

Cognitive Computation 2018
City, University of London
27 February 2018

Thinking beyond deep learning? Neural-Symbolic Computing

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

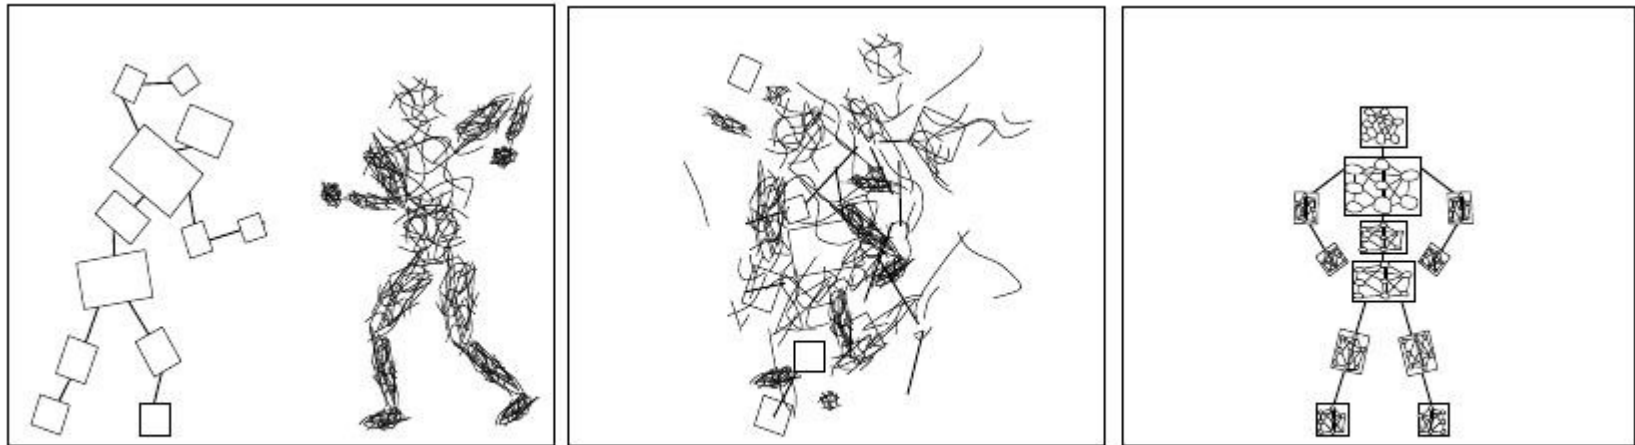


Figure 1. Conflict between theoretical extremes.

Attributed to Marvin Minsky's
Logical Versus Analogical or Symbolic Versus Connectionist or Neat Versus Scruffy, AI Magazine, 1991

The AI revolution...

US\$40B investment in AI (mostly ML) in 2016 and growing, but...

AI adoption still low in 2017 (McKinsey)

Education (active learning)

Finance (time series prediction)

Security (image and speech recognition)

Health (companions, drug design)

Telecom and Tech (infrastructure data analysis)

Gaming (online learning)

Transport (logistics, car industry)

Manufacturing, Retail, Marketing, Energy, etc.

AI revolution mainly due to...

