Welcome to

CSC384 Introduction to Artificial Intelligence Winter 2020

Bahar Aameri & Sonya Allin, University of Toronto, Winter 2020

CSC384: Intro to Artificial Intelligence

Instructor: Bahar Aameri
Office Hour: Tentatively, M 11 am

Instructor: Sonya Allin
Office Hour: Tentatively, F 4 pm

Email for both instructors:
csc384-2020-01@cs.toronto.edu

This course is being co-taught. Lectures & evaluation will be identical in the sections. Profs. Allin & Aameri will each teach certain topics in both sections.

CSC384: Intro to Artificial Intelligence

Lectures/Tutorials:

LEC0101: Mon, Wed, Fri 1:00 - 2:00 pm
BA 1190 (Bahen Centre, 40 St. George Street)

LEC0201/2001: Mon, Wed, Fri 3:00 - 4:00 pm
KP 108 (Koffler House, 569 Spadina Crescent)

LEC5101: Wed 6:00 - 9:00 pm
BA 1190 (Bahen Centre, 40 St. George Street)

The last hour will be a continuation of the lecture period and/or time to go over extra examples and questions. Don't plan to miss it!

CSC384: Textbook

Artificial Intelligence: A Modern Approach, Stuart Russell and Peter Norvig. 3rd Edition, 2010.

- 2 copies of are on 24hr reserve in the Engineering and Computer Science Library.
- Recommended but not required.
- Lecture notes cover much of the course material and will be available online before class.
- Electronic version available online at a reduced price.

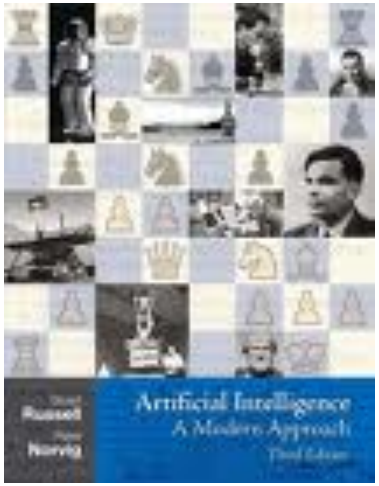

3rd edition:

Other Recommended Books and Resources:

- Knowledge Representation and Reasoning. Brachman & Levesque. 2004.
- Computational Intelligence: A Logical Approach. D. Poole, A. Mackworth, R. Goebel, Oxford Press, 1998.
- Artificial Intelligence Foundations of Computational Agents (2017), Poole & Macworth. The website for this book is http://artint.info/. Note: the text is on-line and free!

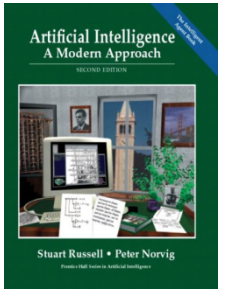

2nd edition

CSC384: Other Resources

Online Courses:

Various materials and lectures are available online, e.g.,

Udacity

- https://www.udacity.com/courses
- Introduction to Artificial Intelligence

Berkeley

- http://ai.berkeley.edu/home.html
- We will be using some of their software

CSC384: Prerequisites

- Some probability (STA 247/STA255/STA257H).
- Good knowledge of **python** (assignments involve python programming).
- Knowledge of basic data structures (stacks, queues, priority queues), graph, familiarity with Big O notation and run time complexity (CSC263, 265).
- If you don't have these prerequisites or don't feel fully comfortable with these ideas you will be responsible for learning any needed background material.

CSC384: Website

- **Course web site**

  http://www.teach.cs.toronto.edu/~csc384h/winter/

  - Primary source of more detailed information, course slides, announcements, etc.

- **Course Communication**

  Piazza: https://piazza.com/utoronto.ca/winter2020/csc384/home

- **Course Assignment Submission and Administration**

  MarkUs: TDB (Check the website and Piazza for uptades)

All announcements will be made on piazza and the course web page. You are responsible for monitoring these regularly.

