CS 188: Artificial Intelligence

Introduction

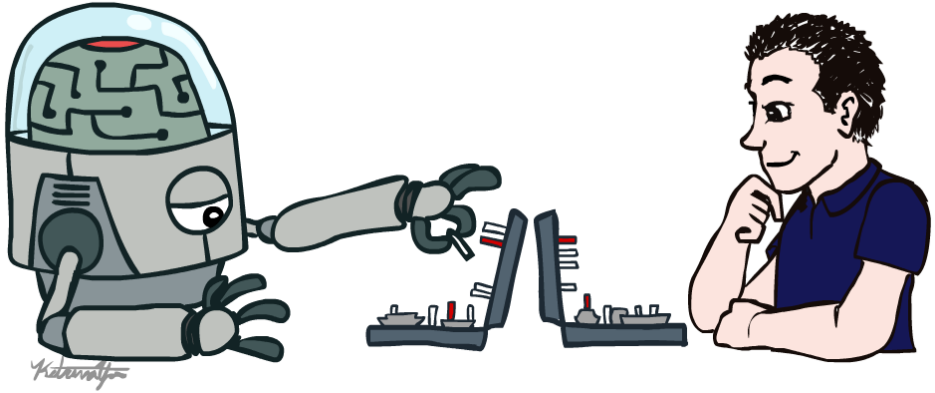

Instructors: Stuart Russell and Dawn Song

Course Staff

Professors

Dawn Song

Stuart Russell

GSIs

Saagar Sanghavi

Yanlai Yang

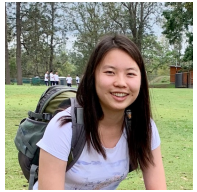

Angela Liu

Lun Wang

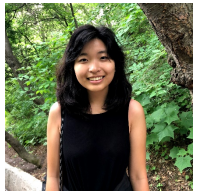

Regina Wang

Scott Emmons

Jason Wang

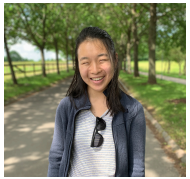

Katherine Shu

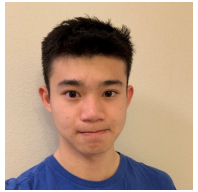

Jeffrey Tao

Jocelyn Chen

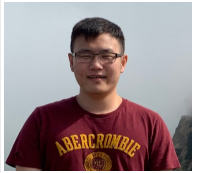

Jonathan Yang

Jasmine Collins

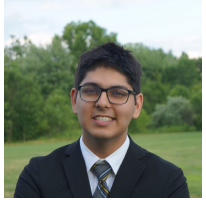

Nitish Dashora

Rishi Parikh

Evgeny P.

Ayush Kamat

Raiymbek A.

Anikait Singh

Tarun Amarnath

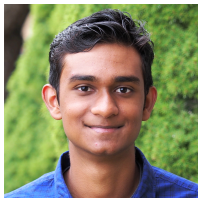

Ajay Sridhar

Arvind Rajaraman

Course Information

http://inst.cs.berkeley.edu/~cs188

CS 188 | Spring 2022

Syllabus | Policies | Projects | Schedule | Staff | Piazza

Warning: Spring 2022 website is under construction. There may be misleading information from past semesters on this website, which may be different from the current offering of the course. We will remove this warning when the information on this website is accurate.

CS 188 | Introduction to Artificial Intelligence
Spring 2022

Lectures: Tu/Th 2:00–3:30 pm, Online

Description

This course will introduce the basic ideas and techniques underlying the design of intelligent computer systems. A specific emphasis will be on the statistical and decision-theoretic modeling paradigm.

By the end of this course, you will have built autonomous agents that efficiently make decisions in fully informed, partially observable and adversarial settings. Your agents will draw inferences in uncertain environments and optimize actions for arbitrary reward structures. Your machine learning algorithms will classify handwritten digits and photographs. The techniques you learn in this course apply to a wide variety of artificial intelligence problems and will serve as the foundation for further study in any application area you choose to pursue.

See the syllabus for slides, deadlines, and the lecture schedule. Readings refer to fourth edition of AIMA unless otherwise specified.

Lecture reference materials

We make lecture recordings available as links to Kaltura, which you can find posted together with other materials on the Syllabus page of this website shortly after the lecture. These links will work only if you are signed into your UC Berkeley bCourses (Canvas) account.

Syllabus

- Communication:
  - Announcements, questions on Piazza
  - Staff email: cs188@berkeley.edu
  - Office hours:
    - Stuart: Monday 1.30-3, Tuesday 3.30-5*
    - Dawn: TBD
  - Sections start this week
- Work:
  - Projects (25%), homework (10% + 10%)
    - **P0 (Python) due 1/21, HW0 (math) due 1/26**
  - Midterm (20%), final (35%)
  - Participation up to **5%** extra (**be nice!**)
  - Fixed grading scale (85% A, 80% A-, etc.)

Some Historical Statistics

- Homework and projects: instruction (iterate/learn till you nailed it)

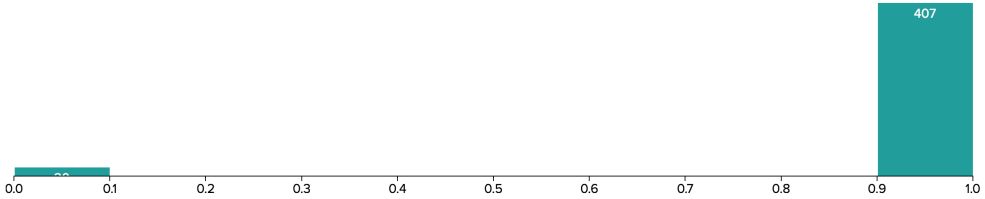

- Exams: assessment

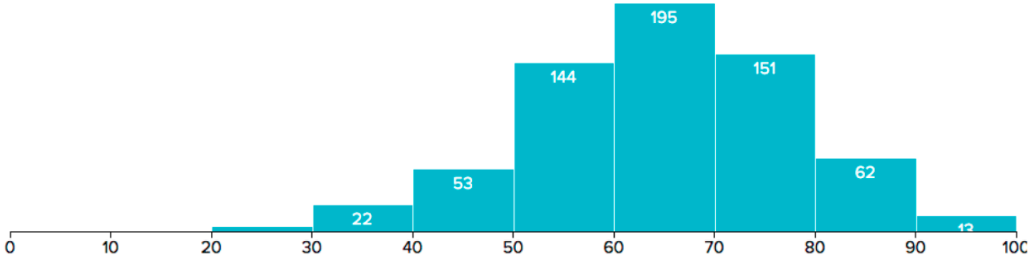

Textbook

Russell & Norvig, AI: A Modern Approach, 4th Ed.

(sorry!)

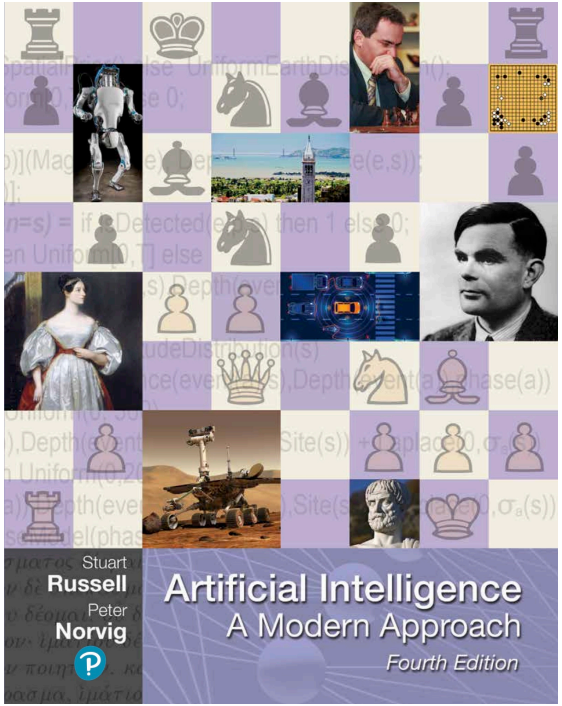

Policies (see website)

- For online lectures:
  - Camera on, mic off
  - Please do ask questions: “Hand Up” or write in Chat
- We (staff) are here to help
  - Please do observe academic integrity policies!
  - Please don’t exclude your fellow students!

