

Artificial Intelligence

WS 15/16

Introduction

Anna Fensel
anna.fensel@sti2.at



Where are we?



#	Title
1	Introduction
2	Propositional Logic
3	Predicate Logic
4	Reasoning
5	Search Methods
6	CommonKADS
7	Problem-Solving Methods
8	Planning
9	Software Agents
10	Rule Learning
11	Inductive Logic Programming
12	Formal Concept Analysis
13	Neural Networks
14	Semantic Web and Services

Overview

- Course home page: <http://bit.ly/1hy77vZ>
(schedule, lecture notes, exercises, etc.)
- Textbooks:
 - G. Görz, C.-R. Rollinger, J. Schneeberger (Hrsg.) “**Handbuch der künstlichen Intelligenz**” Oldenbourg Verlag, 2003, Fourth edition
 - G. Luger “**Artificial Intelligence – Structures and Strategies for Complex Problem Solving**” Addison-Wesley, 2005, Fifth edition
- Lecturers: Ass.-Prof. Dr. Anna Fensel (anna.fensel@sti2.at) and Dr. Ioan Toma (ioan.toma@sti2.at)
- Tutor: Zaenal Akbar (zaenal.akbar@sti2.at).
- Lectures every week and tutorials every two weeks
- Attendance of the tutorials is obligatory!



GEORGE F LUGER
ARTIFICIAL INTELLIGENCE
Structures and Strategies for Complex Problem Solving



- Written exam at the end, with no literature
- Exam grade:

Score in % of the total	Grade
89-100	1
76-88	2
63-75	3
50-62	4
0-49	5



-
1. Introduction
 2. Propositional logic
 3. Predicate logic
 4. Reasoning
 5. Search methods
 6. CommonKADS
 7. Problem-solving methods
 8. Planning
 9. Software Agents
 10. Rule learning
 11. Inductive logic programming
 12. Formal concept analysis
 13. Neural networks
 14. Semantic Web and Services

- **Motivation**
 - What is “Intelligence”?
 - What is “Artificial Intelligence” (AI)?
 - Strong AI vs. Weak AI
- **Technical Solution**
 - Symbolic AI vs. Subsymbolic AI
 - Knowledge-based systems
- **Popular AI systems**
- **Subdomains of AI**
- **Some relevant people in AI**
- **Summary**



Introduction to Artificial Intelligence


MOTIVATION

What is “Intelligence”?

- "Intelligence denotes the ability of an individual to adapt his thinking to new demands; it is the common mental adaptability to new tasks and conditions of life" (William Stern, 1912)
- Being "intelligent" means to be able to cognitively grasp phenomena, being able to judge, to trade off between different possibilities, or to be able to learn.
- An important aspect of "Intelligence" is the way and efficiency how humans are able to adapt to their environment or assimilate their environment for solving problems.
- Intelligence manifests itself in logical thinking, computations, the memory capabilities of the brain, through the application of words and language rules or through the recognition of things and events.
- The combination of information, creativity, and new problem solutions is crucial for acting "intelligent".

Testing “Intelligence” with the Turing Test

- Turing test is a proposal to test a machine’s ability to demonstrate “intelligence”



Source: http://en.wikipedia.org/wiki/Turing_test

- Turing test proceeds as follows:
 - A human judge **C** engages in a natural language conversation with one human **B** and one machine **A**, each of which tries to appear human.
 - All participants are placed in isolated locations.
 - If the judge **C** cannot reliably tell the machine **A** from the human **B**, the machine is said to have passed the test.
 - In order to test the machine's intelligence rather than its ability to render words into audio, the conversation is limited to a text-only channel such as a computer keyboard or screen.
- Turing test is an operational test for intelligent behaviour. For more details see [2].

- The “Chinese room” experiment developed by John Searle in 1980 attempts to show that a symbol-processing machine like a computer can never be properly described as having a “mind” or “understanding”, regardless of how intelligently it may behave.
- With the “Chinese room” John Searle argues that it is possible to pass the Turing Test, yet not (really) think.

