The background features a large, faint watermark of the Brown University crest. The crest includes a shield with a red cross, a crest above it, and a banner at the bottom with the Latin motto "IN DEO SPERAMUS".

Artificial Intelligence

George Konidakis
gdk@cs.brown.edu

Fall 2021

I410 Team

Instructor: George Konidakis

HTAs: Rachel Ma and Jinwoo Choi

TAs:

Bowen He

Edward Tan

Jason Crowley

Jason Gong

Kevin Lu

Kunj Padh

Neil Ramaswamy

Paul Krupski

Rena Jiang

Thien Nguyen



Major Topics Covered

Syllabus:

1. Agents and Agenthood
2. Search
 - (a) Uninformed
 - (b) Informed
 - (c) Game Theory and Adversarial Search
3. Knowledge Representation and Reasoning
 - (a) Logical Representations: Reasoning and Inference
 - (b) Uncertain Knowledge
 - i. Bayes' Rule
 - ii. Probabilistic Reasoning
 - iii. Bayes Nets
 - iv. Hidden Markov Models
4. Planning
 - (a) Classical Planning
 - (b) Robot Motion Planning
 - (c) Planning Under Uncertainty: Markov Decision Processes
5. Learning
 - (a) Reinforcement Learning
 - (b) Supervised Learning
 - (c) Unsupervised Learning
6. Advanced Topics
 - (a) Natural Language Processing
 - (b) Machine Vision
 - (c) Robot Learning
 - (d) Algorithmic Game Theory
7. Philosophy of AI
8. Social and Ethical Issues



On Lectures

The textbook contains everything you need to know.

Lectures contain everything you need to know.

Lecture notes **do not contain everything you need to know.**

Suggested approach:

- Come to lectures and pay attention.
- Revise via textbook (immediately).
- Clarify at office hours.



Logistics

Course webpage:

<http://cs.brown.edu/courses/cs141/>

- Syllabus
- Calendar - **office hours!**
- Assignments etc.

Written assignments and grades etc. via Gradescope
Comms (Q&A, announcements) via EdStem Discussion
Make sure to sign up!



Questions

EdStem: Quick question, or question many people may want to know the answer to.

Office Hours: Assignment and coding questions, material covered in lectures.



Grading

Six assignments

- 75% of grade.
- Python programming + report
- Generally 1-2 weeks long

Extended project: 25%.

Note: basic numpy/Python workshop on Monday from 5-6pm in Motorola.



Academic Honesty

I expect all Brown students to conduct themselves with the highest integrity, according to the *Brown Academic Code*.



It is OK to:

- Have high-level discussions.
- Google for definitions and background.

It is NOT OK TO:

- Hand in anyone else's **code**, or **work**, in part or in whole.
- Google for solutions.

ALWAYS HAND IN YOUR OWN WORK.



Academic Honesty

Consequences of cheating:


- Your case will be reported.
- Possible consequences include zeros on the assignment, suspension, failure to graduate, retraction of job offers.

If I catch you I **will** refer you to the Office of Student Conduct, and I **will** push for a hearing with the Standing Committee.

DO NOT CHEAT.





AI




AI: The Very Idea

For as long as people have made machines, they have wondered whether machines could be made intelligent.





(pictures: Wikipedia)





An epic drama of
adventure and exploration



(pictures: Wikipedia)



Turing



Computing machinery and intelligence. *Mind*, October 1950.

“Can machines think?”








(picture: Wikipedia)



Dartmouth, 1956



Trends



Connectionism I

AI Winter

Deep Learning (C III)

GOF AI

Connectionism II

Bayes

Reinforcement Learning



Modern AI

Subject of intense study:


- Nearly every CS department has at least 1 AI researcher.
- ~ 3000 PhDs a year in the US
- Tens of thousands of research papers written every year.

- Heavily funded (NSF, DARPA, EU, etc.).
 - Pays itself back fast (e.g., DART).

- Most major companies have efforts in this direction
 - Google,
 - Amazon
 - Microsoft, etc.



