



Decision Making for Robots and Autonomous Systems

Fall 2015



George Konidaris
gdk@cs.duke.edu

Admin



Course:

- Tuesdays and Thursdays, 3:05-4:20pm.
- Allen 318
- Two exams (25% each).
- Two assignments (25% each).

Course website:

- <http://www.cs.duke.edu/courses/fall15/compsci590.1/>
- Lecture slides, papers etc. will be uploaded
- Mailing list
- Course content will be decided *at runtime*

Expectations

I expect you to:


- Show up for class (mostly)
- Do the reading (**before class**).

- Ask questions whenever you like:
 - In class
 - By email
 - Make an appointment if you have in-depth questions.

- Please send constructive feedback, at any time.

- THINK about the material, and how it applies to your interests.

Content



Principles and practice of decision-making for autonomous agents (mostly robots).

- Rationality
- Utility theory
- Probabilistic reasoning
- Multi-arm bandits
- Markov decision processes
- Reinforcement Learning
- Partially observable MDPs
- Belief-space learning and planning
- Learning from demonstration
- Inverse RL

AI and Robotics



Agents

Rationality

Robots

Models


Agenthood

AI is about agents:


- “An agent is just something that acts”
- “operate autonomously, perceive their environment, persist over a prolonged time period, adapt to change, and create and pursue goals”

(Russell and Norvig)

