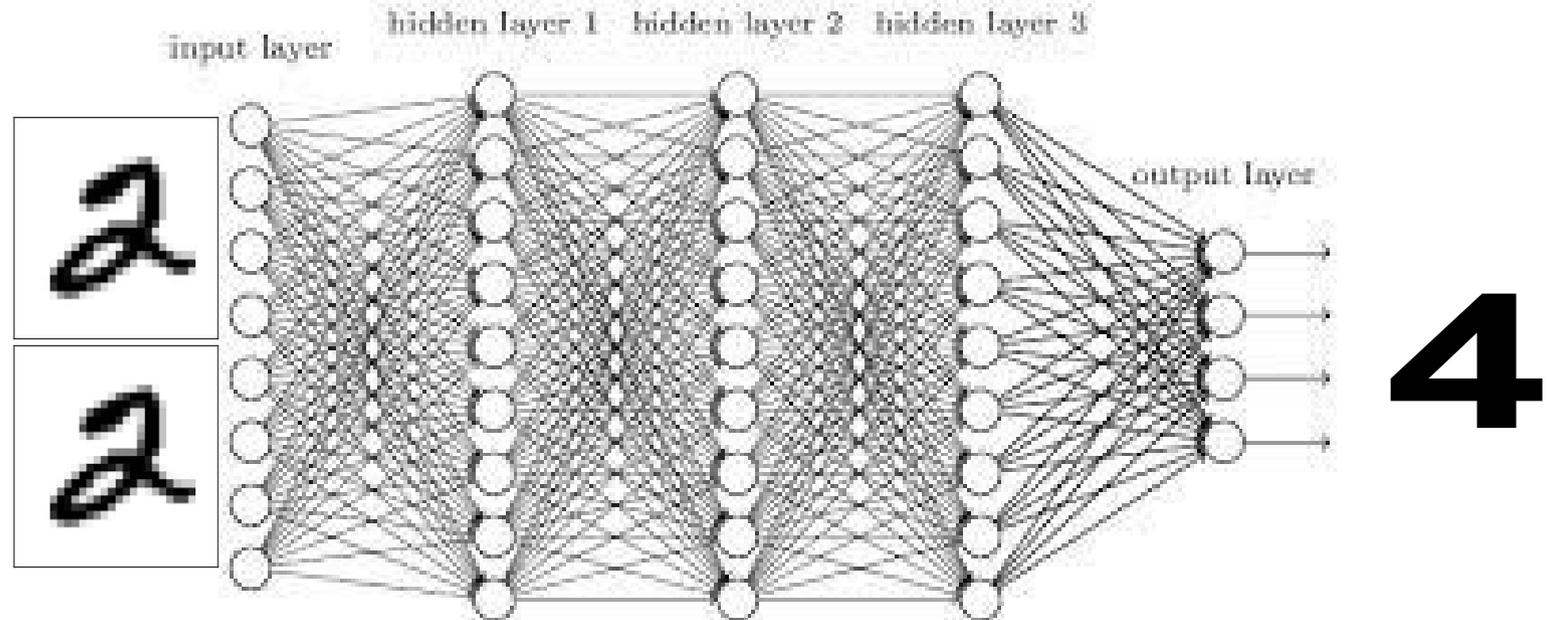
Adversarial perturbations

Ostrich

Knowledge Extraction from Deep Nets



Knowledge Extraction from Deep Nets



$$2+2=4$$

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

A Compiler for Neural Nets

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



translations



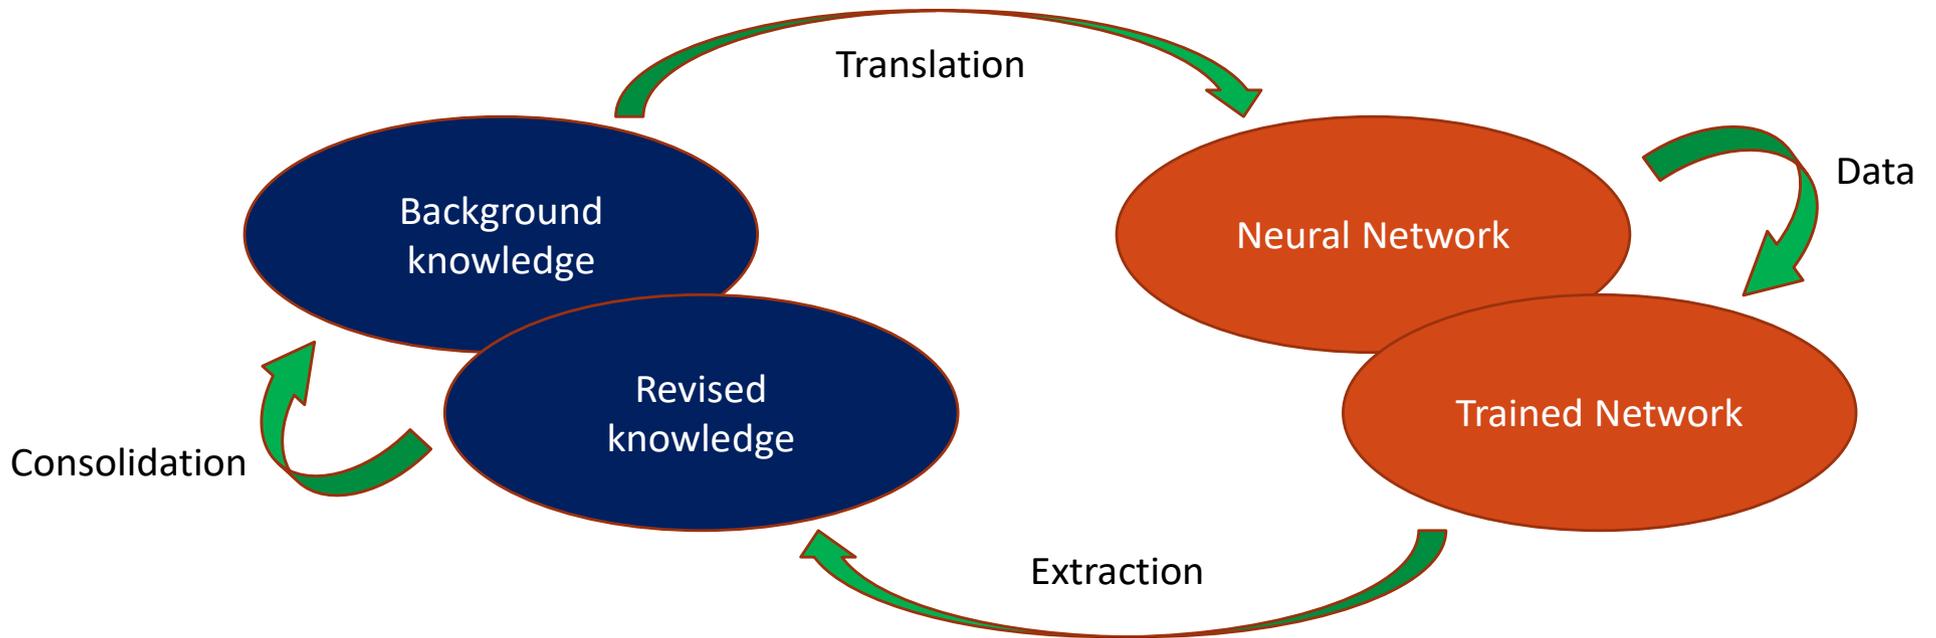
low level, efficient neural structures
(with the same, simple architecture throughout)

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

Neural-symbolic systems were applied to:

- Training in simulators
 - Robotics (robocup)
- Evolution of software models
- Protein classification
- Power systems fault diagnosis
- Semantic web (ontology learning)
- General game playing
- Visual intelligence
- Business compliance

Neural-Symbolic Learning Cycle



Connectionist Inductive Logic Programming (CILP System)