CSC384: E-mail/board policies

Discussion and communication will be via Piazza, unless it's of a personal nature in which case you should email your course instructor. Please be respectful when you post to Piazza (i.e. be careful with your words).

Piazza is a third-party discussion forum with many features that are designed specifically for use with courses. We encourage you to post questions (and answers!) on Piazza, and read what other questions your classmates have posted. However, since Piazza is run by another company, we also encourage you to read the privacy policy carefully and only sign up if you are comfortable with it. If you are not comfortable signing up with your U of T email address, you can access read-only mode. Feel free to discuss further with your instructor.

- For each assignment, a TA(s) will be assigned to answer questions and to monitor Piazza
- Answers that are important to everyone will be posted to the web site.
- For personal email, subject header should start "[CSC384]".
- A silent period will take effect 24 hours before each assignment is due. I.e. no question related to the assignment that is posed within this 24 hour period will be answered.

CSC384: Where to get help

- Piazza
- Office hours
- Help sessions: we'll be running extensive assignment-specific help sessions so that you can get personalized help from TAs

CSC384: How you will be graded

Course work:

- 4 Assignments **(A1: 11%; A2: 11%; A3: 11%; A4 11%)**
- 1 term test **(16%)**
- 1 final exam **(40 %)**

Late Policy/Missing Test:

- You will have 3 grace days. Use them wisely!
- After that, you will be penalized 10% per day for late assignments.
- For some assignments there may be a cut-off date after which assignments will no longer be accepted.

Plagiarism: (submission of work not substantially the student's own)
http://www.cs.toronto.edu/~fpitt/documents/plagiarism.html

CSC384: How you will be graded

Grading Scheme Item	Topic	Weight	Dates (subject to some Uncertainty!)
Assignment 1	Search	11%	Out Jan 20, due Feb 4 @ 10:00pm
Assignment 2	Constraint Satisfaction	11%	Out Feb 3, due Feb 25 @ 10:00pm
Midterm Test		16%	Feb 12 In Class
Assignment 3	Game Tree Search	11%	Out Feb 24, due Mar 10 @ 10:00pm
Assignment 4	Reasoning with Uncertainty	11%	Out Mar 16, due Apr 3 @ 10:00pm
Final Exam		40%	Exam Period

**** Assignment and test dates are tentative and may be updated ****

Grading Summary:

- 4 Assignments **(A1: 11%; A2: 11%; A3: 11%; A4: 11%)**
- 1 term test **(16%)**
- 1 final exam **(40 %)**

WHAT IS AI?

What is Artificial Intelligence (AI)?

How to achieve “intelligent” behaviour through computational means

...but what is “intelligent”?

Are these intelligent?

What about these?

Recent Acceleration

In the last 5 years the number of technologies and devices embodying some form of “machine intelligence” have exploded …

…even surpassing the expectation of those developing the enabling technologies and the basic research

A BRIEF LOOK BACK

Alan Turing (1912 – 1954)

"Computing machinery and intelligence"
Mind, 59, 433-460., 1950

The Imitation Game

"Can machines think?" *I believe to be too meaningless to deserve discussion. Nevertheless I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.*

1956-1997

1956: John McCarthy - Dartmouth Artificial Intelligence Conference

1957: Herbert Simon - “within 10 years a digital computer will be the world’s chess champion”

1967: Marvin Minsky - “within a generation .. the problem of creating ‘artificial intelligence’ will substantially be solved".

1974-80: Slow progress, relative to expectations “AI Winter” 

Early 1980’s: Expert Systems boom

1982-1992: Japanese Fifth-Generation Project

1987-1993: 2nd AI Winter 

1997: Chess

IBM's Deep Blue beats chess grandmaster Garry Kasparov

2011: Watson

IBM's Watson defeats the Jeopardy champions Brad Rutter and Ken Jennings by a significant margin.

AlphaGo

2016: AlphaGo beats 9-Dan pro Go player Lee Sedol
2017: AlphaGo Zero – learns by playing with itself

Autonomous Driving

2005: DARPA Grand Challenge – Stanford autonomous vehicle drives 131 miles along an unrehearsed desert trail.