Source: http://en.wikipedia.org/wiki/Chinese_room

“Chinese Room” (1)

- The “Chinese room” experiment proceeds as follows:
 - Searle, a human, who does not know Chinese, is locked in a room with an enormous batch of Chinese script.
 - Slips of paper with still more Chinese script come through a slot in the wall.
 - Searle has been given a set of rules in English for correlating the Chinese script coming through with the batches of script already in the room.



- Searle is instructed to push back through the slot the Chinese script with which the scripts coming in through the slot are correlated according to the rules.
- Searle identifies the scripts coming in and going out on the basis of their shapes alone. He does not speak Chinese, he does not understand them.
- The scripts going in are called ‘the questions’, the scripts coming out are ‘the answers’, and the rules that Searle follows is ‘the program’.
- Suppose also that the set of rules, the program is so good and Searle gets so good at following it that Searle’s answers are indistinguishable from those of a native Chinese speaker.

- The result:
 - It seems clear that Searle nevertheless does *not* understand the questions or the answers.
 - But Searle is behaving just a computer does, “performing computational operations on formally specified elements”.
- ***Hence, manipulating formal symbols, which is just what a computer running a program does, is not sufficient for understanding or thinking.***

- Many definitions exist, among them:
 - “The study of the computations that make it possible to perceive, reason, and act” (Winston, 1992)
 - “A field of study that seeks to explain and emulate [human] intelligent behaviour in terms of computational processes” (Schalkoff, 1990)
- It is an interdisciplinary field that is based on results from philosophy, psychology, linguistics, or brain sciences
- Difference to “traditional” computer science: Emphasis on cognition, reasoning, and acting
- Generative theory of intelligence:
 - Intelligence emerges from the orchestration of multiple processes
 - Process models of intelligent behaviour can be investigated and simulated on machines



- Two main aspects begin to manifest in the early days of AI
 1. Cognitive modelling, i.e., the simulation of cognitive processes through information processing models.
 2. The construction of “intelligent systems” that make certain aspects of human cognition and reasoning available.

- Strong AI
 - “An artificial intelligence system can *think* and have a *mind*. “ (John Searle 1986)
 - “Machine intelligence with the full range of human intelligence” (Kurzweil 2005)
 - AI that matches or exceeds human intelligence
 - Intelligence can be reduced to information processing
 - “Science Fiction AI”
- Weak AI
 - Intelligence can partially be mapped to computational processes.
 - Intelligence is information processing
 - Intelligence can be simulated



Symbolic vs. Subsymbolic AI; Knowledge-based Systems

TECHNICAL SOLUTIONS



SYMBOLIC AI vs. SUBSYMBOLIC AI

- Research on Information Processing in AI by
 - Exact formulisations.
 - Exemplary realisation via implementations.
- Core aspect: Representation and processing of symbols as a foundation of internal processes.

- Symbols are naming objects which provide access to meaning (Newell, 1958)
- “Spoken words are the symbols of mental experience, and written words are the symbols of spoken words.” (Aristotle) [3]
- Mental abilities of humans can be inspected on a symbolic level independent of neuronal architectures or processes.
- Subject of Symbolic AI is thus the meaning of processes (or their symbolic representations respectively).
- Symbolic AI aims to imitate intelligence via formal models.
- Main persons behind symbolic AI are: Simon, Newell, Minsky

The “(General) Intelligent Agent”

- Core paradigm of symbolic AI is the “Intelligent Agent” [4]:
 - has a memory and the capability to act in his world based on it.
 - has sensors to perceive information from his environment.
 - has actuators to influence the external world.
 - has the capability to probe actions. By that he is able to choose the best possible action.
 - has internal memory for methods and the exploration of the world is guided by knowledge kept in it.