Modern AI



(picture: Wikipedia)






What *is* AI?

Fundamental Assumption

The brain is a computer.



=



(picture: Wikipedia)

What is AI?

This turns out to be a hard question!

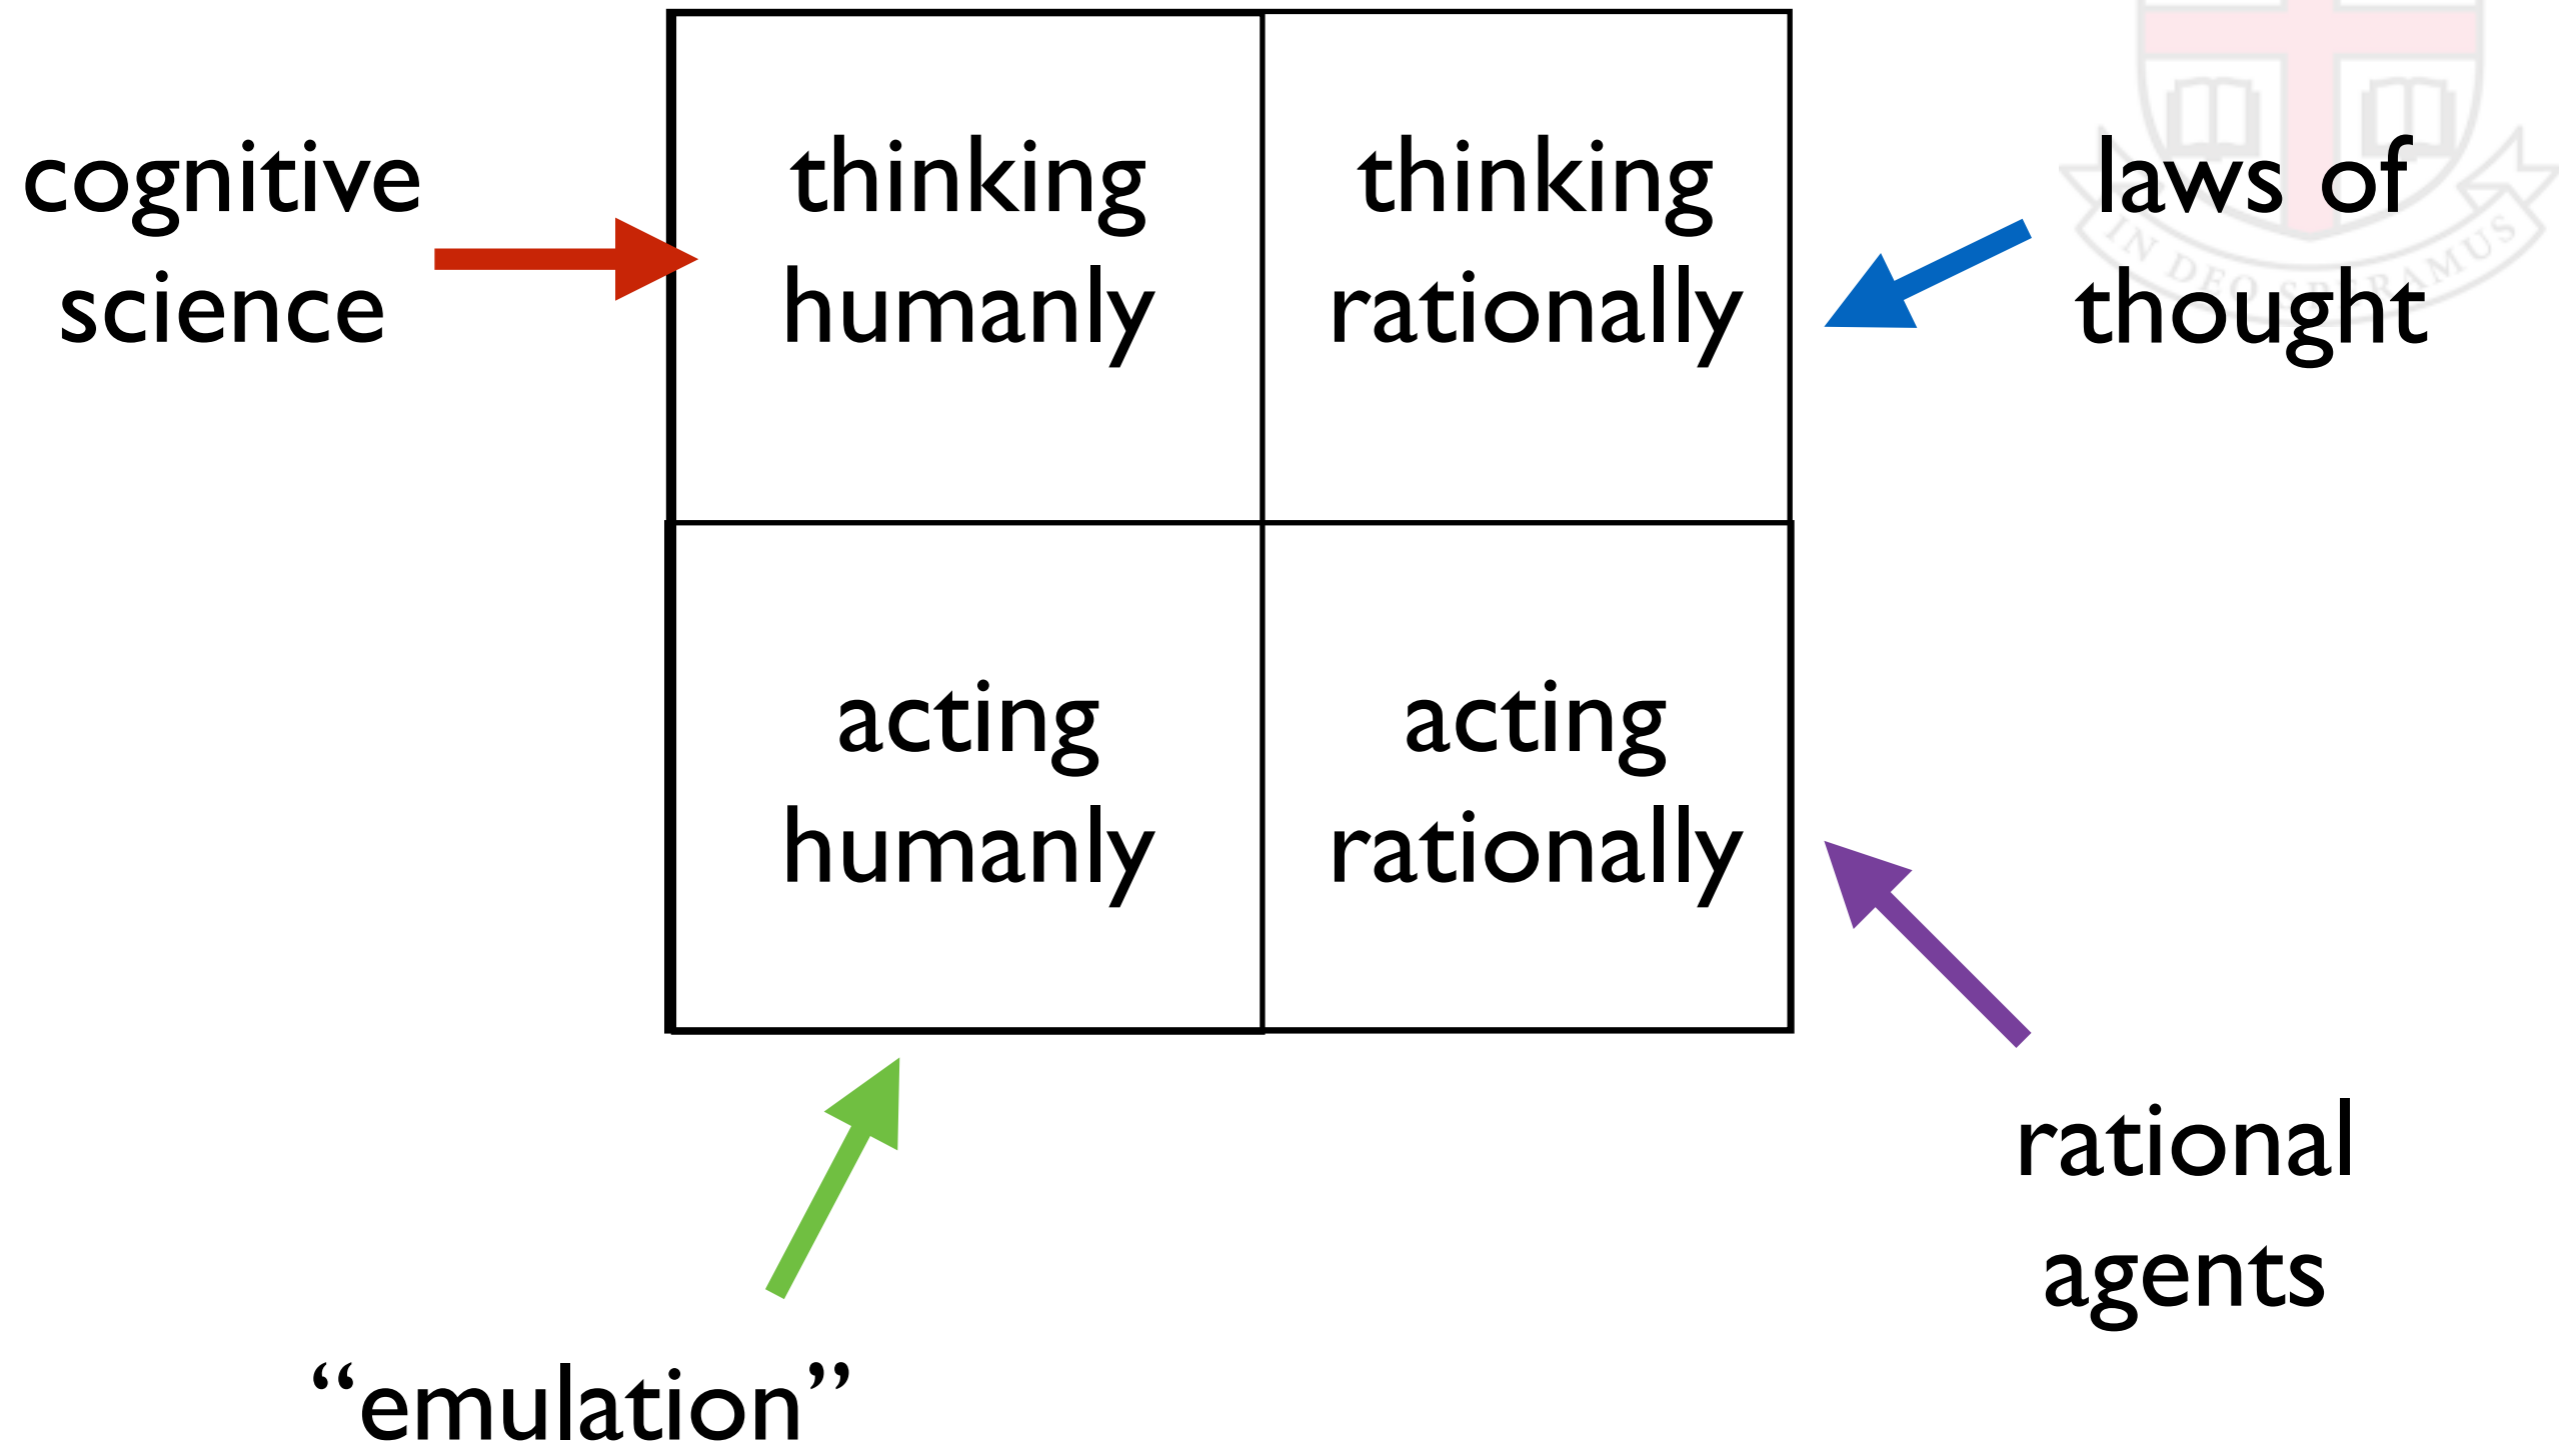


Two dimensions:


- “Humanly” vs “Rationally”
- “Thinking” vs. “Acting”.

thinking humanly	thinking rationally
acting humanly	acting rationally

What is AI?

