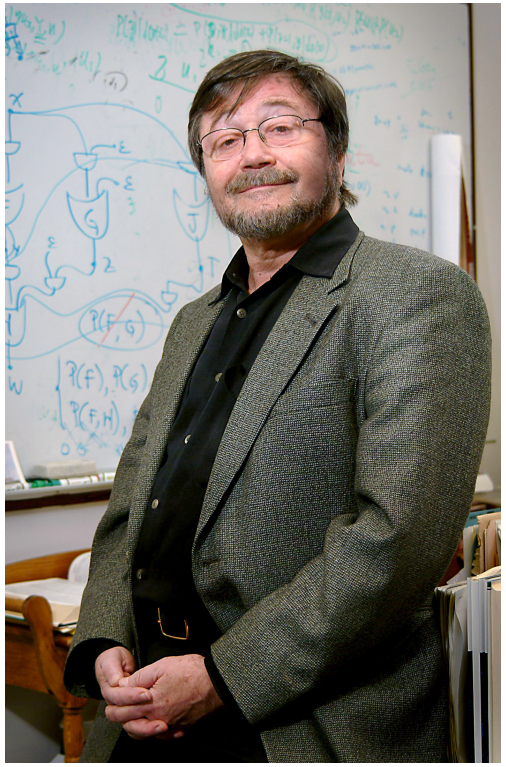

Artificial Intelligence Intro and Agents

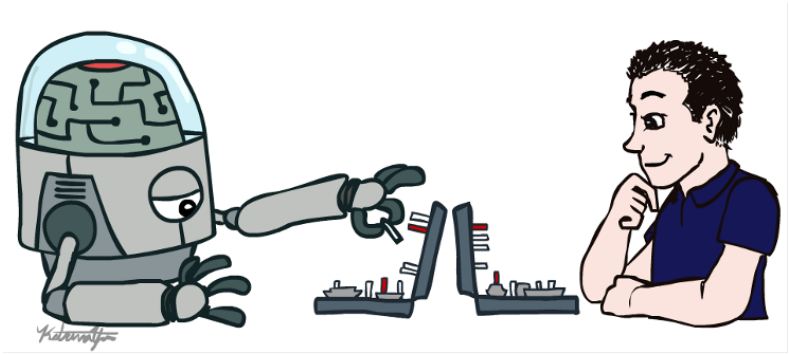

CS 444 – Spring 2021

Dr. Kevin Molloy

Department of Computer Science

James Madison University

JMU JAMES MADISON UNIVERSITY

Some figures and inspiration from
Zoran Duric, Amarda Shehu, Dan Klein, and Pieter Abbeel

Meet and Greet

Who is this person?

- Grew up in Newport News. Last 20 years in Northern Virginia
- Got my PhD in 2015 in computer science with a focus on robotics, artificial intelligence and structural biology (proteins)
- Work/lived in southern France (Toulouse) for ≈ 2 years as a research scientist
- My 3rd year at JMU

Course Information

- Communication
  - Announcements on webpage and Canvas
  - Questions? Discussion on **Piazza**
  - Email: molloykp@jmu.edu
  - Office hours are posted on Canvas.

Course Delivery:

- Lectures: zoom (see canvas for link)
- Review of homework via Youtube channel
- Autograded homework and projects via https://autolab.cs.jmu.edu

- Course Assignment and Topics
  - Class calendar is available at https://w3.cs.jmu.edu/molloykp/cs444/cs444_2021Spring/cs444_Calendar.php
  - All due dates are in Canvas

- Programming Assignments PAs
  - Between 6 and 8
  - Python Programming Language
  - Students will be expected to present their findings/results to the class (at least once in the semester). This may be done live or via video.
  - PA 0 due next Tuesday

Mastery Learning Model

- Quiz each week on material
- If you score well, your done.
- If you are not happy with your score, you can take another quiz on the subject the following week.
- Starting around week 5 or 6, the quiz will consist of the current material plus material from 5 or 6 weeks ago.
- Approximately 12 of these quizzes. No exams.