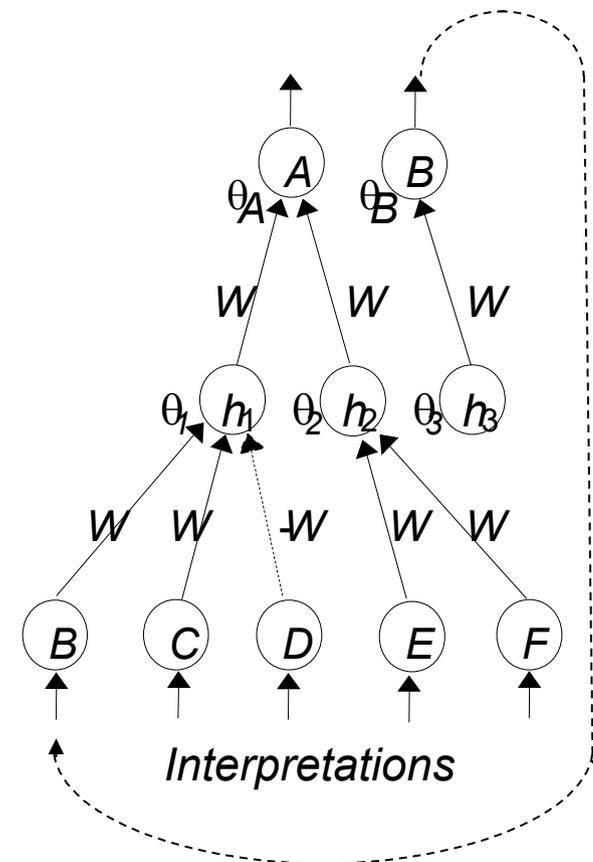
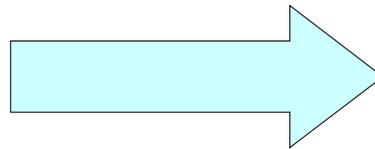
Neural Nets + Logic Programming (rules with exceptions)

Using Background Knowledge
Learning with Backpropagation
Knowledge Extraction

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

$r_2: A \leftarrow E, F;$

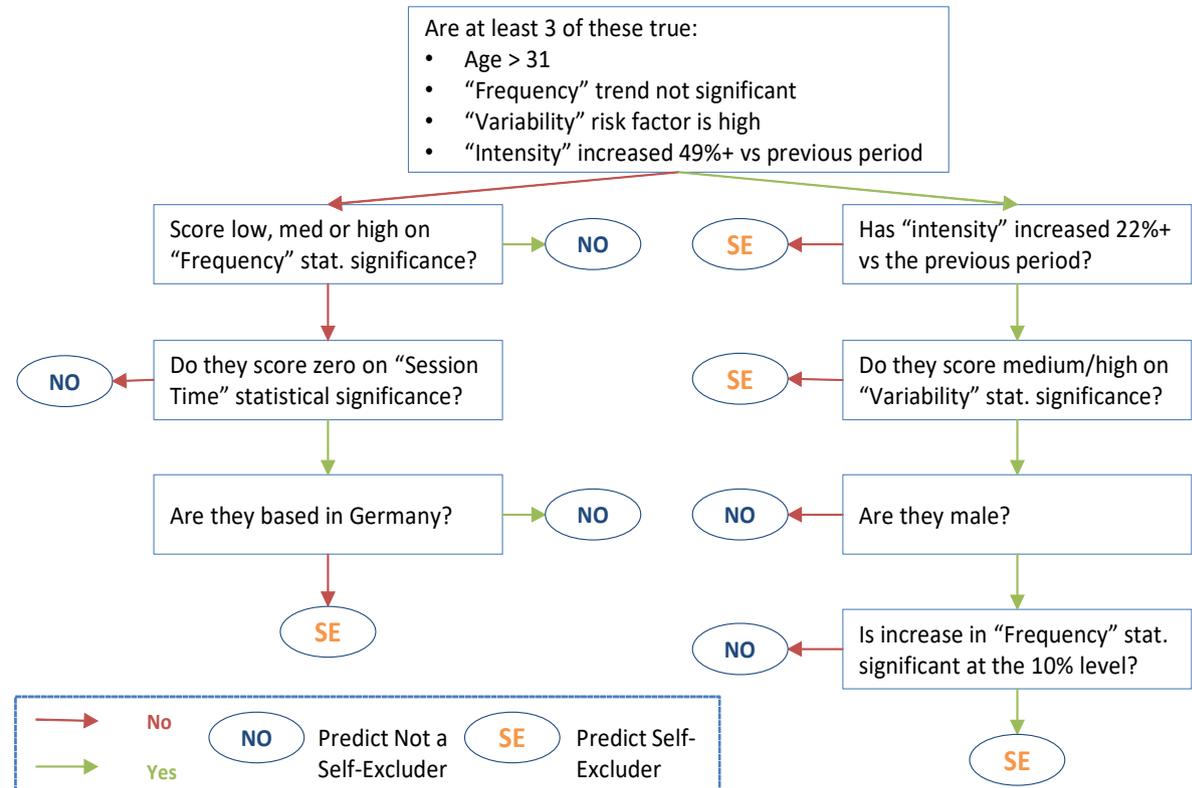
$r_3: B \leftarrow$



Rule Extraction: Neural Net = Black Box?

- Extracted rules can be visualized in the form of a **state transition diagram**

- Alternatively, use TREPAN-like rule extraction 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.

Frosst and Hinton: Distilling a Neural Network Into a Soft Decision Tree, AI-IA CEX workshop, Bari, September 2017.

Relational Learning (CILP++)

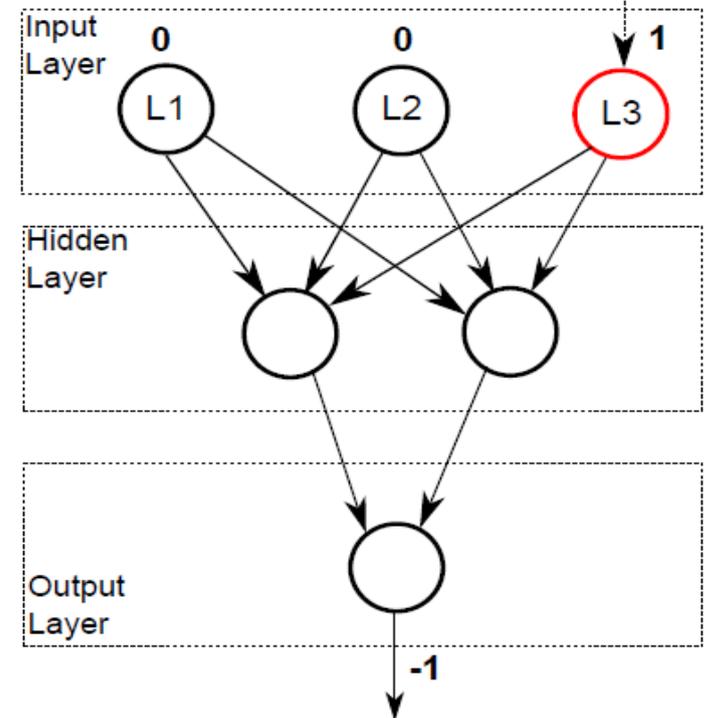
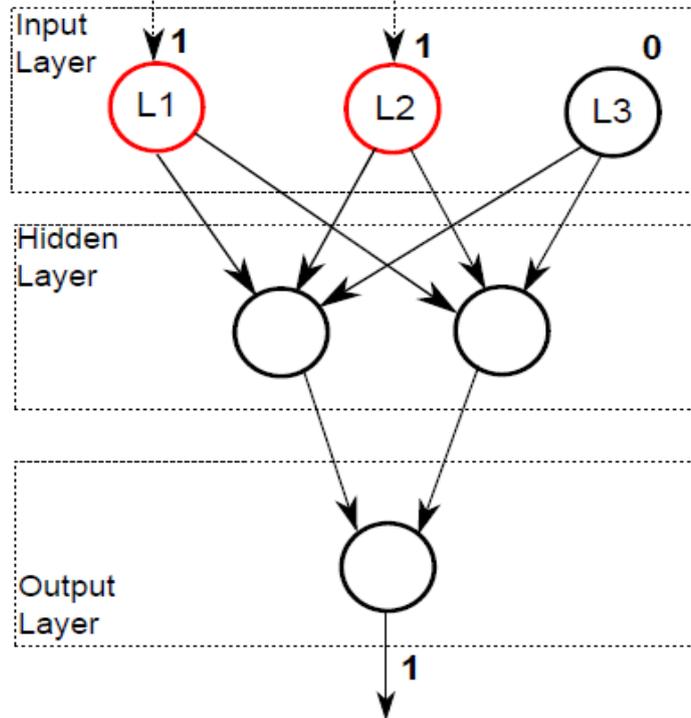
- Extension of CILP allowing the system to be applied directly to ILP problems.
- Each neuron now denotes a first-order literal.
- Data for training the neural network is obtained by **propositionalization**
- Choice of literals to use is important, e.g. consider macro-operators (c.f. Mooney)
- TREPAN-like extraction of first-order rules possible

