Introduction to Artificial Intelligence (AI)

Computer Science cpsc322, Lecture 1

January, 4, 2010

People

Instructor

- **Giuseppe Carenini** (carenini@cs.ubc.ca; office CICSR 129)

Teaching Assistants

MSc

- **Hammad Ali** hammada@cs.ubc.ca

PhD

- **Kenneth Alton** kalton@cs.ubc.ca *(will be starting Jan 18)*

PhD

- **Scott Helmer** shelmer@cs.ubc.ca

MSc

- **Sunjeet Singh** sstatla@cs.ubc.ca

Course Essentials(1)

- **Course web-pages:**

  www.cs.ubc.ca/~carenini/TEACHING/CPSC322-10/index.html

WebSearch: Giuseppe Carenini

  - This is where most information about the course will be posted, most handouts (e.g., slides) will be distributed, etc.
  - CHECK IT OFTEN!

- **Lectures**:
  - Cover basic notions and concepts known to be hard
  - I will try to post the slides in advance (by noon).
  - After class, I will post the same slides inked with the notes I have added in class.
  - Each lecture will end with a set of learning goals:

Student can….

Course Essentials(2)

- **Textbook**: *Artificial Intelligence*, 2nd Edition, by Poole, Mackworth. Under development (here at UBC), but almost domne.
  - It's free!
  - It's available electronically

    http://people.cs.ubc.ca/~poole/aibook/
  - We will cover at least Chapters: 1, 3, 4, 5, 6, 8, 9
  - PDF Available on WebCT

Course Essentials(3)

- **WebCT:** used for textbook, discussion board
  - Use the discussion board for questions about assignments, material covered in lecture, etc. That way others can learn from your questions and comments!
  - Use email for private questions (e.g., grade inquiries or health problems).
- **AIspace** : online tools for learning Artificial Intelligence http://aispace.org/
  - Also under development here at UBC!

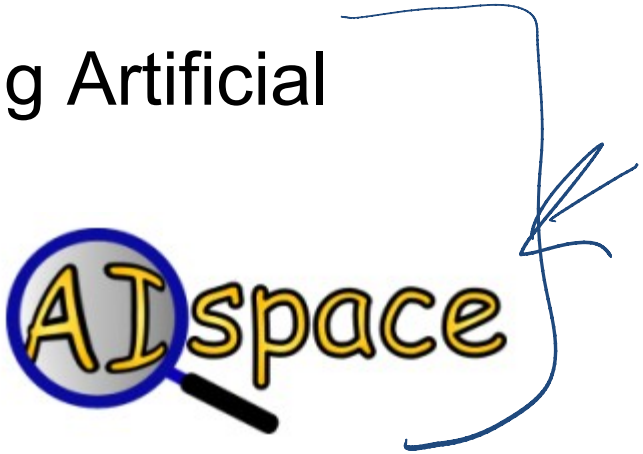

Course Elements

- **Practice Exercises:** 0%
- **Assignments:** 20%
- **Midterm:** 30%
- **Final:** 50%

If your final grade is >= 20% higher than your midterm grade:

- Assignments: 20%
- Midterm: 15% ↓
- Final: 65% ↑

Assignments

- **There will be five assignments in total**
  - Counting “assignment zero”, which you’ll get today
  - They will not necessarily be weighted equally

- **Group work**
  - code questions:
    - ✓ you can work with a partner
    - ✓ always hand in **your own piece of code** (stating who your partner was)
  - written questions:
    - ✓ you may discuss questions with other students
    - ✓ you may not look at or copy each other's written work
    - ✓ you'll be asked to sign an honour code saying you've followed these rules

Assignments: Late Days

- **Hand in by 3PM on due day** (in class or electronically)
- **You get four late days** ☺
  - to allow you the flexibility to manage unexpected issues
  - additional late days will not be granted except under truly exceptional circumstances
- **A day is defined as:** all or part of a 24-hour block of time beginning at 3 PM on the day an assignment is due
- Applicable to assignments 1- 4 **not applicable to assignment 0, midterm, final!**
- if you've used up all your late days, **you lose 20% per day**

Missing Assignments / Midterm / Final

Hopefully late days will cover almost all the reasons you'll be late in submitting assignments.

