Decision Making
in Robots and Autonomous Agents

Introduction

Subramanian Ramamoorthy
School of Informatics

16 January, 2018

Core Concerns of this Course

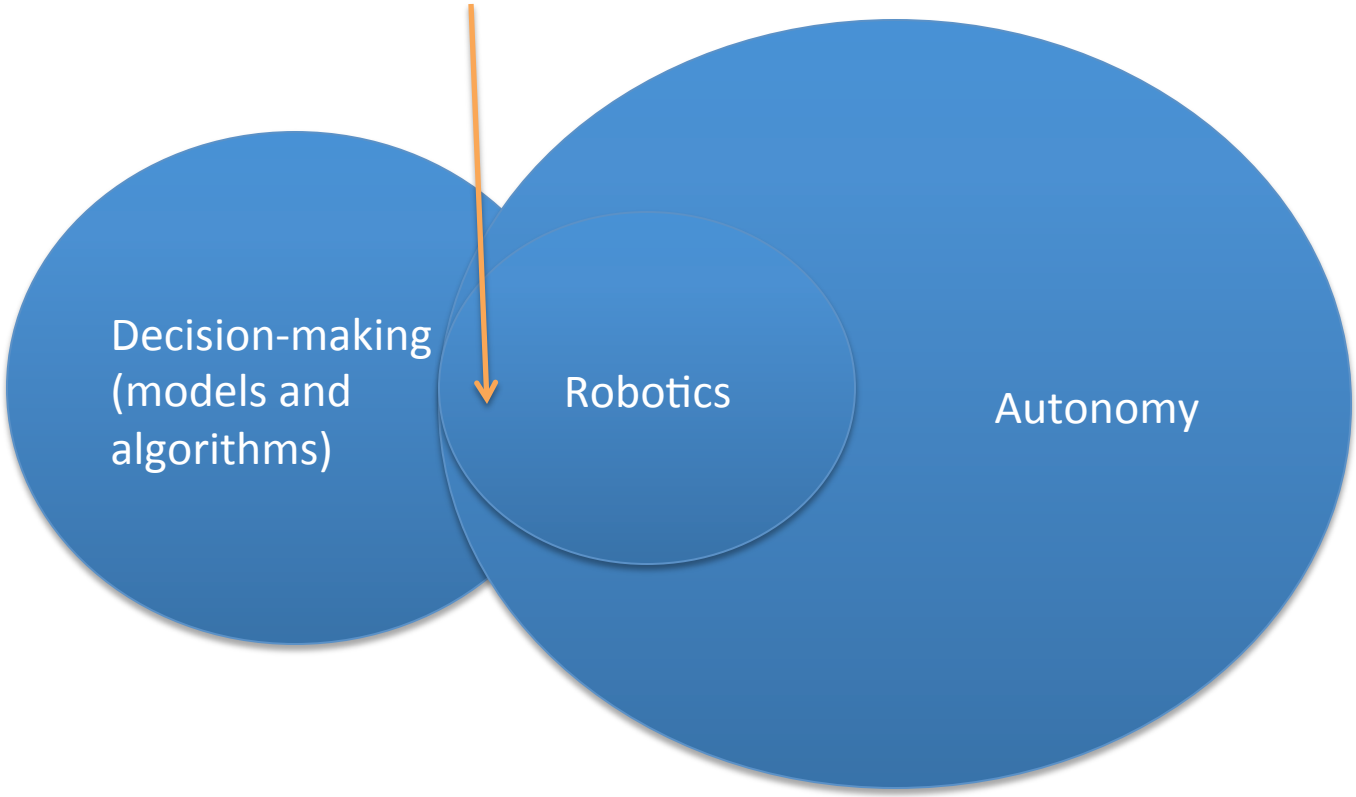

What does each term mean to you?

Early Examples of Autonomy? da Vinci's Mechanical Knight

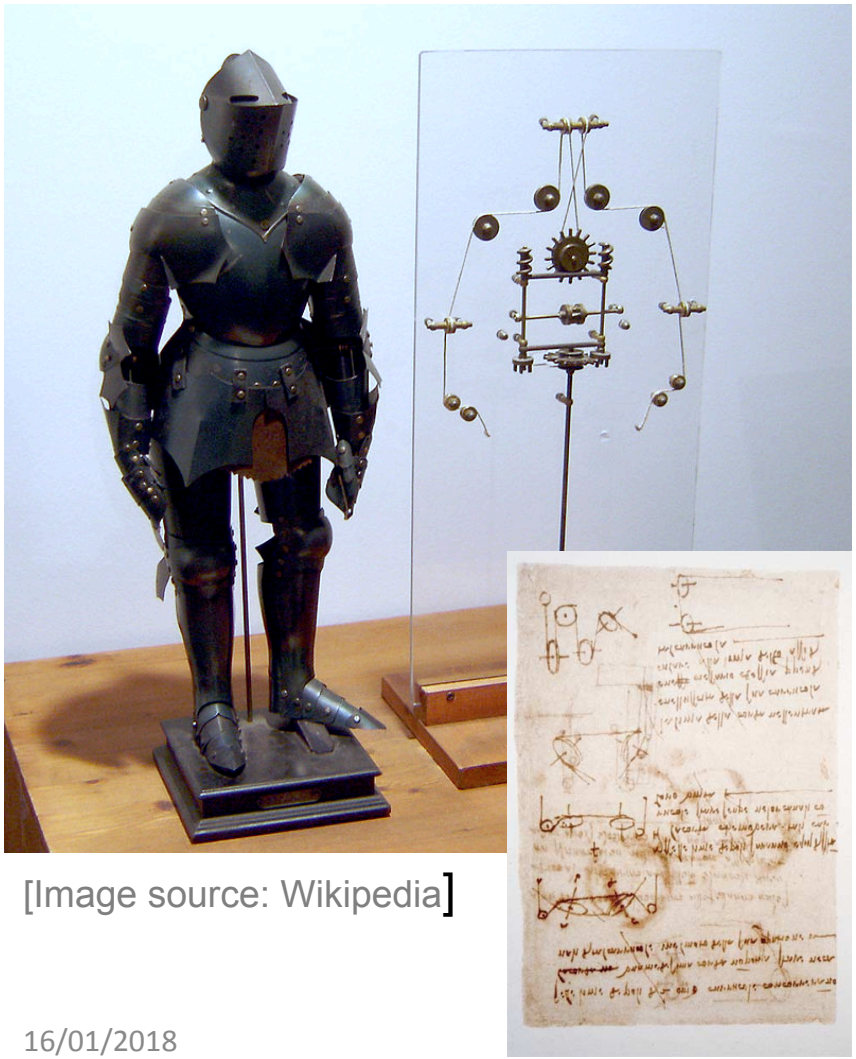

[Image source: Wikipedia]

- Is it a "robot"?
- (What) decisions does this automaton make?

How about this?

How about the Tippe Top?

[Image source: Physics Stack Exchange]

- How would we describe its behaviour?
- Do you know the principle of operation?

- Is this a “robot”?
- This is not a trivial question…

  e.g., passive walkers

[Source: https://www.youtube.com/watch?v=e2Q2Lx8O6Cg]

How about a Marionette?

[Source: https://www.youtube.com/watch?v=bXFPWZSIOs0]

Teleoperation:
"Invisible" Puppet Strings?