CILP++

$\text{motherInLaw}(A,B) :- \underline{\text{mother}(A,C)}, \underline{\text{wife}(C,B)}$

2) $\sim\text{motherInLaw}(A,B) :- \underline{\text{wife}(A,C)}$

L1: mother(A,C)
L2: wife(C,B)
L3: wife(A,C)



Franca, Zaverucha and d'Avila Garcez. Fast relational learning using bottom clause propositionalization with artificial neural networks. Machine Learning 94(1):81–104, 2014

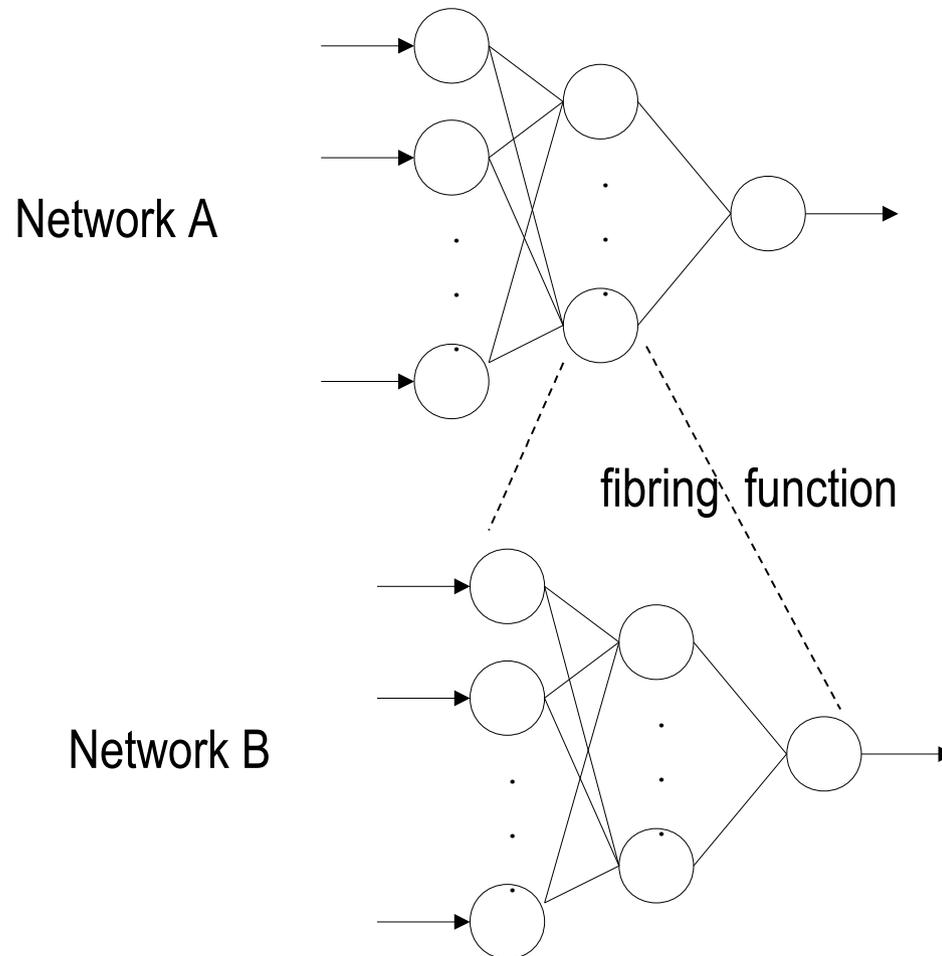
Muggleton and Tamaddoni-Nezhad. QG/GA: A Stochastic Search for Progol. Machine Learning 70(2-3):121-133, 2008.

Richer structures: Fibring of Networks

A neuron that is a network!

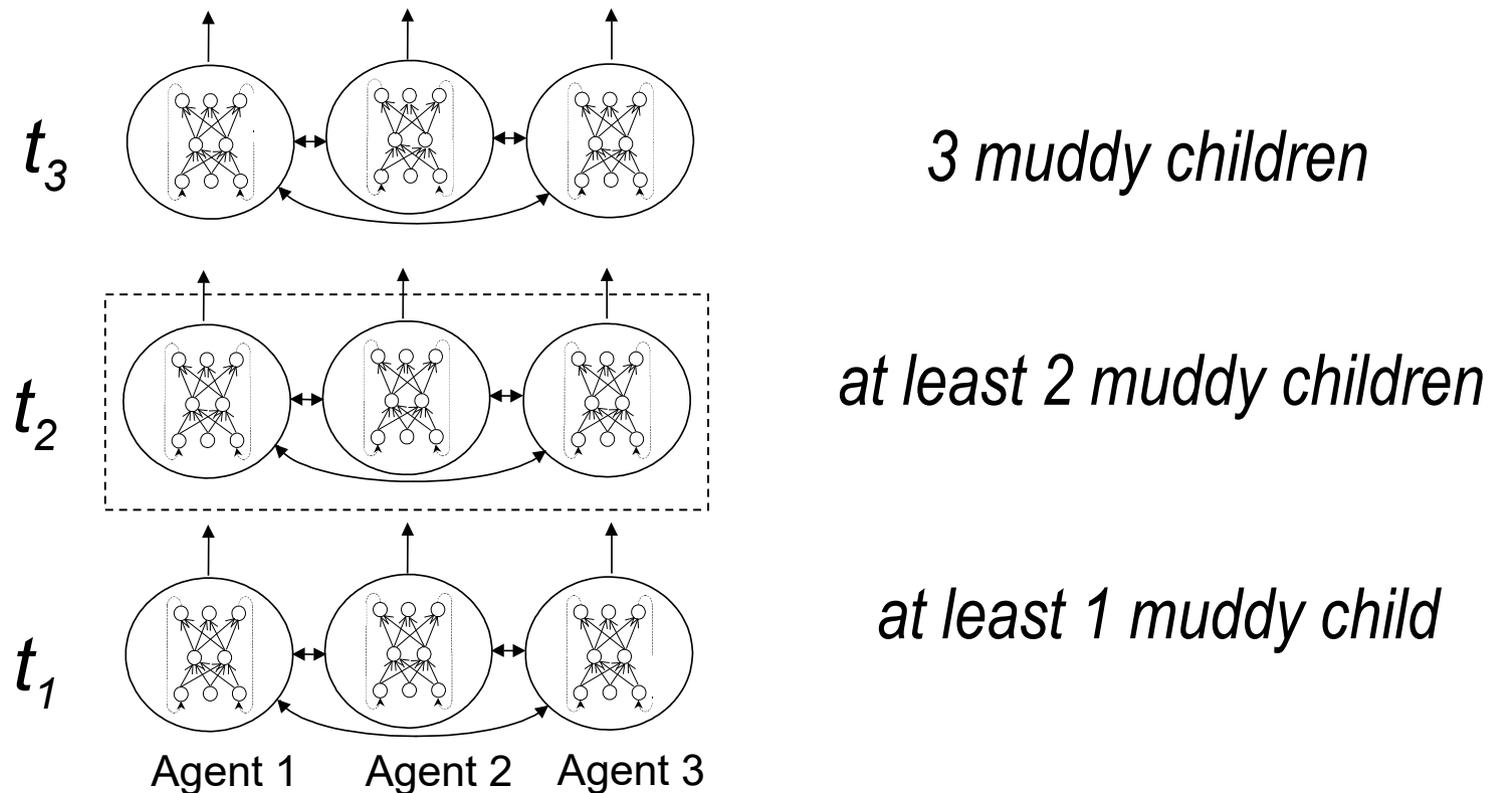
Early form of (**modular**) deep network

