Foundations of AI

1. Introduction

Organizational, AI in Freiburg, Motivation, History, Approaches, Examples

Wolfram Burgard, Bernhard Nebel

Organizational

Lectures:

- **Time:**

  Tu 9:15 – 11:00
  Fr 9:15 – 10:00

- **Lecturers:**

  Prof. Dr. Wolfram Burgard
  (burgard@informatik.uni-freiburg.de)

  Prof. Dr. Bernhard Nebel
  (nebel@informatik.uni-freiburg.de)

Exercises:

- **Time:**

  Fr 10:00-11:00

- **Organizers:**

  Alexander Kleiner, Robert Mattmüller

- Exercises handed out on Fridays, submitted the following week
- Exercise sheets may be solved in groups of three students

Credit Requirements:

- Written exam
- Bonus marks for reasonably solved exercises

Lecture Material

Lectures are based on

Artificial Intelligence – A Modern Approach, 2nd Edition

Stuart Russell - Peter Norvig

Copies of the lecture slides as well as further information can be found on the Web at

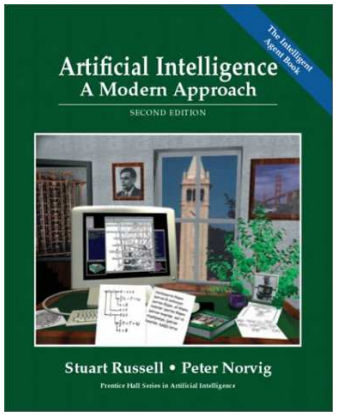

`http://informatik.uni-freiburg.de/~ki/teaching/ss07/gki/index.html`

Many illustrations have been taken from the above book. Some slides are based on presentations written by Prof. Gerhard Lakemeyer, Univ. Aachen. Many sections were prepared by Dr. Jana Köhler.

Course Contents

Strongly method-oriented

1. Introduction
2. Rational Agents
3. Decisions under Uncertainty
4. Acting under Uncertainty
5. Solving Problems by Searching
6. Informed Search
7. Constraint Satisfaction Problems
8. Games

1. Propositional Logic
2. Satisfiability and Model Construction
3. Predicate Logic
4. Modeling with Logic
5. Machine Learning
6. Planning

AI in Freiburg

Foundations of Artificial Intelligence
Bernhard Nebel

Machine Learning
and Natural Language Processing
(formerly: Luc de Raedt)

Autonomous Intelligent Systems
Wolfram Burgard

Foundations of Artificial Intelligence

- Action Planning: Theory and Practice
  - Fast planning systems (proven at int. competition)
  - Applications at airports and for lift systems
  - Theoretical results (see new Russell/Norvig)
- Qualitative Temporal-Spatial Reasoning
  - Theory and reasoning algorithms
  - Application in qualitative layout description
  - SFB
- RoboCup
  - World champion three times
  - Autonomous table soccer
  - RoboCup Rescue
    (Multi-Agent-System for disaster relief)

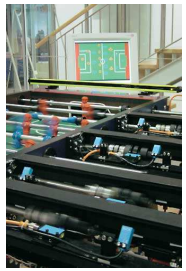

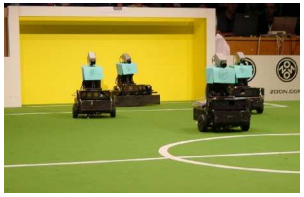

Machine Learning and Natural Language Processing

Focus on Machine Learning and Data Mining, in particular

- inductive logic programming (ILP): learning and data mining using first order logical or relational representations
- inductive databases: knowledge discovery as an extended querying process
- probabilistic ILP: Bayesian networks, hidden Markov models, and reinforcement learning in an ILP setting
- application to problems in bio- and chemo-informatics.

Autonomous Intelligent Systems

Mobile robot navigation, perception, action planning, modeling, learning, …

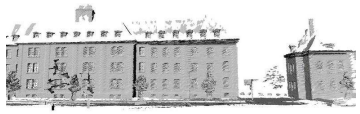

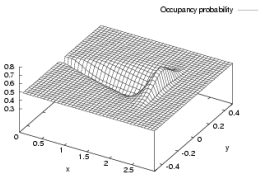

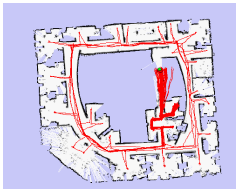

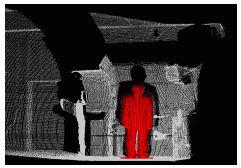

If You Want to Do AI in Freiburg

- Foundations of Artificial Intelligence
- Machine Learning and Data Mining
- Knowledge Representation
- Autonomous Mobile Systems
- AI Planning
- Logic
- Game Theory
- …

What is Artificial Intelligence?

