

CogSysI Lecture 1: Introduction

Intelligent Agents

Ute Schmid (lecture)

Emanuel Kitzelmann (practice)

Applied Computer Science, Bamberg University

last change: April 16, 2008

Intelligent Agents

- Natural or artificial systems which act in an intelligent way
- What is “intelligence”?
What is “artificial intelligence” (AI)?
What is “cognitive AI”?
 - topic of this lecture
- Planning is the reasoning side of acting
 - focus topic of this course
 - covering many topics, approaches, methods of general AI

Subject Matter of AI

- Term “Artificial Intelligence” is problematic
 - no satisfying definition of “intelligence”
 - Operational definition in psychology: *Intelligence is what is measured by an intelligence test*
 - Analytical definition: *Intelligence is the ability to acquire and apply knowledge*
 - “artificial”: problematic connotation, creation of machines which have human intelligence (“strong AI” in the 60ies)
 - Modern term “intelligent systems”

Definitions of AI

A typical definition: *Artificial intelligence is the science of making machines do things that would require intelligence if done by men.* (Minsky, 1963, pp. 23)

Problems:

- term intelligence is used but not defined
- a lot of problems where we ascribe high intelligence, if a human can solve them, are relatively easy to solve by a computer program (solving mathematical equations, playing chess, some classes of mathematical proves ...)
- a lot of problems which can be solved by every child are very hard to realize with a computer program (object recognition, building a tower of blocks, language understanding, ...)

Example: Context Effects



(human intelligence makes heavy and easy use of context,
Selfridge, 1955)

Example: Context Effects

T /- \ E C /- \ T

(human intelligence makes heavily and easy use of context,
Selfridge, 1955)

Example: Context Effects



Intelligence as Ascription

The extent to which we regard something as behaving in an intelligent manner is determined as much by our own state of mind and training as by the properties of the object under consideration. If we are able to explain and predict its behavior we have little temptation to imagine intelligence. With the same object, therefore, it is possible that one man would consider it as intelligent and another would not; the second man would have found out the rules of its behavior.

(Alan Turing, 1947)

Further Definitions of AI

There is no really satisfying definition of AI, two which I like are:

AI researches how one could make a computer do what humans currently can do better. (Elaine Rich)

AI research is concerned with computer problems yet unsolved. (Marvin Minsky)

- ↪ For many concepts/ scientific fields (not only for AI) there is no single satisfying definition.
- ↪ Pragmatic solution: Enumerate the topics, give examples (getting “concrete”)

AI and Computer Science

- AI is an integral part of computer science, the beginnings of computer science research are also the beginnings of AI research, but:
- *AI is that sub-discipline of computer science, which is concerned with analysis and formalization of thought processes which are **not yet formally understood**. In the case of success, that is if these processes got “**de-mystified**”, these areas become part of standard computer science and AI research is moving on to new problems. (Christian Freksa)*

Examples:

problem solving → efficient search algorithms

automatic programming → compiler design

reasoning → theorem proving

Methods of AI

AI as formal science

- **Analytical, formal:** Analysis of problems, formalization, design and evaluation of algorithms (completeness, correctness, optimality; empirical evaluation of error rates and performance times)
↪ *Precise description of problems and algorithms; formal, normative constraints for performance*

Methods of AI cont.

AI as natural/epistemological science

- **Empirical:** Recourse to biological and cognitive principles of organization and processing and their simulation
↳ *Evidence for the existence of abilities and skills, their constraints and characteristics of performance*

Methods of AI cont.

AI as engineering science

- **Engineering:** (Efficient) implementations of algorithms, adaptation of algorithms to requirements of specific applications

↪ *Application is one possibility to evaluate the scope, functionality and relevance of AI-algorithms*

↪ As an AI researcher you should have knowledge of formal methods and of cognitive theories and empirical methods!

Two Perspectives

- AI as **engineering discipline**: make algorithms for solving AI problems
- AI as **cognitive modeling**: make algorithms which solve problems in a similar fashion as humans
↳ **AI as theoretical psychology** (cognitive AI)

Everything is vague to a degree you do not realise till you have tried to make it precise. (Bertrand Russell)

Topics of AI

- **Problem solving and planning**
 - basic search techniques for many AI topics
 - Applications: Scheduling, configuration, manipulation of formula, games
- **Inference and theorem proving**
 - deduction of information which is given implicitly in knowledge bases
 - in contrast to other areas of AI: algorithms which are provably complete and correct
 - non-monotonic reasoning, fuzzy/probabilistic approaches
 - Applications: Program verification, knowledge based/expert systems

Topics of AI cont.

