

CPSC 522 — Artificial Intelligence 2

David Poole

Term 1 — Tuesdays and Thursdays 11:00-12:20

Prerequisite: previous AI and/or ML course(s)

Kahneman's two systems

[Thinking Fast and Slow, 2011]

- **System 2**, the higher level, is slow, deliberate, serial, open to introspection, and based on reasoning. “Cognition / Reasoning”
- **System 1**, the lower level, is fast, automatic, parallel, intuitive, instinctive, emotional, and not open to introspection. “perception”.

Kahneman's two systems

[Thinking Fast and Slow, 2011]

- **System 2**, the higher level, is slow, deliberate, serial, open to introspection, and based on reasoning. “Cognition / Reasoning”
- **System 1**, the lower level, is fast, automatic, parallel, intuitive, instinctive, emotional, and not open to introspection. “perception”.

“Pretty much anything that a normal person can do in < 1 sec, we can now automate with AI.” Andrew Ng, @AndrewYNg, Twitter, 18 Oct 2016.

Kahneman's two systems

[Thinking Fast and Slow, 2011]

- **System 2**, the higher level, is slow, deliberate, serial, open to introspection, and based on reasoning. “Cognition / Reasoning” **The Next Frontier**
- **System 1**, the lower level, is fast, automatic, parallel, intuitive, instinctive, emotional, and not open to introspection. “perception”. **✓ SOLVED!**


“Pretty much anything that a normal person can do in < 1 sec, we can now automate with AI.” Andrew Ng, @AndrewYNg, Twitter, 18 Oct 2016.

What is AI?

Artificial intelligence is the synthesis and analysis of computational agents that act intelligently.

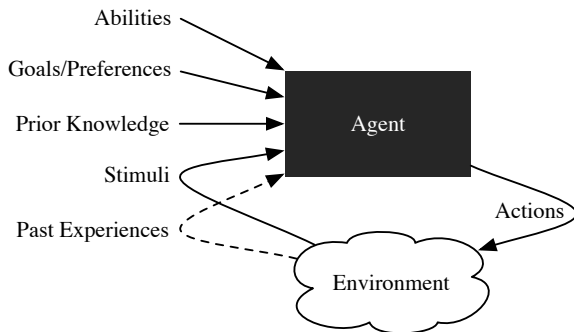
What is AI?

Artificial intelligence is the synthesis and analysis of computational agents that act intelligently.



What is AI?

Artificial intelligence is the synthesis and analysis of computational agents that act intelligently.



- Prior + observations (stimuli) \rightarrow Belief
- Belief + preferences + abilities \rightarrow Action

Topics covered

- Prior + observations \longrightarrow Belief
 - Directed and undirected graphical models
 - Representations, inference, learning
 - Causality and actions
 - Time
 - Relational models
 - Probabilistic relational learning
 - Relational embeddings
 - Ontologies
 - Identity, number, existence uncertainty
- Belief + preferences \longrightarrow Action
 - Preferences: utility and reward
 - Decision-theoretic planning
 - Reinforcement learning
 - Multiagent systems

— and **how they are all interconnected**

Clicker question

What is the real world made of?

- A Features or random variables
- B Words, pixels, phonemes . . .
- C Entities and events (e.g., people, computers, courses, presentations)
- D Huh? There is a real world?

Role of preferences

In designing a system to recommend or carry out a medical procedure, whose preferences should be used?

- the doctor
- whoever is paying
- the individual patient
- a generic patient in the database
- society
- academic ethics researchers

Role of preferences

In designing a system to recommend or carry out a medical procedure, whose preferences should be used?

- the doctor
- whoever is paying
- the individual patient
- a generic patient in the database
- society
- academic ethics researchers

How are these determined?

Social Impact: Moral Machines

“ ...self-driving cars ... There are gonna be a bunch of choices that you have to make ... It's a moral decision, and who's setting up those rules?”

– Barack Obama, 2016

We can embrace the challenges of machine ethics as a unique opportunity to decide, as a community, what we believe to be right or wrong; and to make sure that machines, unlike humans, unerringly follow these moral preferences.

– Awad et al., [2018]

<https://www.cs.ubc.ca/~poole/cs522/2023W1/>