CPSC 322 — Introduction to Artificial Intelligence

- Professor: David Poole
- URL: `http://www.cs.ubc.ca/~poole/cs322/2020/`
  Assignment 0 is due on Tuesday!

CPSC 322 — Introduction to Artificial Intelligence

- Professor: David Poole
- URL: `http://www.cs.ubc.ca/~poole/cs322/2020/`
  Assignment 0 is due on Tuesday!
- Textbook: `http://artint.info` (full text online)

CPSC 322 — Introduction to Artificial Intelligence

- Professor: David Poole
- URL: `http://www.cs.ubc.ca/~poole/cs322/2020/`
  Assignment 0 is due on Tuesday!
- Textbook: `http://artint.info` (full text online)
- Pre-lectures, assignments, etc. on web site.
  Discussions, grades, zoom recordings on Canvas.

CPSC 322 — Introduction to Artificial Intelligence

- Professor: David Poole
- URL: `http://www.cs.ubc.ca/~poole/cs322/2020/`
  Assignment 0 is due on Tuesday!
- Textbook: `http://artint.info` (full text online)
- Pre-lectures, assignments, etc. on web site.
  Discussions, grades, zoom recordings on Canvas.
- Every student should expect to struggle, but can succeed!
  Students learn by doing.
- Ask questions!

CPSC 322 — Introduction to Artificial Intelligence

- Professor: David Poole
- URL: `http://www.cs.ubc.ca/~poole/cs322/2020/`
  Assignment 0 is due on Tuesday!
- Textbook: `http://artint.info` (full text online)
- Pre-lectures, assignments, etc. on web site.
  Discussions, grades, zoom recordings on Canvas.
- Every student should expect to struggle, but can succeed!
  Students learn by doing.
- Ask questions!
- Marks:
  - 26%: assignments.
  - $2 \times 20\%$: midterm.
  - 30%: final exam
  - 4% participation; in-class (2%), postings (2%).

Clicker Questions (no correct answer)

Do you know Python?

A Yes

B No

C Some

Clicker Questions (no correct answer)

Do you know Python?

- A Yes
- B No
- C Some

What is the output of the Python code:

```
>>> [(x+7, x-1) for x in range(10) if x < 3]
```

- A This is obvious
- B I could probably work it out by myself
- C I have no idea, I'd appreciate a tutorial

What is Artificial Intelligence?

- Artificial Intelligence is the synthesis and analysis of computational agents that act intelligently.
- An agent is something that acts in an environment.
- An agent acts intelligently if:
  - its actions are appropriate for its goals and circumstances
  - it is flexible to changing environments and goals
  - it learns from experience
  - it makes appropriate choices given perceptual and computational limitations

Clicker Questions (no correct answer)

A Yes

B No

C Umm, I'm not sure.... what is an agent again?

- Can a book inspire or explain or enlighten?

Clicker Questions (no correct answer)

A Yes
B No
C Umm, I'm not sure.... what is an agent again?

- Can a book inspire or explain or enlighten?
- Can a book be an agent?

Clicker Questions (no correct answer)

A Yes

B No

C Umm, I'm not sure.... what is an agent again?

- Can a book inspire or explain or enlighten?
- Can a book be an agent?
- Can a paragraph be an agent?

Clicker Questions (no correct answer)

- Can a book inspire or explain or enlighten?
- Can a book be an agent?
- Can a paragraph be an agent?
- Can a sentence be an agent?

A Yes

B No

C Umm, I'm not sure.... what is an agent again?

Clicker Questions (no correct answer)

A Yes

B No

C Umm, I'm not sure.... what is an agent again?

- Can a book inspire or explain or enlighten?
- Can a book be an agent?
- Can a paragraph be an agent?
- Can a sentence be an agent?
- Can a word be an agent?

Clicker Questions (no correct answer)

A Yes
B No
C Umm, I'm not sure.... what is an agent again?

- Can a book inspire or explain or enlighten?
- Can a book be an agent?
- Can a paragraph be an agent?
- Can a sentence be an agent?
- Can a word be an agent?
- Can a letter of the alphabet be an agent?