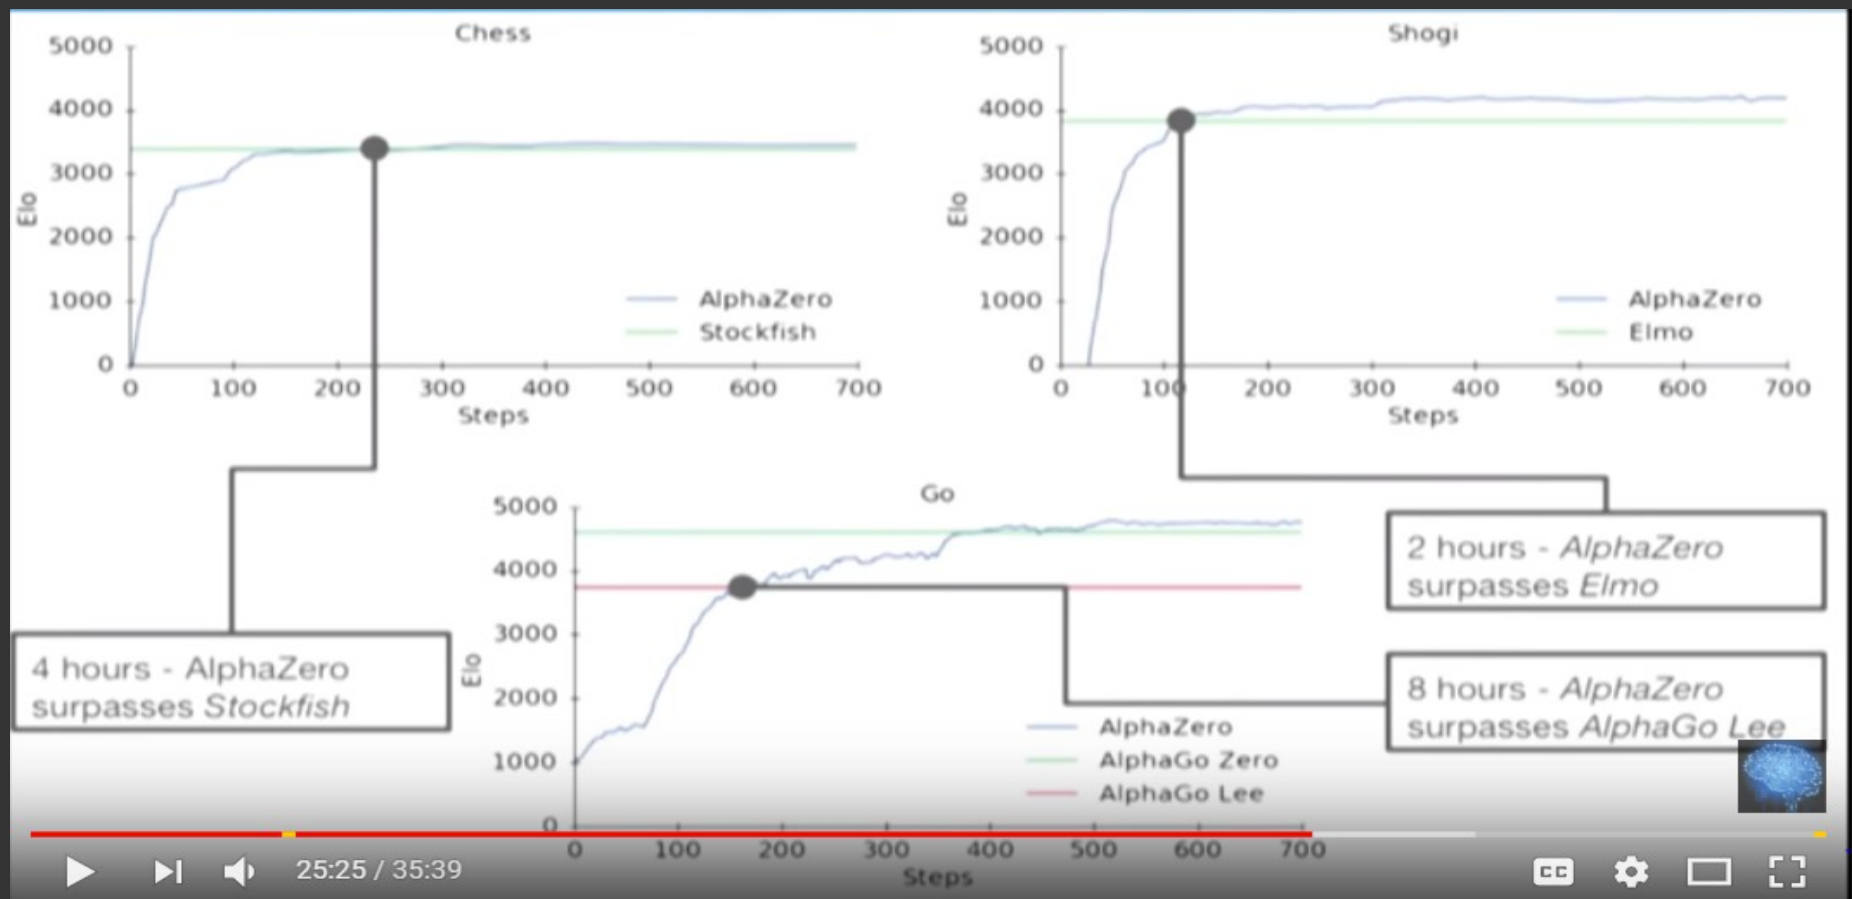
... deep learning

Very nice original idea (deep belief nets; semi-supervised learning) then turned/engineered into systems that work in practice using backprop