Strictly more expressive than shallow nets

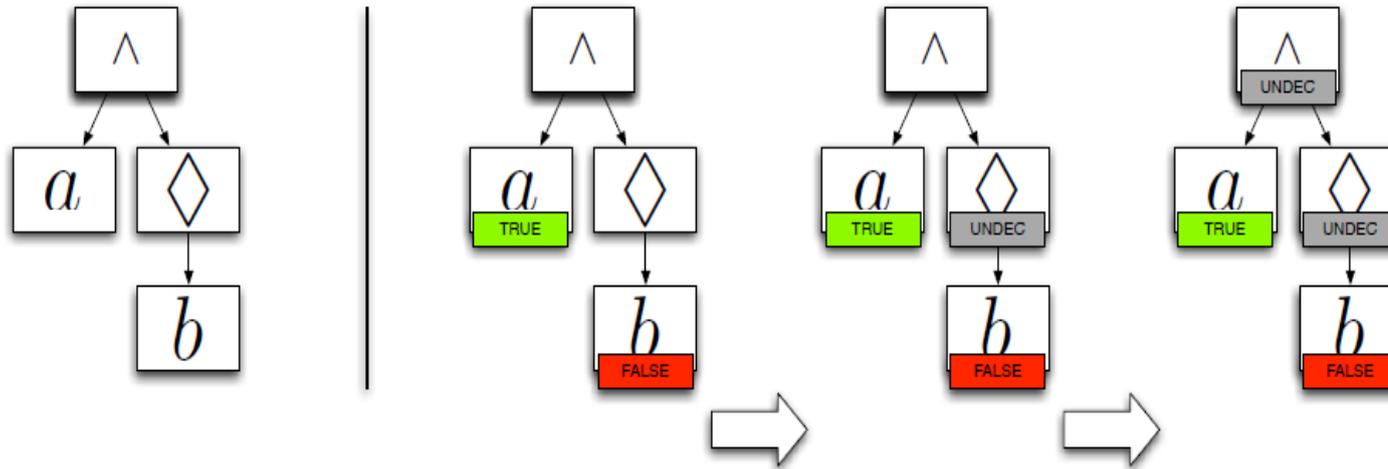
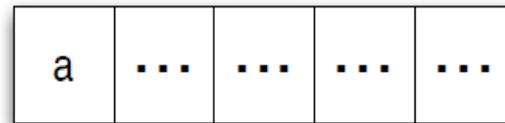


Richer structures: Temporal Reasoning

E.g. Muddy children puzzle (full solution)

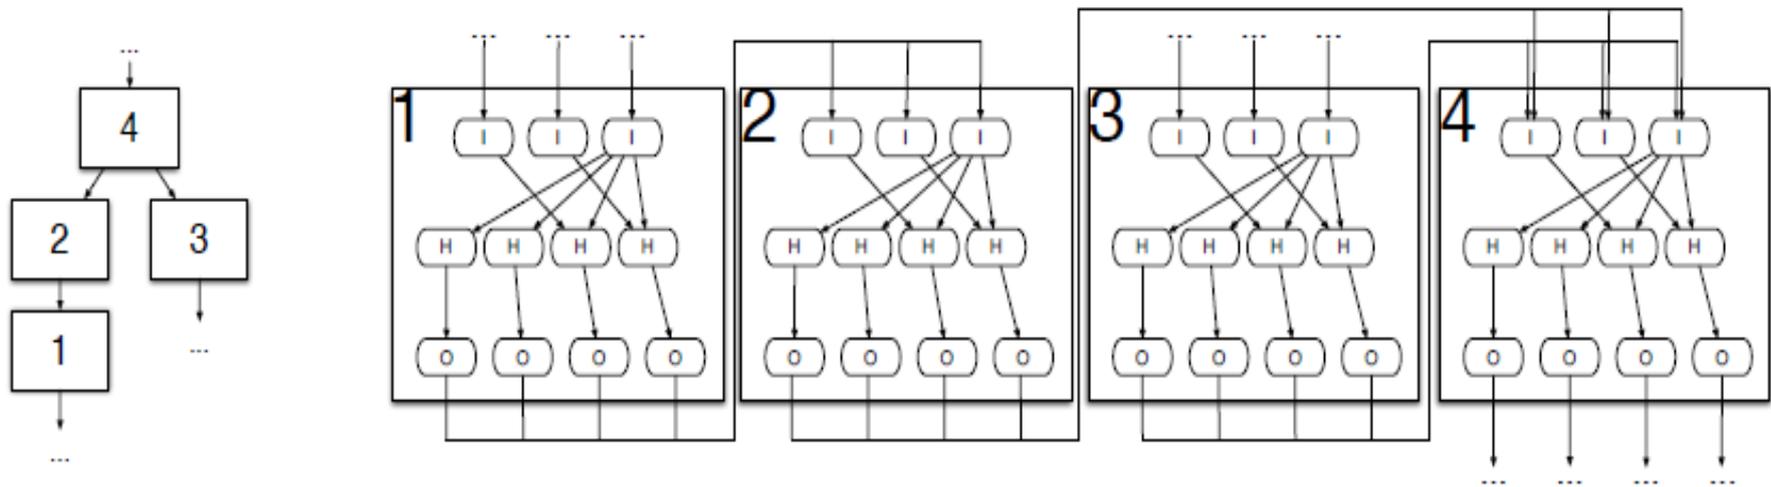


Run-time Monitoring (Propositional Modal Logic)