● Knowledge representation

- can be crucial for performance of AI systems (easy retrieval of stored knowledge)
- logic, semantic nets, frames/schemes, graphs, ...

● Machine learning

- Concept learning, classification learning
- strategy learning
- Applications: data mining, object recognition, process control

Areas of Application

- Natural language processing
- Computer vision (object recognition)
- Multi-Agent-Systems
- Robotics (action planning, navigation)
- ...

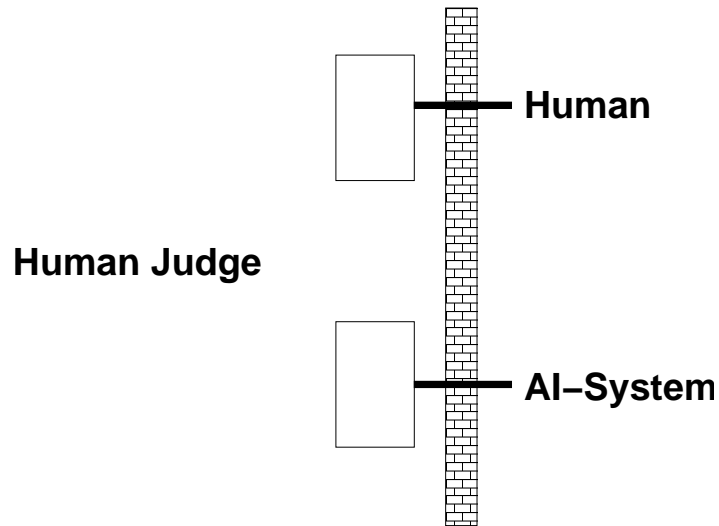
AI programming

- Declarative programming languages (5th generation languages): Prolog, Lisp; today there are no longer typical AI languages, modern planners and most machine learning algorithms are written in C

Relations to Others Disciplines

- **Philosophy:** Early AI made strong promises, philosophers started to ask whether human cognition has inherent aspects which cannot be transferred to a computer ([Dreyfus](#), [Searle](#)). AI is an interesting area for philosophical studies of mind.
- **Linguistics:** Chomsky had very strong influence to computer science ([Chomsky hierarchy](#)) and had a major impact on the beginning of cognitive psychology ([Chomsky-Skinner debate](#)).
Interaction of AI and Linguistics: Computer Linguistic
- **Psychology:** AI models influence psychological models (Minsky's computer metaphor Mind/Brain, Software/Hardware).

Turing Test



- Kasparov vs. Deep Blue, Match 1996: Kasparov suspected human intervention
- Searle critique: Thought experiment “Chinese Room”
- Cognitive Science approach: do not judge superficial performance; investigate process characteristics of the program (correspondence of number of rule applications to reaction times, similar errors etc.)

A Short History of AI

Prehistory

- 1943** McCulloch & Pitts “Logical Calculus of the Ideas Immanent in Nervous Activity” (Architecture for an intelligent system based on a neural net)
- 1948** Wiener “Cybernetics” (information theory)
- 1950** Turing: “Computing Machinery and Intelligence” (Turing-Test)
- 1955** Selfridge: Pattern-matching program

A Short History of AI cont.

Early AI (1956 to mid 60ies)

- Dartmouth conference of 1956: Term AI comes into use (McCarthy)
- Focus on models of cognitive processes and general principles of intelligent behavior
 - Marvin Minsky (MIT): Perception
 - A. Newell and H. Simon (CMU): Inference: theorem prover for propositional logic “Logic Theorist” (1956)
Problem Solving: GPS (1958)
 - John McCarthy (Stanford): Lisp (1958)

A Short History of AI cont.