Autonomous Driving

2007: DARPA Urban Challenge – CMU autonomous vehicle drives 55 miles in an urban environment while adhering to traffic hazards and traffic laws.

Autonomous Driving

Now: Google Car …

and efforts by all the major car manufacturers, Mobileye, Uber, …

Watch also for autonomous wheelchairs, boats, ...

Autonomous Robots

Darpa Grand Challenge
Winner Team KAIST

https://www.youtube.com/watch?v=BGOUSvaQcBs

Where are we in 2020?

More and more devices are appearing that seem
…a little smarter,
…a little more intuitive,
…a little easier to use,
…and that know your preferences.

There may be no “ah ha” moment in which we say now we have “achieved artificial intelligence”. Instead, AI may creep up on us.

Unsolved problems remain …

DARPA Robotics Challenge:

Robot “fails”

including plenty of ethical ones.

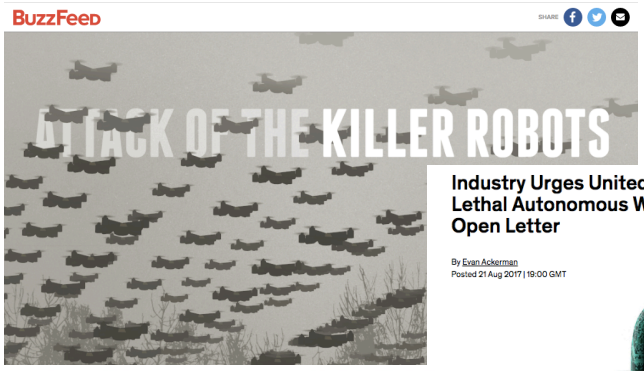

Industry Urges United Nations to Ban Lethal Autonomous Weapons in New Open Letter

By Evan Ackerman
Posted 21 Aug 2017 | 19:00 GMT

More from The Economist | My Subscription

The Economist | World politics | Business & finance | Economics | Science & technology

Special report: Artificial intelligence

The impact on jobs

Automation and anxiety

Will smarter machines cause mass u

Jun 25th 2016

SITTING IN AN office in San Francis
Igor Barani calls up some medical sc
his screen. He is the chief executive
Enlitic, one of a host of startups appl
deep learning to medicine, starting w
analysis of images such as X-rays a

Results of bias scan on COMPAS

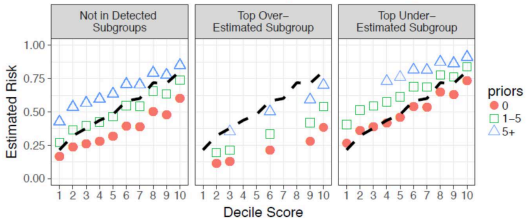

from D. Niell and Z Zhang, 2016

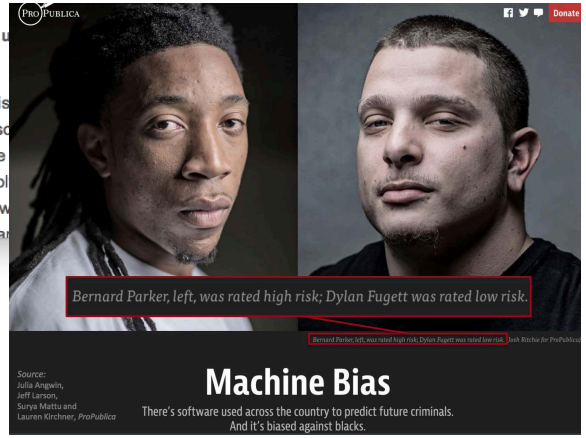

Lots of work in Canada …

theReview

How Canada became a hotspot for artificial intelligence research

TRANSPORTATION UBER SELF-DRIVING CARS

Uber hired noted AI researcher Raquel Urtasun to lead its self-driving expansion into Canada

Urtasun will continue to teach at the University of Toronto part-time and will be joined by eight of her students.

