



EDA132: Applied Artificial Intelligence or TAI: Tillämpad Artificiell Intelligens

Jacek Malec

Dept. of Computer Science, Lund University, Sweden

January 21st, 2014



Plan for today

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





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.



In 2011:



www.elbot.com

A big hello to my new visitor in Sweden!

Powered by **artificial-solutions.com**

Done



A couple of agents more ...

Billiga flygbiljetter och flygresor - sök och boka flyg - SAS

Most Visited - Conf Current - Edu@LU - Software - News - Banki - Travel - Domeny - CL stuff - Måje - roads - EU - ELIN - Slowiki - Sctopia - LUPP - Robotics - Palm

Logga in i Ny användare
Hjäp & kontakt
Fråga Eva

Boka resa
 Checks in & väp plats
 Hitta din bokning
 Är flygget i ord?
 Tilstäbel
 Nyhetsbrev
 Boka hyrtel
 Boka hotell

Boka i lägerskalendern
 Boka bonusresa
 Boka SAS Corporate Booking & Travel Pass
 Boka flyg + hotell

Boka flyg
 Boka ungdoms/student
 Boka ensamtresande barn

Från: Slev stad, flygplats, flygplats eller land
 Till: Slev stad, flygplats, flygplats eller land

Tur och retur Enkel resa

Utresa: januari 2011
 Hämresa: januari 2011

M	T	O	T	F	L	S	M	T	O	T	F	L	S
13	14	15	16	17	18	19	20	21	22	23	24	25	26
27	28	29	30	31									

Visa månadskalender Direkt till valda dagar

1 vecka
 0 barn (2-11 år)
 0 spelbarn (0-23 mån)

Asterisk **Sök**

Lägerskalendern - flyg fr. 495:- Sök

Äre enkel resa fr. 495:-

Riseforsäkring

Trygghet på resa
 Det kan bli dyrt att bli sjuk utomlands - boka försäkring!

Angelholm
 enkel resa fr. 495:-

Flyrail
 Slipp kombinera resan själv!

fr. 395:-
 Sista minuten
 Ungdom

EuroBonus
 Bli medlem idag och tjäna poäng på dina resor!

Fråga Eva
 http://eva.cs.uioa.artificial-
 Firefox prevented this site from opening a pop-up window.
 Fråga Eva
 Du skriver jag vill resa till Thailand
 Din webbläsare ställer inte mig ett äpple
 sidor på SAS webbplatsen. Klicka för att
 öppna sidan i ett nytt fönster.
 Du undrar något om resor och
 destinationer.
 De frågor jag har svar på som handlar om
 resor och destinationer är:
 • Vilka resor kan jag boka på nätet?
 • Var hittar jag historik över mina resor?
 • Kommande flygplan
 • Har många olika destinationer flyger SAS
 till?
 • Var hittar jag SAS citygater?

SAS ger 1:- per beräknat till **Reklam** **För en på**

Tjänst **Reklam** **All om resan** **På EuroBonus**

http://www.sas.se/sv/erbjudanden/Nyhetsbrev/



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 (IBM Watson) still can not answer this question!



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, 15–17 (M:D) and Fridays, 13–15 (E:A)
- Programming assignments (4)
- Home reading (textbook)
- S. Russell, P. Norvig, Artificial Intelligence, a Modern Approach, 3rd 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 Reasoning
- 3 Learning (tentative list)
 - 1 Decision Trees
 - 2 Logistic Regression
 - 3 Bayesian Learning
- 4 Natural Language Processing



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. 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



AI prehistory

<i>Philosophy</i>	logic, methods of reasoning mind as physical system
<i>Mathematics</i>	foundations of learning, language, rationality formal representation and proof algorithms, computation, (un)decidability, (in)tractability, also probability
<i>Psychology</i>	adaptation phenomena of perception and motor control experimental techniques (psychophysics, etc.)
<i>Economics</i>	formal theory of rational decisions
<i>Linguistics</i>	knowledge representation, grammar
<i>Neuroscience</i>	plastic physical substrate for mental activity
<i>Control theory</i>	homeostatic systems, stability, optimal control



History of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1950s Early AI programs, including Samuel's checkers program
Newell & Simon's Logic Theorist, Gelernter's Geometry E
- 1956 Dartmouth meeting: "**Artificial Intelligence**" adopted
- 1965 Robinson's complete algorithm for logical reasoning
- 1966–74 AI discovers computational complexity
- 1969–79 Early development of knowledge-based systems
- 1980–88 Expert systems industry booms
- 1988–93 Expert systems industry busts: "AI Winter"
- 1985–95 Neural networks return to popularity
- 1988– Resurgence of probability; general increase in technical depth, "Nouvelle AI": ALife, GAs, soft computing
- 1995– Agents, agents, everywhere . . .




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 (98% of time driven autonomously)
- Logistics for Gulf, Iraq and Afghanistan
- Warfare for Iraq and Afghanistan
- 2011: Watson defeats humans in Jeopardy
- 2011: Siri



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?