... Early AI (1956 to mid 60ies)

- Games: Samuels (1959), Checkers
- Learning: Winston (1970), Learning by analogy
- Analogical reasoning: Evans (1959)
- Simon, 1965: *“by 1985 machines will be capable of doing any work a man can do”*
- Minsky (1968) “Semantic Information Processing” (an important collection of early work)
- Drawbacks, e.g. in machine translation \hookrightarrow AI critique


Analogy Problems

Evans (1968), grammar inference method



A




B




C




1




2



3



4



5

A Short History of AI cont.

The middle years (70ies)

- no interest in interdisciplinarity and cognition
- focus on knowledge: representation and inferences
- Question-Answering systems, Expert systems
- SHRDLU, Winograd (1970): natural language system for manipulation and verification of statements in a blocksworld (procedural semantics)
- DENDRAL, Feigenbaum (since 1965): Analysis of molecular structures
- MACSYMA (1971): manipulation of formulas
- MYCIN, Shortliff (since 1974): medical diagnosis

A Short History of AI cont.

... The middle years (70ies)

- PROLOG (1973) Colmerauer
- Development of languages for production systems (OPS, McDermott 1977) and shells for expert systems
- Planning: STRIPS (Fikes & Nilsson 1971), NOAH (Sacerdoti 1975)
- non-classic logic (Reiter, 1980), fuzzy logic (Zadeh, 1965)
- hierarchical nets, Quinlan (TLC, 1968); frames, Minsky (1975); Scripts, Schank & Abelson (1973)

A Short History of AI cont.

The 80ies

- Search for real world applications; intelligent tutor systems
- Language research: functional-logic programming
- Re-invention of backpropagation, new interest in neural/distributed information processing (Feldmann & Ballard 1982; Rumelhart & McClelland 1986)
- Big battels of symbolic vs. sub-symbolic AI

A Short History of AI cont.

The 90ies and now

- new interest in interdisciplinary research, birth of cognitive science
- AI-methods are used when helpful:
“Some of the most successful applications of AI are those in which the artificial intelligence is spread like raisins in a loaf of raisin bread: the raisins do not occupy much space, but they often provide the principal source of nutrition.” (Esther Dyson, industrial analyst)
- Focus on learning
- Relation to robotics: embodied intelligence, situated action
- Multi-Agent Systems

Textbooks

Stuart Russell and Peter Norvig (2002). Artificial Intelligence: A Modern Approach. Prentice Hall. (2nd Edition) *The mostly used text book in the field.*

Patrick Winston (1992). Artificial Intelligence (3rd Ed.). Addison-Wesley. (deutsch, 2. Auflage 1987) *A not very formal book. It's advantage is that it addresses relations to cognitive science research.*

Nils Nilsson (1980). Principles of Artificial Intelligence. Springer. (1998). Artificial Intelligence: A new synthesis. Kaufman. *Since Nilsson is the inventor of the A*-algorithm, his chapters on heuristic search algorithms are written very clearly and detailed. His most comprehensive text on search methods is:*

Nils Nilsson (1971). Problem-Solving Methods in Artificial Intelligence, McGraw-Hill.

Textbooks cont.

M. Ghallab, D. Nau, & P. Traverso (2004). Automated Planning. Theory and Practice. Elsevier. *The text book this course will be mainly based on.*

Günther Görz, Claus Rollinger & Josef Schneeberger (2000). Handbuch der Künstlichen Intelligenz. (3. Auflage) Oldenbourg.
A detailed overview of topics of AI. Gives an overview of AI-research in Germany.

Raymond Kurzweil (1992). The Age of Intelligent Machines. MIT Press.
This is not a scientific text book but a most motivating, comprehensive book for a broad readership.

... many others, see course web page

AI Organizations

- AAI: American Association of AI
Most important international conference: IJCAI
Important journal: Artificial Intelligence
- ECCAI: European Coordinating Committee on AI
- Fachbereich 1, KI, der Gesellschaft für Informatik (GI)
Annual conference: KI, interdisciplinary spring school:
IK (Interdisziplinäres Kolleg)

Running Gag

Question: How many AI people does it take to change a lightbulb?

Answer: At least 81.

First part of the solution: **The Psychological Group (5)**

- One to build an apparatus which will time lightbulb-changing performance
- One to gather and run subjects
- One to mathematically model the behavior
- One to call the expert systems group
- One to adjust the resulting system, so that it drops the right number of bulbs