Run-Time Neural Monitor

The tree structure is “flattened” into an ensemble of CILP networks



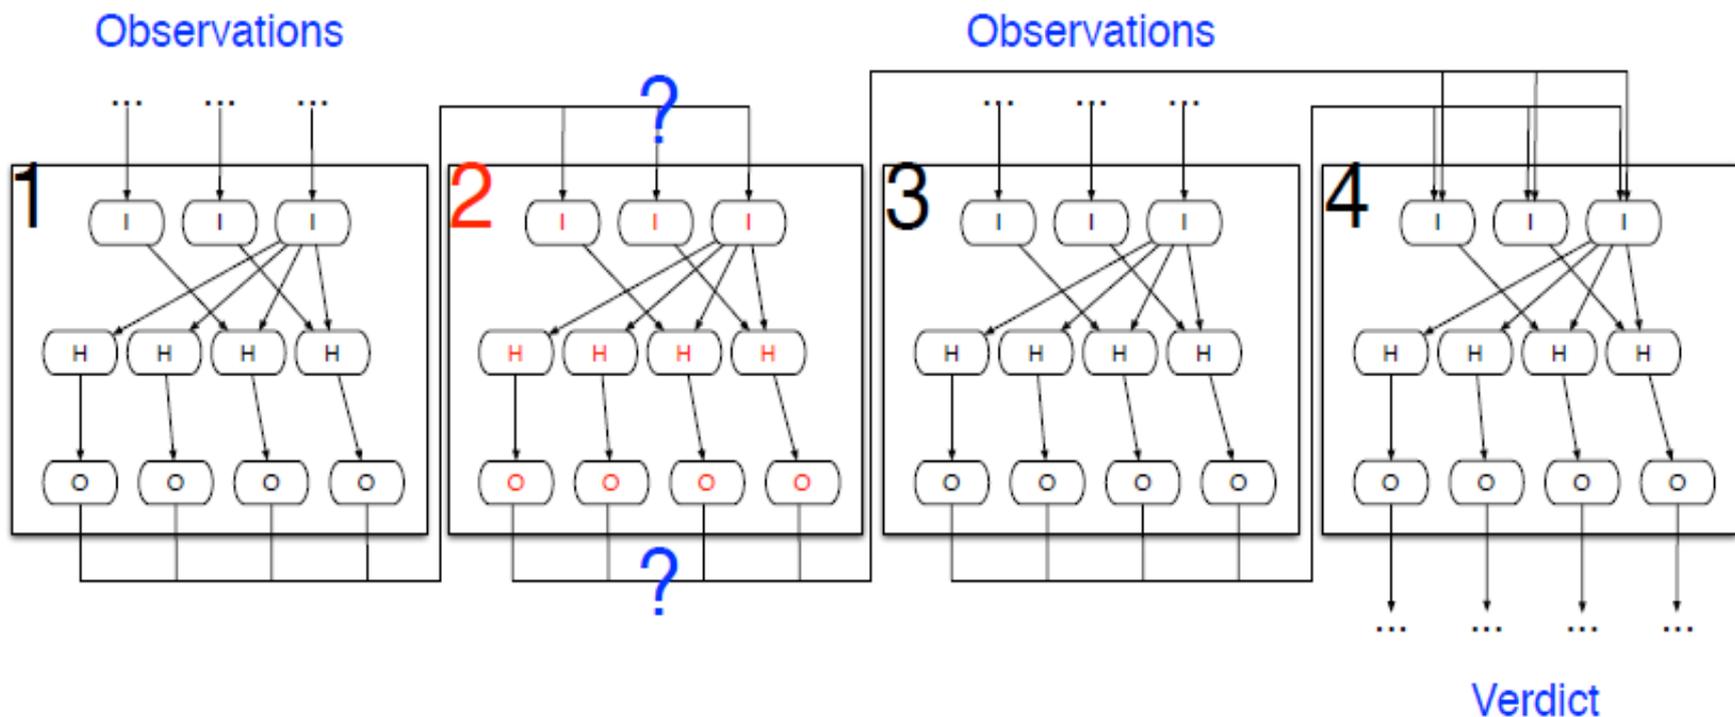
Learning = property adaptation



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

Learning is local (modularity)

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



Transfer Learning (1)



