Logics and Languages for Representation Learning

Vardifest 2022 (!!!)

Hector Geffner
ICREA & Universitat Pompeu Fabra
Barcelona, Spain

Wallenberg Guest Professor
Linköping University, Sweden

Collaborations with Blai Bonet and RLeap Team

Model-Free Learners and Model-Based Solvers in AI

$$\textit{Input } x \Longrightarrow \boxed{\textsc{Function } f} \Longrightarrow \textit{Output } f(x)$$

- **Learners** require **experience over related problems** x but then fast
  - ▷ They compute function f from training, then apply it
- **Solvers** deal with **completely new problems** x but need **models/thinking**
  - ▷ They compute $f(x)$ for each input x from scratch

Learners and Solvers: System 1 and System 2?

Dual process accounts of the human mind assume two processes (D. Kahneman: Thinking, Fast and Slow, 2011; K. Stanovich: The Robot's Rebellion, 2005)

System 1 (Intuitive Mind)	**System 2** (Analytical Mind)
fast	slow
associative	deliberative
unconscious	conscious
effortless	effortful
parallel	serial
specialized	general
...	...
Learners?	Solvers?

Key Challenge in AI

- General **two-way integration** of System 1 and System 2 inference in AI systems
  - ▷ **Learn representations** that support reasoning/planning, general/reusable

- **Yoshua Bengio**'s challenges reflected in title of his IJCAI 2021 talk:
  - ▷ System 2 Deep Learning: Higher-level cognition, agency, out-of-distribution generalization and causality

- **Yann LeCun**'s three challenges, AAAI 2020:
  - ▷ AI must learn to represent the world
  - ▷ AI must think and plan in ways compatible with gradient-based learning
  - ▷ AI must learn hierarchical representation of action plans

Bottom-Up vs. Top-Down Representation Learning

- **Bottom-up approach** *(most common)*
  - ▷ Representations emerge from **architecture**, loss function, and "right" bias
- **Top-down approach** *(logic)*
  - ▷ Representations learned over **language** with "right" syntax and semantics
  - ▷ Meaningful learning bias, transparency, reasoning, **what** vs. **how**

Top-down approach in line with **"traditional AI": just learn from data the representations that have traditionally been handcrafted**

Related but different than **neuro-symbolic AI** where representation languages used mainly to encode **background knowledge**

Example: Learning representations over FO-STRIPS

- **Planning problems** P specified as **instances** $P = \langle D, I \rangle$ of **general domain** D
  - ▷ **Domain** D specified in terms of **action schemas** and **predicates**
  - ▷ **Instance** is $P = \langle D, I \rangle$ where I details **objects**, **init**, **goal**

$move(c, c')$
Preconds: $atRobot(c)$, $adjacent(c, c')$
Effects: $atRobot(c')$, $\neg atRobot(c)$

$pick(o, c)$:
Preconds: $atRobot(c)$, $at(o, c)$, $emptyhand$
Effects: $held(o)$, $\neg at(o, c)$, $\neg emptyhand$

$drop(o, c)$:
Preconds: $atRobot(c)$, $held(o)$
Effects: $at(o, c)$, $\neg held(o)$, $emptyhand$

- Can **symbolic, first-order representations** like this be learned?

Example: Learning FO-STRIPS from State Graph

Input: State graph G of agent in $1{\times}3$ grid, moving/picking/dropping 2 pkgs

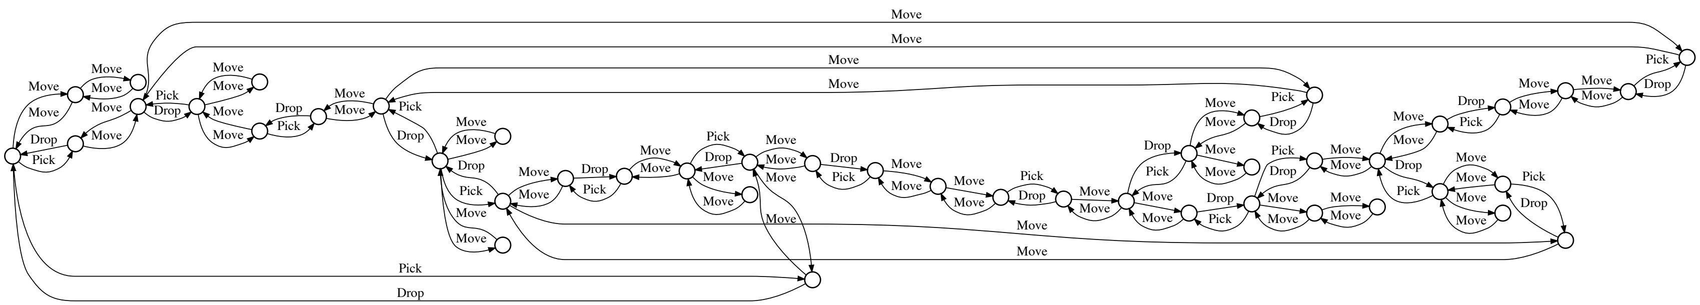

Output: **Simplest** STRIPS representation $P = \langle D, I \rangle$ that **generates** G

```
Move(?to,?from):
  Pre: neq(?to,?from), p5(?to,?from)
  Pre:  p2(?from), -p2(?to)
  Eff: -p2(?from),  p2(?to)

Pick(?p,?x):
  Pre:  p2(?x),  p1, -p3(?p),  p4(?p,?x)
  Eff:          -p1,  p3(?p), -p4(?p,?x)

Drop(?p,?x):
  Pre:  p2(?x), -p1,  p3(?p), -p4(?p,?x)
  Eff:           p1, -p3(?p),  p4(?p,?x)
```

Interpretation of learned predicates:

– p_1: gripper empty

– $p_2(x)$: agent at cell x,

– $p_3(p)$: agent holds pkg p,

– $p_4(p, x)$: pkg p in cell x

– $p_5(x, y)$: cell x adj to y

- Domain D learned from 1×3 grid, 2 pkgs, correct for **any** grid, **any** $\#$ of pkgs

Summary: The unusual scope of logic in AI/ML

- **Learning representations** that support reasoning/planning is central in AI/ML
- **Logic** has **key role** to play: representations learned over **languages** with **known structure** and **semantics**
- Examples from own recent work:
  - ▷ Learning **FO-STRIPS** representations for planning from state graphs
  - ▷ Learning **general policies** using **C2** features or **Graph Neural Nets**
  - ▷ Learning **sketches** for decomposing in subproblems of **bounded width**

AI and Social Impact

- **System 2** not only necessary for AI systems; essential for people and **societies**
- AI far from human-level intelligence, yet it can be used for **good** or **ill**
- **Ethical committees** and **AI principles** good but not sufficient (**Moshe**)
- **Markets and politics** play our **System 1**, focused on the **bottom line**
- If we want **good AI**, we need a **good and decent society** . . .

"Need artificial intelligence for social good because natural intelligence is busy in other pursuits"

:-)