Neuro Symbolic Taxonomies

Kautz's Taxonomy

1. Symbolic Neuro symbolic
2. Symbolic [Neuro]
3. Neuro | Symbolic
4. Neuro:Symbolic → Neuro
5. Neuro_{Symbolic}
6. Neuro [Symbolic]

1. Symbolic Neuro symbolic

- Current deep learning approach for NLP/NLU

- Symbolic input converted into vector representations
- Neural network reasoning
- Symbolic output

- Generally not topical in this course

2. Symbolic [Neuro]

- Symbolic system with a neural sub-routine

- Examples:
  - AlphaGo (MC search with a neural network used as a heuristic function
  - Autonomous vehicles

3. Neuro | Symbolic

- Neural network converts non-symbolic input (e.g., pixels) into symbolic input that is manipulated by a symbolic reasoning systems
- Differs in that the neural component is a “co-routine” instead of a sub-routine

- Example:
  - NeurASP

4. Neuro:Symbolic → Neuro

- Uses architecture of #1 but with symbolic rules to guide the neural training process
- Note: this does not provide a symbolic derivation of the result
- Kautz's taxonomy does not specify if such frameworks should have guarantees, so LNN's fall into this category as well (at least according to Garcez and Lamb)
- Examples vary widely:
  - Symbolic information is normally compiled into the training data (not covered in this course)
  - LNN's
  - Differentiable ILP

5. Neuro_{Symbolic}

- Direct encoding of logical statements into neural structures (through the use of embeddings), they act as a form of regularization
- Is explainability sacrificed?

- Examples:
  - LTN
  - Deep ontological networks

6. Neuro [Symbolic]

- Inspired by System 1 / System 2 (fast / slow) thinking from cognitive science
- System 1 is neural, system 2 is symbolic
- System 1 conducts initial reasoning, when it puts attention on a certain part of the problem, it triggers System 2 which performs a combinatorial search

Garcez and Lamb's Options for NSR

- Option 1: Symbols translated into the neural network, reasoning performed within the network
- Option 2: Network interreacts with a symbolic reasoning system
- Option 3 (non-NSR): Expert knowledge is made available, rather than learned from data and sound reasoning is performed

Shortcomings of the taxonomy w.r.t. results over the past 5 years?

- Should we really consider the use of symbolic knowledge to create more training examples a true NSR approach? And if so, should they be in the same taxonomic category of LNN's?
- Should LNN's be in type 4 or type 5? If in type 5, shouldn't there be a divide between systems that are explainable and those that are not?
- Should there be a division between systems that do / do not support inductive rule learning?
- Can reasoning about combinatorial cases be performed without necessitating a type 6 system?
- Are type 6 systems truly superior to types 3, 4, and 5? Are type 6 systems just a special case of type 3 systems by definition?

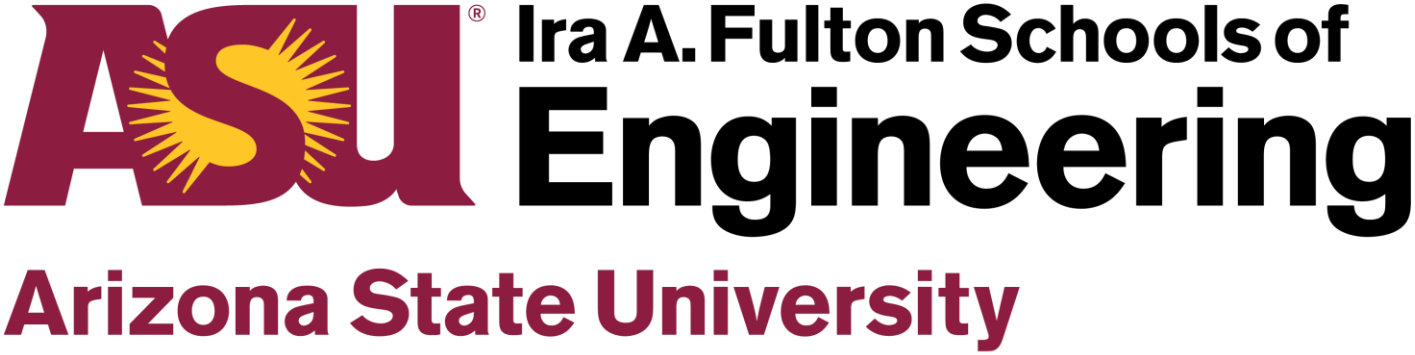
Ira A. Fulton Schools of
Engineering
Arizona State University