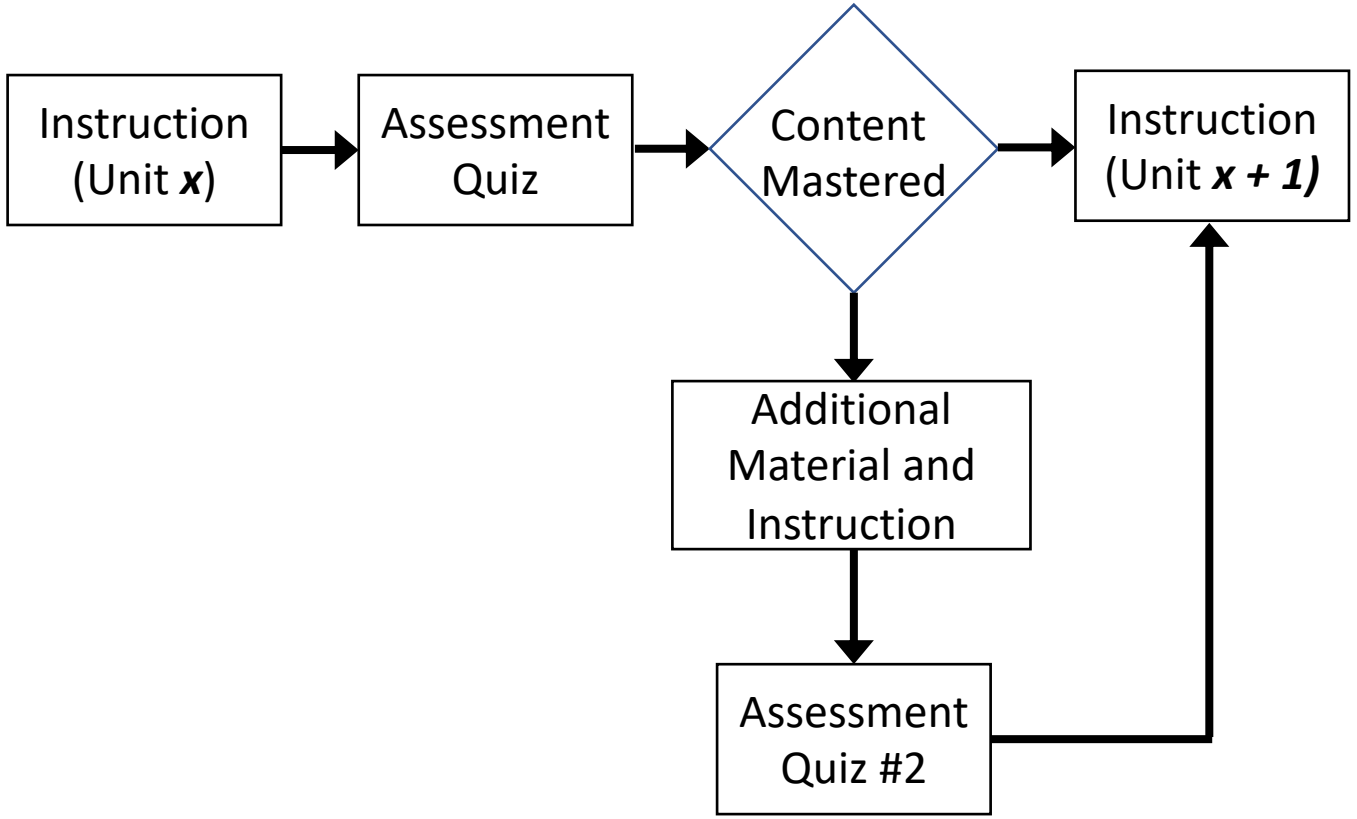

Course Information

- Prerequisites:
  - We will be using some basic statistics and linear algebra. Review material will be provided.
  - We will use a little calculus in this book for analysis.
- Textbook:
  - Artificial Intelligence, Forth Edition by Stuart Russell and Peter Norvig
  - The lecture slides and other reading material will be posted.
- Grading

Task	Number	Percent
Weekly Quizzes	12	50%
Programming Assignments	6 to 8	35%
Homework, Labs, Participation	6 to 8	15%

Why Take This Class?