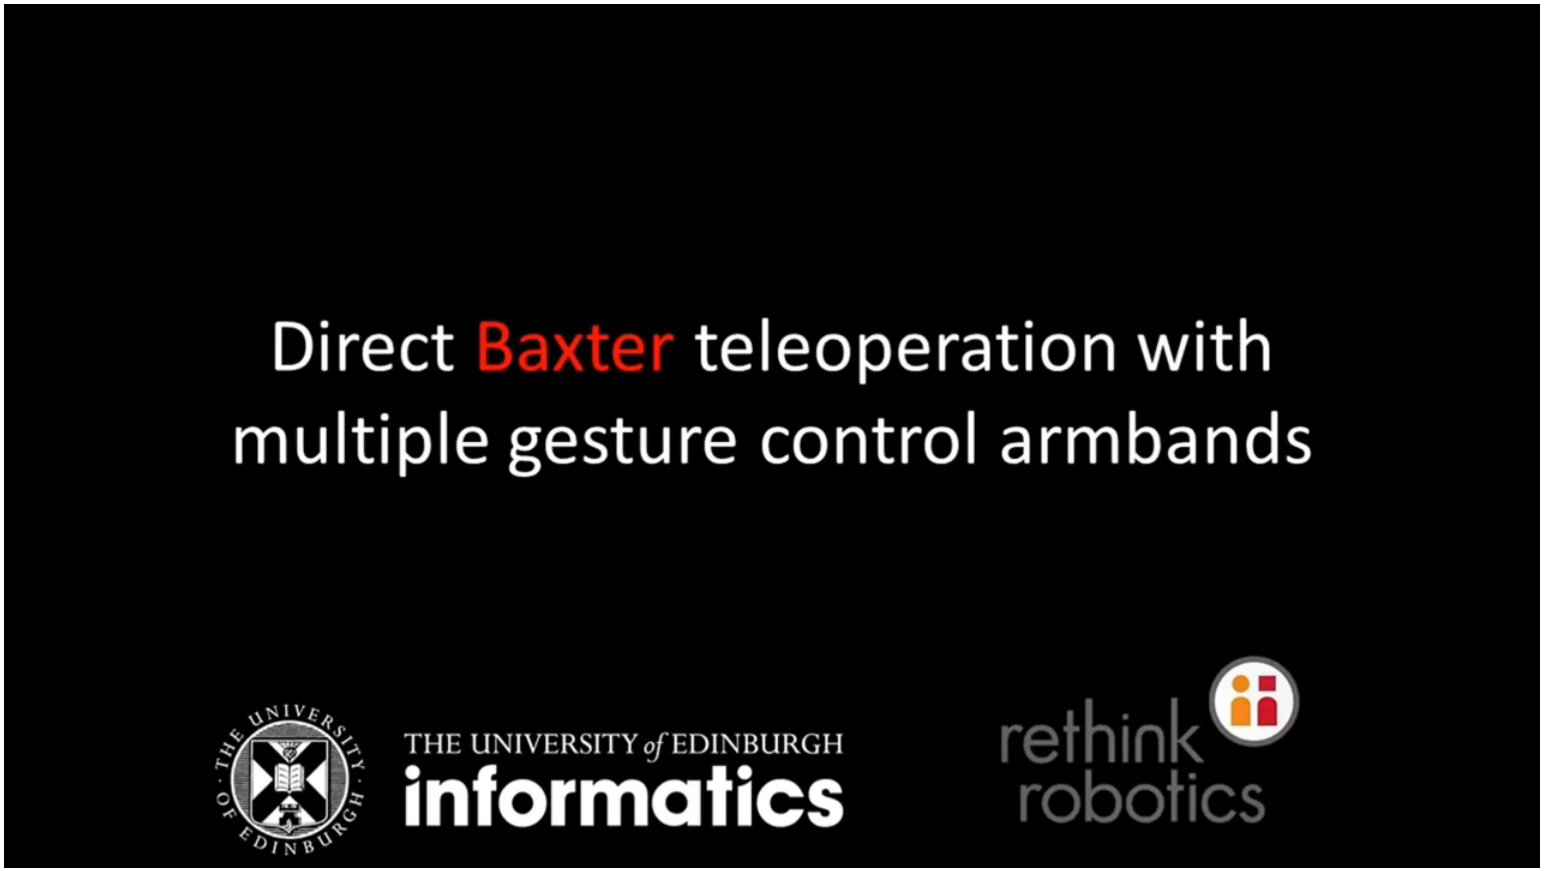

[Source: https://www.youtube.com/watch?v=fSskyIaWkMk]

On Robotic Paradigms

- Questions so far may seem like pedantic nitpicking, but they have been at the heart of discussions regarding *paradigms*
- Paradigm: Philosophy or set of assumptions and/or techniques which characterize an approach to a class or problems
  - Rarely is any one paradigm uniquely best for all problems (bit like cartesian vs. polar coordinates in calculus)
- Robotic paradigms can be described in terms of:
  - Relationship between commonly accepted primitives: SENSE, PLAN and ACT
  - Ways in which sensory data is processed and distributed throughout the system

Robot Primitives in terms of I/O

ROBOT PRIMITIVES	INPUT	OUTPUT
SENSE	Sensor data	Sensed information
PLAN	Information(sensed or cognitive)	Directives
ACT	Sensed information or directives	Actuator commands

The *Hierarchical* Paradigm

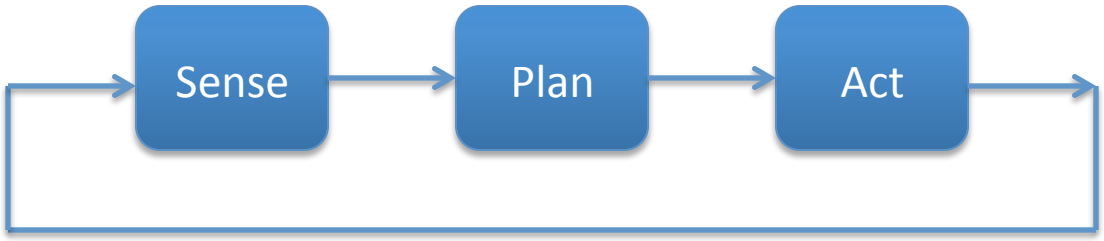

- One of the oldest approaches (1967 – 1990)
- Top down, sensed data is compiled into world model and planner operates on this global model
- Can be hard and brittle due to *closed world assumption* and the so-called *frame problem*

The Hierarchical Paradigm

ROBOT PRIMITIVES	INPUT	OUTPUT
SENSE	Sensor data	Sensed information
PLAN	Information(sensed or cognitive)	Directives
ACT	~~Sensed information or~~ directives	Actuator commands

Example: Shakey and STRIPS

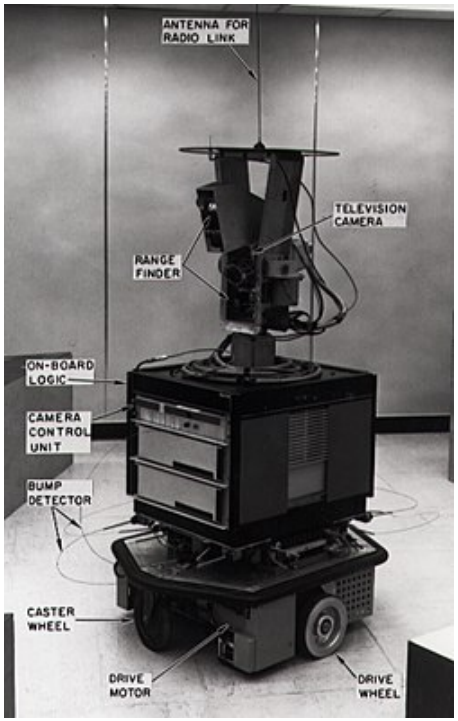

[Source: Wikipedia]