- However, something more serious like an extended illness may occur ☹

- **For all such cases:** you'll need to provide a note from your doctor, psychiatrist, academic advisor, etc.
- **If you miss:**
  - **an assignment,** your score will be reweighted to exclude that assignment
  - **the midterm,** those grades will be shifted to the final. (Thus, your total grade = 80% final, 20% assignments)
  - **the final,** you'll have to write a make-up final as soon as possible.

How to Get Help?

- Use the course **discussion board** on WebCT for questions on course material (so keep reading from it)

- Go to office hours (newsgroup is NOT a good substitute for this) – times will be finalized next week
  - **Giuseppe:** TBA (CICSR #129)
  - **Hammad** TBA (learning Center)
  - **Ken :** TBA (learning Center)
  - **Scott:** TBA (learning Center)
  - **Sunjeet:** TBA (learning Center)

Can schedule by appointment if you can document a conflict with the official office hours

Getting Help from Other Students? (Plagiarism)

- **It is OK to talk with your classmates about assignments; learning from each other is good**
- **But you must:**
  - Not copy from others (with or without the consent of the authors)
  - Write/present your work completely on your own (code questions exception)

- See **UBC official regulations** on what constitutes plagiarism (pointer in course Web-page)
- Ignorance of the rules will not be a sufficient excuse for breaking them

Getting Help from Other Students? (Plagiarism)

When you are in doubt whether the line is crossed:

- Talk to me or the TA's

Any unjustified cases will be **severely dealt with** **by the Dean's Office** (that's the official procedure)

- My advice: better to skip an assignment than to have "academic misconduct" recorded on your transcript and additional penalties as serious as expulsion from the university!

To Summarize

- All the course logistics are described in the course Webpage

www.cs.ubc.ca/~carenini/TEACHING/CPSC322-10/index.html

WebSearch: Giuseppe Carenini

(And summarized in these slides)

- Make sure you carefully read and understand them!

What is Intelligence?