BY JOHANA BHUIYAN | @JMBOOYAH | MAY 8, 2017, 10:32AM EDT

TWEET SHARE LINKEDIN

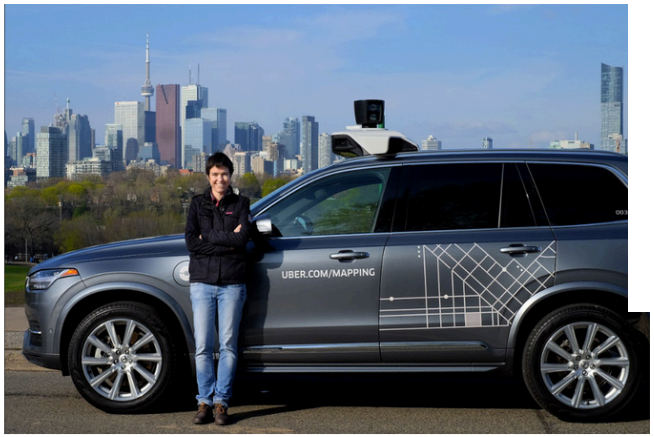

Raquel Urtason | uber

Vector Institute is just the latest in Canada's AI expansion

By Jessica Murphy
BBC News, Toronto

29 March 2017 Share

Watch this 'SNL' skit about an Amazon Echo Silver for senior citizens

Does all this success mean we've machines with human-level intelligence?

Back to our definition

AI studies … how to achieve “intelligent” behaviour

through computational means

What is Intelligence?

- Webster says:
  - The capacity to acquire and apply knowledge.
  - The faculty of thought and reason.
  - …

- What features/abilities do humans (animals/animate objects) have that you think are indicative or characteristic of intelligence?

- *Abstract concepts, mathematics, language, problem solving, memory, logical reasoning, planning ahead, emotions, morality, ability to learn/adapt, etc…*

Human Intelligence

- Turing provided some very persuasive arguments that a system passing the Turing test ***is intelligent***.
  - We can only really say it ***behaves like a human***
  - Nothing guarantees that it thinks like a human

- The Turing test does not provide much traction on the question of how to actually build an intelligent system.

Human Intelligence

- Recently some claims have been made of AI systems that can pass the Turing Test.
- However, these systems operate on subterfuge, and were able to convince a rather naïve jury that they were human like.
- The main technique used is obfuscation…rather than answering questions the system changed the topic!
- This is not what Turing described in his Turing Test.

Computational Intelligence

- *AI tries to understand and model intelligence as a computational process.*
- Thus we try to construct systems whose computation achieves or approximates a desired notion of intelligence.
- Hence AI is part of Computer Science.
  - Other areas interested in the study of intelligence lie in other areas or study, e.g., cognitive science which focuses on human intelligence. Such areas are very related, but their central focus tends to be different.

Computational Intelligence != Human Intelligence

- In general there are various reasons why trying to mimic humans might **not** be the best approach to AI:
  - Computers and Humans have a very different architecture with quite different abilities.
  - Numerical computations
  - Visual and sensory processing
  - Slow parallel vs. fast serial

	Computer	Human Brain
Computational Units	8 CPUs, 10^{10} gates	10^{11} neurons
Storage Units	10^{10} bits RAM 10^{13} bits disk	10^{11} neurons 10^{14} synapses
Cycle time	10^{-9} sec	10^{-3} sec
Bandwidth	10^{10} bits/sec	10^{14} bits/sec
Memory updates/sec	10^{10}	10^{14}

AlphaGo vs. Lee Sedol

	Lee Sedol	AlphaGo
Processing	1 brain (10^{11} neurons)	1920 CPUs + 280 GPUs
Power	20 watts	1 MW (50,000 times more)

Human Intelligence

- But more importantly, we know very little about how the human brain performs higher level processes. Hence, this point of view provides little information from which a scientific understanding of these processes can be built.