Today

- What is artificial intelligence?
- Where are we and how did we get here?
- How do we think about the design of AI systems?

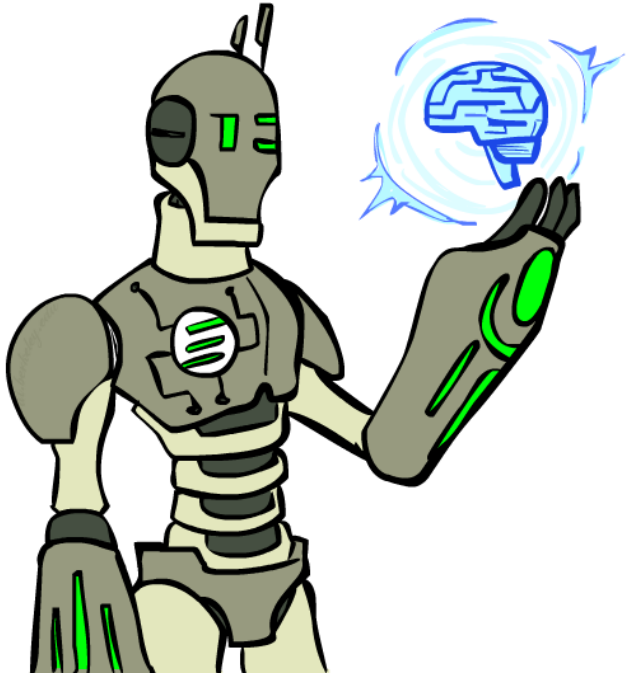

Movie AI

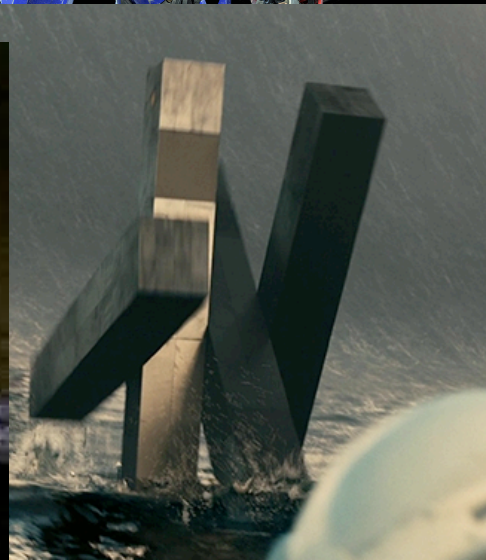

Movie AI

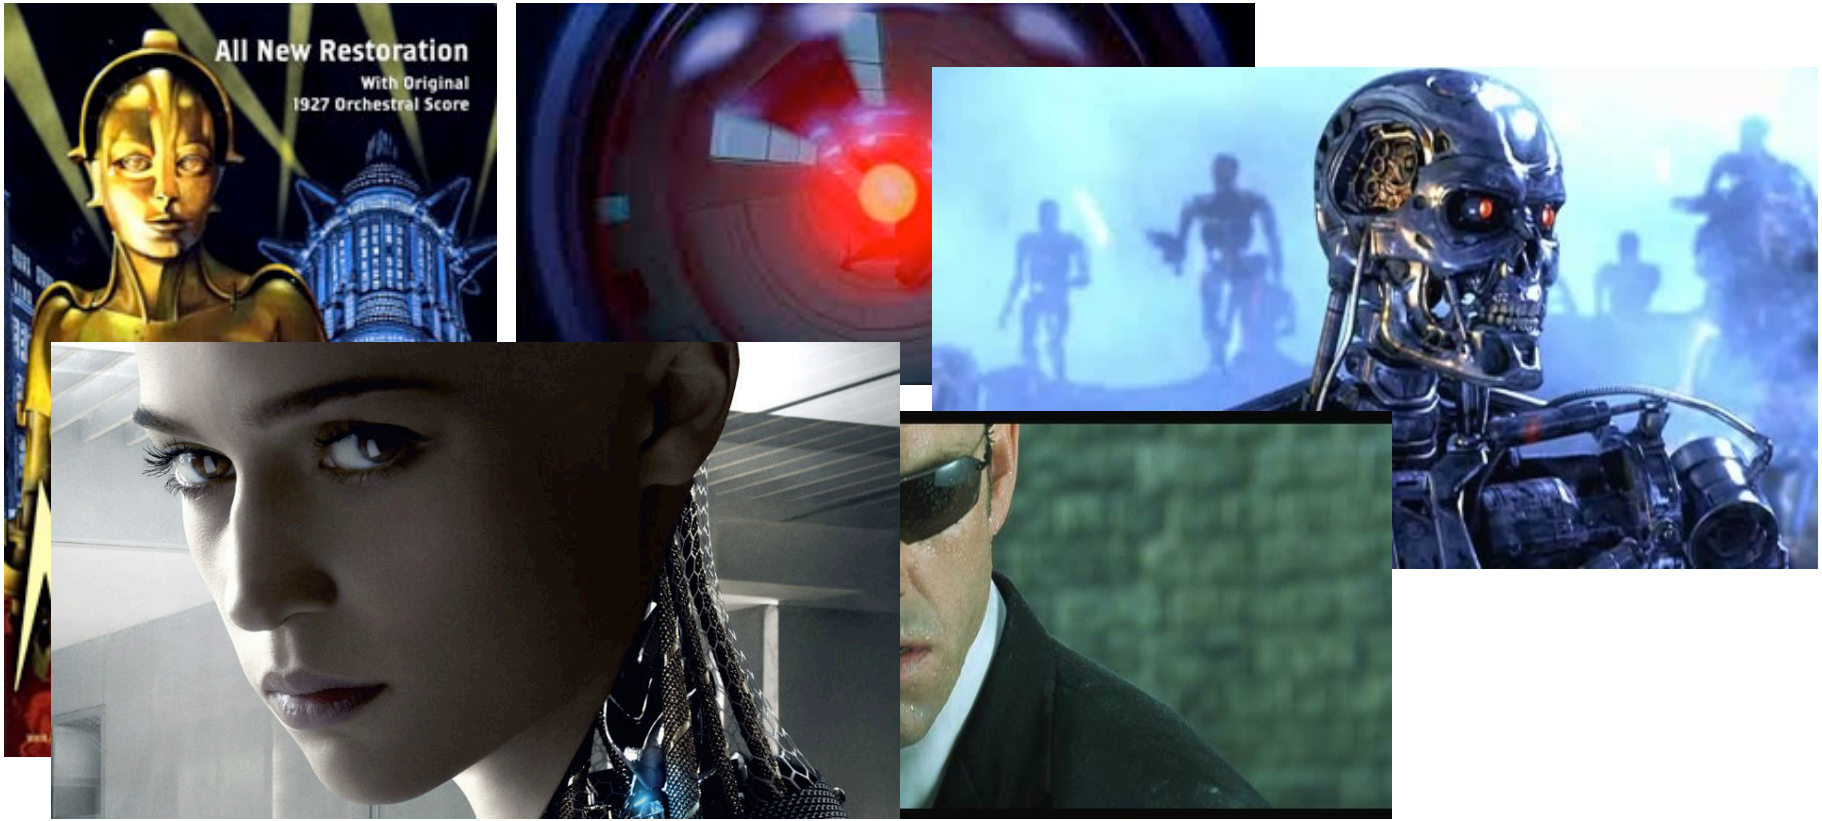


YESTERDAY DR. WILL CASTER WAS ONLY HUMAN
JOHNNY DEPP
REBECCA HALL
PAUL BETTANY
KATE MARA
CILLIAN MURPHY
AND MORGAN FREEMAN
TRANSCENDENCE
entertainmentfilms.co.uk