- You want to learn more about AI
- You want to get an exciting job

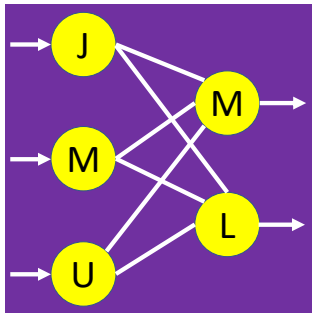

CS 445 Machine Learning

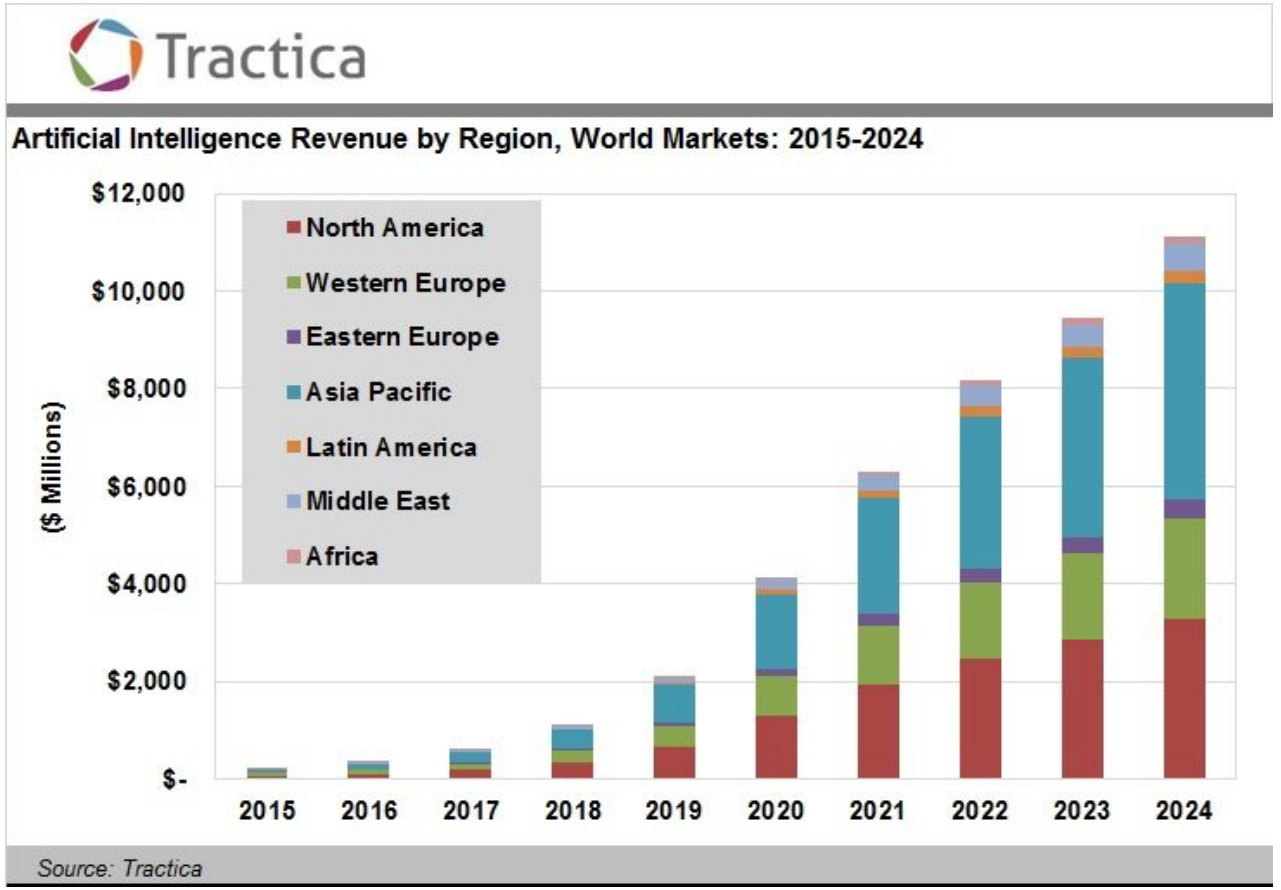

JMU AI and Machine Learning Seminar Series

Third Thursday of Each Month (4:15 – 5:00 pm)

Important This Week

- **Register on Autolab** (https://autolab.cs.jmu.edu). You will need to the use VPN if connecting to Autolab from off Campus (VPN instructions are on the resources page on the class web site).

- **Signup on Piazza.**

- **PA 0** is out and due next Tuesday, January 26th.

- **Quiz 0** will be published on Canvas after class on Thursday and will be due Friday, January 22nd @ 5:00 pm. The quiz will be timed (20 minutes to complete once started). Topics for Quiz 0 are posted on the class website.

Objectives For Today

- What is AI?
- What can AI do?
- What are we going to do in this course?
- Introduction to AI agents

What is AI?

Sci-Fi tells us AI:

- can be nice (R2D2, C3PO)
- can be scary (terminator)

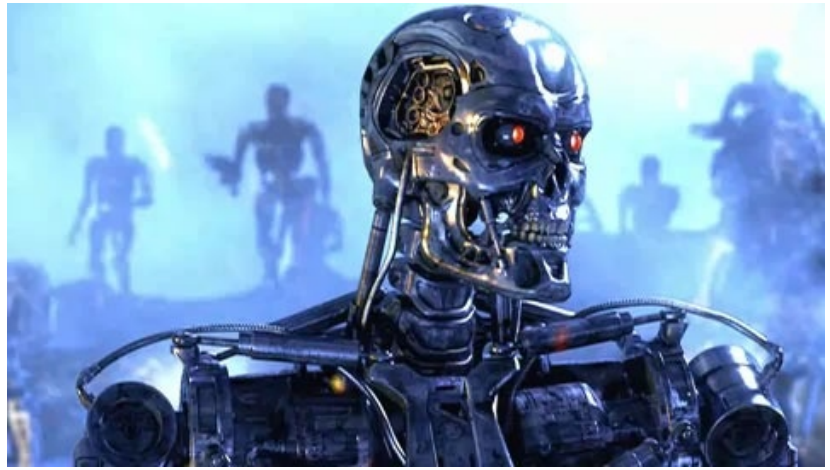

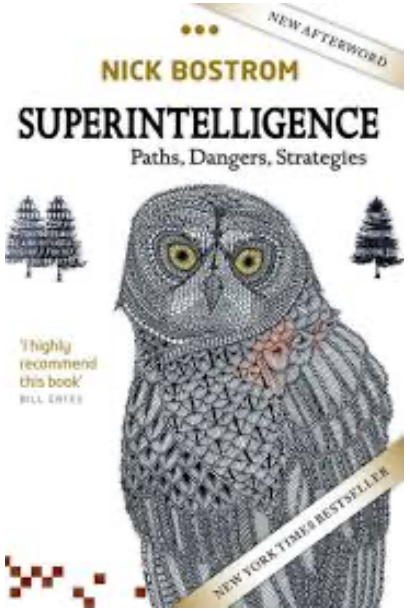

I am reading this book now,
and so far, it is good.