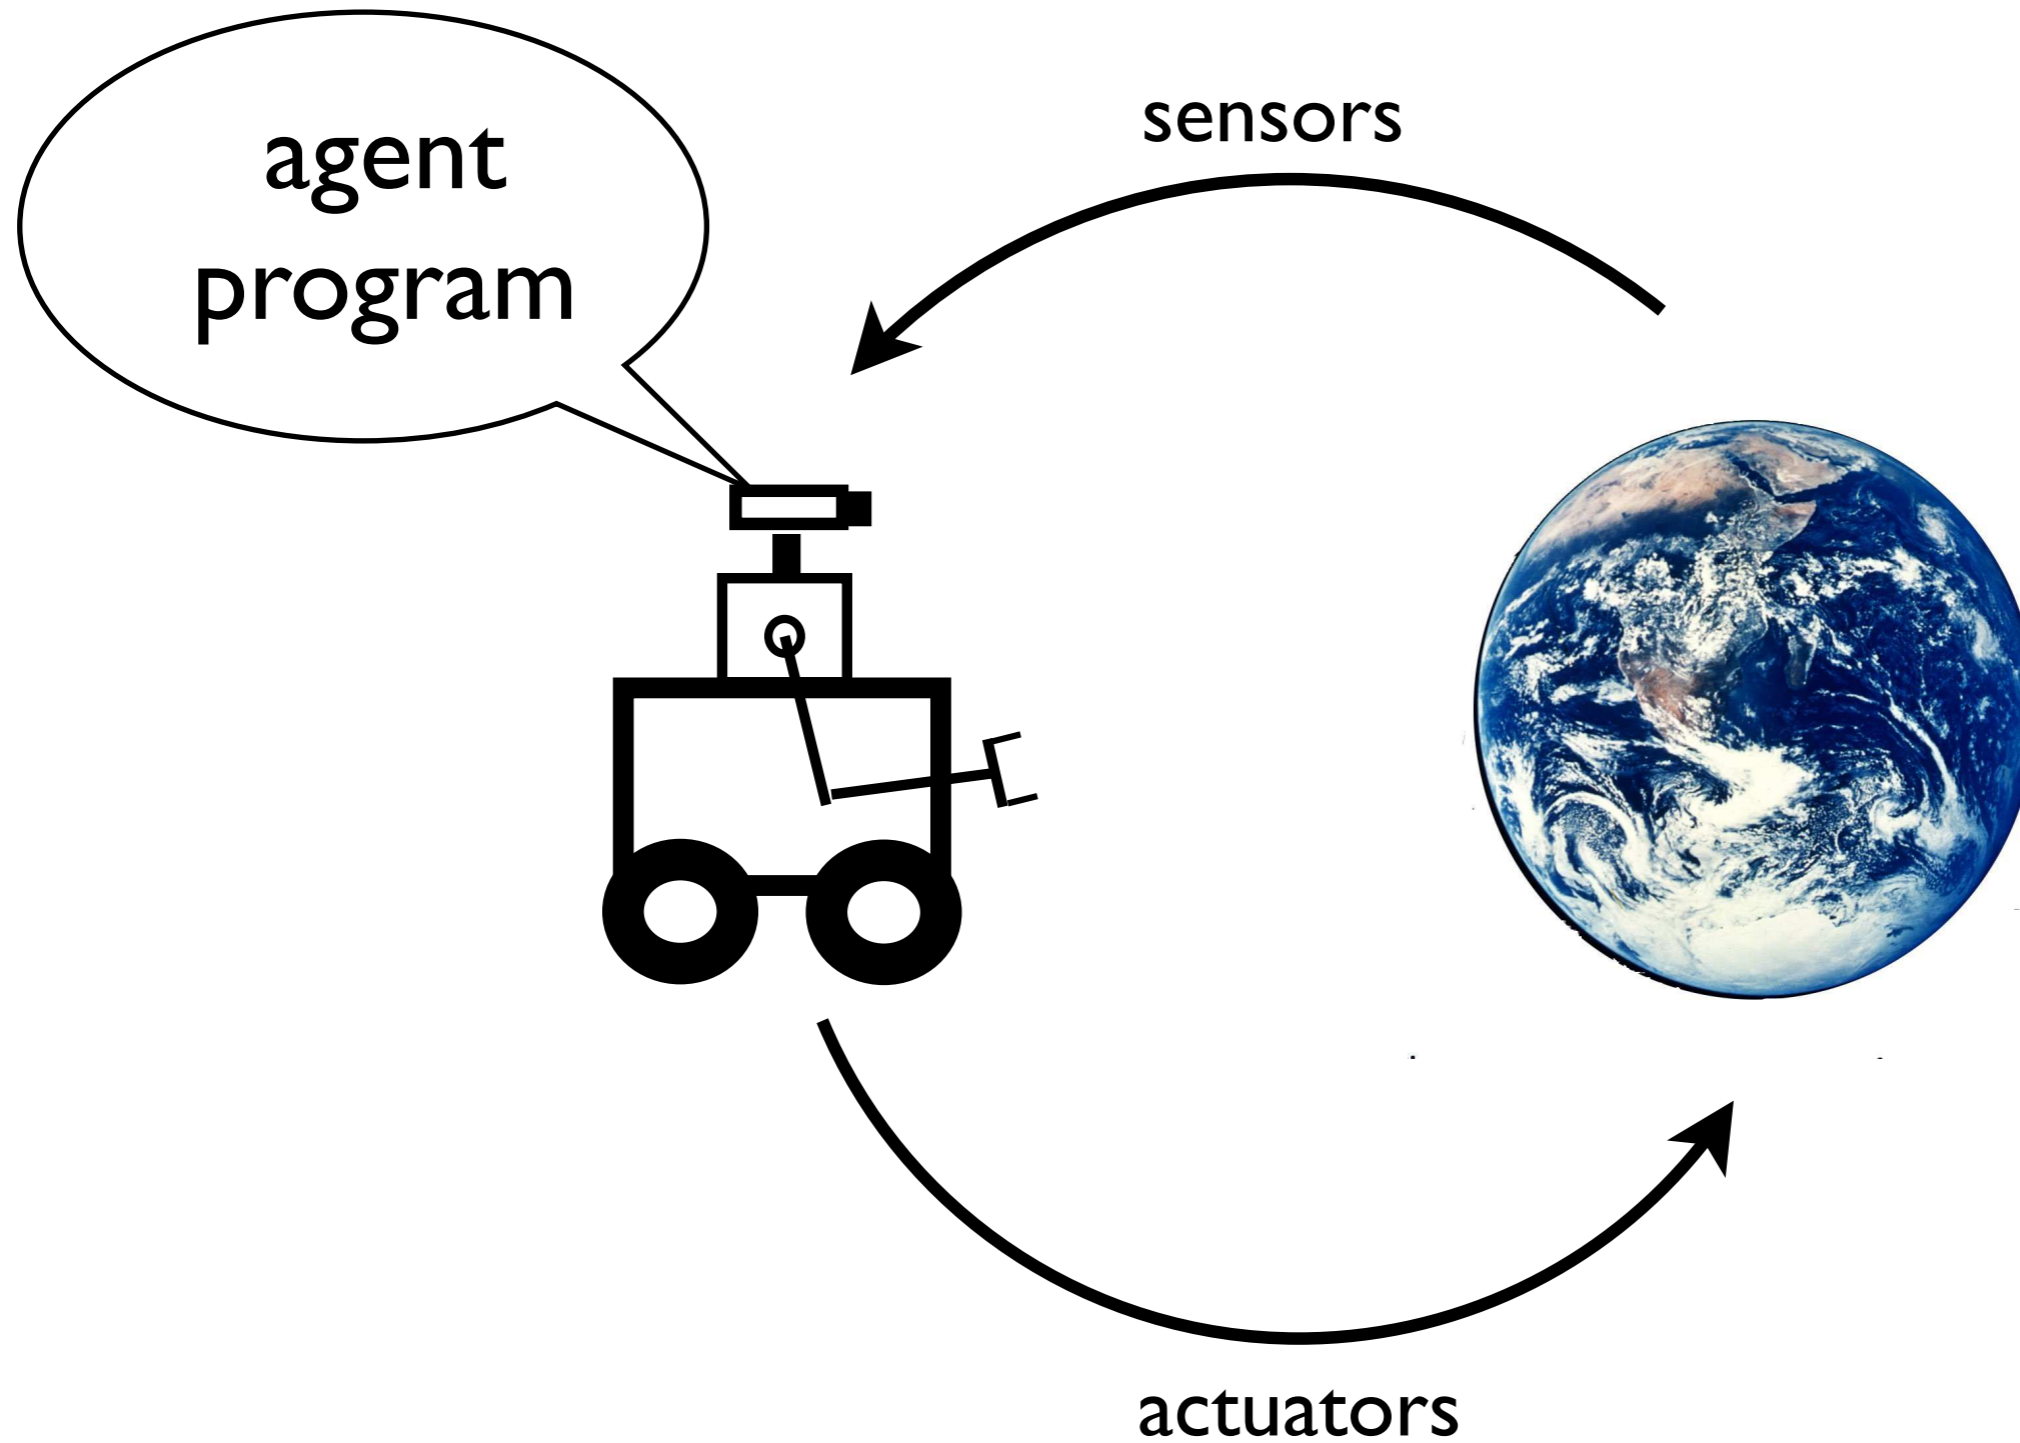
Agenthood



Agenthood



Agenthood



Performance measure.

Rational Agents

AI is about constructing **rational agents**.


- Performance measure = objective function.
- Design can be tricky.



- Rational agents maximize their performance measure.

“For each possible percept sequence, a rational agent should select an action that is expected to maximize its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.” (Russell and Norvig)

Satisficing vs. Optimizing



Herb Simon: *satisficing*

- Threshold level of performance
- “Good enough”

How does this compare with optimizing?

- Still need a performance metric.
- Still trying to improve “score”.
- Don’t insist on achieving optimal performance.

Optimization-based approaches have made major progress in AI in the last 20 years, because they allow us to apply a well-understood set of mathematical tools.

Robots




humans are robots



Central Metaphor



Central Metaphor



Central Metaphor




Robots




robots are *fully realized* agents

Models



Even though robots are *the one true way*, most of this course will be about **models**.

Models




A model is a formal specification of a *class* of problems.

- Not a particular problem.
- Captures (abstractly) the essential components of class.
- *Generalizes across* robots, environments, utility functions.

A model is *not* an algorithm.

- *Harder to get model right than algorithm right.*
- Algorithm often follows from model.



model



robot

Models

Properties of models:

- They are always wrong.
- They are sometimes useful.
- They are never “real”.
 - AI communities often make this mistake. Models are not real. Only robots are real.
- They make assumptions.
 - These assumptions are often wrong.
 - Sometimes we make them anyway.


Models

This course is going to examine a series of models.

For each model, we will:

- Write a **formal definition**.
- Make its **assumptions explicit**.
- Discuss the circumstances in which it is **applicable**.
- Consider its **theoretical properties**.
- Describe **algorithms** for solving it.
- Consider a few “**real-world**” **examples** (usually robots).

Doing Research



When faced with a hard (robot) research problem:


FIRST formalize the problem in the most widely applicable way.

- Capture the essential properties.
- Get the model right.
 - In full knowledge that it is wrong/approximate.
- NEVER be idiosyncratic.
 - This specific problem may never occur again.

THEN gain insight into the structure that the model exposes.

THEN develop/apply a general-purpose algorithm.

Thursday



The course starts in earnest!

1. Utility theory
2. Probabilistic reasoning and uncertainty