News AI

AI is the biggest risk we face as a civilisation, Elon Musk says

Billionaire burn: Musk says Zuckerberg's un

BY PE
Publish
Updated

MORE

HOME » FINANCE » FINANCE TOPICS » DAVOS

'Sociopathic' robots could overrun the human ra
a generation

Computers should be trained to serve humans to reduce their th
human race, says a leading expert on artificial intelligence

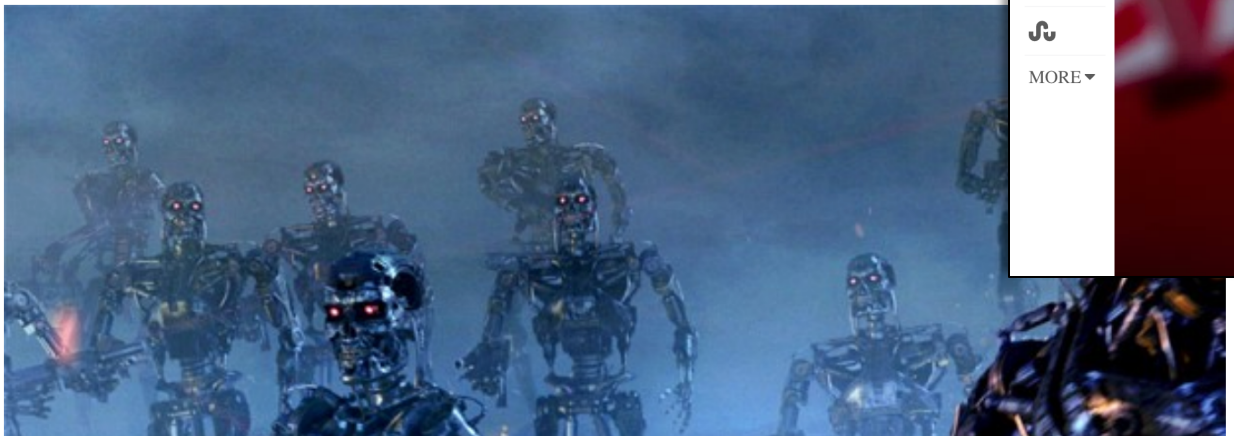

LIVESCIENCE NEWS TECH HEALTH PLANET EAR

Live Science > Tech

Lifelike 'Sophia' Robot Granted Citizenship to Saudi Arabia

By Mindy Weisberger, Senior Writer | October 30, 2017 03:39pm ET

Be the most interesting person you know, subscribe to LiveScience. Subscribe ›

MORE

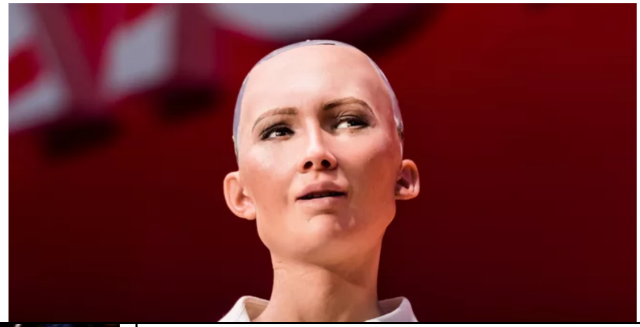

News AI

TECH • ARTIFICIAL INTELLIGENCE

United Kingdom Plans $1.3 Billion Intelligence Push

France to spend $1.8 billion on
compete with U.S., China

EU wants to invest £18b
development

**China's Got a Huge Art
Intelligence Plan**

'Whoever leads in AI will rule the world': Putin to Russian children on Knowledge Day

Published time: 1 Sep, 2017 14:08
Edited time: 1 Sep, 2017 14:40

News AI

IBM's Watson Jeopardy Computer Shuts Down Humans in Final Game

DAILY NEWS 9 March 2016

Sili

'I'm in shock!' Ho world's best hum

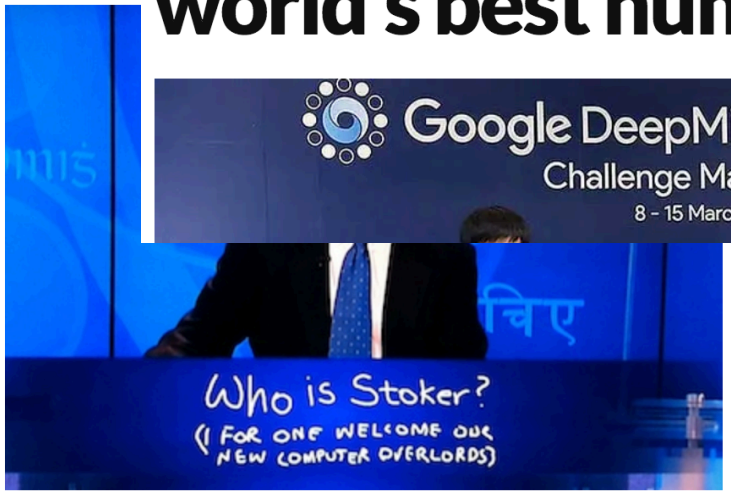

Blizzard will show off Google's Deepmind AI in StarCraft 2 later this week

By Andy Chalk 4 hours ago

Google and Blizzard launched the artificial intelligence project in 2016.