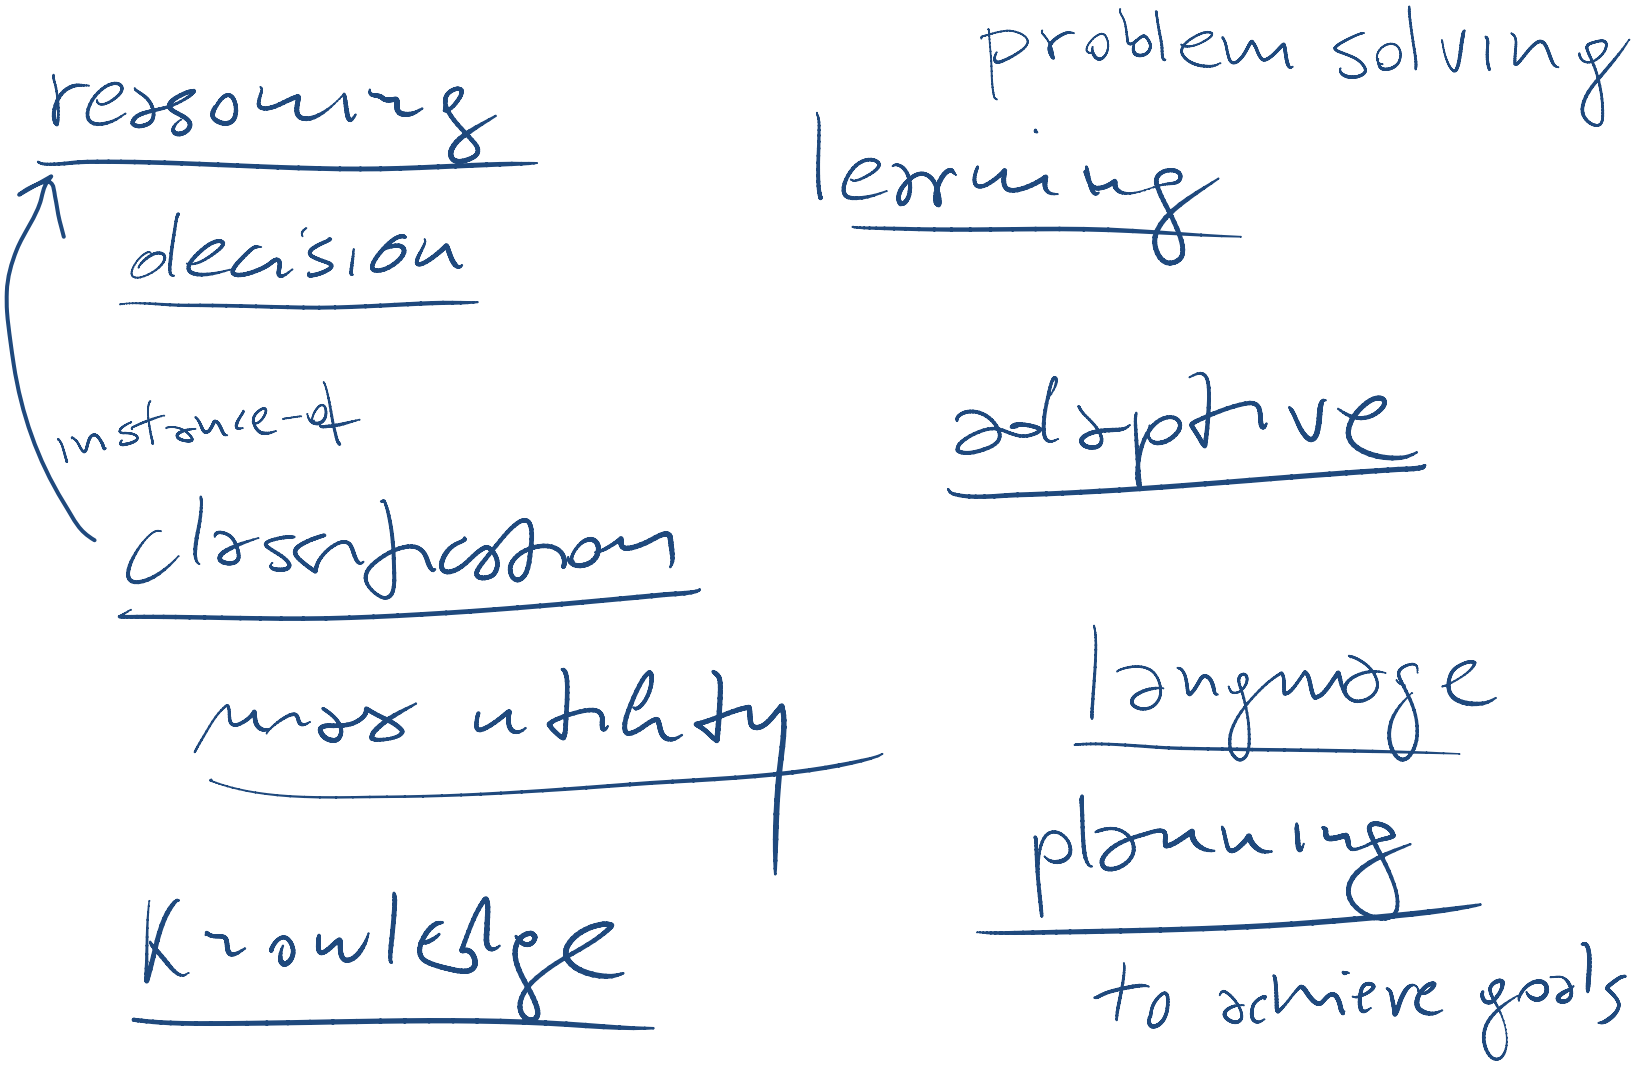

What is Artificial Intelligence?