- Nevertheless, Neuroscience has been very influential in some areas of AI. For example, in robotic sensing, vision processing, etc.

- Humans might not be best comparison …
  - Don’t always make the best decisions
  - Computer intelligence can aid in our decision making

Rationality

- The alternative approach relies on the notion of **rationality**.
- Typically this is a precise mathematical notion of what it means to *do the right thing* in any particular circumstance. Provides
  - A precise mechanism for analyzing and understanding the properties of this ideal behaviour we are trying to achieve.
  - A precise benchmark against which we can measure the behaviour the systems we build.

Rationality

- Mathematical characterizations of rationality have come from diverse areas like logic (laws of thought) and economics (utility theory how best to act under uncertainty, game theory how self-interested agents interact).
- There is no universal agreement about which notion of rationality is best, but since these notions are precise we can study them and give exact characterizations of their properties, good and bad.
- We'll focus on acting rationally
  - this has implications for thinking/reasoning

Overview of an AI System

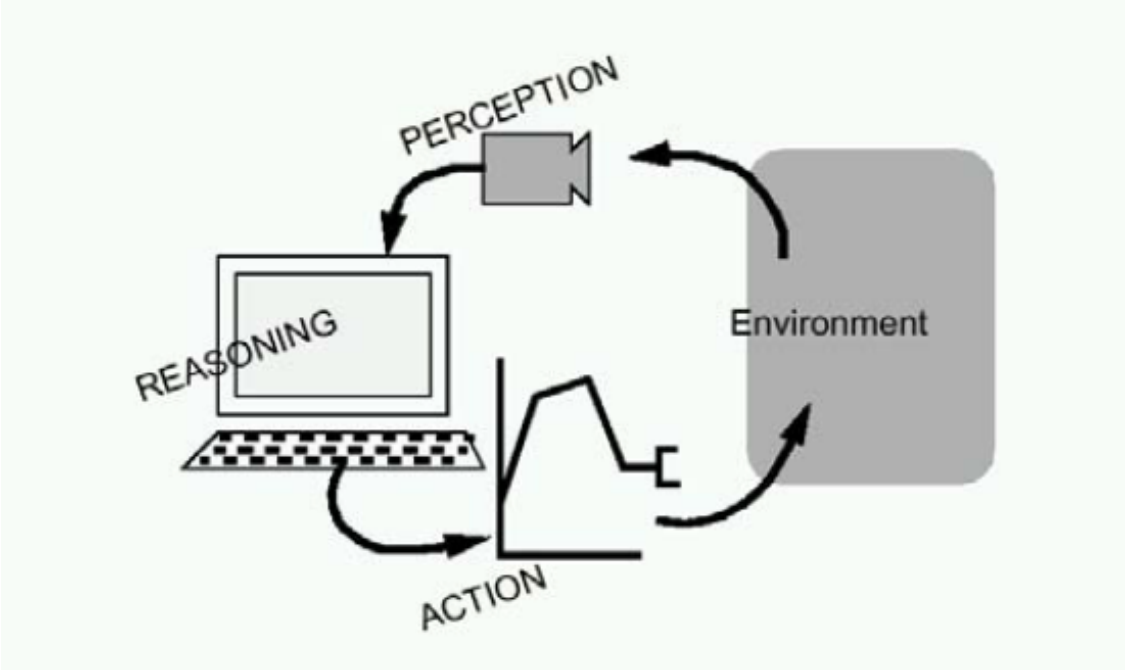

The Perception-Action Cycle

Overview of an AI System

Game Playing Agents

Environment?
Perceptions?
Decisions?
Actions?