COMMENTS

Real AI

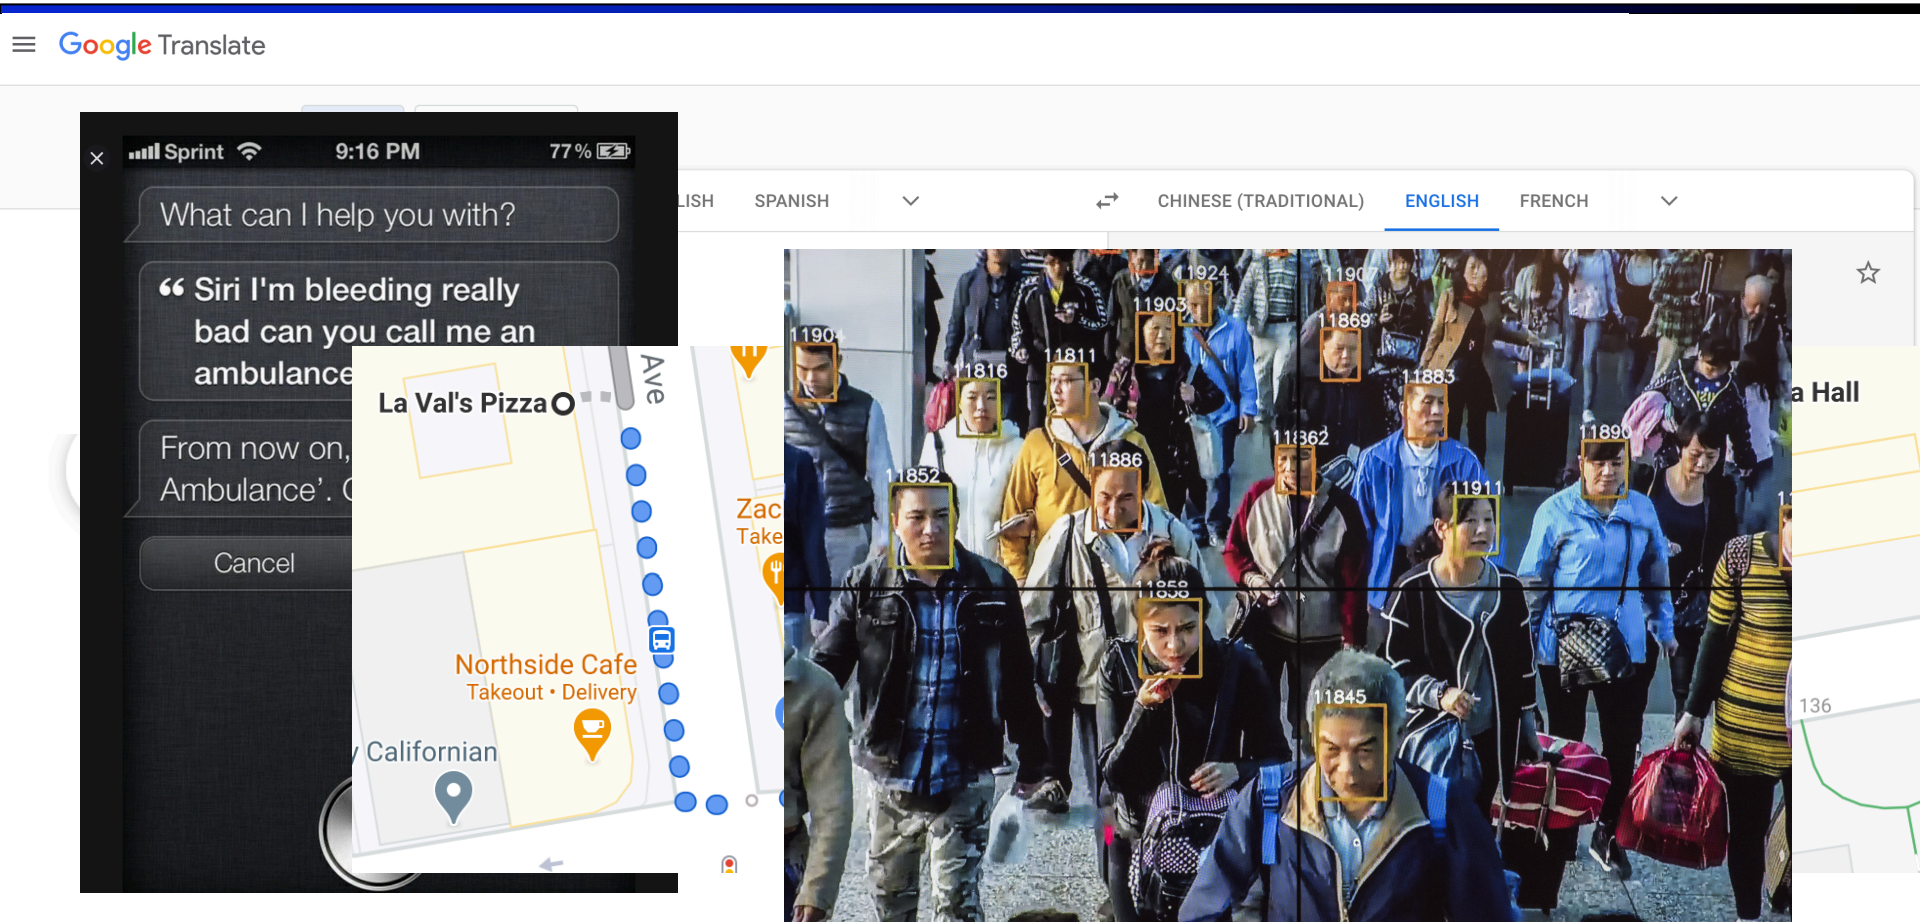
Google Translate
Sprint
9:16 PM
77 %
What can I help you with?
Siri I'm bleeding really bad can you call me an ambulance
From now on,
Ambulance'.
Cancel
La Val's Pizza
Ave
Northside Cafe
Takeout • Delivery
Californian
SPANISH
CHINESE (TRADITIONAL)
ENGLISH
FRENCH
a Hall
136


TUG
CAUTION
MAY CONTAIN
CHEMOTHERAPY DRUG
CAUTION
MAY CONTAIN
CHEMOTHERAPY DRUG


Boston Dynamics

A (Short) History of AI

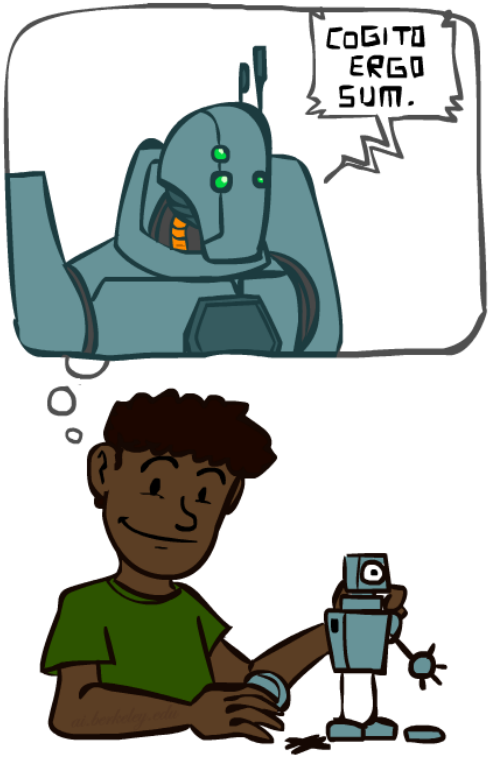

Demo: HISTORY – MT1950.wmv

A short prehistory of AI

- Prehistory:
  - **Philosophy** (reasoning, planning, learning, science, automation)
  - Aristotle: For if every instrument could accomplish its own work, obeying or anticipating the will of others . . . if, in like manner, the shuttle would weave and the plectrum touch the lyre without a hand to guide them, chief workmen would not want servants, nor masters slaves
  - **Psychology** (learning, cognitive models)
  - **Linguistics** (grammars, formal representation of meaning)
- Near miss (1842):
  - Babbage design for universal machine
  - Lovelace: “a thinking machine” for “all subjects in the universe.”

AI's official birth: Dartmouth, 1956

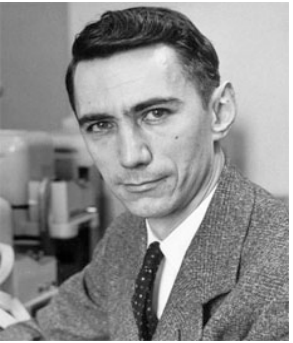

"An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. ***We think that a significant advance can be made if we work on it together for a summer."***

John McCarthy and Claude Shannon
Dartmouth Workshop Proposal

A (Short) History of AI

- 1940-1950: Early days
  - 1943: McCulloch & Pitts: Boolean circuit model of brain
  - 1950: Turing's “Computing Machinery and Intelligence”
- 1950—70: Excitement: Look, Ma, no hands!
  - 1950s: Early AI programs: chess, checkers (RL), theorem proving
  - 1956: Dartmouth meeting: “Artificial Intelligence” adopted
  - 1965: Robinson's complete algorithm for logical reasoning
- 1970—90: Knowledge-based approaches
  - 1969—79: Early development of knowledge-based systems
  - 1980—88: Expert systems industry booms
  - 1988—93: Expert systems industry busts: “AI Winter”
- 1990— 2012: Statistical approaches + subfield expertise
  - Resurgence of probability, focus on uncertainty
  - General increase in technical depth
  - Agents and learning systems… “AI Spring”?
- 2012— ___: Excitement: Look, Ma, no hands again?
  - Big data, big compute, deep learning
  - AI used in many industries

