



EDA132: Applied Artificial Intelligence or TAI: Tillämpad Artificiell Intelligens

Jacek Malec
Dept. of Computer Science, Lund University, Sweden
January 19th, 2016



Plan for today

- Administrative stuff
- Brief intro (AIMA Chapter 1)
- Agents (AIMA Chapter2)



What is (Artificial) Intelligence?



What is (Artificial) Intelligence?

What is Intelligence?



What is (Artificial) Intelligence?

What is Intelligence?

- adaptivity



What is (Artificial) Intelligence?

What is Intelligence?

- adaptivity
- learning



What is (Artificial) Intelligence?

What is Intelligence?

- adaptivity
- learning
- creativity



What is (Artificial) Intelligence?

What is Intelligence?

- adaptivity
- learning
- creativity
- logical reasoning



What is (Artificial) Intelligence?

What is Intelligence?

- adaptivity
- learning
- creativity
- logical reasoning
- problem solving capability
- ...

Can it be compared? Measured?



What is Artificial Intelligence?

Artificial intelligence (AI) is the intelligence of machines and the branch of computer science that aims to create it.

Textbooks define the field as “the study and design of intelligent agents,” where an **intelligent agent** is a system that **perceives** its environment and takes **actions** that maximize its chances of success.



What is Artificial Intelligence?

Artificial intelligence (AI) is the intelligence of machines and the branch of computer science that aims to create it.

Textbooks define the field as “the study and design of intelligent agents,” where an **intelligent agent** is a system that **perceives** its environment and takes **actions** that maximize its chances of success.

Late John McCarthy, who coined the term in 1956, defines it as “*the science and engineering of making intelligent machines.*” (Wikipedia)




In 2005:

Automated agent traders account for over 50% of portfolio trades by value most weeks on the New York Stock Exchange and, in some weeks, as much as 70% of portfolio trades.

Problems: “2010 flash crash”




In 2011:



www.elbot.com



A couple of agents more ...



Movie time

- Watson;
- Geminoid;
- Google car.



But...

“What is the only former Yugoslav republic in the European Union?”
(2012)

2880 POWER7 cores, 16 Terabytes memory, 4 Terabytes clustered storage (i.e., IBM Watson) still could not answer this question!



Subdomains of Artificial Intelligence

- Search, Problem solving
- Reasoning, Logical reasoning, Probabilistic reasoning
- Machine Learning
- Natural Language Processing
- Perception, Computer Vision
- Autonomous Robots
- Knowledge Processing
- ...



About the course

- EDA132: Applied artificial intelligence
- <http://cs.lth.se/EDA132>
- Serves as an announcement board as well!
- Meaning I expect you to read it often!!!
- Teachers: Pierre Nugues, Elin Anna Topp, Jacek Malec
- Administrator: Lena Ohlsson



Contents

- 7,5 hp (ECTS)
- Lectures (14), normally Tuesdays, 13–15 in A:C and Thursdays, 13–15, in E:B
- Three programming assignments
- Home reading (textbook)
- S. Russell, P. Norvig, Artificial Intelligence, a Modern Approach, 3rd int. ed., Prentice Hall



Evaluation

- Exam: worth 4,5p out of 7,5p. Material pointed to in the “reading advice” section.
- Programming assignments: worth 3p. Important: both correctness **and presentation** count.
- Complexity level of programming assignments may vary, although we strive for even division of labour;

but

- We need your feedback ...
- Kursombud (course representatives) need to be chosen



Programming assignments

- 1 Search
- 2 Probabilistic reasoning
- 3 Machine Learning (tentative list)
 - 1 Decision Trees
 - 2 Logistic Regression



Programming assignment submission

- The submission is to be sent to tai@cs.lth.se in the format described on the course web.
- This address works for **assignment submissions only!**. May or may not work for other things, so mail me (or any of the other teachers) **directly** for other purposes.
- All assignments are to be handed in electronically (as pdf documents), **on time!**



End of the admin stuff

Questions? Comments?

Please elect a course representative. Thank you.



What is AI


Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally



Acting humanly: The Turing test

Turing (1950) “Computing machinery and intelligence”:

- Can machines think? → Can machines behave intelligently?
- Operational test for intelligent behavior: the *Imitation Game*



- Loebner prize
- Anticipated all major arguments against AI in last 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning

Problem: Turing test is not *reproducible*, *constructive*, or amenable to *mathematical analysis*



Thinking humanly: cognitive science

1960s “*cognitive revolution*”: information-processing psychology replaced the then prevailing orthodoxy of *behaviorism*

Requires scientific theories of internal activities of the brain

- What level of abstraction? “Knowledge” or “circuits”?
- How to validate? Requires
 - Predicting and testing behavior of human subjects (top-down),
 - or Direct identification from neurological data (bottom-up)

Both approaches (roughly, *Cognitive Science* and *Cognitive Neuroscience*) are now distinct from AI

Both share with AI the following characteristic: *the available theories do not explain (or engender) anything resembling human-level general intelligence*

Hence, all three fields share one principal direction!



Thinking rationally: laws of thought

Aristotle: what are correct arguments/thought processes?

Several Greek schools developed various forms of **logic**:

notation and *rules of derivation* for thoughts;

may or may not have proceeded to the idea of mechanization

Direct line through mathematics and philosophy to modern AI

Problems:

- Not all intelligent behavior is mediated by logical deliberation
- What is the purpose of thinking? What thoughts *should* I have out of all the thoughts (logical or otherwise) that I *could* have?



Acting rationally

Rational behavior: doing the right thing

The right thing: that which is expected to maximize goal achievement,
given the available information

Doesn't necessarily involve thinking—e.g., blinking reflex—but thinking should be in the service of rational action

Aristotle (Nicomachean Ethics):

Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good



Rational agents

An *agent* is an entity that perceives and acts

This course is about designing *rational agents*

Abstractly, an agent is a function from percept histories to actions:

$$f : \mathcal{P}^* \rightarrow \mathcal{A}$$

For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance

Caveat: *computational limitations make perfect rationality unachievable*

→ design best **program** for given machine resources




State of the art

- Human-level AI back on the agenda
- 1997: Deep Blue defeats Kasparov
- Robbins conjecture (mathematics) proven after decades of human attempts
- No hands accross America: autonomous driving, flying, sailing, ...
- Logistics for Gulf, Iraq and Afghanistan
- Warfare for Iraq and Afghanistan
- 2011: Watson defeats humans in Jeopardy
- 2011: Siri
- Medical diagnoses and treatment
- ...



What's in the course



What's missing in the course





Ethics

- enhancements of our capacities (bodies, minds)
 - do we want that?
 - can we afford not having that?
- elderly care, rehabilitation, medicine
vs. war-fighting, sex, socializing
- emotional artificial partners
- large finances come from military sources (e.g., DARPA)
 - defensive
 - preventive attacks
 - robots that kill

Do we have the **right** to create robot servants?