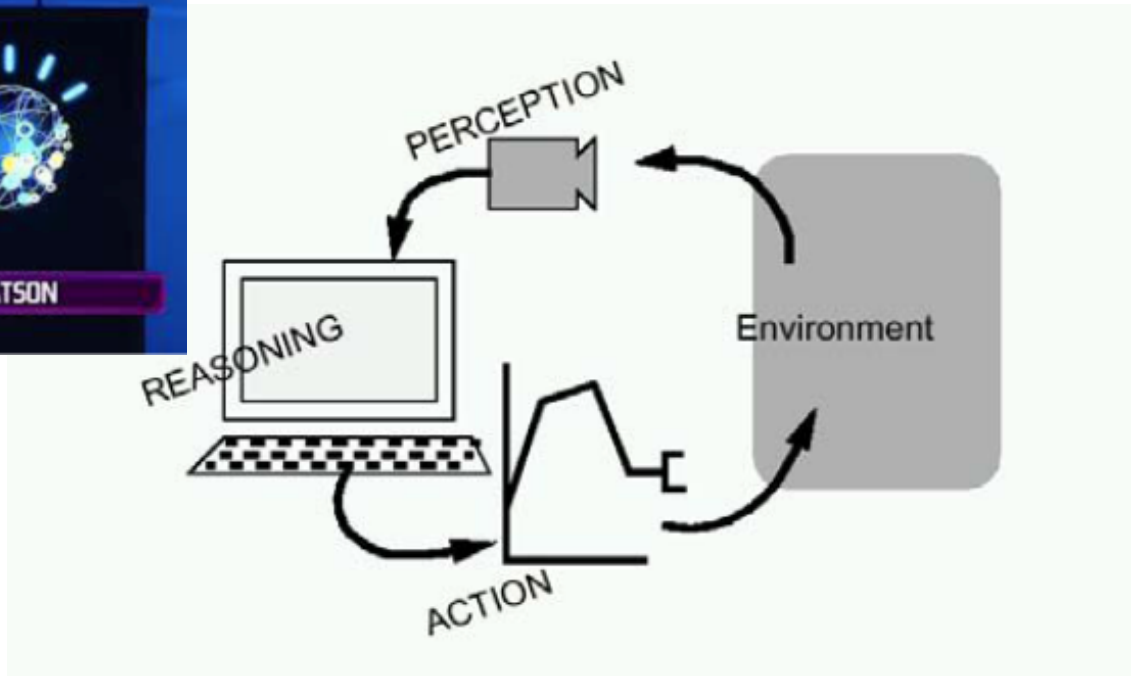

The Perception-Action Cycle

Overview of an AI System

Machine Vision, Speech Understanding, etc.