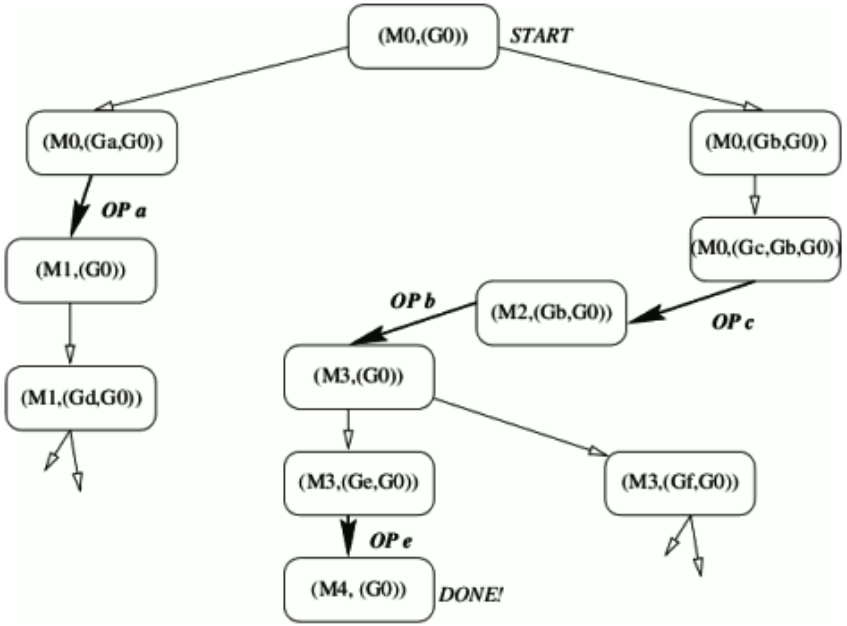

[R.E. Fikes, N.J. Nilsson, STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving, Artificial Intelligence. 2 (3–4): 189–208, 1971.]

The *Reactive* Paradigm

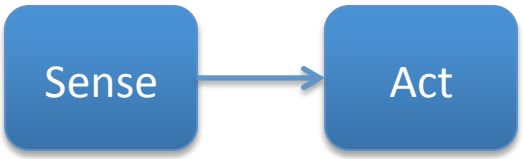

- Started due to disappointment with features of the hierarchical paradigm (1988 – 1992, but older roots in biology and cognitive science)
- Threw out planning altogether! Leveraged availability of low-cost hardware and computing resources
- Several clever robot insect demonstrations, but not sufficiently general purpose for robotics

The *Reactive* Paradigm

ROBOT PRIMITIVES	INPUT	OUTPUT
SENSE	Sensor data	Sensed information
~~PLAN~~	~~Information(sensed or cognitive)~~	~~Directives~~
ACT	Sensed information ~~or directives~~	Actuator commands

Example: Brooks' Insect Robots

[Source: ai.mit.edu]

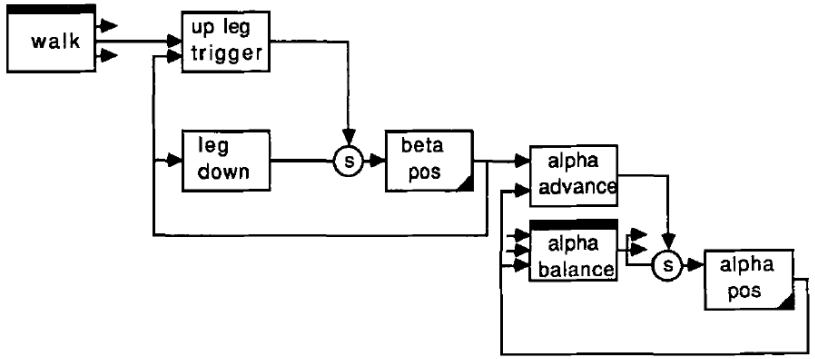

[Brooks, R.A., A robot that walks; emergent behaviors from a carefully evolved network. Neural computation, 1(2), pp.253-262, 1989]

The *Hybrid Deliberative/Reactive* Paradigm

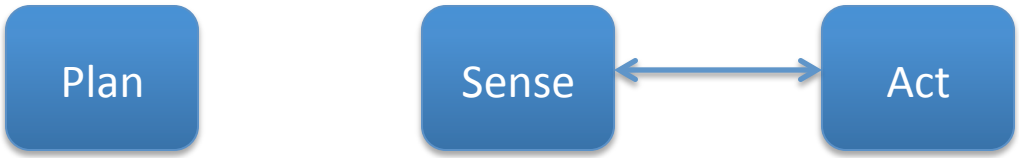

- Many current robots use this approach (1990s onwards)
- First, the robot deliberates how to break down task into subtasks (mission planning)
- Then the individual behaviours are executed as per a fast reactive paradigm
- PLAN, SENSE-ACT (P, S-A)

The Hybrid Paradigm

ROBOT PRIMITIVES	INPUT	OUTPUT
PLAN	Information(sensed or cognitive)	Directives
SENSE-ACT (behaviours)	Sensor data	Actuator commands

Example: “Modern” Mobile Robots

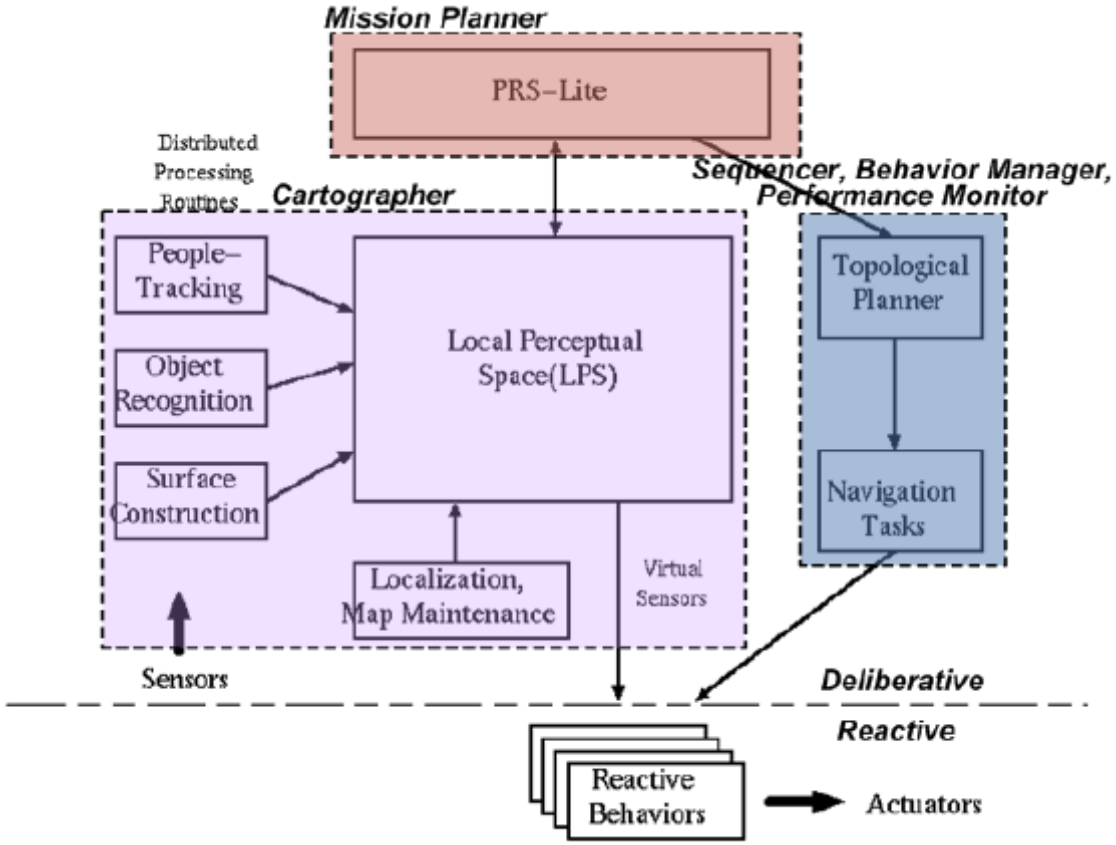

[Konolige, K., Myers, K., Ruspini, E., & Saffiotti, A. The Saphira architecture: A design for autonomy. Journal of experimental & theoretical artificial intelligence, 9(2-3), 215-235, 1997.]