Agents acting in an environment

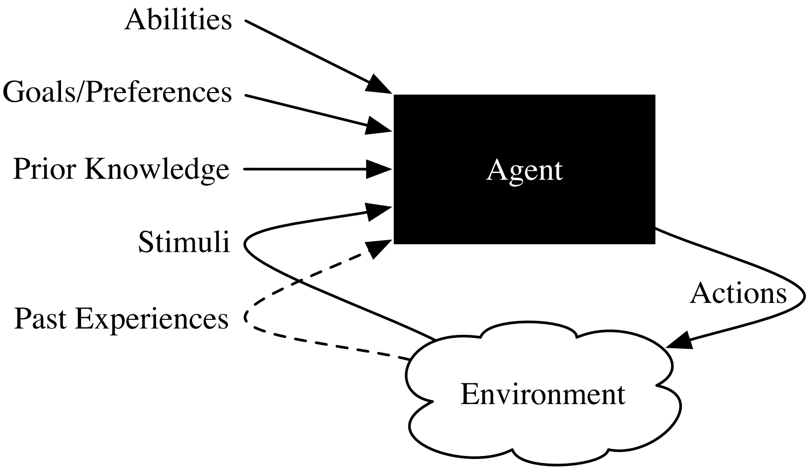

Example agent: autonomous car

- abilities:

Example agent: autonomous car

- abilities: steer, accelerate, brake

Example agent: autonomous car

- abilities: steer, accelerate, brake
- goals:

Example agent: autonomous car

- abilities: steer, accelerate, brake
- goals: safety, get to destination, timeliness,...

Example agent: autonomous car

- abilities: steer, accelerate, brake
- goals: safety, get to destination, timeliness,...
- prior knowledge:

Example agent: autonomous car

- abilities: steer, accelerate, brake
- goals: safety, get to destination, timeliness,...
- prior knowledge: what signs mean, what to stop for

Example agent: autonomous car

- abilities: steer, accelerate, brake
- goals: safety, get to destination, timeliness,...
- prior knowledge: what signs mean, what to stop for
- stimuli:

Example agent: autonomous car

- abilities: steer, accelerate, brake
- goals: safety, get to destination, timeliness,...
- prior knowledge: what signs mean, what to stop for
- stimuli: vision, laser, GPS...

Example agent: autonomous car

- abilities: steer, accelerate, brake
- goals: safety, get to destination, timeliness,...
- prior knowledge: what signs mean, what to stop for
- stimuli: vision, laser, GPS...
- past experiences:

Example agent: autonomous car

- abilities: steer, accelerate, brake
- goals: safety, get to destination, timeliness,. . .
- prior knowledge: what signs mean, what to stop for
- stimuli: vision, laser, GPS. . .
- past experiences: streetmaps, how breaking, steering affects direction..

Example agent: thermostat for heater

- abilities:

Example agent: thermostat for heater

- abilities: turn heater on or off

Example agent: thermostat for heater

- abilities: turn heater on or off
- goals:

Example agent: thermostat for heater

- abilities: turn heater on or off
- goals: conformable temperature, save fuel, save money

Example agent: thermostat for heater

- abilities: turn heater on or off
- goals: conformable temperature, save fuel, save money
- prior knowledge:

Example agent: thermostat for heater

- abilities: turn heater on or off
- goals: conformable temperature, save fuel, save money
- prior knowledge: there is a delay, 24 hour cycle, weekends

Example agent: thermostat for heater

- abilities: turn heater on or off
- goals: conformable temperature, save fuel, save money
- prior knowledge: there is a delay, 24 hour cycle, weekends
- stimuli:

Example agent: thermostat for heater

- abilities: turn heater on or off
- goals: conformable temperature, save fuel, save money
- prior knowledge: there is a delay, 24 hour cycle, weekends
- stimuli: temperature, set temperature, who is home, outside temperature

Example agent: thermostat for heater

- abilities: turn heater on or off
- goals: conformable temperature, save fuel, save money
- prior knowledge: there is a delay, 24 hour cycle, weekends
- stimuli: temperature, set temperature, who is home, outside temperature
- past experiences:

Example agent: thermostat for heater

- abilities: turn heater on or off
- goals: conformable temperature, save fuel, save money
- prior knowledge: there is a delay, 24 hour cycle, weekends
- stimuli: temperature, set temperature, who is home, outside temperature
- past experiences: delay between heater being turned off and the room stopping warming, when people come and go, who likes what temperature

Other Agents — small groups

- Apple Inc.
- user interface
- smart home
- medical doctor (GP)
- surgeon
- teacher

  . . .