Image from Padgham/Winikoff “Developing Intelligent Agents (Wiley 2004)

- Subsymbolic AI (SSAI) aims to model intelligence empirically.
- SSAI was inspired by biological systems: A model which imitates neural nets in the brain is the basis for the creation of artificial intelligence.
- Neural nets consist of a network of neurons which have weighted connections with each other.
- Early work by Rosenblatt (1962): the “Perceptron” [6]
- Advantages of artificial neuronal nets:
 - Distributed representation
 - Representation and processing of fuzziness
 - Highly parallel and distributed action
 - Speed and fault-tolerance




Image: <http://www.neuronaletesnetz.de>




KNOWLEDGE-BASED SYSTEMS

-
1. General Problem Solver
 2. Knowledge-is-power hypothesis
 3. Knowledge levels
 - 3a. Newell's 3 levels of knowledge
 - 3b. Brachman's 5 levels of knowledge
 4. Problem Solving Methods

- The General Problem Solver (GPS) is a universal problem solving approach.
- GPS is the first approach that makes the distinction between *knowledge of problems domains* and *how to solve problems*.
- GPS is domain and task independent approach.
- GPS does not put any restrictions both on the domain knowledge and on the task.
- Examples of GPS are: *automated theorem proving* and *generic search methods*.

- Automatic theorem provers are GPS for which every problem can be expressed as logical inference
- Automated theorem proving is about proving of mathematical theorems by a computer program

More in Lecture 4



- Generic Search Methods are GPS for which every problem can be expressed as search.
- One particular example of a Generic Search Method is the **A*** algorithm.
- **A*** works for problems that can be represented as a state space i.e. a graph of states. Initial conditions of the problem are represented as *start state*, goal conditions are represented as *end state*.
- **A*** is an informed search or heuristic search approach that uses the estimation function:

$$f(n) = g(n) + h(n)$$

- $g(n)$ the cost to get from the *start state* to current state n
- $h(n)$ estimated cost to get from current state n to *end state*
- $f(n)$ estimated total cost from *start state* through current state n to the *end state*

More in Lecture 5

- **However, GPS has a set of limitations:**
 - It works in theory but in practice works only on toy problems (e.g. Tower of Hanoi)
 - Could not solve real-world problems because search was easily lost in the combinatorial explosion of intermediate states

2. Knowledge-is-power hypothesis



Knowledge-is-power hypothesis, also called the **Knowledge Principle** was formulated by E.A. Feigenbaum in 1977:

“knowledge of the specific task domain in which the program is to do its problem solving was more important as a source of power for competent problem solving than the reasoning method employed” [15]

2. Knowledge-is-power hypothesis (1)

- The Knowledge-is-power hypothesis shifted the focus on how to build intelligent systems from *inference* to the *knowledge base*.
- Problem solving is guided by experiential, qualitative, heuristic knowledge.
- The meaning of intelligence as knowledge is the common meaning in English world.
- The Central Intelligence Agency (CIA) defines intelligence as knowledge.
- The Knowledge-is-power hypothesis lead to the emergence of a new field i.e. *expert systems* and a new profession i.e. *knowledge engineer*.

3. Knowledge levels



3a. Newell's 3 levels of knowledge

3b. Brachman's 5 levels of knowledge

-
- In his work from 1981, Newell tried to answer questions such as
 - How can knowledge be characterised?
 - What is the relation of this characterisation and the representation of knowledge?
 - What is characteristic about a system which holds knowledge?
 - Newell distinguished 3 levels in the context of knowledge representation:
 - **Knowledge Level**
 - **Logical Level**
 - **Implementation Level**

- **Knowledge Level**

- The most abstract level of representing knowledge.
- Concerns the total knowledge contained in the Knowledge Base.

- Example:

The automated DB-Information system knows that a trip from Innsbruck to Vienna costs 120€

- **Logical Level**

- Encoding of knowledge in a formal language.