Another “Modern” Issue: Interaction

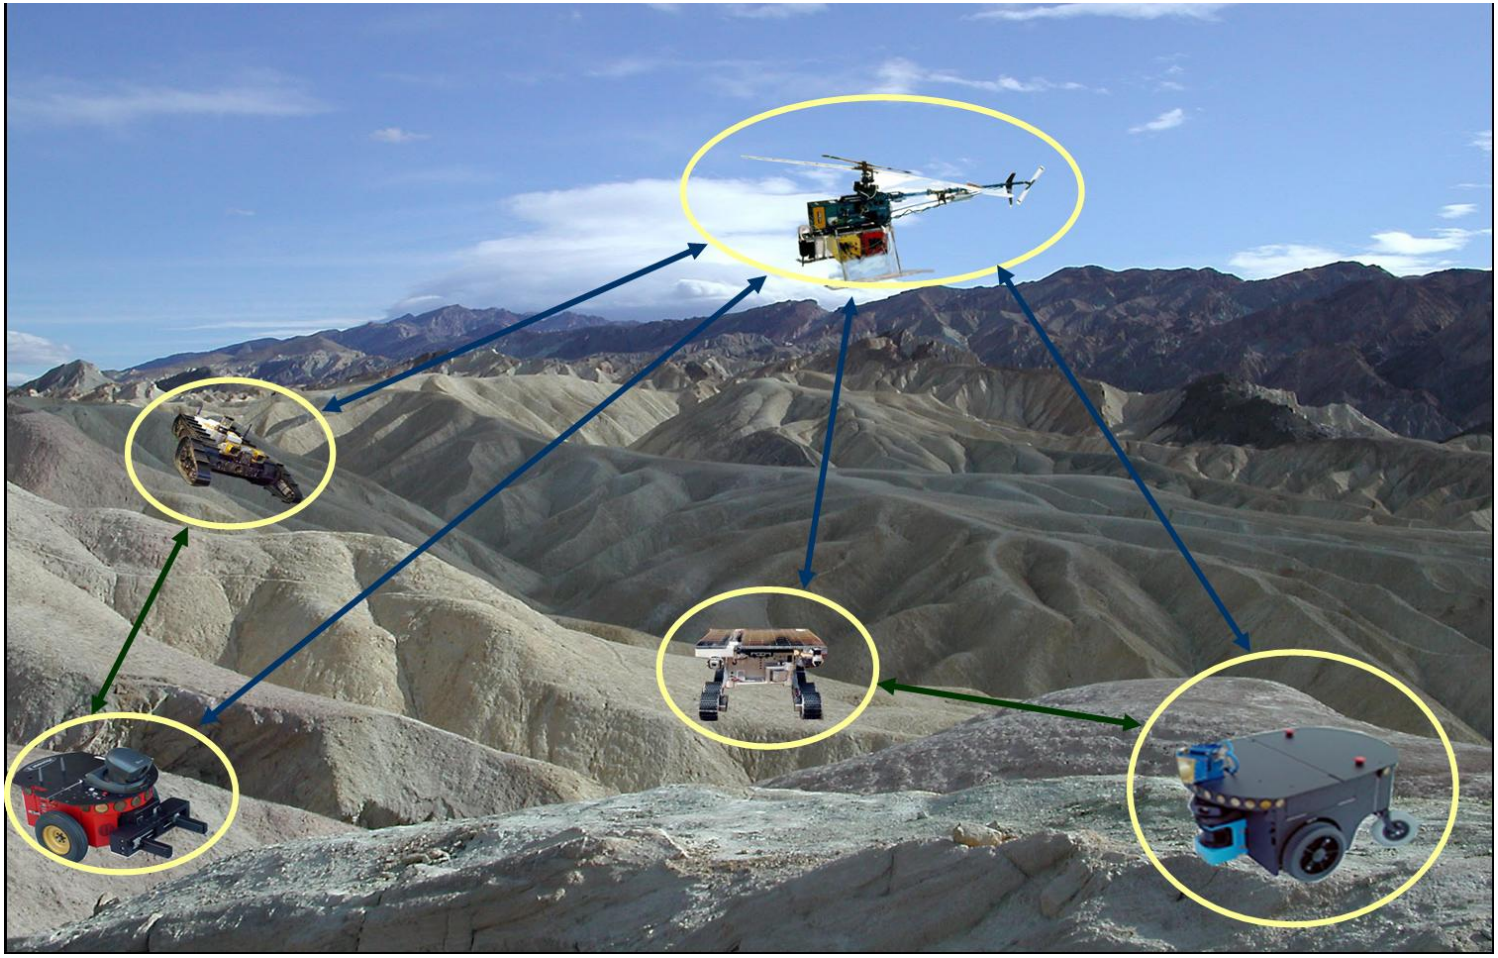

[Source: http://www.ee.ucr.edu/~mourikis/project_pages/images/multi.jpg]

So, What **is** a Robot?