- abilities:
- goals:
- prior knowledge:
- stimuli:
- past experiences:

Dimensions

- Research proceeds by making simplifying assumptions, and gradually reducing them.
- Each simplifying assumption gives a dimension of complexity
  - multiple values in a dimension: from simple to complex
  - simplifying assumptions can be relaxed in various combinations

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

Modularity

- Model at one level of abstraction: flat
- Model with interacting modules that can be understood separately: modular
- Model with modules that are (recursively) decomposed into modules: hierarchical

Modularity

- Model at one level of abstraction: flat
- Model with interacting modules that can be understood separately: modular
- Model with modules that are (recursively) decomposed into modules: hierarchical
- Example: Planning a trip from here to a see the Mona Lisa in Paris.

Modularity

- Model at one level of abstraction: flat
- Model with interacting modules that can be understood separately: modular
- Model with modules that are (recursively) decomposed into modules: hierarchical
- Example: Planning a trip from here to a see the Mona Lisa in Paris.
- Flat representations are adequate for simple systems.
- Complex biological systems, computer systems, organizations are all hierarchical
- A flat description is either continuous or discrete. Hierarchical reasoning is often a hybrid of continuous and discrete.

By a hierarchic system, or hierarchy, I mean a system that is composed of interrelated subsystems, each of the latter being in turn hierarchic in structure until we reach some lowest level of elementary subsystem. In most systems of nature it is somewhat arbitrary as to where we leave off the partitioning and what subsystems we take as elementary. Physics makes much use of the concept of "elementary particle," although the particles have a disconcerting tendency not to remain elementary very long . . .
Empirically a large proportion of the complex systems we observe in nature exhibit hierarchic structure. On theoretical grounds we would expect complex systems to be hierarchies in a world in which complexity had to evolve from simplicity.

– Herbert A. Simon, The Sciences of the Artificial, 1996

Planning horizon

...how far the agent looks into the future when deciding what to do.

- Static: world does not change

Planning horizon

...how far the agent looks into the future when deciding what to do.

- Static: world does not change
- Finite stage: agent reasons about a fixed finite number of time steps

Planning horizon

...how far the agent looks into the future when deciding what to do.

- Static: world does not change
- Finite stage: agent reasons about a fixed finite number of time steps
- Indefinite stage: agent reasons about a finite, but not predetermined, number of time steps

Planning horizon

...how far the agent looks into the future when deciding what to do.

- Static: world does not change
- Finite stage: agent reasons about a fixed finite number of time steps
- Indefinite stage: agent reasons about a finite, but not predetermined, number of time steps
- Infinite stage: the agent plans for going on forever (process oriented)

Representation

Much of modern AI is about finding compact representations and exploiting the compactness for computational gains.
A agent can reason in terms of:

- Explicit states — a state is one way the world could be

Representation

Much of modern AI is about finding compact representations and exploiting the compactness for computational gains.

A agent can reason in terms of:

- Explicit states — a state is one way the world could be
- Features or propositions.
  - States can be described using features.
  - 30 binary features can represent $2^{30} = 1,073,741,824$ states.

Representation

Much of modern AI is about finding compact representations and exploiting the compactness for computational gains.

A agent can reason in terms of:

- Explicit states — a state is one way the world could be
- Features or propositions.
  - States can be described using features.
  - 30 binary features can represent $2^{30} = 1,073,741,824$ states.
- Individuals and relations
  - There is a feature for each relationship on each tuple of individuals.
  - Often an agent can reason without knowing the individuals or when there are infinitely many individuals.

Computational Limits

- Perfect rationality: the agent can determine the best course of action, without taking into account its limited computational resources.

Computational Limits

- Perfect rationality: the agent can determine the best course of action, without taking into account its limited computational resources.
- Bounded rationality: the agent must make good decisions based on its perceptual, computational and memory limitations.

Learning from experience

Whether the model is fully specified a priori:

- Knowledge is given.

Learning from experience

Whether the model is fully specified a priori:

- Knowledge is given.
- Knowledge is learned from data or past experience.

Learning from experience

Whether the model is fully specified a priori:

- Knowledge is given.
- Knowledge is learned from data or past experience.