MNIST

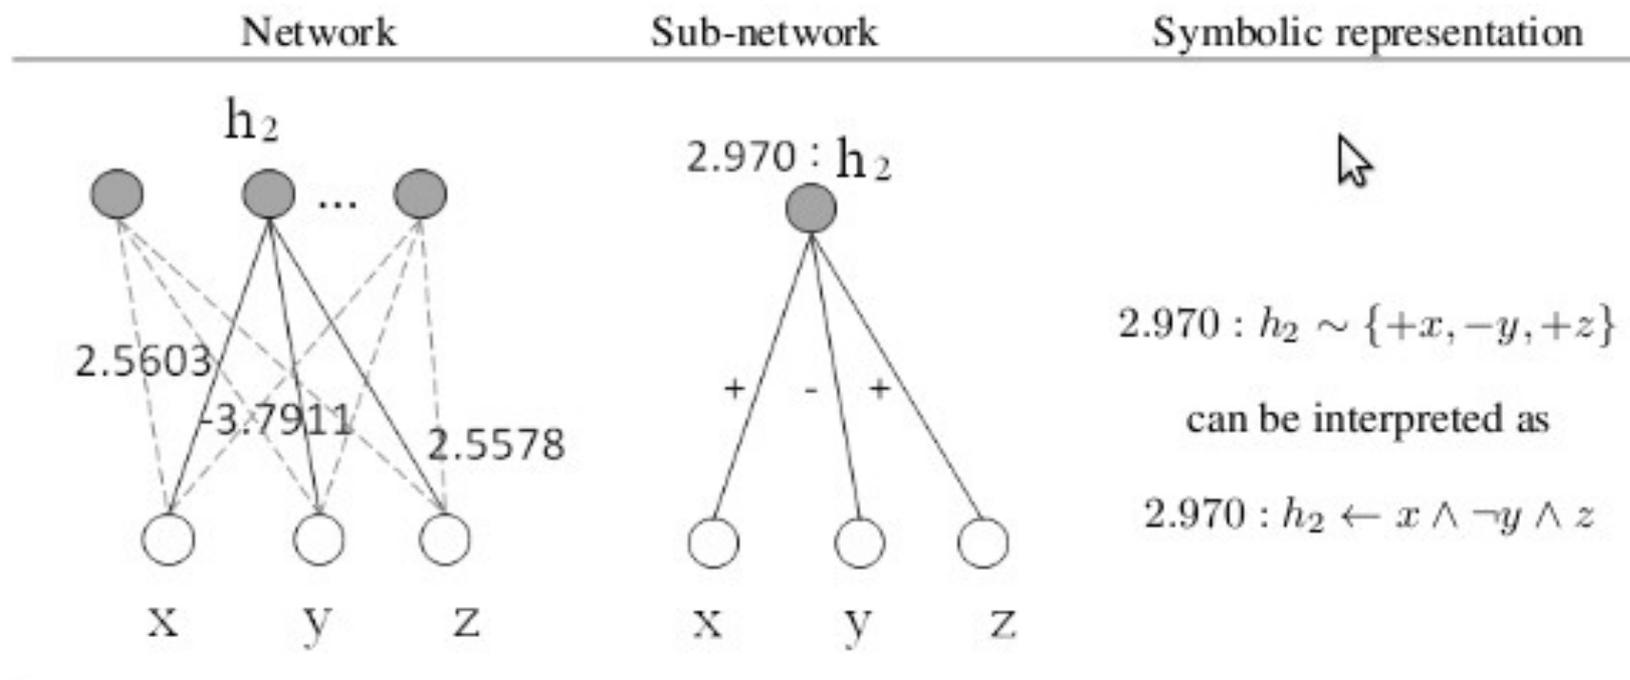


TiCC

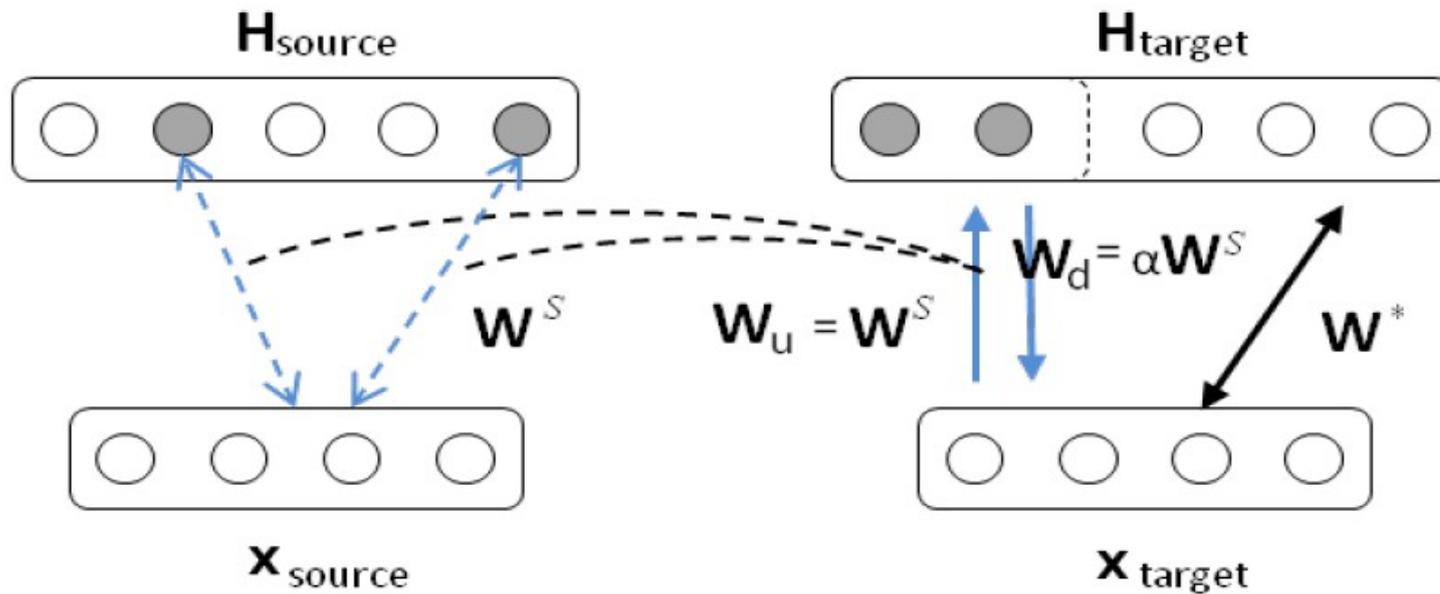


ICDAR

Knowledge extraction from RBMs (originally the building block of (modular) deep nets, c.f. Hinton's DBNs)



Transfer Learning (2)



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

Logic Tensor Networks (LTNs)

- Neural nets with rich structure can represent more than classical propositional logic
- But... neural nets are essentially propositional (John McCarthy was right)
- To take advantage of full FOL, a more **hybrid** approach is needed
- One needs to get the representation right first: the logical statements act as (soft) **constraints** on the neural network...

Semantic Image Interpretation (1)

Given a picture extract a graph that describes its semantic content

Normally, every cat has a tail

Q. Get me the red thing next to the sheep

A. The horse's muzzle? Yes.

$$\forall xy(\text{partOf}(x, y) \rightarrow \neg\text{partOf}(y, x))$$



Make sure your system does not distinguish cats from leopards 99% correctly because of the snow in the background...

Semantic Image Interpretation (2)

In LTN, we build the graph by predicting facts given the bounding boxes, e.g.: Cow(b1), PartOf(b2,b1), Head(b2), etc.

In LTN, an object is described by a vector of features: e.g.
John = (NI number, age, height, 3x4 picture, etc.)

Object detection (bounding box detection and labeling) is performed by an object detector (Fast RCNN)

LTN assigns a **degree of truth** (the grounding G) to atomic formulas: $G(\text{Cow}(b1)) = 0.65$, $G(\text{PartOf}(b2,b1)) = 0.79\dots$

$G(b_i) = \langle \text{score}(\text{Cow}), \text{score}(\text{Leg}) \dots \text{score}(\text{Head}), x, y, x', y' \rangle$



Semantic features: the score of the bounding box detector on b_i for each class of objects



Geometric features: the coordinates of b_i

LTN in action

1. $\forall x(\neg PartOf(x, x))$
2. $\forall xy(PartOf(x, y) \rightarrow \neg PartOf(y, x))$
3. $\forall xy(Cow(x) \wedge PartOf(x, y) \rightarrow Leg(y) \vee Neck(y) \vee Torso(y) \vee Head(y))$
4. $\forall xy(Cow(x) \rightarrow \neg PartOf(x, y))$
5. $\forall xy(Torso(x) \rightarrow \neg PartOf(y, x)).$

- Grounding for PartOf is given by the % of intersection between two bounding boxes
- One can query the knowledge-base (KB) to obtain further groundings for training
- Learning is... **maximizing satisfiability!**

Learning in LTNs...

Given a KB and groundings, LTN calculates a grounding for the entire KB compositionally in the “usual ways” ...

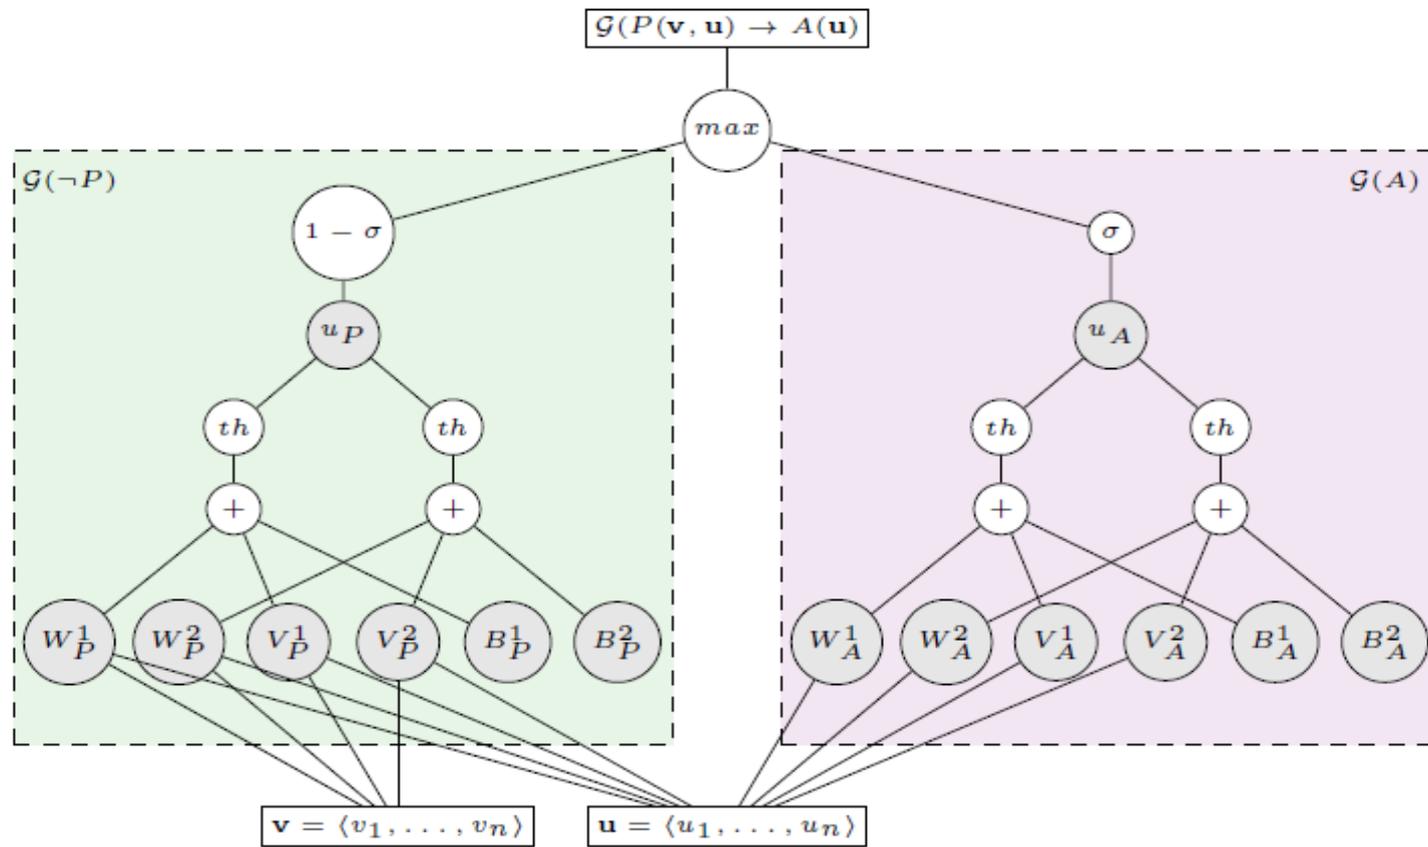


Fig. 1. Tensor net for $P(x, y) \rightarrow A(y)$, with $\mathcal{G}(x) = \mathbf{v}$ and $\mathcal{G}(y) = \mathbf{u}$ and $k = 2$.