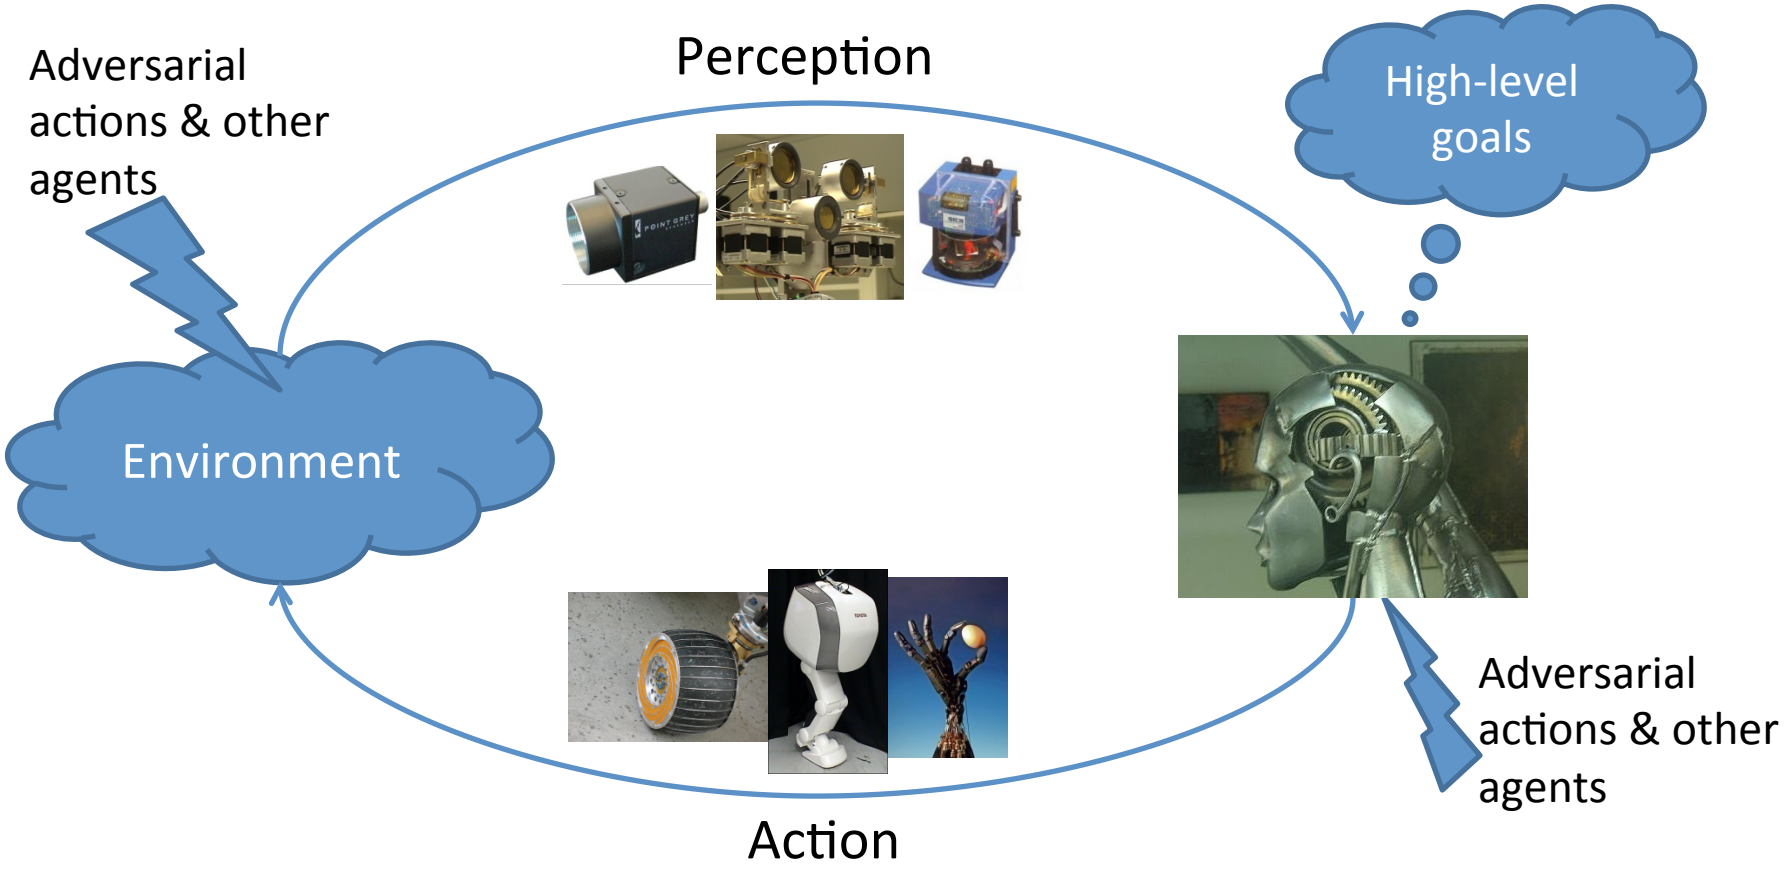

Problem: How to generate actions, to achieve high-level goals, using limited perception and incomplete knowledge of environment & adversarial actions?

Example Application: Autonomous Vehicles

- http://www.youtube.com/watch?gl=GB&v=1W27Q6YvTXc

- What are the various decisions involved?
- What paradigm(s) would you adopt?

Example Application: Rescue Robots

- http://www.youtube.com/watch?v=F7lqriYKsX4

- What are the various decisions involved?
- What paradigm(s) would you adopt?

Example Application: Automated Warehouses

- https://www.youtube.com/watch?v=6KRjuuEVEZs

- What are the various decisions involved?
- What paradigm(s) would you adopt?

Example Application: Humanoid Robots at Work!

- https://www.youtube.com/watch?v=DpTSXeei9zo

- What are the various decisions involved?
- What paradigm(s) would you adopt?

The Designer's Task: Components of the Problem

In each case,

- what are the components? how do you delineate?
- what does one (i.e., your robot) need to know?
- what does a motion strategy consist of?
  - what properties must the strategy satisfy?

What changes?
Who else is around?