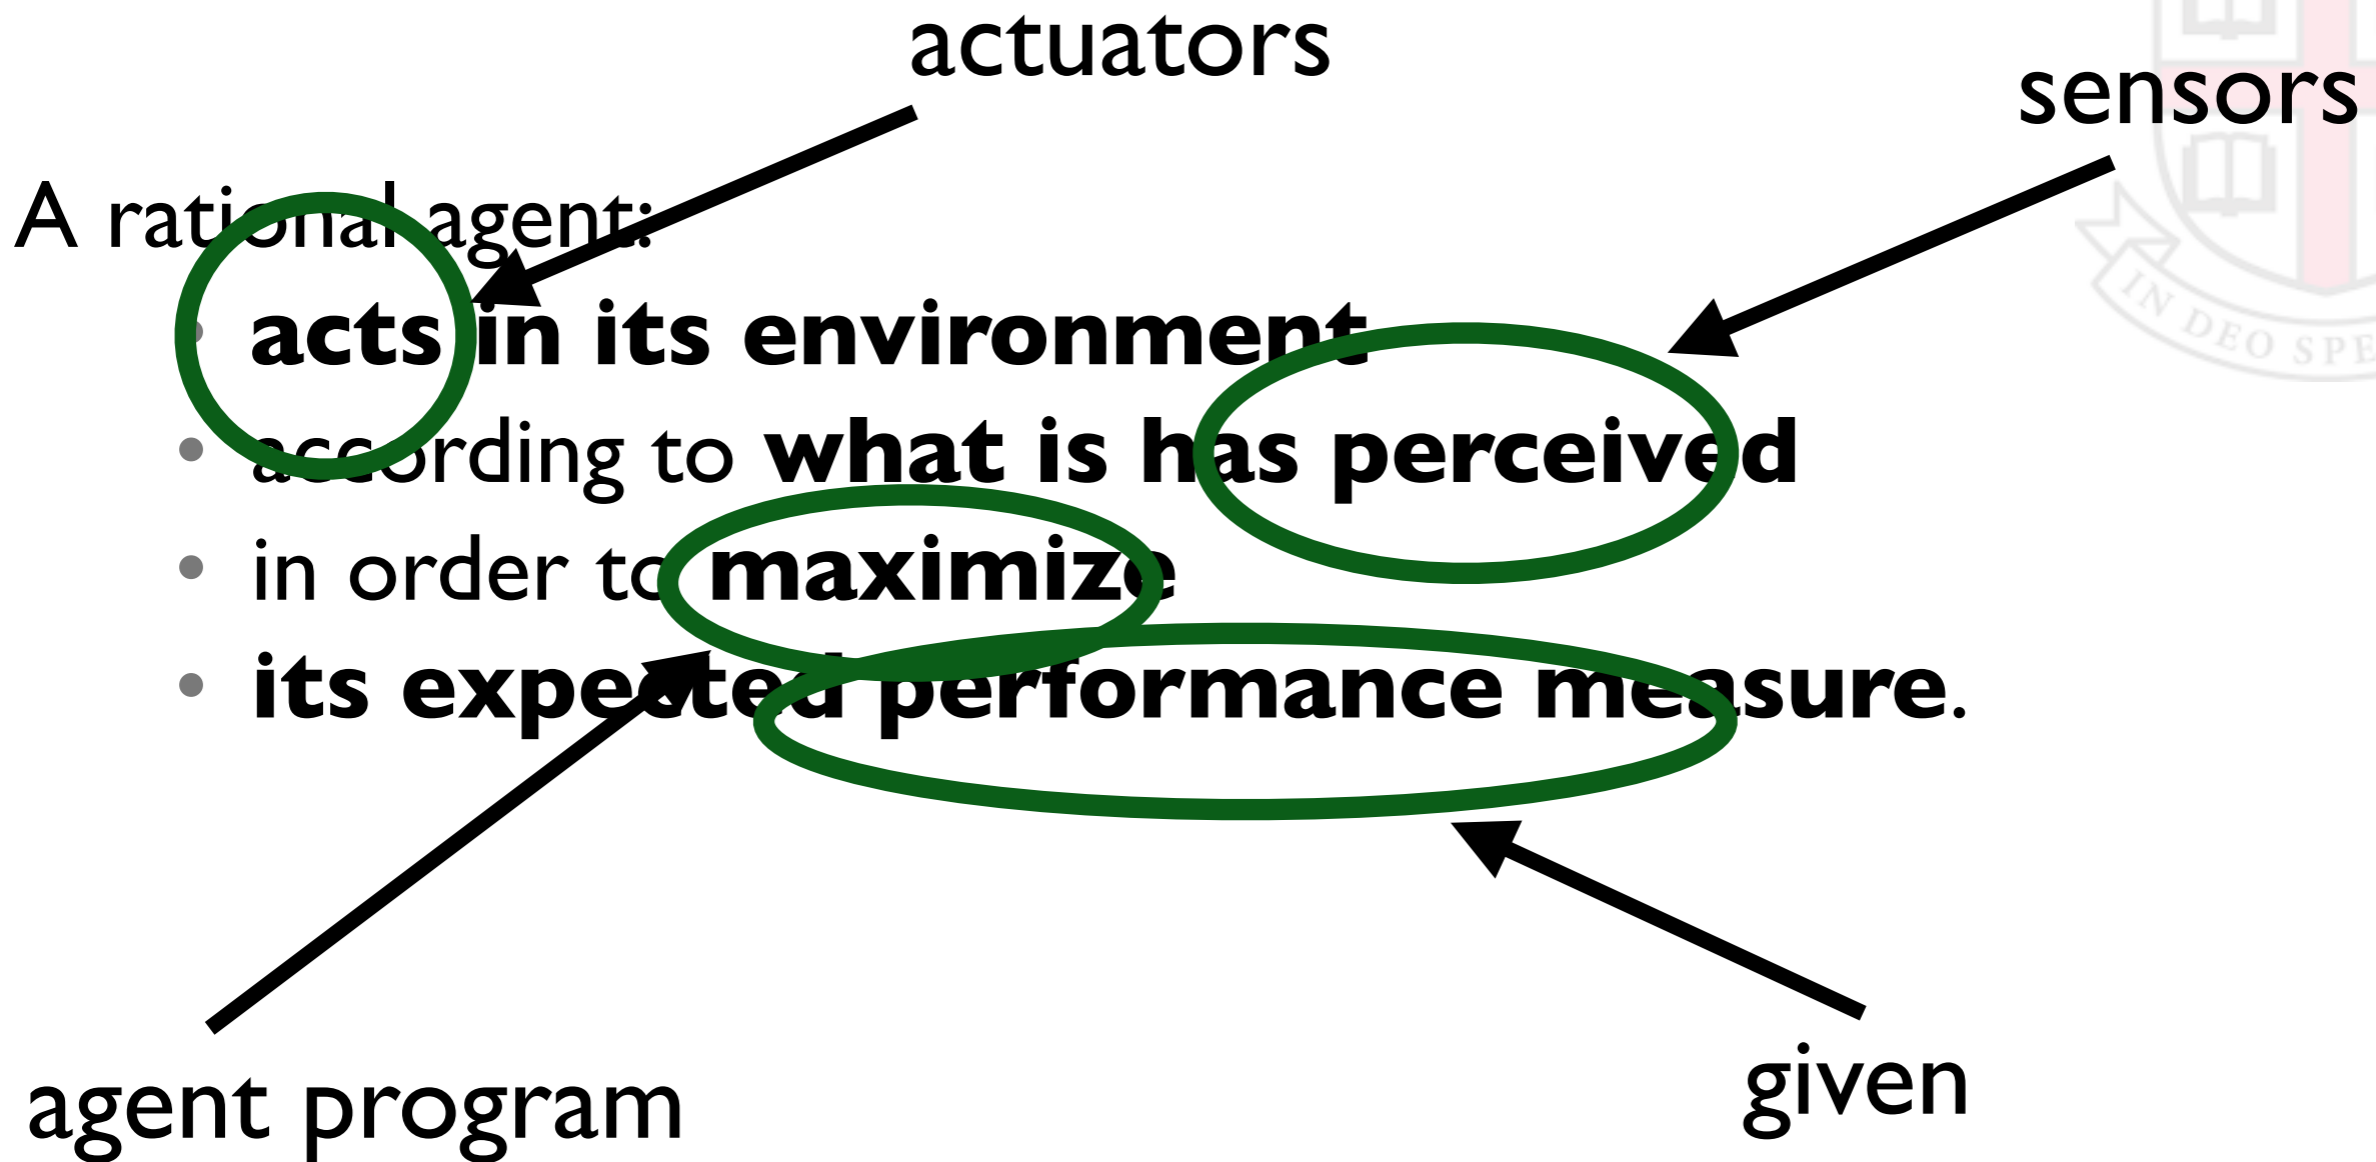


What is a Rational Agent?



Performance measure.

Rational Agents



Example: Chess



Performance measure?

Environment?

Prior knowledge?

Sensing?

Actions?



Chess


The chess environment is:

- Fully observable.
- Deterministic.
- Episodic.
- Static.
- Discrete.
- “Known”.



(picture:Wikipedia)

Example: Mars Rover



Performance measure?

Environment?

Prior knowledge?