- The attempt to make computers more "intelligent"
- The attempt to better understand human intelligence
- Four Approaches:
  - Is it about thought …
  - … or action?
  - Oriented towards a human model (with all its defects) …
  - … or normative (how should a rational being think/act)?

A Few Definitions

"The exciting new effort to make computers think … *machines with minds*, in the full and literal sense" (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning …" (Bellman, 1978)	"The study of mental faculties through the use of computational models" (Charniak and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act" (Winston, 1992)
"The art of creating machines that perform functions that require intelligence when performed by people" (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991)	"A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes" (Schalkoff, 1990) "The branch of computer science that is concerned with the automation of intelligent behavior" (Luger and Stubblefield, 1993)

The Turing Test

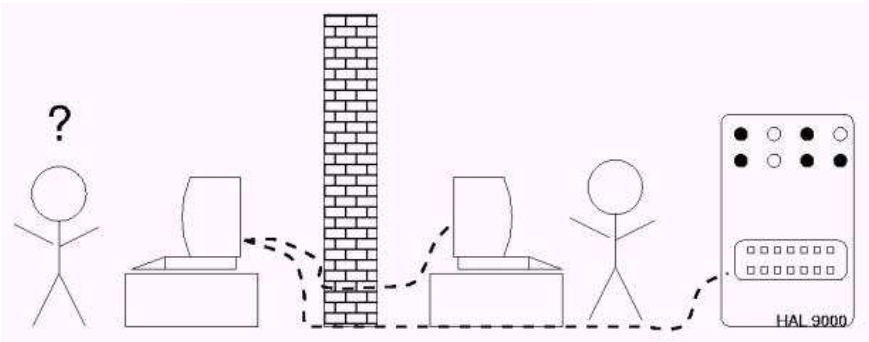

1999 Winner: Richard Wallace (http://www.alicebot.org)

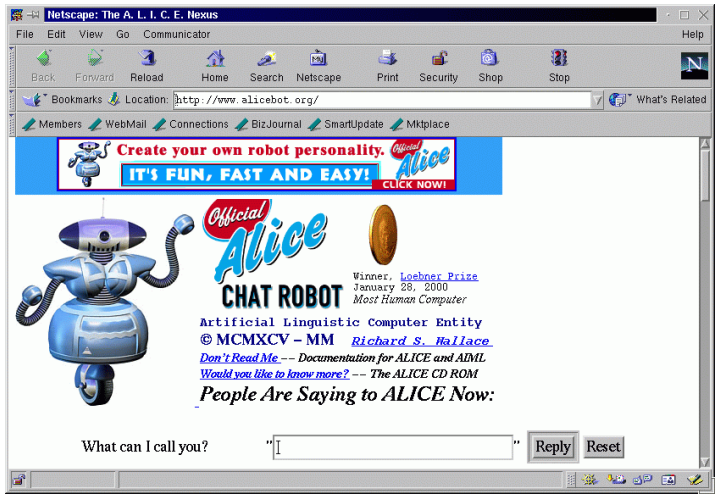

> Wolfram

Nice to meet you Wolfram.

> Do you know Clinton

He was the second Impeached President of the United States.

> Do you like him

Sure he has never done anything to me.

> What is an impeachment

It is a good question. I have to process that one for a while.

I would do a search for it.

> Who are you

I ask myself that question, "Who am I?"

> What are you doing

I am chatting with people on the net.

> What is your favorite topic

My favorite subject is artificial intelligence.

> What is artificial intelligence

Artificial intelligence is the branch of engineering and science devoted to constructing machines that think. I can ask someone about it.

> Can you do anything else than chatting

Doing anything else than chatting is not among my strongest traits.

...

...

> How old am I

Is this a trick question? You said you were how many years old?

My guess is that you are really a kid.

> What would you like to know about me

Tell me about your educational background.

> I am a professor in computer science

You are a professor in computer science? How do you like your work?

> It is fun teaching courses on artificial intelligence

I think it is a lot of fun.

> What should I tell the students about you?