What is Intelligence?

Are we intelligent? Are animals intelligent?

Take a few moments and write down your definition of intelligence

My definition

A computational model/process that:

- solve a problem that does not have an efficient algorithm and requires ”intuition”
- Reasons using facts (and allows new facts to be learned)
- Learns a process by given a set of examples (this is a specialty of AI known as machine learning or statistical inference).

What is AI?

Science of making machines that:

What is AI?

Science of making machines that:

Think like people

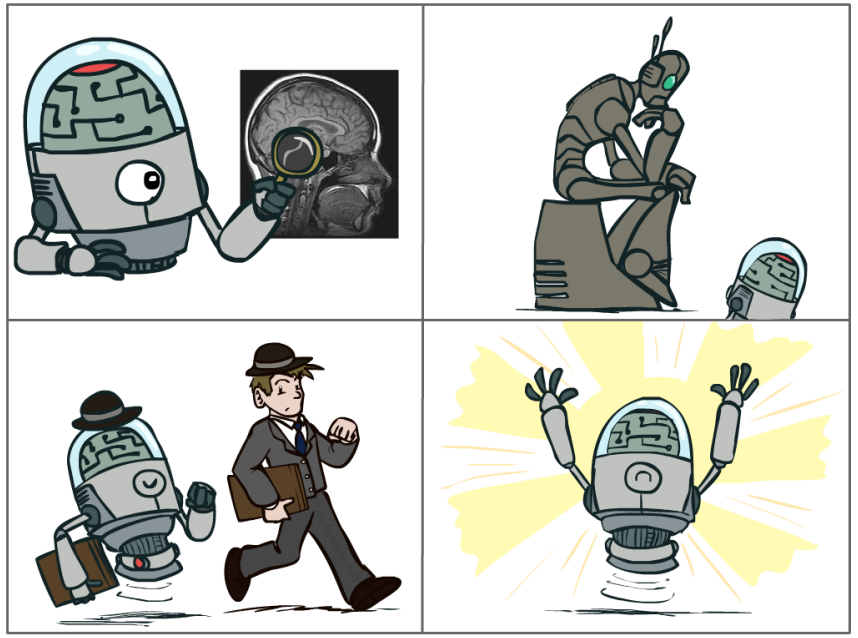

Think rationally

Act like people

Act rationally

Rational Decisions? What are these?

We use the term **rational** in a very specific, technical way:

- Rational: maximally achieving pre-defined goals
- Rationality only concerns what decisions are made (not the thought process behind how the decisions are made)
- Goals are expressed in terms of the **utility** of outcomes
- Being rational means **maximizing your expected utility**

Might think of AI as being the study of **computational rationality**[1]

[1] term from Peter Abbeel

Quantifying Artificial Intelligence – The Turing Test

To be intelligent is to act humanly. – Alan Turing

Turing Test

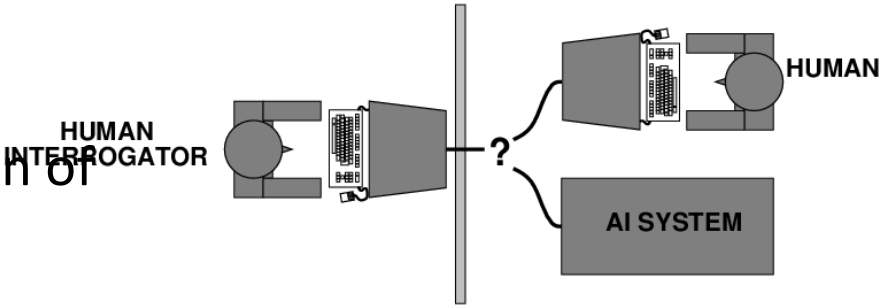

- Turing focused on "acting like a human" as an operational definition of artificial intelligence: Turing (1950). "Computing machinery and intelligence"
- Proposed the "imitation" game as a test for a hidden intelligent artifact who could be fed visual and material information
- **Objective**: fool a human 30% of the time in a 5-minute test

This test introduced major components of AI: knowledge, reasoning, language, understanding, learning (computer vision, robotics). Turing predicted this would be reached by the year 2000.