Very successful/state-of-the-art at object recognition, speech/audio and games, language translation, and some video understanding

Alpha Zero

Deep learning RL system that in 24 hours of training achieved superhuman performance at chess playing, Go and shogi



Next Steps

A few years ago this would have been unthinkable

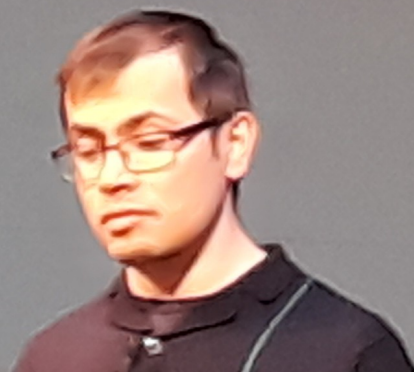
Very general, but *not* claiming this has solved AI!

Lots of huge challenges remain: learning the model

Dealing with partial information and uncertainty

Efficiency, one-shot learning, abstract knowledge

Human efficiency (but evolution, transfer, apprenticeship)



Another important next step...

Making sense of AlphaZero

Long-term positional sacrifices

Grand-masters now watching how it plays

The need for Knowledge Extraction

(<http://www.staff.city.ac.uk/~aag/talks/GarcezDSI.pdf>)

Correctness / soundness

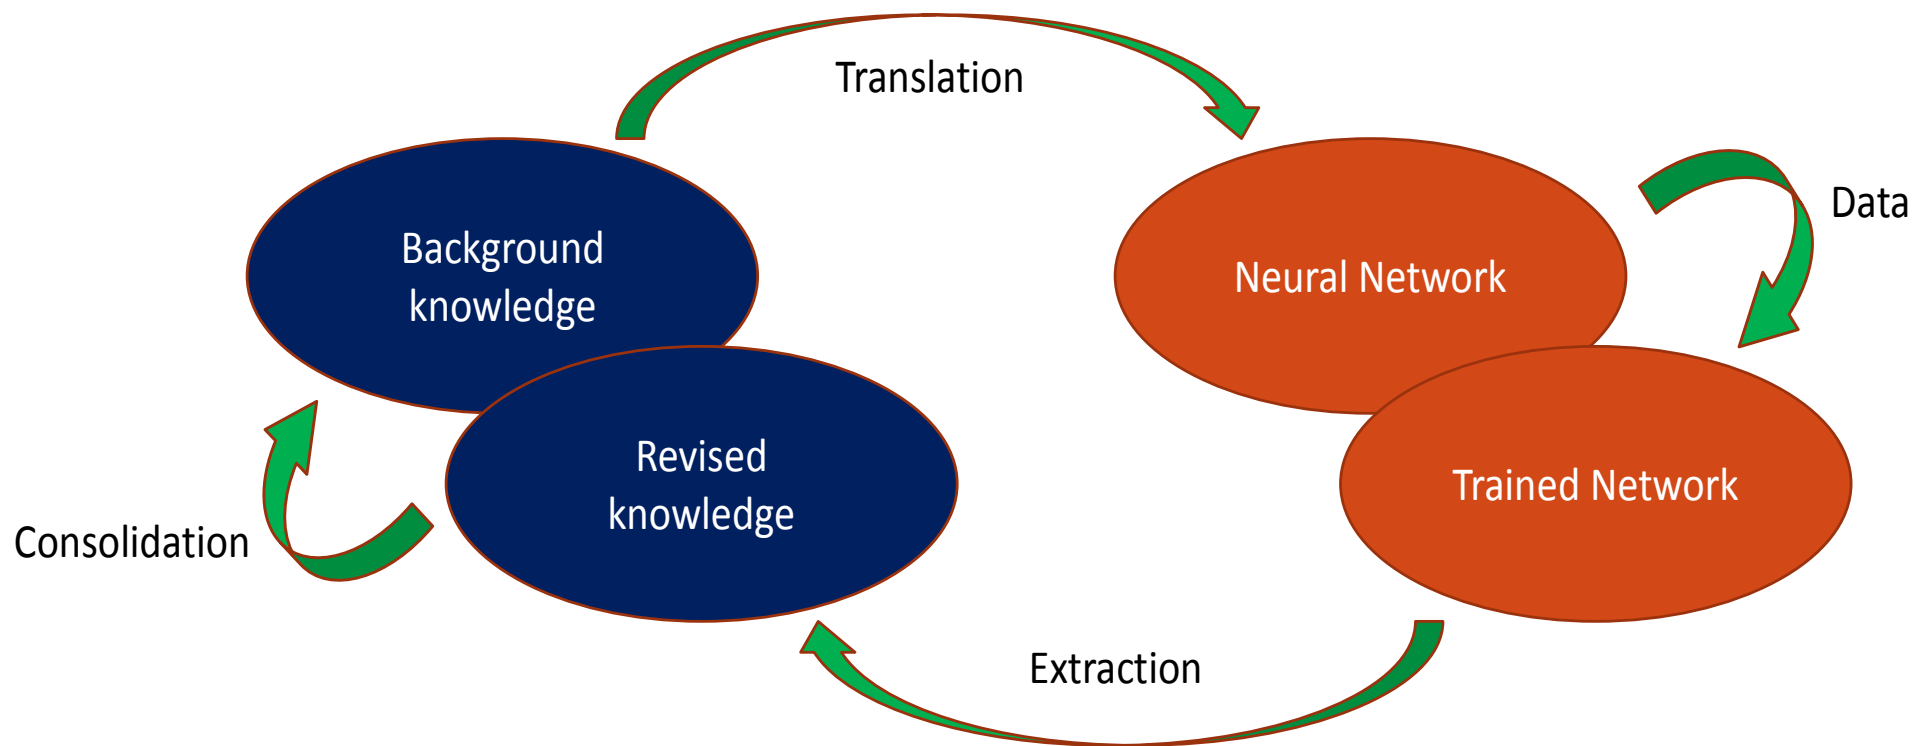
Proof history (goal-directed reasoning)