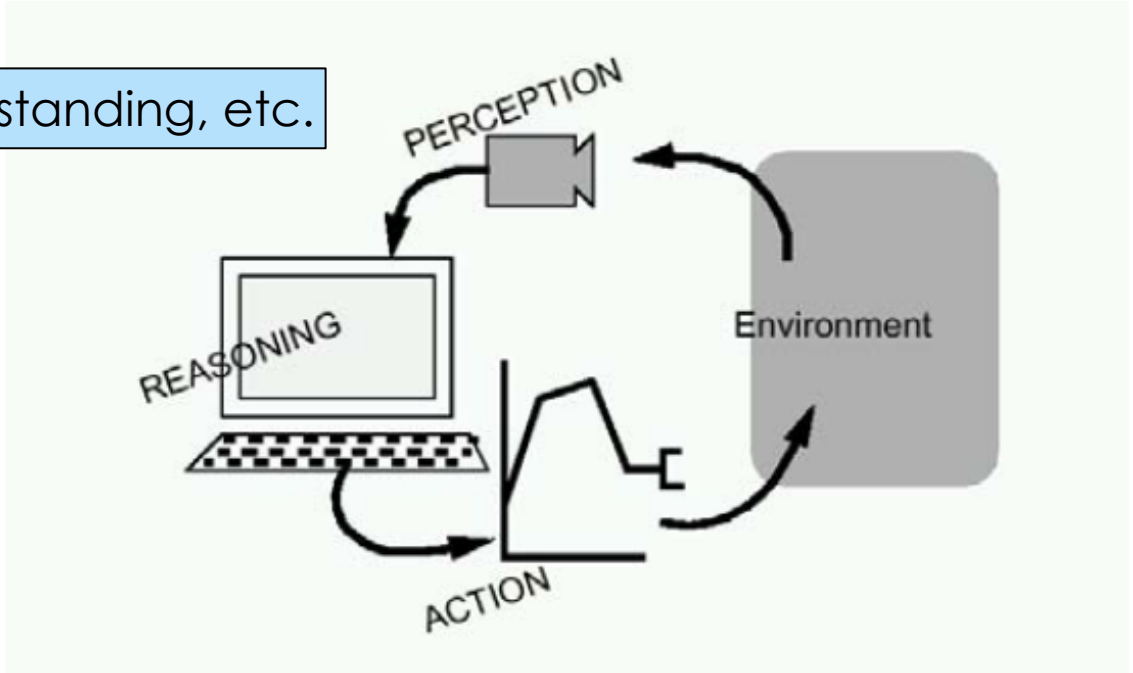

The Perception-Action Cycle

Overview of an AI System

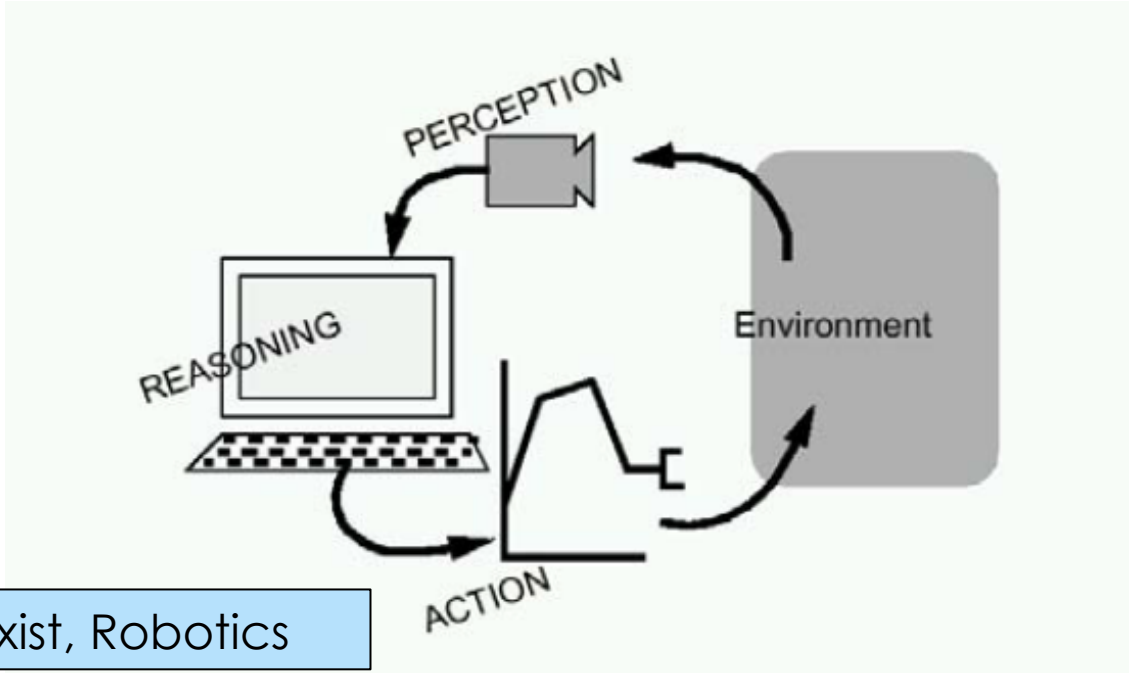

Where mechanical actuators exist, Robotics

The Perception-Action Cycle

Overview of an AI System

Natural Language Processing

Reasoning:

Knowledge Representation

Decision-Making *(search, planning, decision theory)*

Reasoning Processes *(logical, probabilistic)*

Machine Learning, Neural Networks

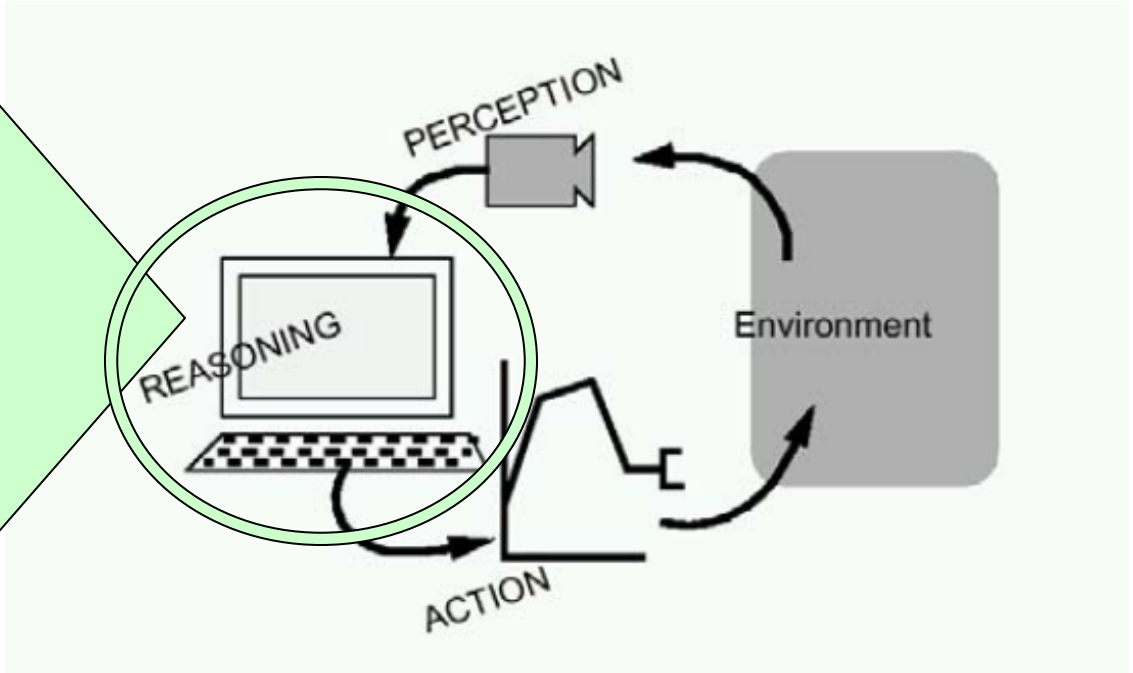

The Perception-Action Cycle

What We Cover in CSC384

- Search
  - Heuristic Search. (Chapter 3,4)
    - Search spaces
    - Heuristic guidance
  - Backtracking Search (Chapter 6)
    - “Vector of features” representation
    - Case analysis search
  - Game tree search (Chapter 5)
    - Working against an opponent

What We Cover in CSC384 (cont.)

- Knowledge Representation (Chapter 7-9,12)
  - First order logic for more general knowledge
  - Knowledge represented in declarative manner
- Planning (Chapter 10-11)
  - Predicate representation of states
  - Planning graph
- Uncertainty (Chapter 13-14 a bit of 15)
  - Probabilistic reasoning, Bayes networks
  - In passing: Utilities and influence diagrams (Chapter 16, 17)

Further Courses in AI

- CSC320H "Introduction to Visual Computing"
- CSC321H "Introduction to Neural Networks and Machine Learning"
- CSC401H1 "Natural Language Computing"
- CSC411H "Machine Learning and Data Mining"
- CSC412H1 "Probabilistic Learning and Reasoning"
- CSC420H1 "Introduction to Image Understanding"
- CSC485H1 "Computational Linguistics"
- CSC486H1 "Knowledge Representation and Reasoning"

Where We'll Start

- **Readings: Russell & Norvig.**
  - **Chapters 1 & 2 – optional but interesting!**
  - **Chapter 3 – topic to be covered over the next week+ and Assignment 1**

Get Involved!

- **Undergraduate AI Group (UAIG)**
- **Undergraduate Summer Research Assistantships (USRAs)**
- **UofT Self-driving Car Team**