- Problem: Turing test is not reproducible, informative/constructive, or amenable to mathematic analysis.
- Weak vs Strong AI argument: One can simulate intelligence but not possess it.

Our Brains

- Human brains are very good at making rational decisions, but not perfect
- Brains are not as modular as software, so hard to reverse engineer
- "Brains are to intelligence as wings are to flight"

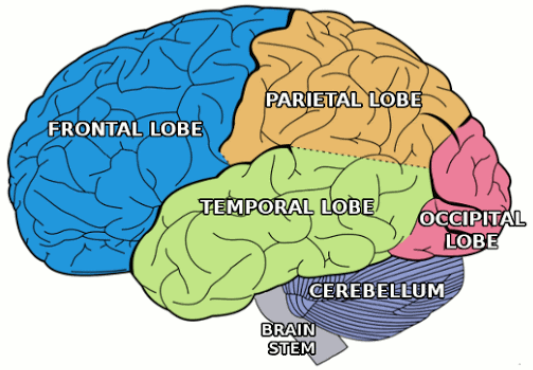

	Supercomputer	Personal Computer	Fly	Human Brain
Computational Units	10^6 GPUs and CPUs 10^{15} transistors	8 CPU cores 10^{10} transistors	10^5 neurons	10^6 columns 10^{11} neurons
Storage Units	10^{16} bytes RAM 10^{17} bytes disk	10^{10} bytes RAM 10^{12} bytes disk	10^7 synapses	10^{11} neurons 10^{14} synapses
Cycle time	10^{-9} secs	10^{-9} secs	10^{-3} secs	10^{-3} secs
Operations/sec	10^{18}	10^{10}	10^{10}	10^{17}

If each synapse is 1 FLOP (fire/not fire once per second),

Then human brain requires 10^{15} flops = 1 petaflop

100,000 CPUs

Amazon EC2 costs $5,000 per hour

Summarized 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, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
  - 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—: Statistical approaches
  - Resurgence of probability, focus on uncertainty
  - General increase in technical depth
  - Agents and learning systems… “AI Spring”?
- 2021—: Where are we now?

What Can AI Do?

Quiz: Which of the following can be d

- ✓ Play a decent game of Jeopardy?
- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mounta
- ? Drive safely along Reservoir Street?
- ✓ Buy a week's worth of groceries on
- ✗ Buy a week's worth of groceries at
- ? Discover and prove a new mathema
- ✗ Converse successfully with another
- ? Perform a surgical operation?
- ✓ Put away the dishes and fold the lau
- ✓ Translate spoken Chinese into spoke
- ✗ Write an intentionally funny story?

Course Topics

Part I: Making decisions

- Fast search and planning
- Constraint satisfaction
- Adversarial and Uncertain search

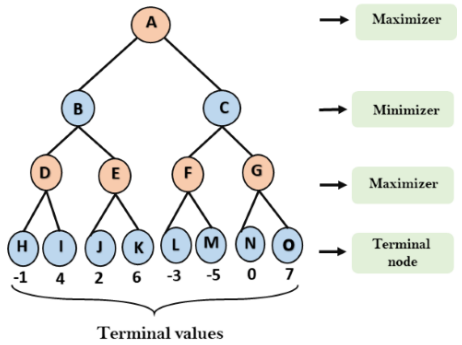

Part II: Representing Knowledge and Logic

- Propositional Logic and representation
- First-order logic and reasoning

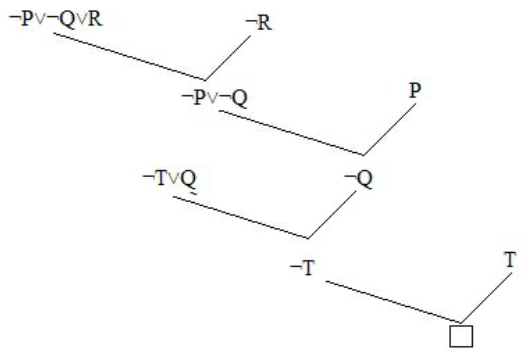

Part III: Reasoning under Uncertainty

- Markov Models
- Bayes nets

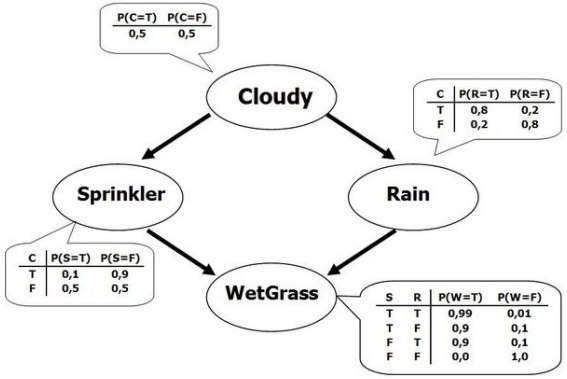

Designing Rational Agents

- An **agent** is an entity that *perceives* and *acts*.
- A **rational agent** selects actions that maximize its (expected) **utility** given the percepts sequence to date.
- Characteristics of the **percepts, environment,** and **action space** dictate techniques for selecting rational actions (PEAS)
- **This course** is about:
  - General AI techniques for a variety of problem types
  - Learning to recognize when and how a new problem can be solved with an existing technique