- Example:

Price(Innsbruck, Vienna, 120)

- **Implementation Level**
 - The internal representation of the sentences.
 - Example:
 - *As a String “Price(Innsbruck, Vienna, 120)”*
 - *As a value in a matrix*

3b. Brachman's 5 Levels of Knowledge [12]



- Brachman defines 5 levels for different types of representations.
- Levels interpret the transition from data to knowledge.
- Each level corresponds to an explicit set of primitives offered to the knowledge engineer.
- Ordering of knowledge levels from simple/abstract to complex/concrete:
 - **Implementational Level**
 - **Logical Level**
 - **Epistemological Level**
 - **Conceptual Level**
 - **Linguistic Level**

- **Implementational Level**
 - The primitives are pointers and memory cells.
 - Allows the construction of data structures with no a priori semantics.

- **Logical Level**

- The primitives are logical predicates, operators, and propositions.
- An index is available to structure primitives.
- A formal semantic is given to primitives in terms of relations among objects in the real world.
- No particular assumption is made however as to the nature of such relations.

- **Epistemological Level**

- The primitives are concept types and structuring relations.
- Structuring relations provide structure in a network of conceptual types or units. (i.e. inheritance: conceptual units, conceptual sub-units).
- The epistemological level links formal structure to conceptual units.
- It contains structural connections in our knowledge needed to justify conceptual inferences.



- **Conceptual Level**


- The primitives are conceptual relations, primitive objects and actions.
- The primitives have a definite cognitive interpretation, corresponding to language-independent concepts like elementary actions or thematic roles.

- **Linguistic Level**

- The primitives are words, and (linguistic) expressions.
- The primitives are associated directly to nouns and verbs of a specific natural language.
- Arbitrary relations and nodes that exist in a domain.


- Problem Solving Methods (PSM) abstract from details of the implementation of the reasoning process.
- Characteristics of PSM [10]:
 - A PSM specifies which inference actions have to be carried out for solving a given task.
 - A PSM determines the sequence in which these actions have to be activated.
 - Knowledge roles determine which role the domain knowledge plays in each inference action.

- Generic inference pattern “Heuristic Classification” describes the problem-solving behaviour of these systems on the Knowledge Level in a generic way.



See [16] for more details

Propose & Revise



- The *propose & revise* method is an efficient method for solving the task of *parametric design*. (see more details in [14])
- The method depends on the following inferences:
 - **propose** – derives an initial design based on the requirements;
 - **C-test** – requires knowledge that describes which possible designs are valid (i.e., the domain constraints);
 - **revise** – tries to improve an incorrect design based on the feedback of the C-test step.

More in Lecture 7

- KBS are realized based on a **knowledge base** (KB).
- KB contains a model represented in a (logical) formalism which can be interpreted by an interpreter (**inference engine**) that is able draw conclusions from it.
- KBs typically capture knowledge of a domain.
- Methodologies for the development of KBS: e.g. CommonKADS
- Examples: CYC
 - One of the first systems that aimed to capture common knowledge in a knowledge base


- Special form of a KBS.
- Definition: An expert system is a software application that stores knowledge about a certain domain. It is able to draw conclusions from that knowledge and offers concrete solutions for problems in that domain.
- ES simulate human experts and thus the knowledge base typically consists of highly specialized expert knowledge.
- Reasoning of human experts vs. reasoning in ES:
 - Human experts are able to master unforeseen effects and situations.
 - Human experts are able to learn from experiences.
 - Human experts expand their knowledge continuously.
 - Human experts derive new knowledge not only based on drawn conclusions but via analogy and intuition.



POPULAR AI SYSTEMS

- Early computer program capable of natural language processing.
- Written by J. Weizenbaum between 1964 and 1966.
- ELIZA simulated a psychotherapist by reformulating questions posed by the user.
- Sample ELIZA conversation:
(Source: Wikipedia)

More information: [9]



```
File Edit Options Buffers Tools Help
I am the psychotherapist. Please, describe your problems.
Each time
you are finished talking, type RET twice.

I have a terrible headache.

Maybe you should consult a doctor of medicine, I am a psychiatrist.

Okay, I'll do that then.

Can you elaborate on that?

What are you talking about? No I cannot elaborate. See you later.

Why do you say that?

Now you're just talking nonsense!


What makes you believe now I am just talking nonsense?

You are not making any sense.

Perhaps your plans have something to do with this.
-u:** *doctor* (Doctor Fill)--L1--Top-----
X
```