Levels of abstraction (modularity)

Transfer learning (analogy)

System maintenance/improvement

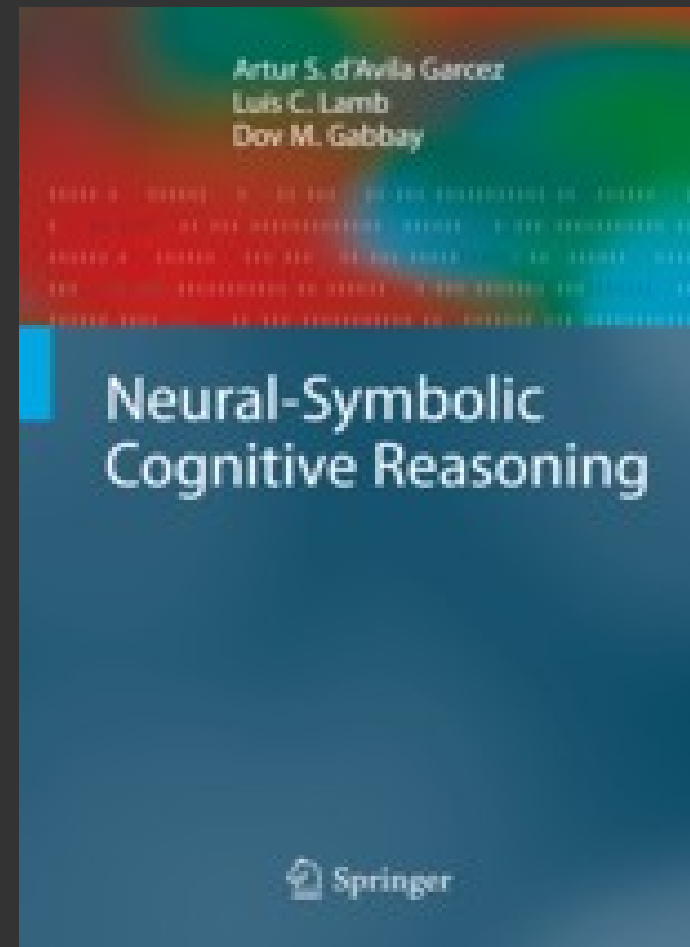
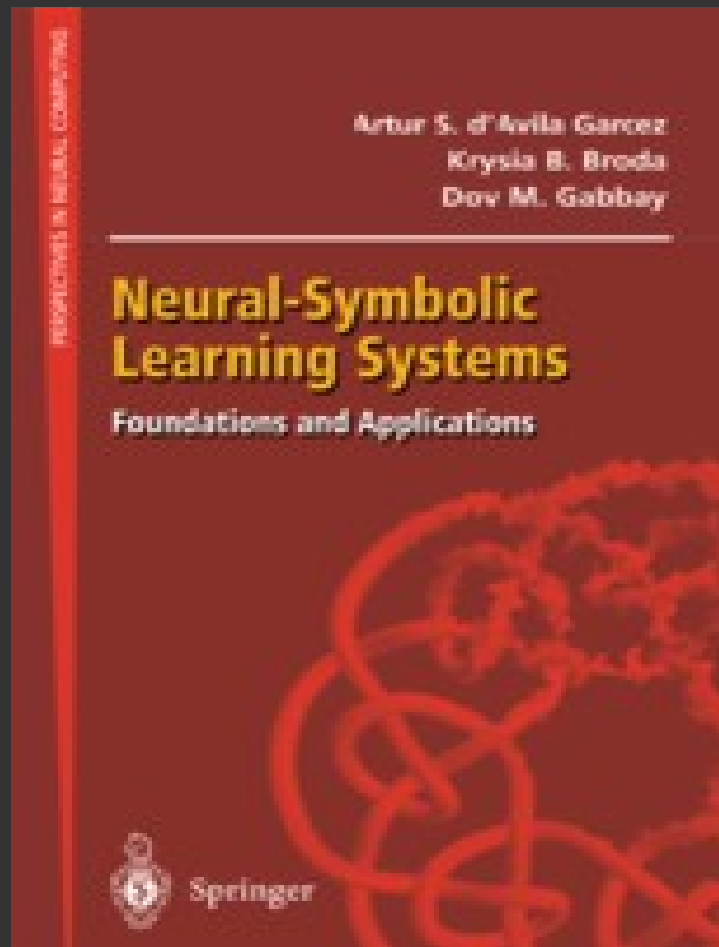
Neural-Symbolic Learning Cycle



Neural-Symbolic Computing

- Neural-symbolic computing = neural networks + logical structure (**compositionality**)
- Neural networks provide the machinery for effective learning and computation
- Perception alone is insufficient: AI needs reasoning, planning, explanation and transfer
- Rich knowledge representation models: nonmonotonic, relational, recursive, temporal reasoning under uncertainty...

For more information...



Challenges for Deep Learning AI

G. Marcus, Deep Learning: A Critical Appraisal,
<https://arxiv.org/abs/1801.00631>, 2018

Identifying challenges is important...

It should not encourage:

Those who don't like neural nets to continue to
dismiss it

Those who don't like logic to continue to ignore it

Let's not go back to 1991!

Challenge 1

Deep learning thus far is data hungry;

BUT... adding knowledge to neural nets
allows for more effective learning when
fewer data is available

Artur d'Avila Garcez and Gerson Zaverucha, The Connectionist Inductive Learning and Logic Programming System, Applied Intelligence 11(1):59-77, 1999