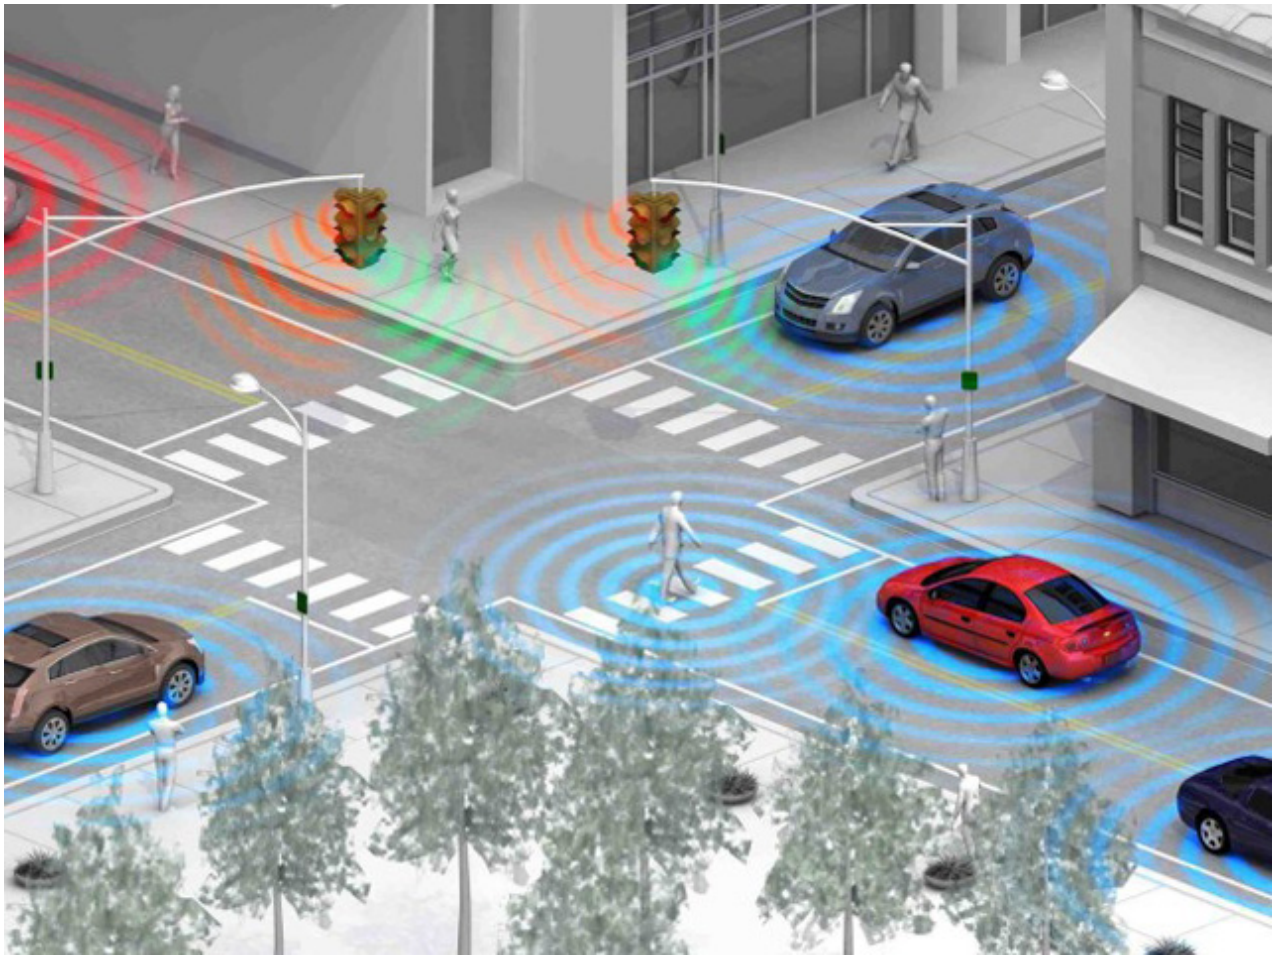

How does the car move?
- Kinematics, Dynamics

Where does the car move?
- World models

What Makes Robotics Problems ***Hard***?

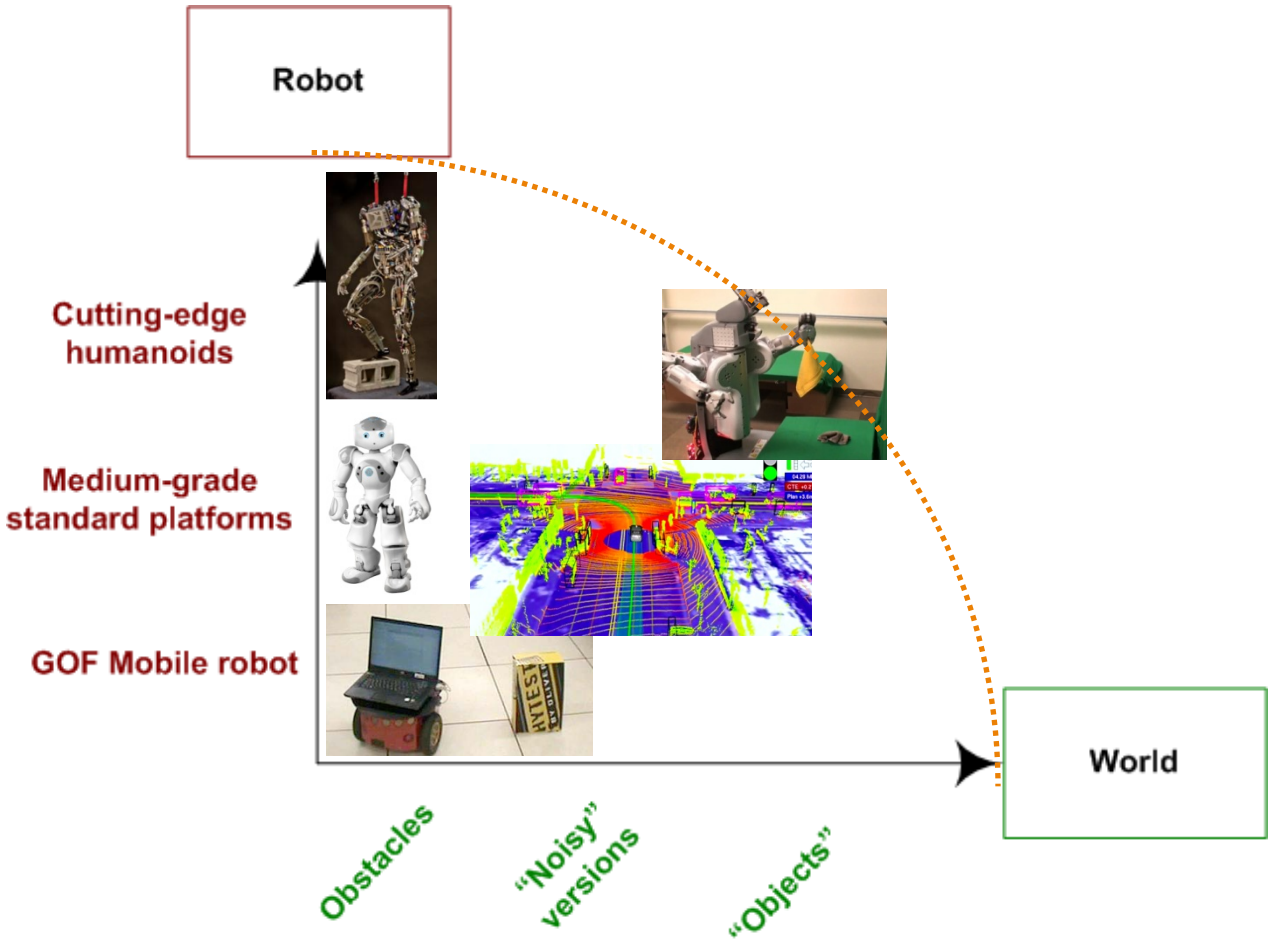

One more thing…

What Makes Robotics Problems Hard?