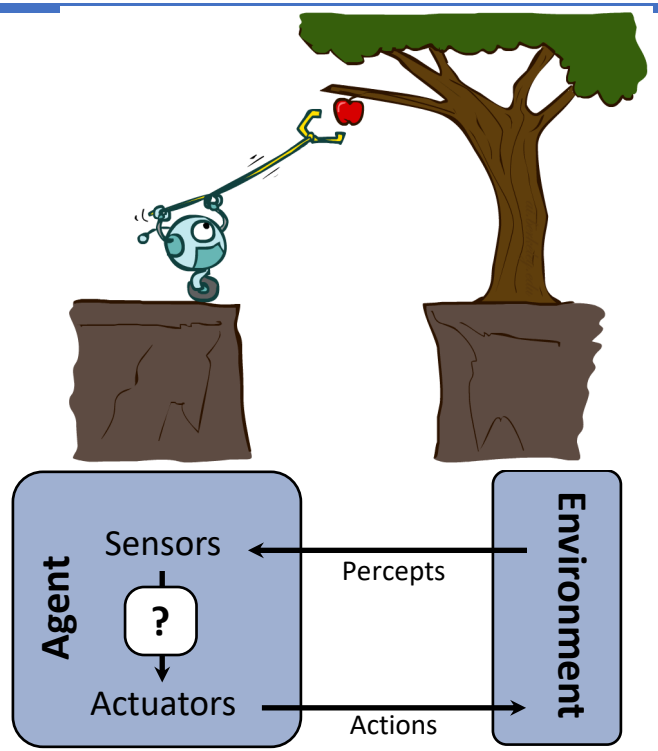

Rational ≠ omniscient
- Percepts may not supply all relevant information

Rational ≠ clairvoyant
- Action outcomes may not be as expected

Hence, rational does not always equal successful.

Agents and Environments

Agents include humans, robots, softbots, thermostats, etc.

The agent function maps from percept histories to actions: $f : \mathcal{P}^* \rightarrow \mathcal{A}$

The agent program runs on the physical architecture to produce f.

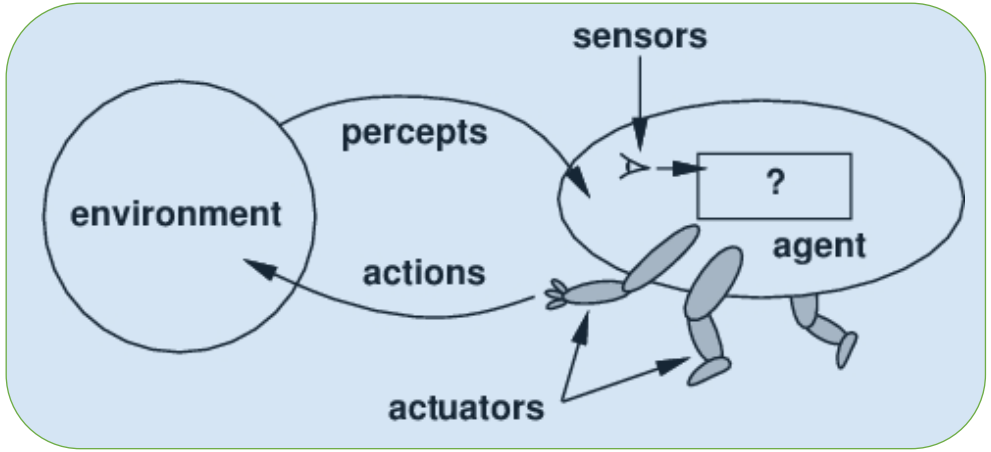

Environment Types

Do the agent's sensors give complete information (relevant to the choice of action) about the estate of the environment at each point in time?

- Fully vs. partially-observable

Does the agent operate in an environment with other agents?

- Single vs. multi-agent (competitive, cooperative)

Is the next state of the environment complete determined by the current state and agent action?

- Episodic vs. sequential

Can the environment change while the agent is deliberating?

- Static vs dynamic

What is the domain of values for variables racking environment state, agent state, and time?

- Discrete vs. continuous

Does the agent know outcomes of all its actions?