SIRI

- Personal assistant - mobile app for the iPhone
"Just like a real personal assistant Siri understands what you say, fulfills tasks and adapts over time to the preferences. Today, Siri can help you find things and plan. "
- Launched in February 2010
- Siri was acquired by Apple in April 2010




- Chess-playing computer developed by IBM that won against world champion Garry Kasparov in 1997.
- Applied a brute force strategy, processing was highly parallel.
- Evaluation of 200 million positions per second.
- Deep Blue's knowledge base contained over 4,000 positions and 700,000 grandmaster games.
- It was fine-tuned by chess grand masters.
- Admission from IBM: „Deep Blue, as it stands.
- today, is not a "learning system." It is therefore not capable of utilizing artificial intelligence to either learn from its opponent or "think" about the current position of the chessboard.“



Link: <http://www.research.ibm.com/deepblue/>

The Humanoid Robot COG

- Project at the MIT Artificial Intelligence Lab
- The goal of the COG project was to build a robot capable of interacting with humans and objects in a human-like way.
- "As I pondered [this] and thought about HAL, I decided to try to build the first serious attempt at a robot with human-level capabilities, the first serious attempt at a HAL-class being." (Rodney Brooks, Inventor of COG)



Link: <http://groups.csail.mit.edu/lbr/humanoid-robotics-group/cog/>


CALO („Cognitive Assistant that Learns and Organizes“)

- DARPA funded project, “Personal assistant that learns” – program
- Involves 25 partners, 300+ researchers, including top researchers in AI
- 500+ publications in first four years
- “The goal of the project is to create cognitive software systems, that is, systems that can reason, learn from experience, be told what to do, explain what they are doing, reflect on their experience, and respond robustly to surprise. “ (calosystem.org)
- CALO assists its user with six high-level functions:
 - *Organizing and Prioritizing Information*
 - *Preparing Information Artifacts*
 - *Mediating Human Communications*
 - *Task Management*
 - *Scheduling and Reasoning in Time*
 - *Resource allocation*



Link: <http://www.calosystem.org/>

- An advanced device capable of performing a variety of tasks and interacting with its human users (companions?).
- The HAL9000 communicates by voice and can control auxiliary devices on a spaceship.
- It (he?) has an unfortunate tendency towards obsessing over minor details or inconsistencies in the instructions given it, however.
- In the events described in Arthur C. Clarke's "2001: A Space Odyssey," HAL's tendency toward obsessive literalism led to the unfortunate death of most of its spaceship's human crew.



- SEAS (“Synthetic Environment for Analysis and Simulation”)
 - Can be used to simulate realistic events; has a world model
 - <http://www.krannert.purdue.edu/centers/perc/html/aboutperc/seaslabs/seaslabs.htm>
- SYSTRAN
 - Early machine translation system
 - Foundation for Yahoo’s Babelfish or Google Translator
 - <http://www.systransoft.com/>
- VirtualWoman
 - Virtual-reality based chatbot
 - <http://virtualwoman.net/>
- For further references, see http://en.wikipedia.org/wiki/List_of_notable_artificial_intelligence_projects




SUBDOMAINS OF AI

- Cognition as information processing
- Artificial neuronal networks
- Heuristic search methods
- Knowledge representation and logic
- Automatic theorem proving
- Non-monotonic reasoning
- Case-based reasoning
- Planning
- Machine Learning
- Knowledge Engineering
- Natural Language Processing
- Image Understanding
- Cognitive Robotics
- Software Agents

- Deals with complex software systems that directly interact and communicate with human users.
- Characteristics of cognitive systems (CS):
 - CS are directly integrated in their environment, act in it, and are able to communicate with it.
 - CS are able to direct and adapt their actions based on the environment they are situated in.
 - CS typically represent system-relevant aspects of the environment.
 - Their information processing capabilities are characterized through learning aptitude and anticipation
- Examples of cognitive system:
 - Organisms / biological cognitive systems
 - Technical systems such as robots or agents
 - Mixed human-machine systems