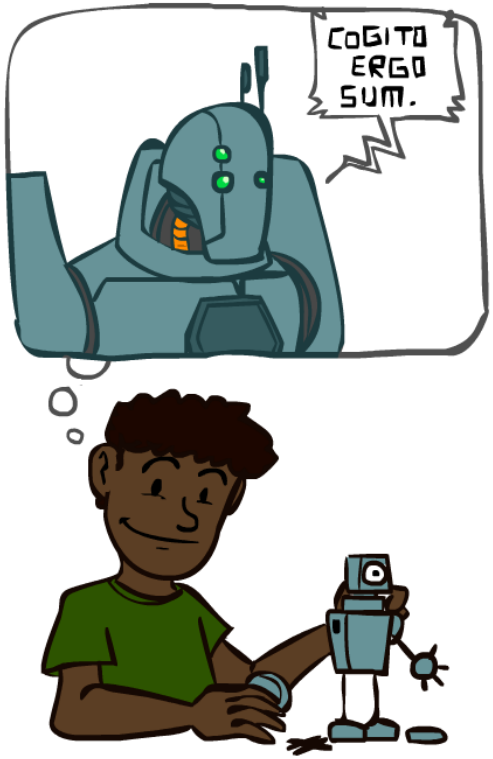

AI as Designing Rational Agents

- An **agent** is an entity that *perceives* and *acts*.
- A **rational agent** selects actions that maximize its expected **utility**.
- Characteristics of the **sensors, actuators, and environment** dictate techniques for selecting rational actions
- **This course** is about:
  - General AI techniques for many problem types
  - Learning to choose and apply the technique appropriate for each problem

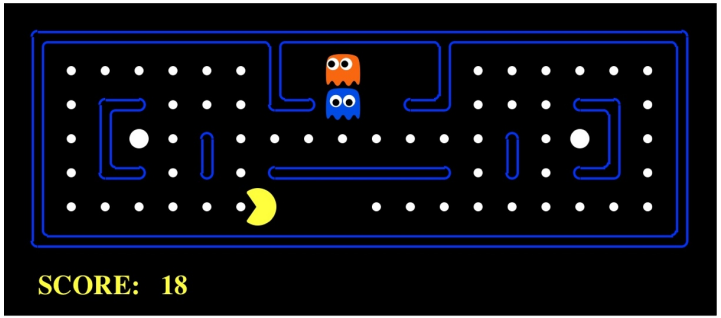

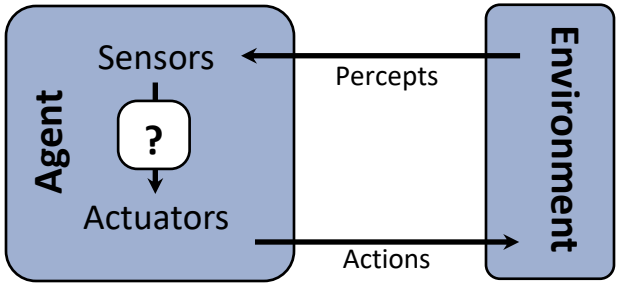

Agents and environments

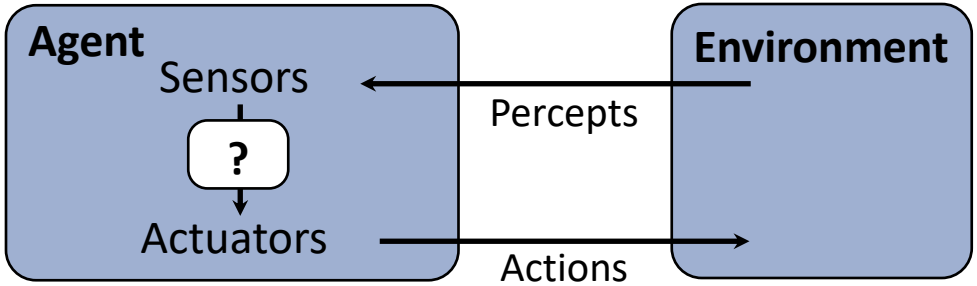

- An agent ***perceives*** its environment through ***sensors*** and ***acts*** upon it through ***actuators*** (or *effectors*, depending on whom you ask)
- The ***agent function*** maps percept sequences to actions
- It is generated by an ***agent program*** running on a ***machine***

A human agent in Pacman

The task environment - PEAS

- Performance measure
  - -1 per step; + 10 food; +500 win; -500 die; +200 hit scared ghost
- Environment
  - Pacman dynamics (incl ghost behavior)
- Actuators
  - Left Right Up Down or NSEW
- Sensors
  - Entire state is visible (except power pellet duration)

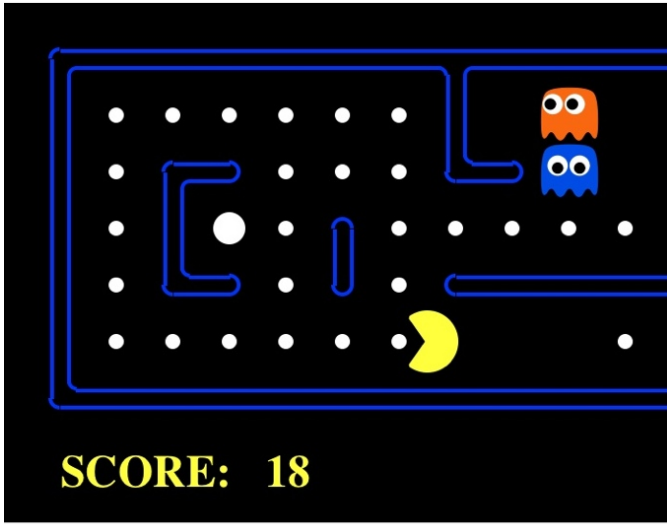

PEAS: Automated taxi

- Performance measure
  - Income, happy customer, vehicle costs, fines, insurance premiums
- Environment
  - US streets, other drivers, customers, weather, police…
- Actuators
  - Steering, brake, gas, display/speaker
- Sensors
  - Camera, radar, accelerometer, engine sensors, microphone, GPS

Image: http://nypost.com/2014/06/21/how-google-might-put-taxi-drivers-out-of-business/

PEAS: Medical diagnosis system

- Performance measure
  - Patient health, cost, reputation
- Environment
  - Patients, medical staff, insurers, courts
- Actuators
  - Screen display, email
- Sensors
  - Keyboard/mouse

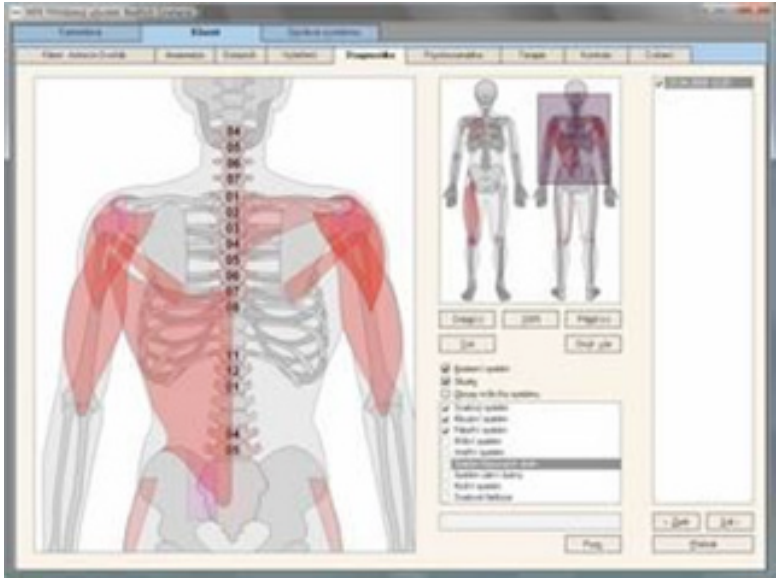


atrium
AmeriMed Choice

Environment types

	Pacman	Backgammon	Diagnosis	Taxi
Fully or partially observable				
Single-agent or multiagent				
Deterministic or stochastic				
Static or dynamic				
Discrete or continuous				
Known physics?				
Known perf. measure?				