. . . always some mix of prior (innate, programmed) knowledge and learning (nature vs nurture).

Learning from experience

Whether the model is fully specified a priori:

- Knowledge is given.
- Knowledge is learned from data or past experience.

. . . always some mix of prior (innate, programmed) knowledge and learning (nature vs nurture).

- Learning is impossible without prior knowledge (bias).

Uncertainty

There are two dimensions for uncertainty. In each dimension an agent can have

- No uncertainty: the agent knows what is true
- Disjunctive uncertainty: there is a set of states that are possible
- Probabilistic uncertainty: a probability distribution over the worlds.

Why Probability?

- Agents need to act even if they are uncertain.
- Predictions are needed to decide what to do:
  - definitive predictions: you will be run over tomorrow
  - disjunctions: be careful or you will be run over
  - point probabilities: probability you will be run over tomorrow is 0.002 if you are careful and 0.05 if you are not careful

Why Probability?

- Agents need to act even if they are uncertain.
- Predictions are needed to decide what to do:
  - definitive predictions: you will be run over tomorrow
  - disjunctions: be careful or you will be run over
  - point probabilities: probability you will be run over tomorrow is 0.002 if you are careful and 0.05 if you are not careful
- Acting is gambling: agents who don't use probabilities will lose to those who do.
- Probabilities can be learned from data and prior knowledge.

Sensing Uncertainty

Whether an agent can determine the state from its stimuli:

- Fully-observable: the agent can observe the state of the world.

Sensing Uncertainty

Whether an agent can determine the state from its stimuli:

- Fully-observable: the agent can observe the state of the world.
- Partially-observable: there can be a number states that are possible given the agent's stimuli.

Effect Uncertainty

If an agent knew the initial state and its action, could it predict the resulting state?

Effect Uncertainty

If an agent knew the initial state and its action, could it predict the resulting state?

The dynamics can be:

- Deterministic: the resulting state is determined from the action and the state

Effect Uncertainty

If an agent knew the initial state and its action, could it predict the resulting state?
The dynamics can be:

- Deterministic: the resulting state is determined from the action and the state
- Stochastic: there is uncertainty about the resulting state.

Preference

What does the agent try to achieve?

- achievement goal is a goal to achieve. This can be a complex logical formula.

Preference

What does the agent try to achieve?

- achievement goal is a goal to achieve. This can be a complex logical formula.
- complex preferences may involve tradeoffs between various desiderata, perhaps at different times.

Preference

What does the agent try to achieve?

- achievement goal is a goal to achieve. This can be a complex logical formula.
- complex preferences may involve tradeoffs between various desiderata, perhaps at different times.
  - ordinal only the order matters

Preference

What does the agent try to achieve?

- achievement goal is a goal to achieve. This can be a complex logical formula.
- complex preferences may involve tradeoffs between various desiderata, perhaps at different times.
  - ordinal only the order matters
  - cardinal absolute values also matter

Preference

What does the agent try to achieve?

- achievement goal is a goal to achieve. This can be a complex logical formula.
- complex preferences may involve tradeoffs between various desiderata, perhaps at different times.
  - ordinal only the order matters
  - cardinal absolute values also matter

Examples: coffee delivery robot, medical doctor

Number of agents

Are there multiple reasoning agents that need to be taken into account?

- Single agent reasoning: any other agents are part of the environment.
- Multiple agent reasoning: an agent reasons strategically about the reasoning of other agents.

Agents can have their own goals: cooperative, competitive, or goals can be independent of each other

Interaction

When does the agent reason to determine what to do?

- reason offline: before acting
- reason online: while interacting with environment

Dimensions of Complexity

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

State-space Search

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

Deterministic Planning

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

Decision Networks

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

Markov Decision Processes (MDPs)

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

Decision-theoretic Planning

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

Reinforcement Learning

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

Classical Game Theory

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

Humans

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

The Dimensions Interact in Complex Ways

- Partial observability makes multi-agent and indefinite horizon reasoning more complex
- Modularity interacts with uncertainty and succinctness: some levels may be fully observable, some may be partially observable
- Three values of dimensions promise to make reasoning simpler for the agent:
  - Hierarchical reasoning
  - Individuals and relations
  - Bounded rationality