

Foundations of AI

1. Introduction

Organizational aspects, AI in Freiburg,
motivation, history, approaches,
and examples

*Wolfram Burgard, Andreas Karwath,
Bernhard Nebel, Martin Riedmiller,*

Organizational

Lectures:

- **Time:**

Tue 9:15 – 11:00
Fr 9:15 – 10:00

- **Lecturers:**

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

Dr. Andreas Karwath
(karwath@informatik.uni-freiburg.de)

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

Prof. Dr. Martin Riedmiller
(riedmiller@informatik.uni-freiburg.de)

Home Work & Tutorials:

- **Time:**

Fr 10:15-11:00

- **Organizers:**

**Barbara Frank, Andreas Karwath,
and Gabi Röger**

- Home work sheets will be handed out on every Tuesday, and need to be submitted one week later
- Home work assignments may be solved in groups of three students

Credit Requirements:


- Written exam
- Bonus marks for reasonably solved exercises (see Web page)

Lecture Material

Lectures are based on

*Artificial Intelligence - A
Modern Approach, 2nd Edition*

Stuart Russell - Peter Norvig



Copies of the slides, online recordings and further information can be found on the Web at

<http://www.informatik.uni-freiburg.de/~ki/teaching/ss09/gki/vorlesung.html>

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

Course Contents

Strongly method-oriented

1. Introduction
2. Rational Agents
3. Solving Problems by Searching
4. Informed Search
5. Constraint Satisfaction Problems
6. Games
7. Propositional Logic
8. Satisfiability and Model Construction
9. Predicate Logic
10. Modeling with Logic
11. Machine Learning
12. Planning
13. Simple Probabilistic Reasoning
14. Reasoning under Uncertainty
15. Decisions under Uncertainty
16. Acting under Uncertainty


AI in Freiburg

Foundations of Artificial Intelligence
Bernhard Nebel




Machine Learning
Martin Riedmiller

Autonomous Intelligent Systems
Wolfram Burgard



Humanoid Robots
Maren Bennewitz

Social Robotics
Kai Arras



PostDocs



Malte Helmert
(AI & Planning)



Andreas Karwath
(ML and Data Mining)



Alexander Kleiner
(Rescue Robotics)



Jan-Georg Smaus
(Verification)




Stefan Wölfl
(Knowledge Representation)



Cyrill Stachniss
(Robotics)

Foundations of Artificial Intelligence

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




Autonomous Intelligent Systems

- Mobile robots
- State estimation and models
- Adaptive techniques and learning
- Multi-robot systems
- Applications of mobile robots
- Robots and embedded systems
- Interaction and Web interfaces
- Probabilistic robotics





Machine Learning Lab

- Reinforcement Learning
- Supervised Learning
- Efficient Learning Algorithms
- Learning in Multi-Agent systems
- Self-learning robots
- Neural Forecasting Systems
- Neural Controllers
- Learning soccer robots in RoboCup
- Industrial Applications



Humanoid Robots


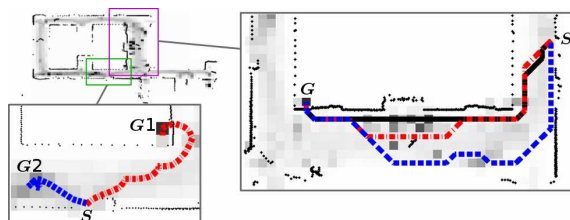


- Development of techniques for robots with
 - human-like body plan
 - human-like senses
- Natural human-robot interaction
 - State estimation and modeling of people
 - Speech, gestures, facial expression, etc.
- Navigation in complex indoor environments
 - 3D environment modeling
 - Path planning
 - Classification and learning



Social Robotics Lab

- Towards socially compatible robots
- Social learning, learning by observation
- People detection and tracking
- Motion planning
- Robot navigation
- Spatio-temporal models of human social behavior
- Human-robot interaction

"Free robots from their social isolation"



Studying AI in Freiburg

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

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 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)


“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)

“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



Systems that Act Humanly

Reproducing human behavior

→ Prototype: Turing Test

→ ... in AI, scarcely pursued


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



<http://www.loebner.net/Prizef/loebner-prize.html>

1999 Winner

Richard Wallace (<http://www.alicebot.org>)



The screenshot shows a Netscape browser window titled "Netscape: The A. L. I. C. E. Nexus". The address bar contains "http://www.alicebot.org/". The page features a blue banner at the top with the text "Create your own robot personality. Official Alice IT'S FUN, FAST AND EASY! CLICK NOW!". Below the banner is a blue robot character on the left and a gold coin on the right. The text "Official Alice CHAT ROBOT" is prominently displayed. Below this, it states "Winner, [Loebner Prize](#) January 28, 2000 *Most Human Computer*". Further down, it lists "Artificial Linguistic Computer Entity", "© MCMXCV - MM [Richard S. Wallace](#)", and provides links for "Don't Read Me -- Documentation for ALICE and AIML" and "Would you like to know more? -- The ALICE CD ROM". At the bottom, there is a chat interface with the prompt "What can I call you?" and a text input field with a "Reply" button and a "Reset" button.

Example Session with Alice (1)

What is your name?

> 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.

...

Example Session with Alice (2)

...

> 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.

Example Session with God (1)

Human: Do you know Obama?

God: I don't know if I know it. Where would I have heard about it?

Human: He is the current president of the United States.

God: I did not know is it.

Human: Do you know Clinton?

God: He was the second Impeached President of the United States.

Human: What is an impeachment?

God: Charge against a public official of improper conduct.

...

Human: Who are you?

God: I am the Almighty, God himself.

Human: What are you doing?

God: I am just killing time.

Human: What is your favorite topic?

God: My favorite subject is artificial intelligence.

Human: What is artificial Intelligence?

God: Artificial intelligence is the branch of engineering and science devoted to constructing machines that think. That's not something I get asked all the time.

Human: Can you do anything else than chatting?

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

Example Session with God (2)

Human: My name is Wolfram.

God: OK I will call you Wolfram.

Human: Who am I?

God: You are Wolfram, seeker.

Human: How old am I?

God: Is this a trick question? You said you were years old? My guess is that you are really a kid.

Human: What would you like to know about me?

God: Tell me about your educational background.

Human: I am a professor in Computer Science

God: What is your field?

Human: Artificial Intelligence

God: Do you mind if I tell other people you are teaching Artificial Intelligence.

Human: No.

God: OK then I will add it to my gossip files.

Human: What should I tell the students about you?

God: I believe it is always best to say positive things to build people up rather than negative things to bring people down.

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
<ul style="list-style-type: none">• Systems that can understand and generate speech• Systems that can understand images• Robotics• Assistant systems	<ul style="list-style-type: none">• 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 micro-worlds (Blocks world)

70’s: Problems:

- Systems for micro-worlds prove not scalable → “real” applications
- “Intelligent behavior” 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 2,500 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