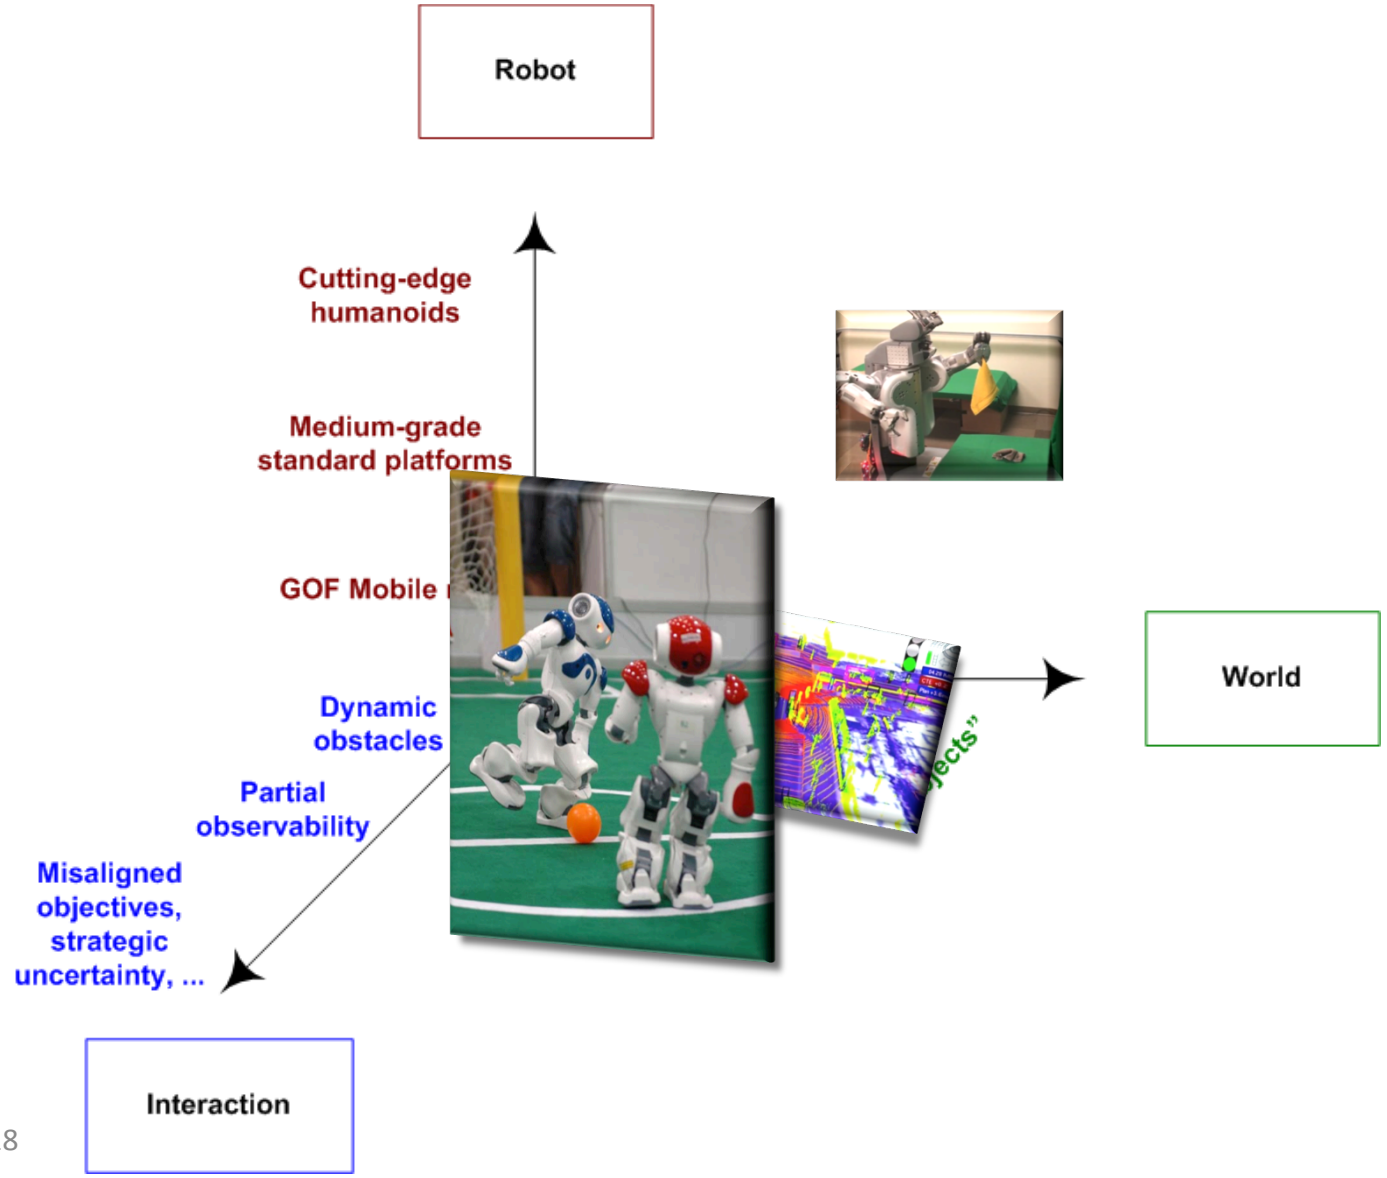

What Happens if You Plug in *Real* People?