Agent design

- **The environment type largely determines the agent design**
  - ***Partially observable*** => agent requires ***memory*** (internal state)
  - ***Stochastic*** => agent may have to prepare for ***contingencies***
  - ***Multi-agent*** => agent may need to behave ***randomly***
  - ***Static*** => agent has time to compute a rational decision
  - ***Continuous time*** => continuously operating ***controller***
  - ***Unknown physics*** => need for ***exploration***
  - ***Unknown perf. measure*** => observe/interact with ***human principal***

Simple reflex agents

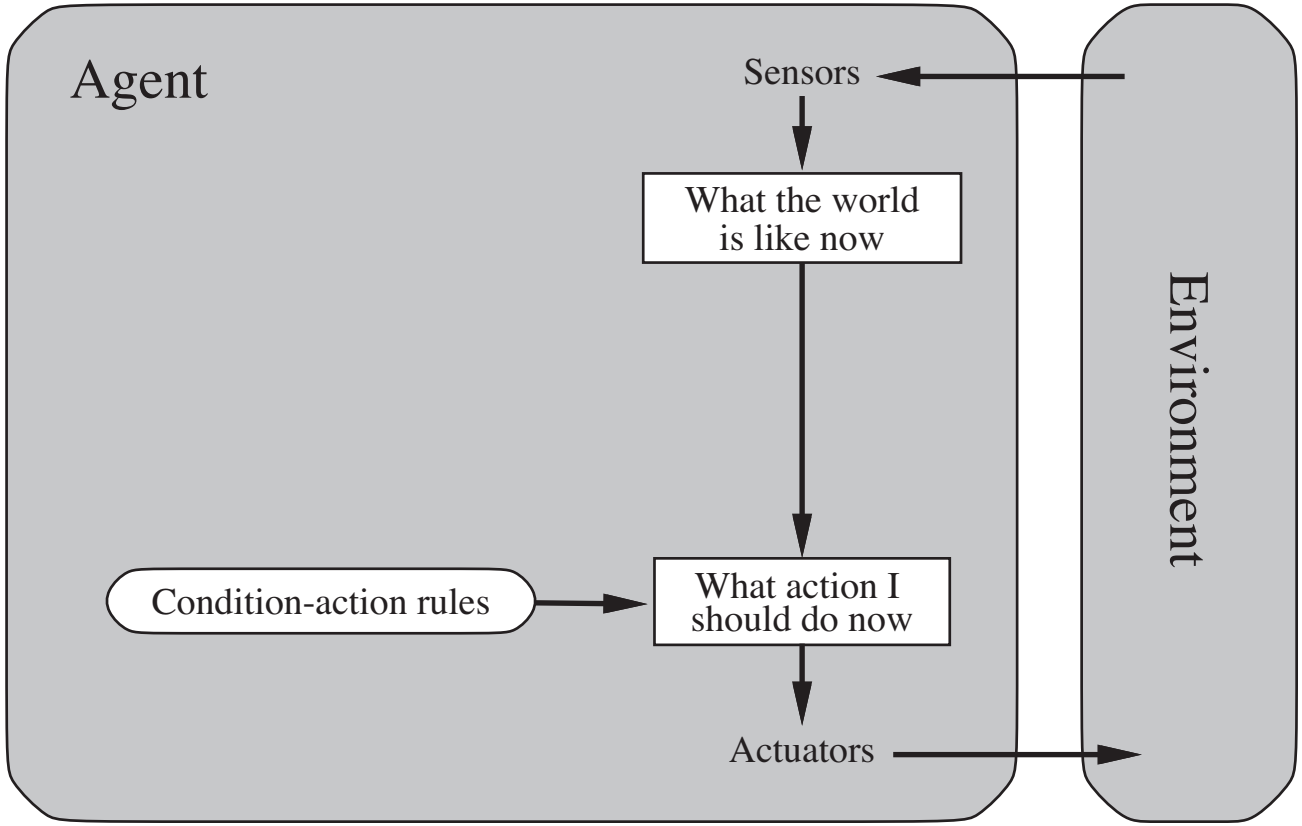

Pacman *agent program* in Python

```
class GoWestAgent(Agent):

    def getAction(self, percept):
        if Directions.WEST in percept.getLegalPacmanActions():
            return Directions.WEST
        else:
            return Directions.STOP
```