- Neural networks are networks of neurons as in the real biological brain.
- **Neurons** are highly specialized cells that transmit impulses within animals to cause a change in a target cell such as a muscle effector cell or glandular cell.
- The **axon**, is the primary conduit through which the neuron transmits impulses to neurons downstream in the signal chain
- Humans: 10^{11} neurons of > 20 types, 10^{14} synapses, 1ms-10ms cycle time
- Signals are noisy “spike trains” of electrical potential




- What we refer to as Neural Networks in the course are mostly Artificial Neural Networks (ANN).
- ANN are approximation of biological neural networks and are built of physical devices, or simulated on computers.
- ANN are parallel computational entities that consist of multiple simple processing units that are connected in specific ways in order to perform the desired tasks.
- Remember: **ANN are computationally primitive approximations of the real biological brains.**
- **Application examples:** e.g., handwriting recognition, time series prediction, kernel machines (support vector machines, data compression, financial predication, speech recognition, computer vision, protein structures.

- Search Methods are typically helping humans to solve complex tasks by generating (optimal) plans (i.e. a set of operations / states) that includes sequences / actions to reach a goal state.

- Example problem: Tower of Hanoi

- Initial status: ((123)())
- Goal status: (())(123)



- Definition: A search method is defined by picking the order of node expansion.
- Search strategies are evaluated according to completeness, time complexity, space complexity, optimality.
- Time and space complexity are measured in terms of maximum branching, depth of the least-cost solution, maximum depth of the state space.
- Distinction between informed / uninformed search techniques

- The term *knowledge representation* describes the design and implementation of formalisms, to model a part of the reality (a domain).
- A model represented using formalisms and implemented by an interpreter is often called a knowledge base.
- A knowledge base is a collection of facts and beliefs.
- Modelling of knowledge bases happens on a conceptual level.
- Intention: To model a domain of discourse and to draw inferences about the objects in the domain (reasoning)
- Logic studies the principles of reasoning and offers
 - Formal languages for expressing knowledge
 - Well understood formal semantics
 - Reasoning methods to make implicit knowledge explicit

- Automatic theorem proving deals with the design and implementation of computer programmes that are capable of making mathematical proofs.
- Theorem provers deduce new formulas from given formulas via logical deduction rules until the target formula is found.
- Theoretical foundation of automated theorem proving: mathematical logic; typically first-order-logic.
- Formulas are mathematically precisely defined via interpretations (provide semantics for function and predicate symbols via mappings and relations respectively).

- Classical logic is monotonic in the following sense: whenever a sentence A is a logical consequence of a set of sentences T , then A is also a consequence of an arbitrary superset of T [13].
- Non-monotonic reasoning:
 - Additional information may invalidate conclusions.
 - Non-monotonic reasoning is closer to (human) common-sense reasoning.
 - Most rules in common-sense reasoning only hold with exceptions (i.e. `university_professors_teach`)
- Important approaches to formalise non-monotonic reasoning:
 - Default-Logics: Non-classical inference rules are used to represent defaults.
 - The modal approach: Modal operators are used to explicitly declare if something is believed in or is consistent.
 - Circumscription: Validity can be restricted to specific models.
 - Conditional approaches: A conditional junctor is used to represent defaults in a logical language.

- Definition: A “case” is an experience made during the solving of a problem.
- A case is typically informally given and covers the problem and the solution.
- Experiences (resp. cases) are used to solve newly occurring problems.
- Cases are collected in a so-called case-base (analogous to a knowledge base in KBS).
- Case-based reasoning is inspired by human problem solving capabilities.
- Application scenarios are characterized through:
 - A considerable amount of cases has to be available
 - Using the cases to solve the problem has to be easier than solving the problem directly.
 - Available information is incomplete or unsecure and imprecise.
 - The construction of a KBS and the modelling of cases is not easily possible.
- Typical application scenarios can be found in the area of diagnostics, electronic sales, configuration, or planning.