Two definitions that have been proposed:

- Systems that **think** and **act** like **humans**
- Systems that **think** and **act rationally**

Thinking and Acting Humanly

Model the cognitive functions of human beings

- Humans are our only example of intelligence: we should use that example!

Problems:

- But... humans often think/act in ways that we don't consider intelligent **(why?)**

emotions

cognitive limitations

memory

tired

incorrect

missing knowledge

- And... detailed model of how people's minds operate not yet available

Thinking Rationally

Rationality: an **abstract "ideal" of intelligence**, rather than ``whatever humans think/do''

- Ancient Greeks invented *syllogisms*: argument structures that always yield correct conclusions given correct premises
  - This led to **logic**, and **probabilistic reasoning** which we'll discuss in this course
- But correct sound reasoning is not always enough "to survive" "to be useful"…

Acting (&thinking) Rationally

This course will emphasize a view of AI as building **agents**: artifacts that are able to think and act rationally in their environments

Rationality is **more cleanly defined** than human behavior, so it's a better design objective

(Eg: "intelligent" vacuum cleaner: maximize area cleaned, minimize noise and electricity consumption)

Agents that can answer queries, plan actions and solve complex problems

And when you have a rational agent you can always tweak it to make it irrational!

Why do we need intelligent agents?

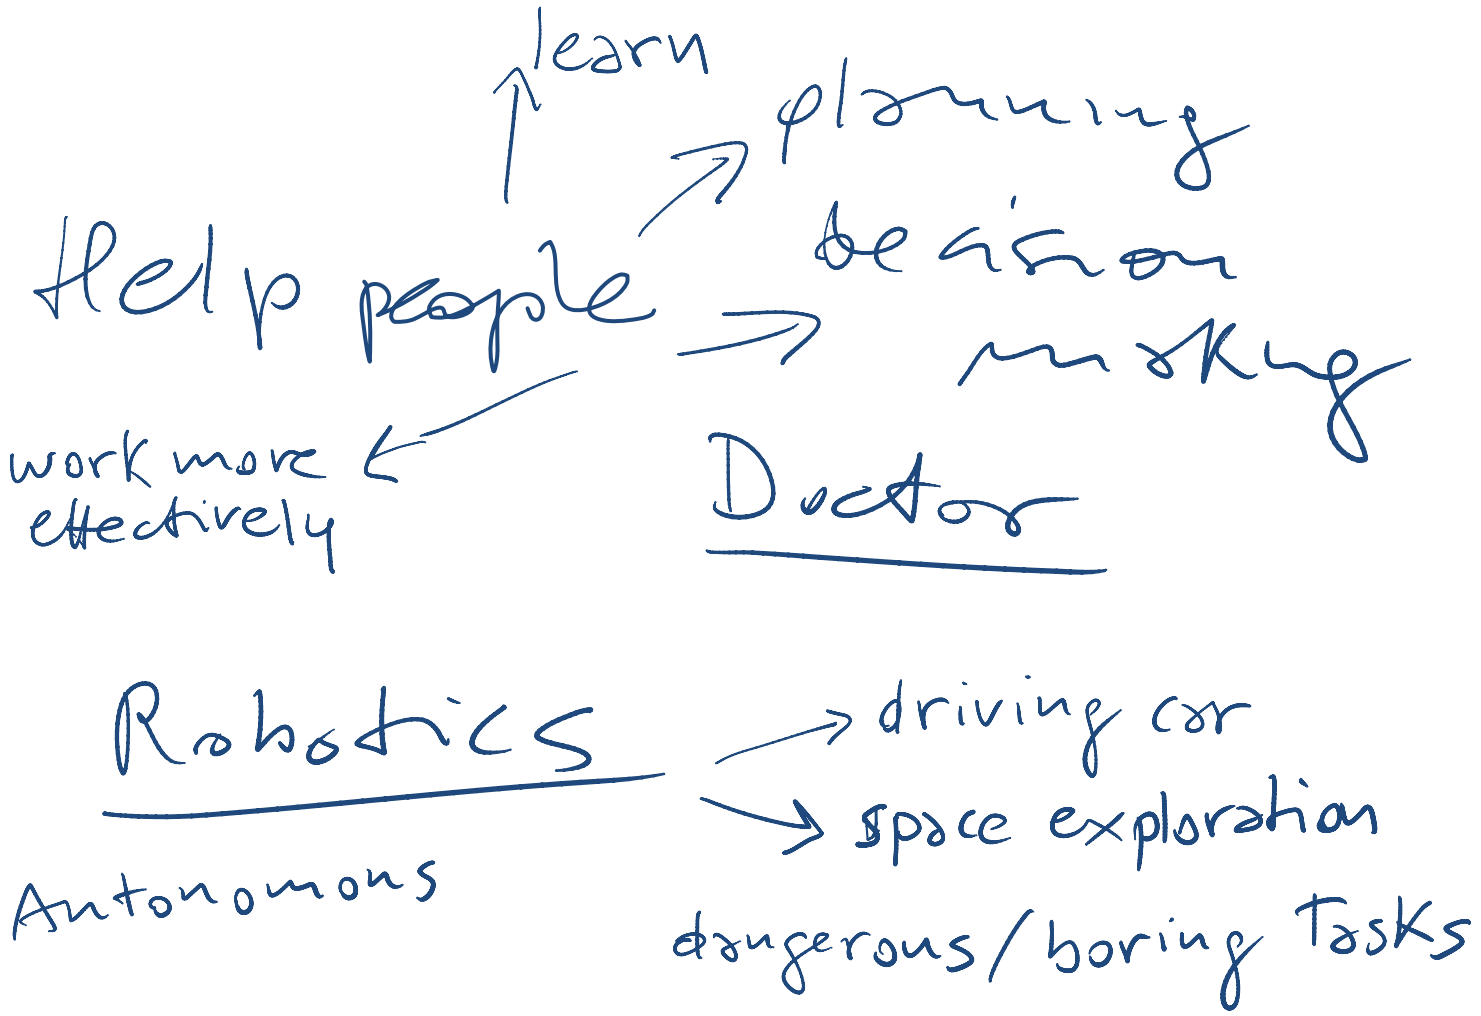

Agents acting in an environment

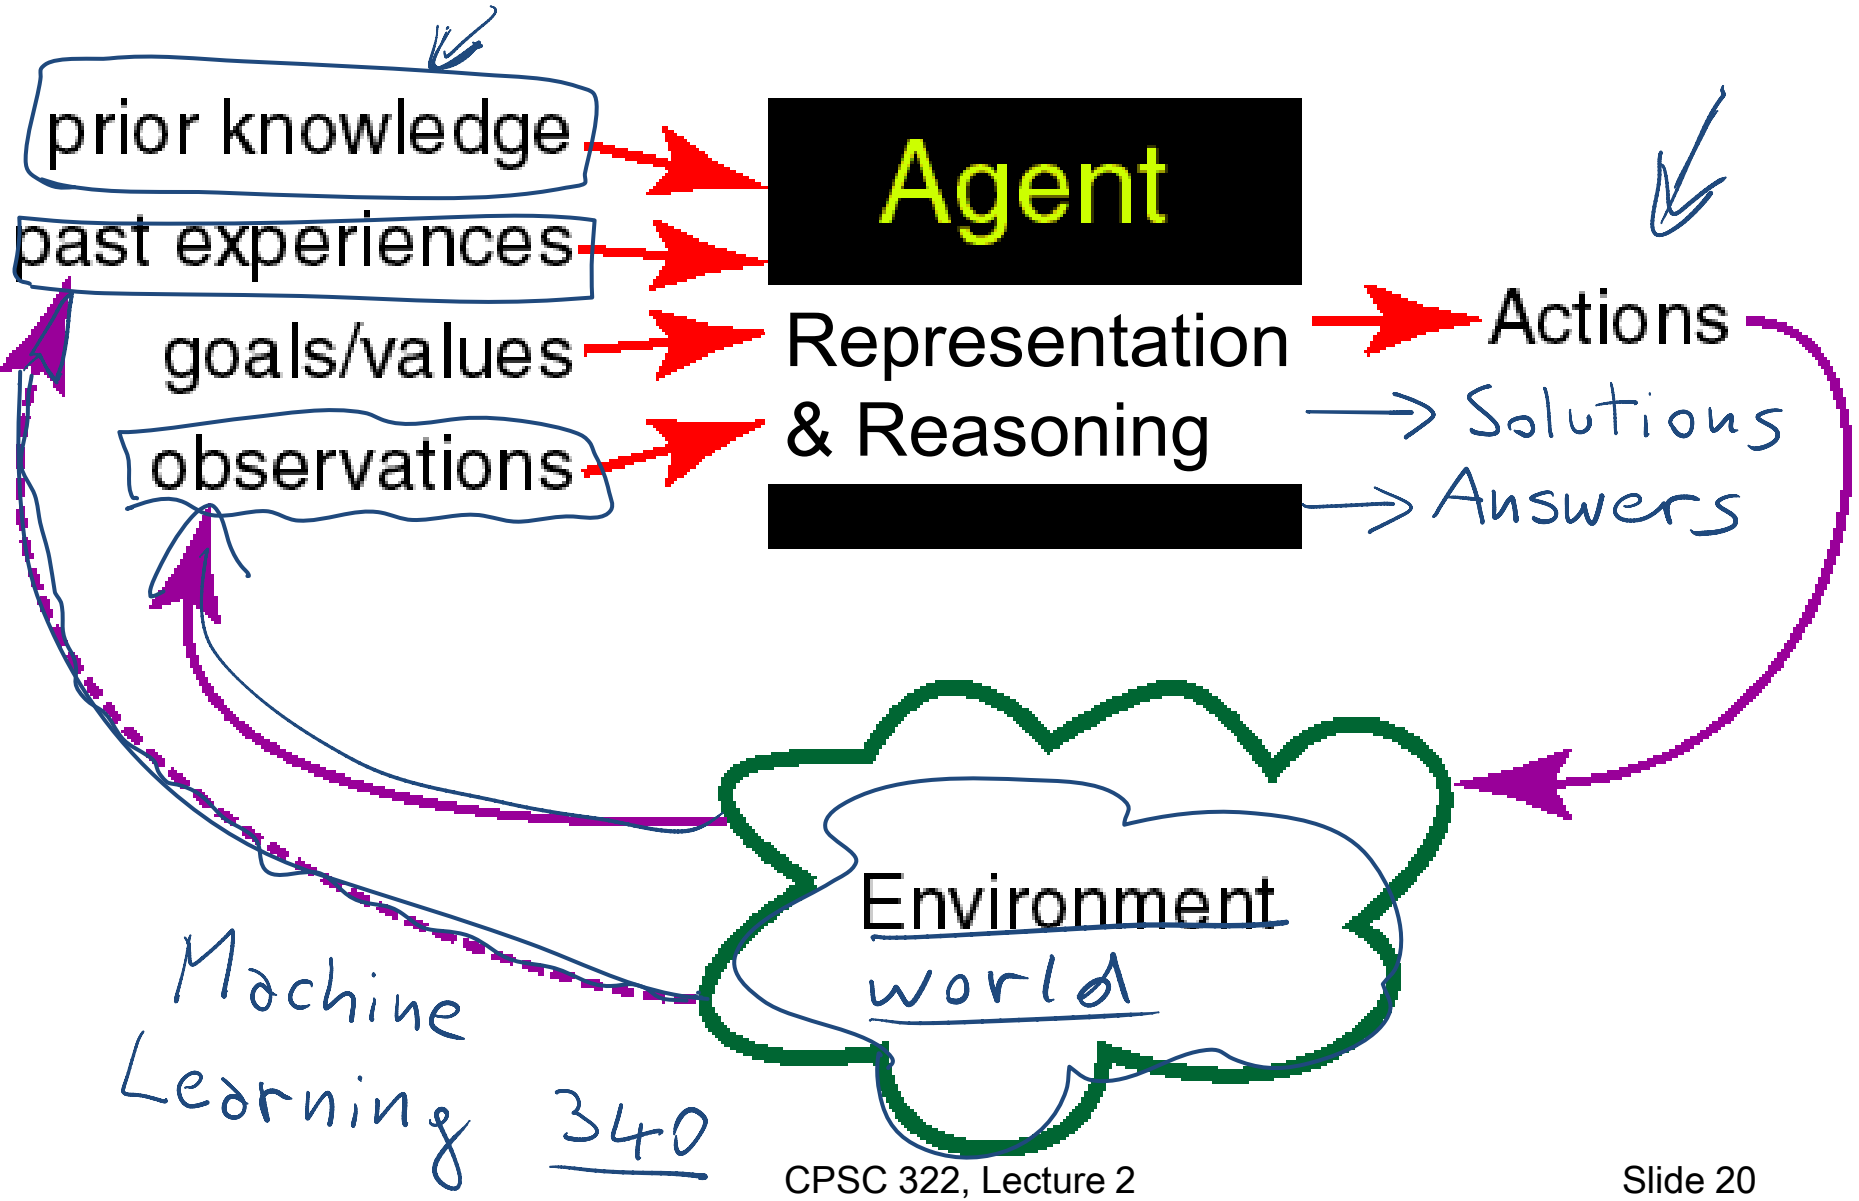

What is an agent?

It has the following characteristics:

- It is situated in some **environment**
  - does not have to be the real world---can be an abstracted electronic environment

Medical test / Eyetracking

- It can make **observations** *(perhaps imperfectly)*