Sensing?

Actions?

(picture: Wikipedia)



Mars Rover


The Mars Rover environment is:

- Partially observable.
- Stochastic.
- Continuing.
- Dynamic.
- Continuous.
- Partially known.



Are We Making Progress?

Specific vs. General



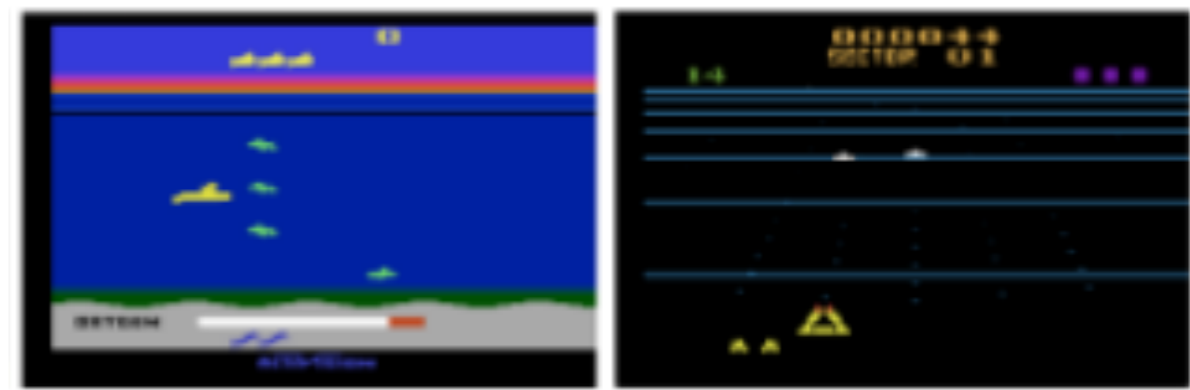


Progress



Starting out - 10 minutes of training

**The algorithm tries to hit the ball back, but
it is yet too clumsy to manage.**


Atari



Structure of the Field

AI is fragmented:

- Learning
- Planning
- Vision
- Language
- Robotics
- Reasoning
- Knowledge Representation
- Search



Progress

Progress in AI:

- Clear, precise models of a *class of problems*
- Powerful, general-purpose tools

- A clear understanding of what each model and tool can and cannot do

- Occasionally: *vividly illustrative applications.*
- *Arduous and slow*

- *Huge opportunity*