- What is Planning?
 - “*Planning is the process of thinking about the activities required to create a desired goal on some scale*” [Wikipedia]
- We take a more pragmatic view – *planning is a flexible approach for taking complex decisions*:
 - decide about the schedule of a production line;
 - decide about the movements of an elevator;
 - decide about the flow of paper through a copy machine;
 - decide about robot actions.
- By “flexible” we mean:
 - the problem is described to the planning system in some generic language;
 - a (good) solution is found fully automatically;
 - if the problem changes, all that needs to be done is to change the description.
- Planning looks at methods to solve any problem that can be described in the language chosen for the particular planning system.
- Approaches for the generation of action sequences: action planning and situated activity planning.

- Machine Learning (ML) is a central research area in AI to acquire knowledge.
- ML deals with the computer-aided design and realisation of learning problems.
- Learning is defined as the process that enables a systems to perform better during solving of the same or similar tasks in the future (Simon, 1983)
- Reduction of learning to mathematical theories: Deduction, Induction, Abduction.
- Learning task is typically characterized through the description of inputs, expected outputs, and environmental conditions.
- Typical machine learning applications: Data mining, Speech recognition, text analysis, control learning, hidden markov networks, etc.

- Knowledge engineering is concerned with the acquisition, management, use and transformation of knowledge.
- Goals are similar to software engineering, i.e. to systematically develop expert systems using existing methods and tools.
- Core process in knowledge engineering: knowledge acquisition; During knowledge acquisition knowledge is formalised, i.e. transformed from a natural language representation to a formal representation.
- Process models for knowledge acquisition: Model by Puppe; model by Buchanan; Harmon/Mauss/Morrissey; Waterman; or MIKE.
- Methodical approaches and tools: D3; CommonKADS; MIKE; Protégé-II; RASSI.
- Application cases include the development of expert systems, workflow systems or knowledge management.

-
- Goal: Processing and understanding of speech or written language.
 - Early applications include question-answer systems, natural-language based access to databases or speech-based control of robots.
 - Challenges include information re-construction from spoken words or information selection and reduction during speech production.
 - Application areas: Tools for inter-human communication, tools for text generation or text correction (i.e. identification of grammatical errors based on language models), information classification or filtering, or human-machine communication.

- Image Understanding (IU) deals with the analysis and interpretation of visual information. IU denotes the reconstruction and interpretation of scenes based on images.
- Early approaches based on pattern recognition (still one of the most important foundations of this field).
- Prominent application: object recognition of still and moving objects



- Application areas: symbol recognition, medical image analysis, vehicle navigation, image archiving, gesture recognition,...

-
- AI deals with the development of robots as autonomous and intelligent systems.
 - Robotic covers many sub-areas of AI and involves interdisciplinary work including mechanic and electrical design and cognitive areas.
 - Types of robots: static robots, mobile robots, and humanoid robots.
 - Application areas: construction, planning, or observation.

- Core paradigm in AI.
- Definition: A software agent is a long-term operating program whose function can be described as autonomous execution of tasks or tracing of goals via interaction with his environment.
- Agent (see earlier slide)
 - has a memory and the capability to act in his world based on it.
 - has sensors to perceive information from his environment.
 - has actuators to influence the external world.
 - has the capability to probe actions
 - has internal memory for methods and the exploration of the world is guided by knowledge kept in it.
- Applications: Data collection and filtering, event notification, planning and optimization in various application areas (commerce, production, military, education),...