Challenge 2

Deep learning thus far has limited capacity for transfer;

BUT... knowledge extraction enables transfer

Son Tran and Artur d'Avila Garcez. Deep Logic Networks: Inserting and Extracting Knowledge From Deep Belief Networks, IEEE TNNLS, Feb 2018

Challenge 3

Deep learning thus far has no natural way to deal with hierarchical structure?

Artur d'Avila Garcez, Luis C. Lamb and Dov M. Gabbay. Connectionist Modal Logic: Representing modalities in neural networks, TCS:371(1-2):34-53, 2007

Artur d'Avila Garcez and Dov M. Gabbay. Fibring Neural Networks, AAAI 2004

Challenge 4

Deep learning thus far has struggled with open-ended inference;

Various forms of reasoning have been shown to be implementable in neural nets

Artur d'Avila Garcez, Luis C. Lamb and Dov M. Gabbay.
Neural-Symbolic Cognitive Reasoning, Springer, 2009

Although goal-directed reasoning is still a challenge

Challenge 5

Deep learning thus far is not sufficiently transparent;

BUT... knowledge extraction with modularity and reasoning can solve this problem

Artur d'Avila Garcez, K. Broda and D. M. Gabbay. Symbolic Knowledge Extraction from Trained Neural Networks: A Sound Approach. *AIJ*, 125(1-2):153-205, 2001.