- It is able to **act** *(provide an answer, buy a ticket)*
- It has **goals or preferences** *(possibly of its user)*

realestate advisor

- It may have **prior knowledge or beliefs**, and some way of **updating beliefs** based on new experiences (to reason, to make inferences)

TODO for this week

For Wed: Read Chp 1

For Fri: Assignment 0

- Your first assignment asks you to find **two examples of fielded or experimental AI agents**, and to explain some high-level details about how they work.

- The assignment is available from the **course web page**
- **submit electronically** and **you can't use late days**

Examples

Which of these things is an **agent**,

and why or why not?

- A soccer-playing robot?
- A rock?
- Machine Translator?
- A thermostat?
- A dog?
- A car?

Which of these things is an **intelligent agent**,

and why or why not?

Acting (&thinking) Rationally

This course will emphasize a view of AI as building **agents**: artifacts that are able to think and act rationally in their environments

- they act appropriately given goals and circumstances
- they are **flexible** to changing environments and goals
- they **learn** from experience
- they make appropriate choices given perceptual and computational limitations (sometimes they act without thinking!)
- They **gather information** (if cost less than expected gain)

Acting Humanly

The **Turing Test**

- Don't try to come up with a list of characteristics that computers must satisfy to be considered intelligent
- Instead, use an operational definition: consider it **intelligent when people can't tell a computer apart from other people**

The original test involved typing back and forth; the `**Total Turing Test** includes a video signal to test perception too

- But... is acting just like a person what we really want?
- For example, again, don't people often do things that we **don't** consider intelligent?