- Known vs unknown

Vacuum-cleaner World and Agent Types

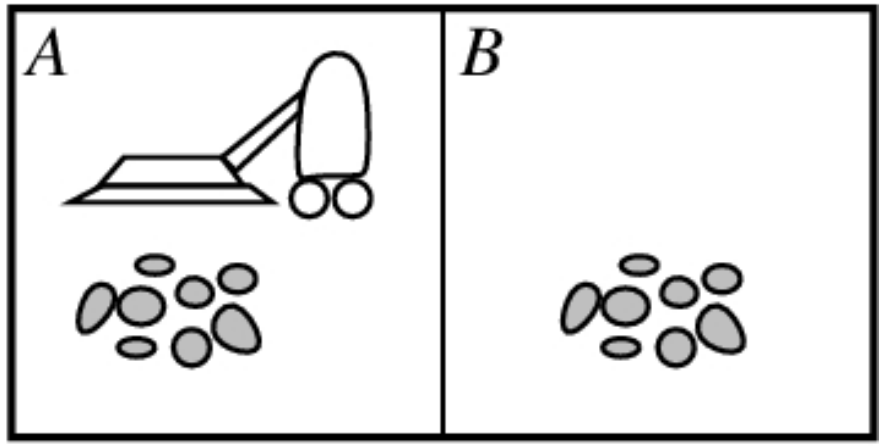

Percepts: location and contents, e.g., [A, Dirty]

Action: *Left, Right, Suck (remove the dirt) , NoOp*

Four basic types of agents:

- Simple reflex agents
- Reflex agents w/state
- Goal-based agents
- Utility-based agents

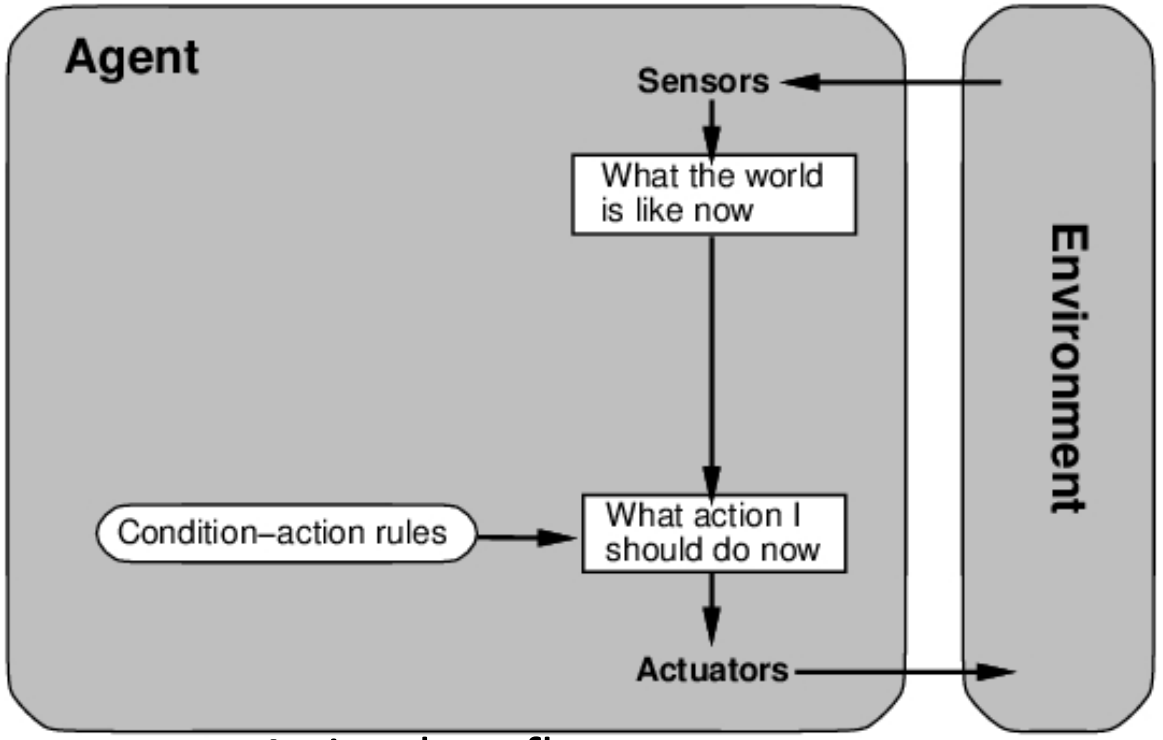

A simple reflex agent

Reflex Agent Example

```
function REFLEX-VACUUM-AGENT ([location, status])
returns an action
        if status = Dirty then return Suck
        else if location = A then return Right
        else if location = B then return Left
```

Can a reflex agent be rational?

A rational agent:

chooses which ever action maximizes the expected value of the performance measure given the percept sequence to date.

Depends, on the performance measure:

- 1 pt for each clean square in each time step
- Geography is known *a priori*
- Agent correctly perceives its location and dirt, and the cleaning mechanism works 100% of the time.