SOME RELEVANT PEOPLE IN AI

Some relevant people in AI

- Isaac Asimov (<http://www.asimovonline.com/>)
- Arthur C. Clark (<http://www.clarkefoundation.org/>)
- John McCarthy (<http://www-formal.stanford.edu/jmc/>)
- Marvin Minsky (<http://web.media.mit.edu/~minsky/>)
- Donald Michie (<http://www.aiai.ed.ac.uk/~dm/dm.html>)
- Allen Newell (<http://www.princeton.edu/~hos/frs122/newellobit.html>)
- Herbert A. Simon (<http://www.psy.cmu.edu/psy/faculty/hsimon/hsimon.html>)
- Alan Turing (<http://www.turing.org.uk/turing/>)



Asimov



Clark




McCarthy



Minsky



Michie



Newell



Simon



Turing



SUMMARY

- Birth of AI in the 1950s
- Broad spectrum of subdomains and combination of disciplines
- Distinction between
 - Weak and strong AI
 - Symbolic and subsymbolic AI
- Central role: symbols and knowledge representation
- Knowledge-based systems and intelligent agents are core concepts in AI



REFERENCES

- Mandatory reading:
 - [1] G. Görz, C.-R. Rollinger, J. Schneeberger (Hrsg.) “Handbuch der künstlichen Intelligenz” Oldenbourg Verlag, 2003, Fourth edition
- Further reading:
 - [2] A. Turing. "Computing Machinery and Intelligence", Mind LIX (236): 433–460, October, 1950.
 - [3] Aristotle “On Interpretation”, 350 B.C.E, see: <http://classics.mit.edu/Aristotle/interpretation.html>
 - [4] A. Newell, H.A. Simon, “Human Problem Solving” Englewood Cliffs, N.J.: Prentice Hall, 1972
 - [5] A. Newell. “The Knowledge Level”, AI Magazine 2 (2), 1981, p. 1-20.
 - [6] F. Rosenblatt. “Strategic Approaches to the Study of Brain Models” In: Förster, H.: Principles of Self-Organization. Elmsford, N.Y.: Pergamon Press, 1962.
 - [7] S. Russell, E.H. Wefald. "Do the Right Thing: Studies in Limited Rationality" MIT Press, 1991.
 - [8] C. Beierle and G. Kern-Isberner "Methoden wissensbasierter Systeme. Grundlagen, Algorithmen, Anwendungen" Vieweg, 2005.

- [9] J. Weizenbaum. "ELIZA - A Computer Program For the Study of Natural Language Communication Between Man And Machine", *Communications of the ACM* 9 (1): p. 36–45, 1966.
- [10] W. Birmingham and G. Klinker "Knowledge Acquisition Tools with Explicit Problem-Solving Methods" *The Knowledge Engineering Review* 8, 1 (1993), 5-25
- [11] A. Newell and H. Simon "GPS, a program that simulates human thought" In: *Computation & intelligence: collected readings*, pp. 415 - 428, 1995.
- [12] R. J. Brachman "On the Epistemological Status of Semantic Networks" In: N.V. Findler (ed.): *Associative Networks: Representation and Use of Knowledge by Computers*. New York: Academic Press, 1979, 3-50.
- [13] G. Brewka, I. Niemelä, M. Truszczyński "Nonmonotonic Reasoning" In: V. Lifschitz, B. Porter, F. van Harmelen (eds.), *Handbook of Knowledge Representation*, Elsevier, 2007, 239-284
- [14] D. Fensel "Problem-Solving Methods: Understanding, Description, Development and Reuse", Springer LNAI 1791, 2000
- [15] E.A. Feigenbaum. "The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering," *Proceedings of the International Joint Conference on Artificial Intelligence, Cambridge, MA, 1977*
- [16] W.J. Clancey. "Heuristic Classification", *Artificial Intelligence*, 27:289-350, 1985

- Wikipedia links:
 - http://en.wikipedia.org/wiki/List_of_notable_artificial_intelligence_projects
 - http://en.wikipedia.org/wiki/Turing_test
 - http://en.wikipedia.org/wiki/General_Problem_Solver

Next Lecture

#	Title
1	Introduction
 2	Propositional Logic
3	Predicate Logic
4	Reasoning
5	Search Methods
6	CommonKADS
7	Problem-Solving Methods
8	Planning
9	Software Agents
10	Rule Learning
11	Inductive Logic Programming
12	Formal Concept Analysis
13	Neural Networks
14	Semantic Web and Services

Questions?