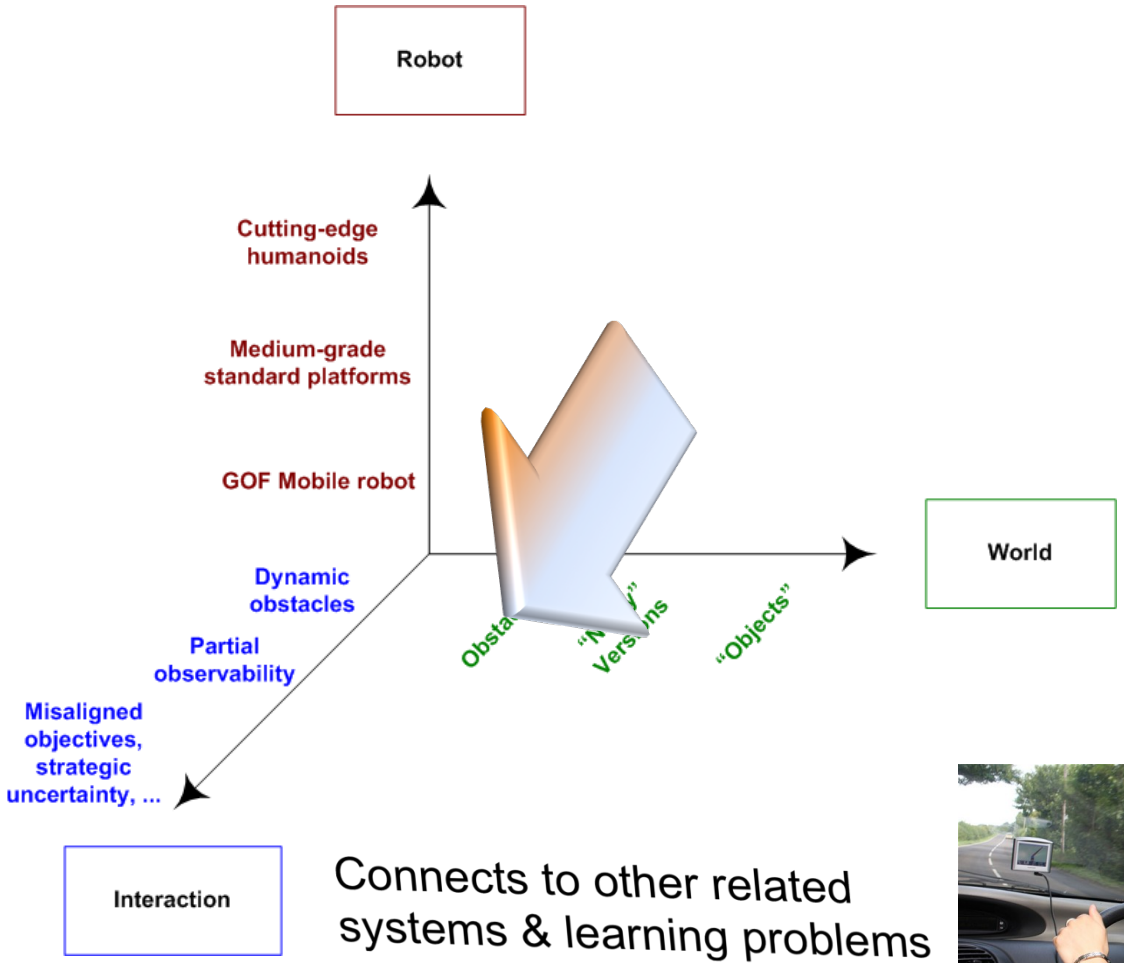

Connects to other related systems & learning problems

Computational Issues: Toy Example

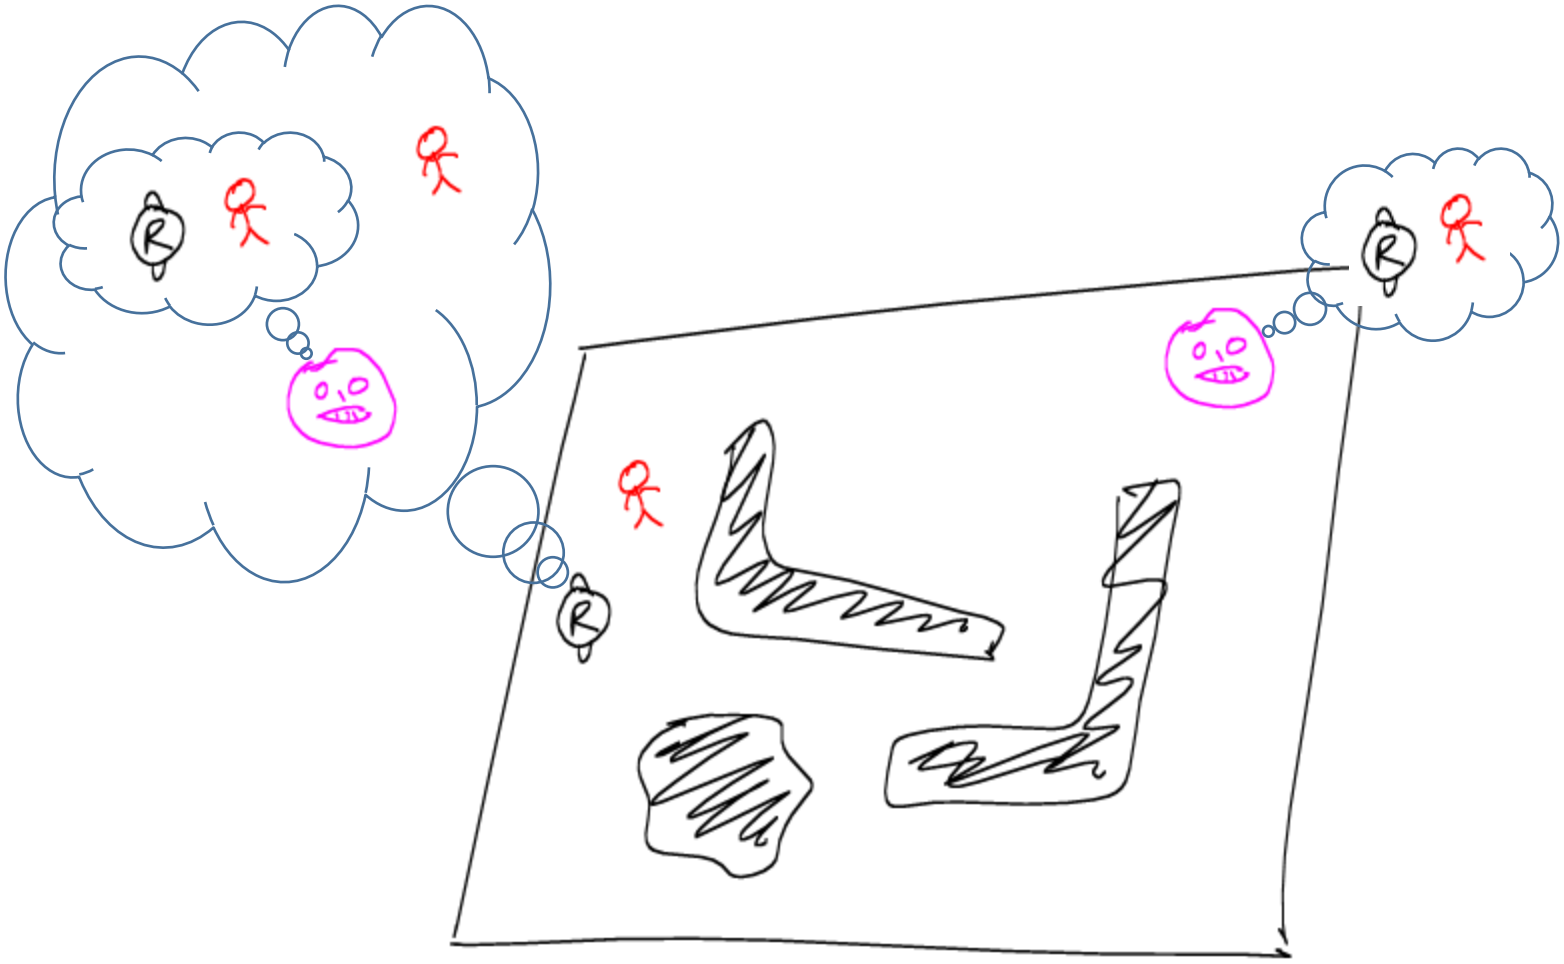

Non-stationarity, plan recognition, personalisation, incentives, strategic coordination

Levels of Difficulty in Interaction

Consequences for *hardness of learning*:

1. Base case: spatial asymmetry
   - Learn a vector field

2. Next level: deal with reactive behaviour
   - 'Inverse' planning, plan recognition

3. Harder case: recursive exchange of beliefs (e.g., signaling, implicit coordination, trust, persuasion)
   - Need to model as a game?

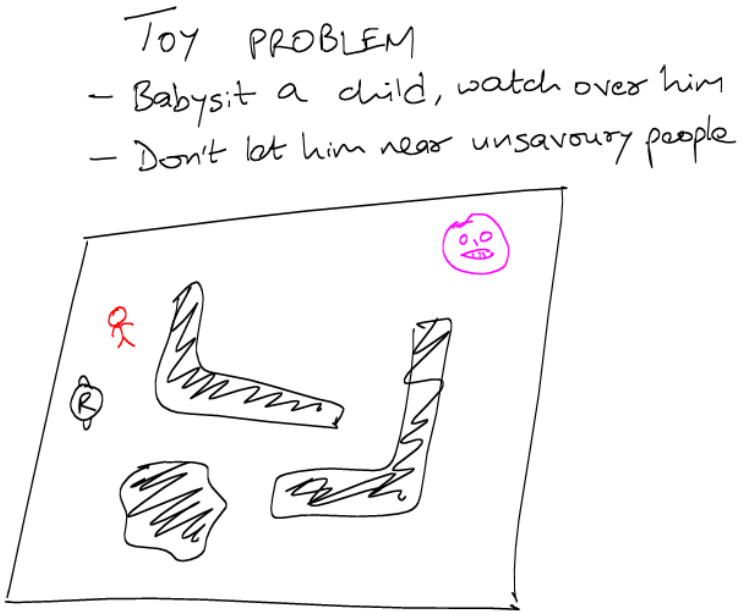

In this course…

We will focus on how to model and compute decisions (*choices*),

- over time, under uncertainty, with incompleteness in models
- emphasizing difficulties involving hidden causality, interaction, etc.
- possibly needing methods for learning from experience and data.
- also, we'll think a bit about how ***real*** people make choices!

Major Themes:

1. Different models of decision making
2. Understand issues, sometimes through case studies
   - What to model, what to analyse?
3. Special issues: safety, security, explainability, bounded rationality

Course Structure

- Schedule of lectures is available at the course web site

  http://www.inf.ed.ac.uk/teaching/courses/dmr/

  I will attempt to upload slides by day before (except in first week)
- Two homework assignments
  - Pen-and-paper exercise on models, concepts, methods (10%)
  - Practical programming exercise in a mock-up domain (20%)
- Term Paper
  - 4 page conference-style review of your chosen topic
- Final Exam (60% of final mark)
- Resources:
  - No prescribed textbook
  - Suggested readings assigned with lecture slides

Ask Questions!

- During the lecture
- After class, if your questions are brief
- After hours, by prior appointment *only* (arranged via email)

- You could also approach TA and Demonstrator:
  - Emmanuel Kahembwe E.Kahembwe@ed.ac.uk
  - Yordan Hristov yordan.hristov@ed.ac.uk

Acknowledgements

The material regarding robotic paradigms is from R.R. Murphy, Introduction to AI Robotics, MIT Press.