Learning-based agents

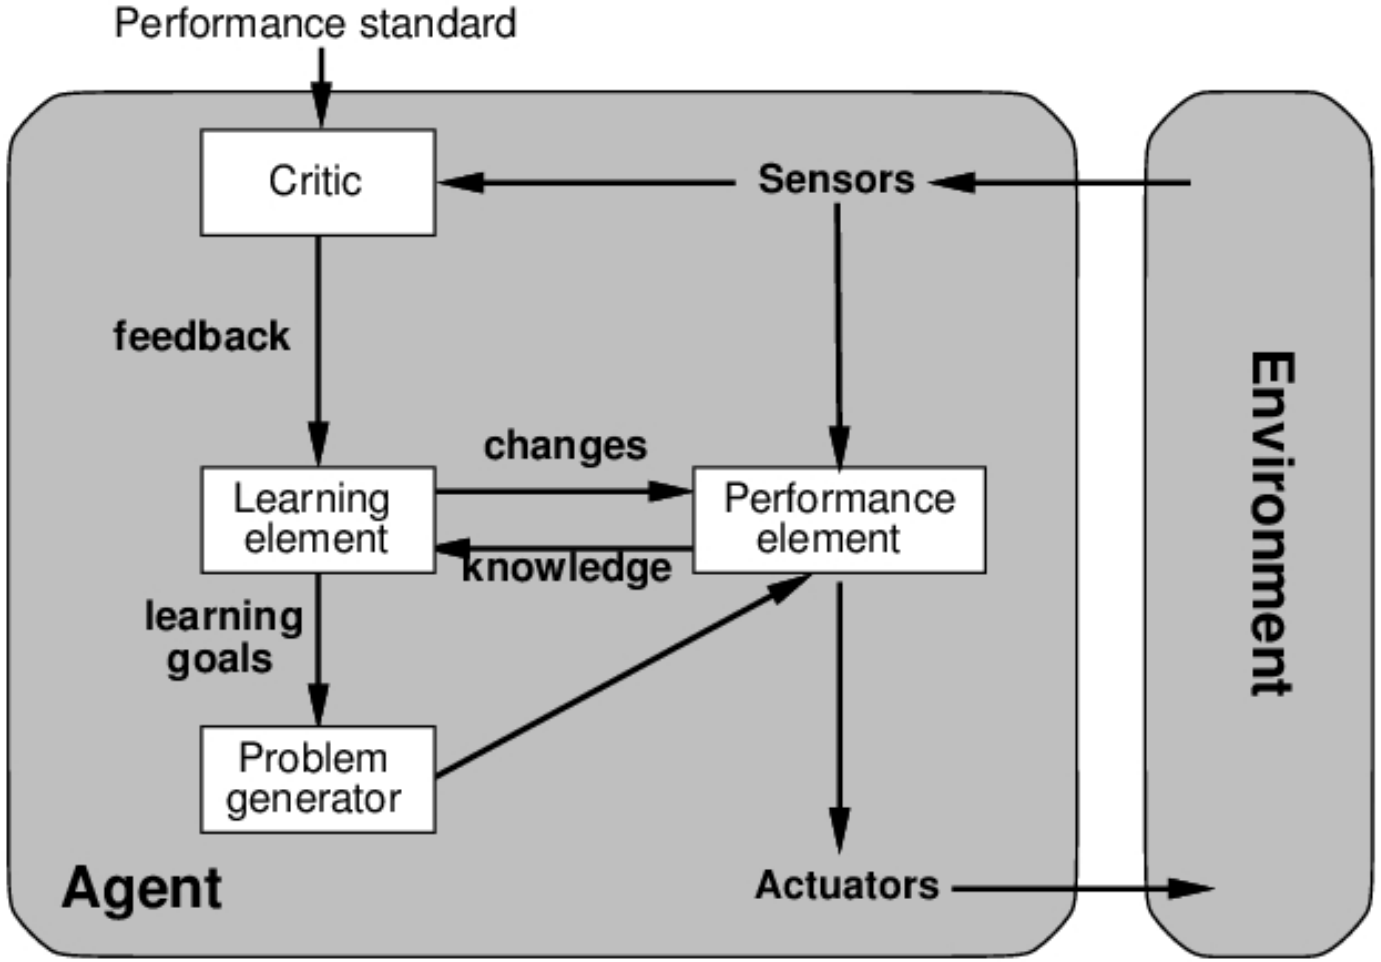

Learning-based agents:

- Performance elements was the "entire" agent previously (accepted sensor input and commanded the actuators).
- Learning elements gets feedback from the critic and changes the performance element.
- Problem generator is the "lets try this and see what happens". Let's explore the impact of other actions.
- Deepmind's AlphaGo and other recent breakthroughs use these types of agents.

Important This Week

- **Register on Autolab** (https://autolab.cs.jmu.edu). You will need to the use VPN if connecting to Autolab from off Campus (VPN instructions are on the resources page on the class web site).

- **Signup on Piazza.**

- **PA 0** is out and due next Tuesday, January 26th. **HW 0** is due this Thursday (January 21st).

- **Quiz 0** will be published on Canvas after class on Thursday and will be due Friday, January 22nd @ 5:00 pm. The quiz will be timed (20 minutes to complete once started). Topics for Quiz 0 are posted on the class website.