The Tensor Network...

$$\mathcal{G}(f)(\mathbf{v}_1, \dots, \mathbf{v}_m) = M_f \mathbf{v} + N_f$$

$$\mathcal{G}(P) = \sigma \left(u_P^T \tanh \left(\mathbf{v}^T W_P^{[1:k]} \mathbf{v} + V_P \mathbf{v} + B_P \right) \right)$$

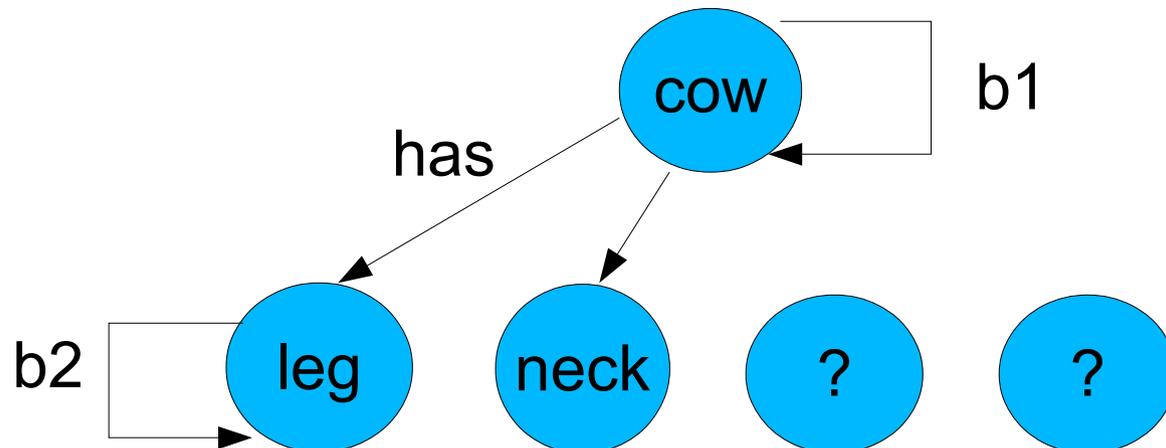
$$\mathcal{G}^* = \operatorname{argmin}_{\hat{\mathcal{G}} \subseteq \mathcal{G} \in \mathbb{G}} \sum_{\langle [v, w], \phi(\mathbf{t}) \rangle \in \mathcal{K}_0} \operatorname{Loss}(\mathcal{G}, \langle [v, w], \phi(\mathbf{t}) \rangle)$$

Fast RCNN + LTN improves on Fast RCNN (state of the art at the time) at object type classification:

I. Donadello, L. Serafini and A. S. d'Avila Garcez. Logic Tensor Networks for Semantic Image Interpretation. In Proc. IJCAI'17, Melbourne, Australia, Aug 2017.

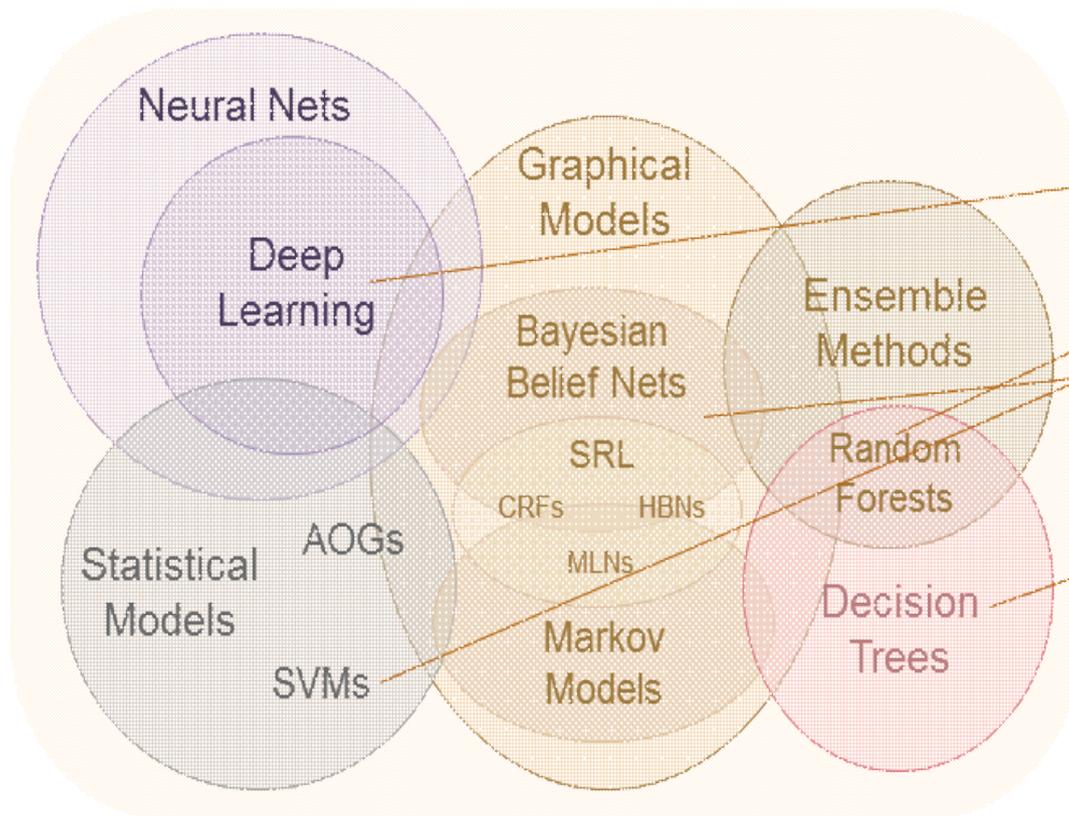
And finally, the knowledge graph...

- Given a trained LTN, start with an unlabeled graph.
- For every bounding box b_i ask the LTN for the set of facts $\{\text{Cow}(b_i), \text{Leg}(b_i), \text{Neck}(b_i), \text{Torso}(b_i), \dots\}$ and select the facts with grounding larger than a threshold.
- For every bounding box b_i ask the LTN for the set of facts $\{\text{PartOf}(b_i, b_j)\}$ with $j = 1, \dots, n$. Then, select the facts with grounding larger than a threshold.

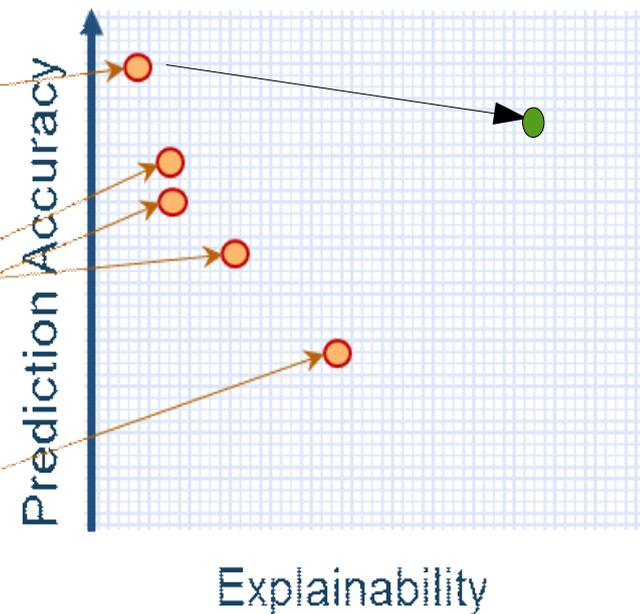


Explainable AI = ML + KR

Learning Techniques (today)



Explainability (notional)



Source: DARPA

- What do I need to change in order to have my credit application accepted the next time?

Verification of Neural Nets

Whose fault is it when a self-driving car gets into an accident?