Eat adjacent dot, if any

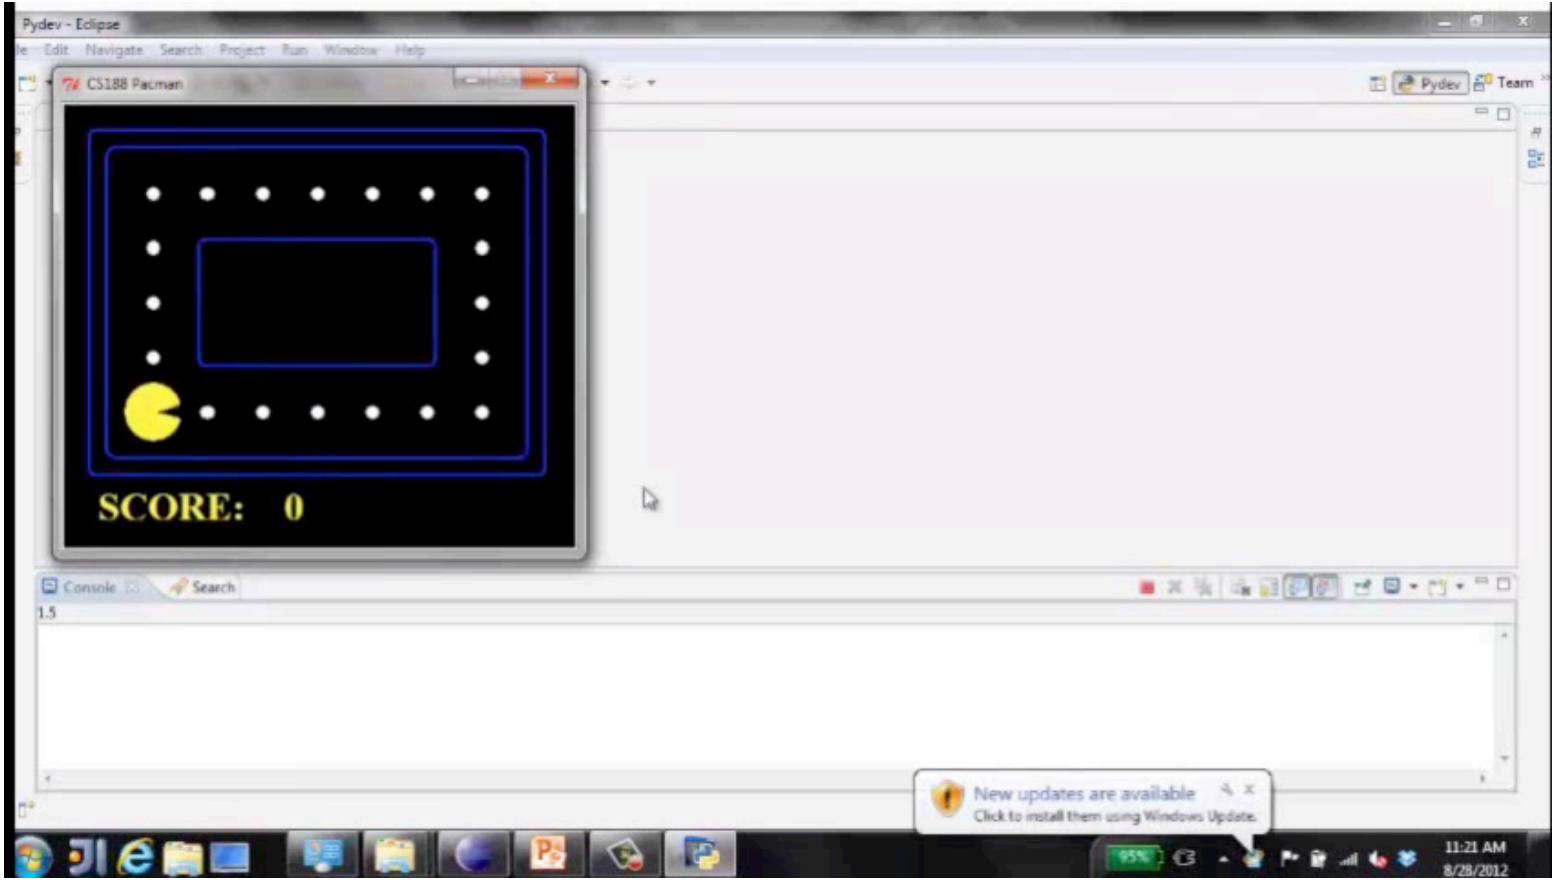

Eat adjacent dot, if any

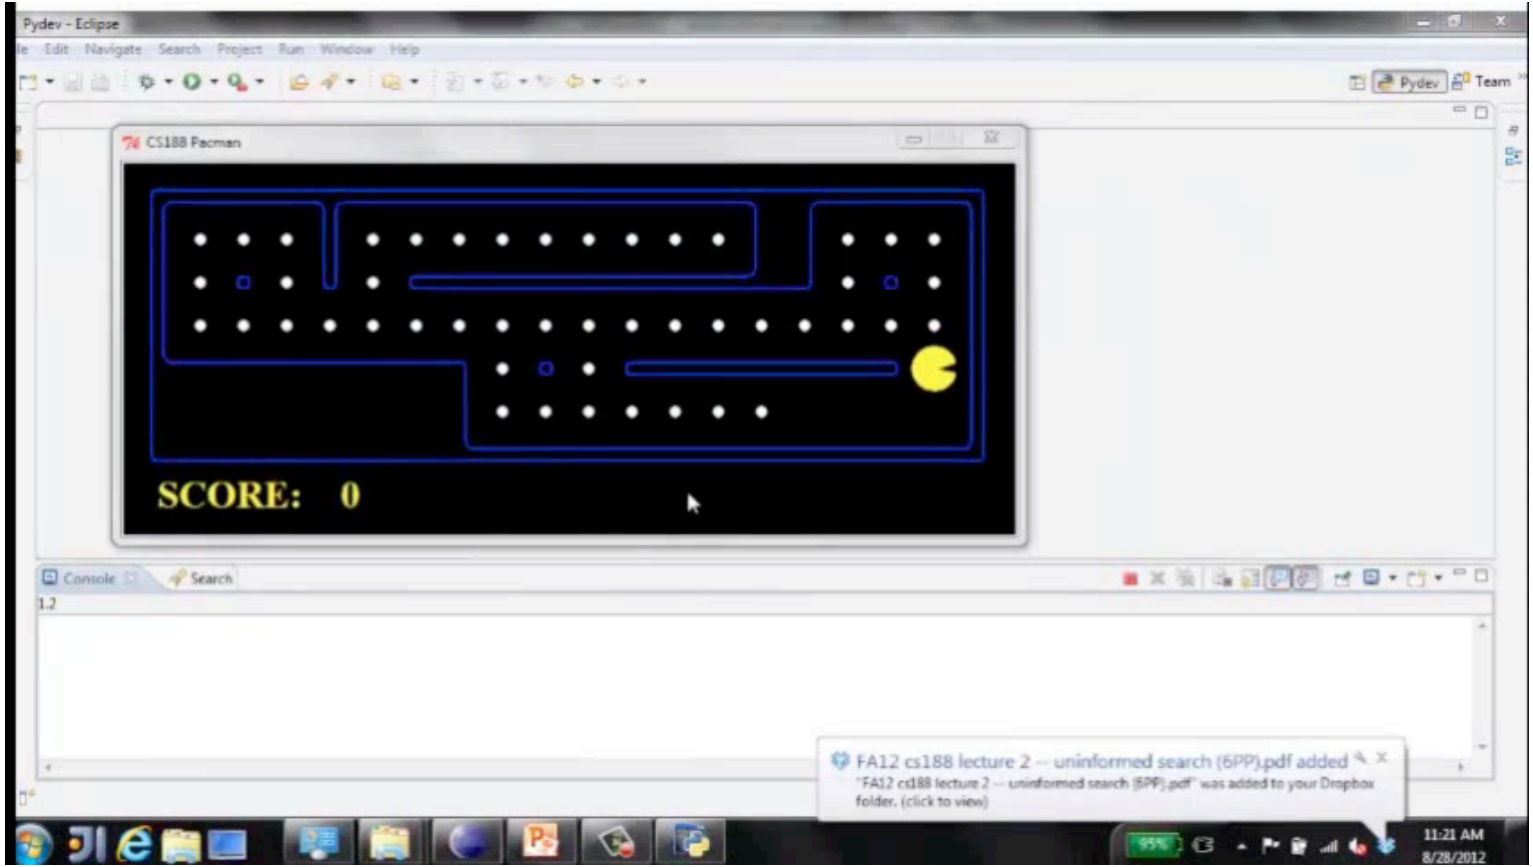

Pacman agent contd.

- Can we (in principle) extend this reflex agent to behave well in all standard Pacman environments?
  - No – Pacman is not quite fully observable (power pellet duration)
  - Otherwise, yes – we can (*in principle*) make a lookup table…..
  - *How large would it be?*