Challenge 6

Deep learning thus far has not been well integrated with prior knowledge?

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

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

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

Etc. etc.

Challenge 7

Deep learning thus far cannot inherently distinguish causation from correlation...

Is this distinction really important?

A. d'Avila Garcez, L. C. Lamb and D. M. Gabbay. A Connectionist Model for Constructive Modal Reasoning. In Proc. NIPS 18, Vancouver, 2005

Challenge 8

Deep learning presumes a largely stable world, in ways that may be problematic;

Doesn't the whole of ML assume this!?

Commonsense (nonmonotonic) reasoning needed on top of deep learning/ML
c.f. CILP system, “rules with exceptions”

Artur d'Avila Garcez, K. Broda and Dov M. Gabbay,
Neural-Symbolic Learning Systems, Springer, 2002

Challenge 9

Deep learning thus far works well as an approximation, but its answers often cannot be fully trusted;

c.f. recent work on verification and system monitoring using neural networks, as well as knowledge extraction (again)

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

Challenge 10

Deep learning thus far is difficult to engineer with;

The problems that AI seeks to solve are generally difficult (c.f. CYC project)

Two conclusions

- "Deep learning is... not a universal solvent, but simply one tool among many, a power screwdriver in a world in which we also need hammers, wrenches..." Gary Marcus
- Deep learning is universal as long as it is combined with a logical interpretation. The question is how to make it explainable, to make it generally applicable, compositional, and to allow it to **reason about what has been learned**, as put forward by Les Valiant.

Logic Tensor Networks

- Suppose your system has learned from millions of images how to recognise horses extremely well
- Now it needs to recognise zebras; which of the following is better?
 1. Collect millions of images with zebras to train together with the other million of images again, or
 2. Tell the system that a zebra is a horse with stripes (may need to learn “stripes” from millions of images)

Logic Tensor Networks (<https://arxiv.org/abs/1705.08968>, IJCAI 2017) does the latter; it uses full first-order logic as a constraint on the learning of a deep network...

DeJuergenating

Lots of relevant work by others...

A Neural-Symbolic Approach to Natural Language Tasks.
Huang, Smolensky, He, Deng, Wu, 2017

Jude Shavlik's team Knowledge-based Artificial Neural Networks and the TREPAN method for extracting decision trees from neural nets

Pascal Hitzler's and co. work on the representation of first-order logic by neural nets, and more recently on how to make CNNs more explainable

End-to-end differentiable ILP (Evans and Grefenstette), lifted relational neural nets (Zelezny et al)

- Work on Neural Tensor Networks, including learning and reasoning by Bowman and Socher et al.
- Work on learning to reason by Dan Roth and others, and seminal work on robust logics and a neuroidal architecture for cognitive computation by Les Valiant
- Gori et al. Graph Neural Nets, leading to Interaction Networks (Battaglia et al.) applied to objects, relations and naive (intuitive) physics
- Bordes et al. Memory Networks and their application to question answering (agents who ask questions!)
- Distilling a Neural Network Into a Soft Decision Tree, by Nicholas Frosst and Geoffrey Hinton, 2017
- Garnelo, Arulkumaran and Shanahan. Towards Deep Symbolic Reinforcement Learning, 2016

I'd say there are **three** main challenges:

First-order logic knowledge extraction from very large networks (sound and efficient, explaining entire model)

Goal-directed commonsense reasoning about what has been learned by a deep network

Human-network communication; an agent who asks questions and checks her understanding!

Good news: multiple PhD theses in each of the above with the possibility of being area defining...

Specific challenging tasks for human-like machines:

Based on: Building Machines That Learn and Think Like People. Brenden M. Lake et. al.. BBS, 2016
<https://arxiv.org/pdf/1604.00289.pdf>

PhD 1: Learning of a novel handwritten character a la MNIST but measured w.r.t. human-like, not human-level performance

PhD 2: Games where agent can refer to objects and their relations to perform goal-directed reasoning, e.g. agent creates sub-goals to solve first

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!