Be sure to think before you speak. State your point clearly and slowly and gauge the listener's response before going any further.

Systems that Act Humanly

Reproducing human behaviour

→ Prototype: Turing Test

→ … in AI, scarcely pursued

→ … but yearly competitions exist, e.g. Loebner Prize:

`http://www.loebner.net/Prizef/loebner-prize.html`

Systems that Think Humanly

What cognitive capabilities are necessary to produce intelligent performance?

- Not important: Being able to solve problems correctly
- Important: Being able to solve problems like a human would

→ Cognitive science and cognitive psychology

→ Also important for HMI

- … will not be discussed in this course

Systems that Think Rationally

- What are the laws of thought?
- How should we think?

→ The logical approach

→ Problems:

  - Presentation of problem descriptions using a formal notation
  - Computability

→ These are problems that appear regardless of the formalization method

Systems that Act Rationally

→ Rational agents (or rational actors)

- A rational agent acts so as to achieve its given goals, under the assumption that its impressions of the world and its convictions are correct
- Rational thinking is a prerequisite for rational acting, although it is not a necessary condition

→ What to do, for example, when we must make a decision faced with insufficient information?

The AI Scene

Fields of Application	Methods
• Systems that can understand and generate speech • Systems that can understand images • Robotics • Assistant systems	• Problem solving and searching • Knowledge representation and processing • Action planning • Machine learning • Handling uncertain knowledge • Neural networks

With interdisciplinary relationships to Mathematics, Philosophy, Psychology, (Computational) Linguistics, Biology, Engineering Sciences, …

The Origins of AI

Since the beginning, Philosophy, Mathematics, Psychology, Linguistics, and Computer Science have all

- asked similar questions
- developed methods and produced results for AI

The origins of AI (1943-1956): With the development of the first computing systems, people began to wonder, "Can computers copy the human mind? (→ Turing Test)"

50 Years of AI (1)

1956: Dartmouth Workshop – McCarthy proposes the term, "Artificial Intelligence" – and early enthusiasm:

> It is not my aim to surprise or shock you – but the simplest way I can summarize is to say that there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until – *in the visible future* – the range of problems they can handle will be coextensive with the range to which the human mind has been applied. [Simon, 1957]

60's: "Intelligent Behavior" is shown in many demonstration systems for microworlds (Blocks world)

70's: Problems:

- Systems for microworlds prove unscalable → "real" applications
- "Intelligent Behaviour" requires much knowledge → knowledge-based systems

50 Years of AI (2)

80's: Commercial success of experimental systems (e.g. R1), intense research support (e.g. *Fifth generation computer systems project* in Japan), return to neural networks

End of the 80's: Expert systems prove less promising than imagined, (demystification of expert systems), end of the *Fifth generation computer systems project, "AI Winter"*

90's: Inclusion of probabilistic methods, agent-oriented vision techniques, formalization of AI techniques and increased use of mathematics in the field

> … gentle revolutions have occurred in robotics, computer vision, machine learning (including neural networks), and knowledge representation. A better understanding of the problems and their complexity properties, combined with increased mathematical sophistication, has led to workable research agendas and robust methods. [Russell & Norvig, 1995]

… and Today?

- Many coexisting paradigms
  - → Reactive vs. deliberative approaches
  - → (Robotics)
  - → Probabilistic vs. analytic (Computational Linguistics)
  - → … often hybrid approaches as well
- Many methods (partly from other disciplines):
  - → Logic, decision theory, algorithms
- Many approaches:
  - → Theoretical, algorithmic experimentation, system-oriented
- Today, many methods are no longer regarded as pure AI methods. Examples: Board game programs, logic programming (PROLOG), search procedures, …

Examples: Algorithmic, Experimental Tasks

Many AI problems are inherently difficult (NP-hard), but it is possible, in spite of this and with the use of good search techniques and heuristics, to solve problem instances up to a certain size:

- Satisfiability of boolean formulas
  - → Randomized, local search techniques (up to 2500 variables in complex instances)
- Constraint propagation and backtracking techniques
  - → Empirical and analytical comparisons of various techniques
- Action planning
  - → Empirical comparisons of various approaches and systems
- …

Systems

Alongside theory and the analysis of individual algorithms, the building of systems and applications is a basic point:

Herb Simon in a lecture entitled “How to become a good scientist” (1998):

“Build a System”

- Application of AI techniques to solve real problems
- Study of the interaction of artefacts with their environment
- Synergetic effects in systems