In the news recently: Tesla self-driving system cleared in deadly crash

There will be fewer deaths from self-driving cars!
Is this argument good enough?



Deep Learning + Symbolic ML

- We want good classification and prediction but also useful descriptions... e.g.: learning factorial ($n!$)

$$0! = 1$$

$$1! = 1$$

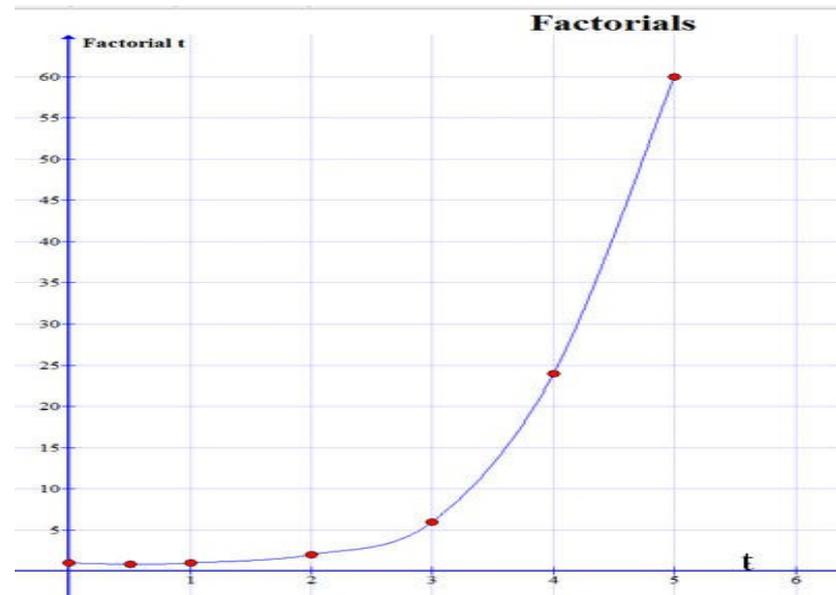
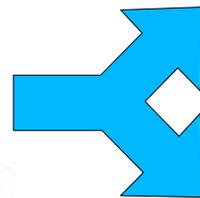
$$2! = 1 \cdot 2 = 2$$

$$3! = 1 \cdot 2 \cdot 3 = 6$$

$$4! = 1 \cdot 2 \cdot 3 \cdot 4 = 24$$

$$5! = 1 \cdot 2 \cdot 3 \cdot 4 \cdot 5 = 120$$

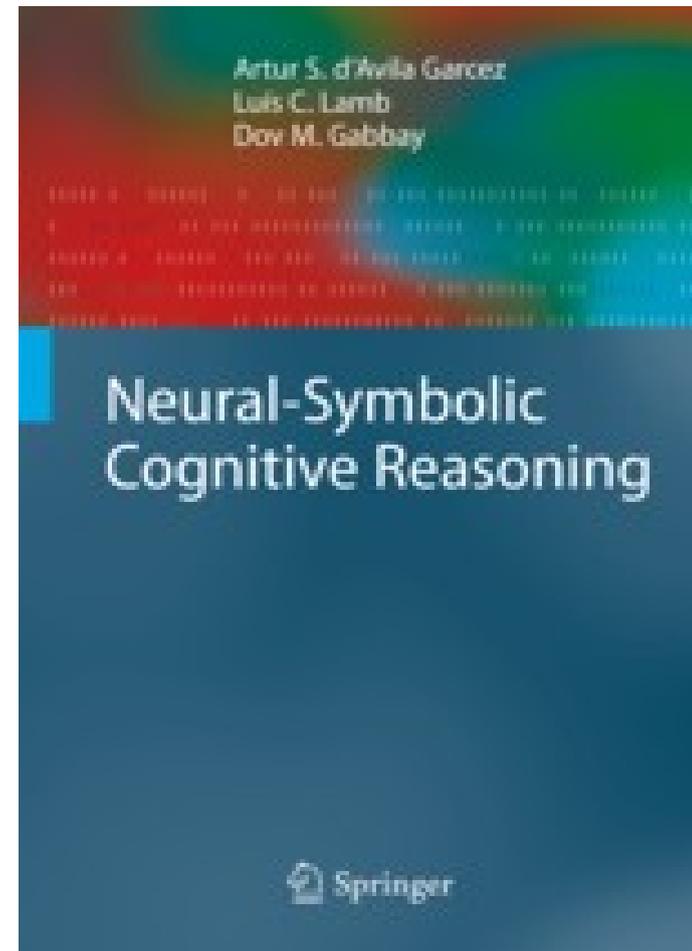
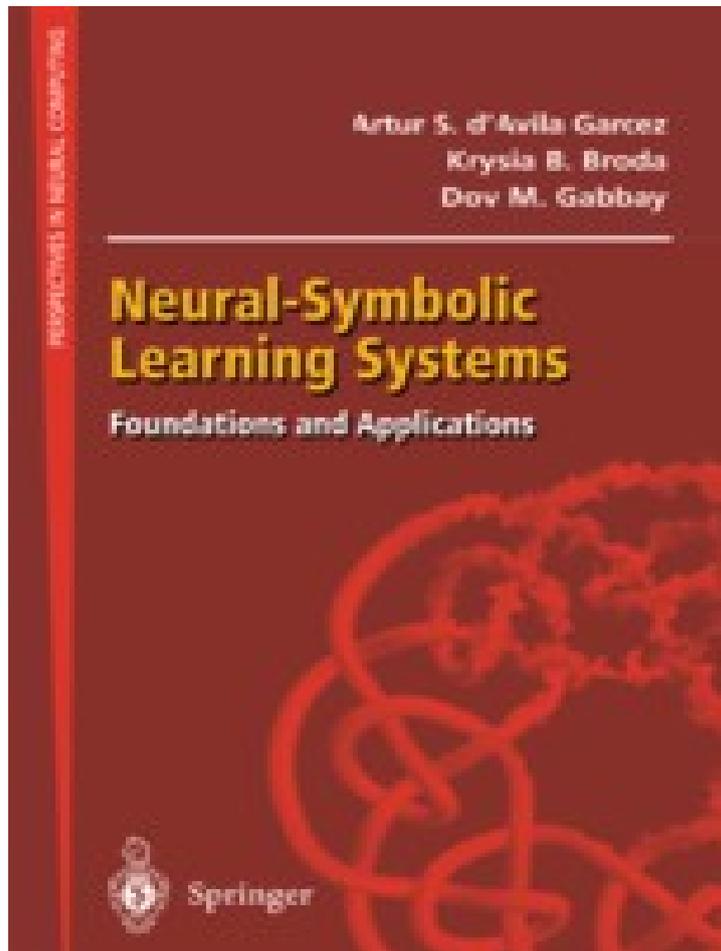
$$6! = 1 \cdot 2 \cdot 3 \cdot 4 \cdot 5 \cdot 6 = 720$$



`factorial(0,1).`

```
factorial(N,F) :- N>0, N1 is N-1,
                 factorial(N1,F1),
                 F is N * F1.
```

For more information...



Cognitive Computation Symposium: Thinking Beyond Deep Learning

27 Feb 2018, 9am – 5pm

City, University of London, Northampton Square, EC1V 0HB

Speakers:

Antoine Bordes (Facebook AI Research)

Edward Grefenstette (Google DeepMind)

Barbara Hammer (Bielefeld University)

Frank van Harmelen (VU Amsterdam)

Kristian Kersting (TU Darmstadt)

Alessio Lomuscio (Imperial College London)

Stephen Muggleton (Imperial College London)

Michael Spranger (SONY CSL)

Geraint Wiggins (Queen Mary, University of London)

Michael Witbrock (IBM Research AI)

Michael Wooldridge (University of Oxford)

Willem Zuidema (University of Amsterdam)

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 AI applications.

Thank you!

Throw away your paradigm...

neurons



symbols



The future is neural-symbolic

Paraphrased from
Murray Shanahan