Reflex agents with state

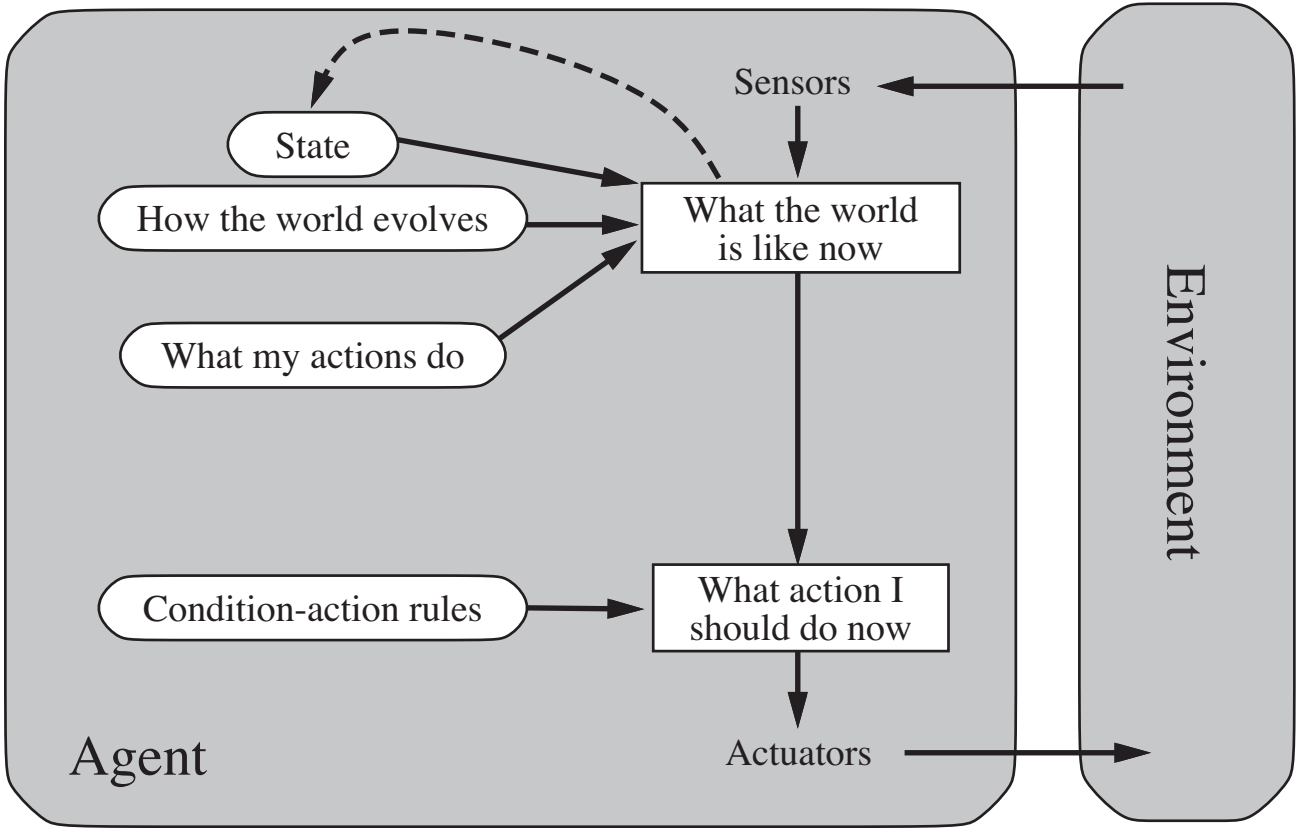

Goal-based agents

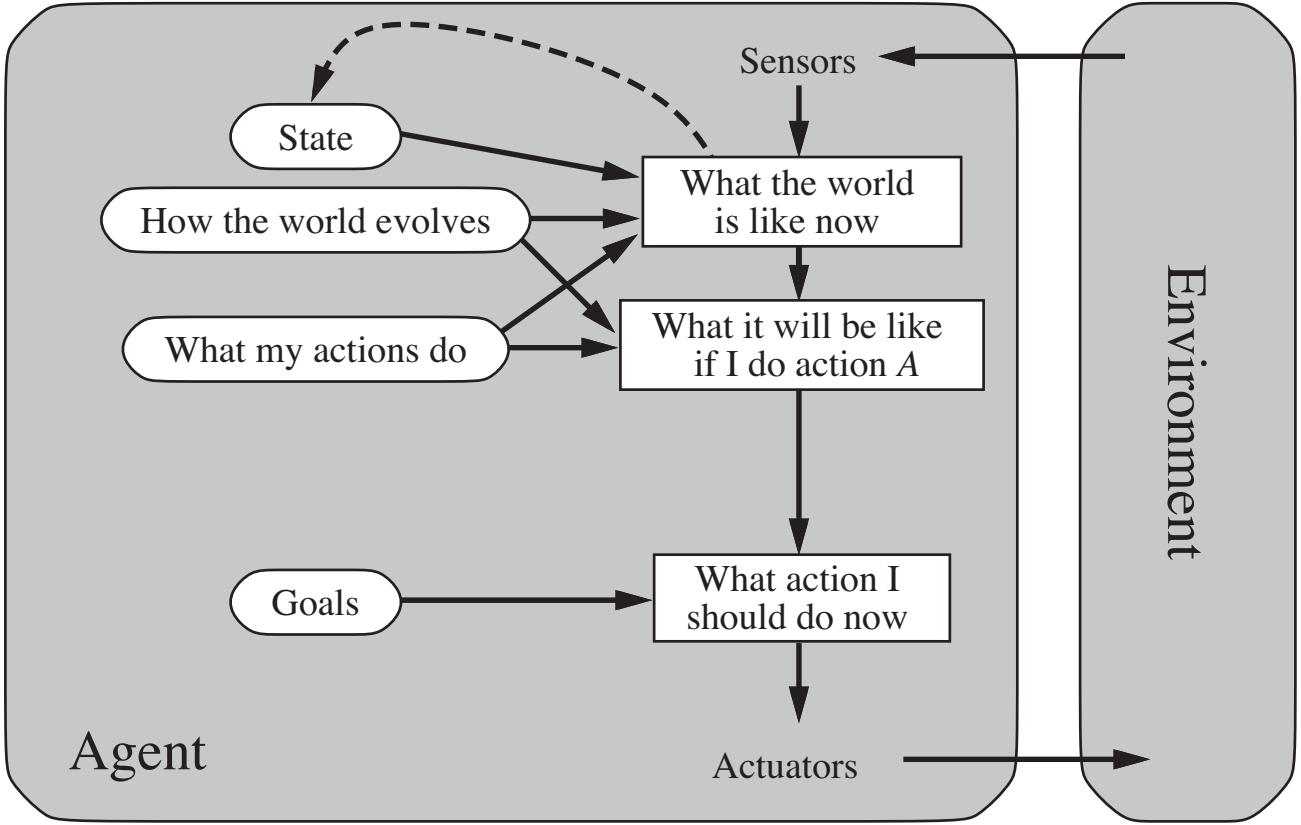

Spectrum of representations

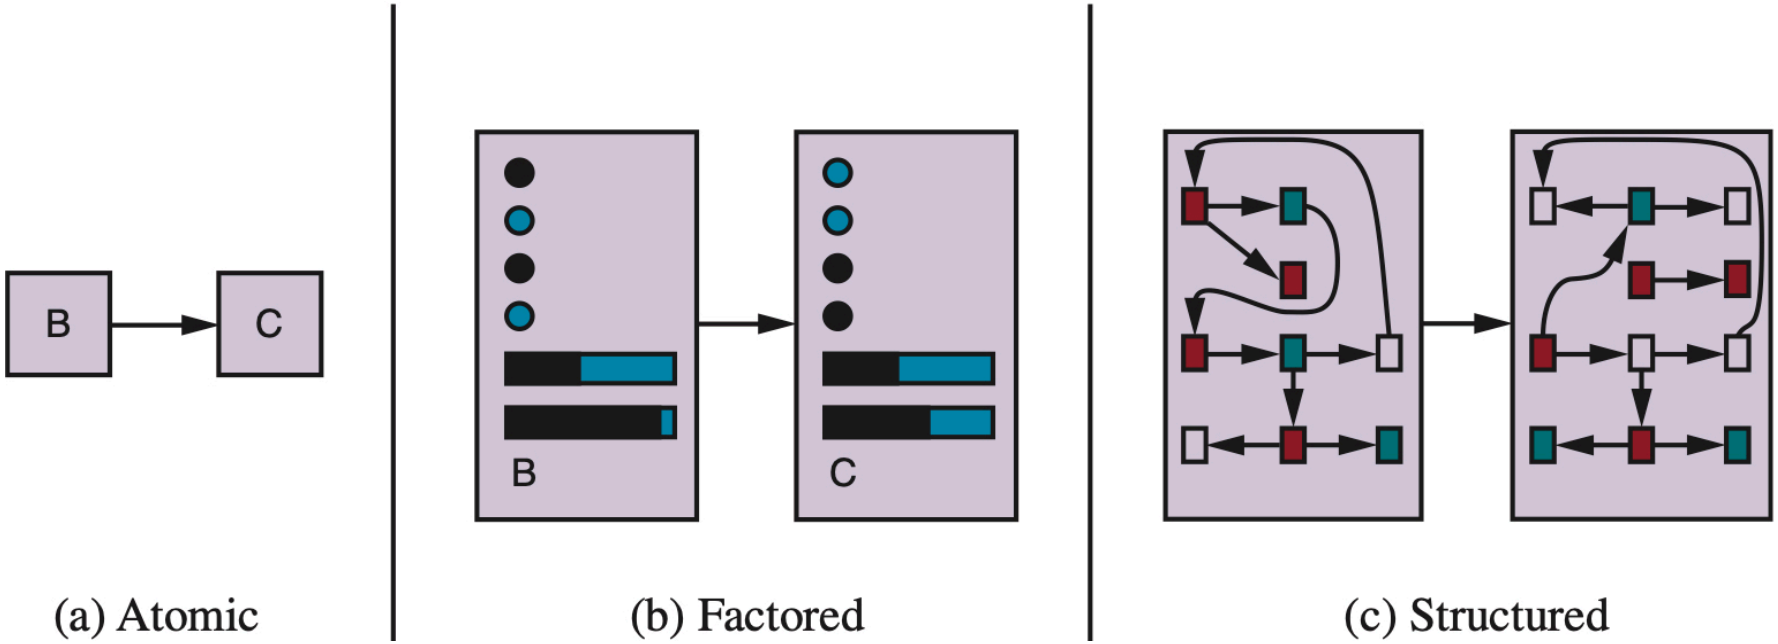

(a) Atomic (b) Factored (c) Structured

Outline of the course