Artificial Intelligence 1
Winter Semester 2023/24

– Lecture Notes –

Prof. Dr. Michael Kohlhase

Professur für Wissensrepräsentation und -verarbeitung
Informatik, FAU Erlangen-Nürnberg
Michael.Kohlhase@FAU.de

2023-10-17

0.1 Preface

0.1.1 Course Concept

Objective: The course aims at giving students a solid (and often somewhat theoretically oriented) foundation of the basic concepts and practices of artificial intelligence. The course will predominantly cover symbolic AI – also sometimes called "good old-fashioned AI (GofAI)" – in the first semester and offers the very foundations of statistical approaches in the second. Indeed, a full account sub symbolic, machine learning based AI deserves its own specialization courses and needs much more mathematical prerequisites than we can assume in this course.

Context: The course "Artificial Intelligence" (AI 1 & 2) at FAU Erlangen is a two-semester course in the "Wahlpflichtbereich" (specialization phase) in semesters 5/6 of the Bachelor program "Computer Science" at FAU Erlangen. It is also available as a (somewhat remedial) course in the "Vertiefungsmodul Künstliche Intelligenz" in the Computer Science Master's program.

Prerequisites: AI-1 & 2 builds on the mandatory courses in the FAU Bachelor's program, in particular the course "Grundlagen der Logik in der Informatik" [Glo], which already covers a lot of the materials usually presented in the "knowledge and reasoning" part of an introductory AI course. The AI 1& 2 course also minimizes overlap with the course.

The course is relatively elementary, we expect that any student who attended the mandatory CS courses at FAU Erlangen can follow it.

Open to external students:

Other Bachelor programs are increasingly co-opting the course as specialization option. There is no inherent restriction to computer science students in this course. Students with other study biographies – e.g. students from other Bachelor programs our external Master's students should be able to pick up the prerequisites when needed.

0.1.2 Course Contents

Goal: To give students a solid foundation of the basic concepts and practices of the field of Artificial Intelligence. The course will be based on Russell/Norvig's book "*Artificial Intelligence; A modern Approach*" [RN09]

Artificial Intelligence I (the first semester): introduces AI as an area of study, discusses "rational agents" as a unifying conceptual paradigm for AI and covers problem solving, search, constraint propagation, logic, knowledge representation, and planning.

Artificial Intelligence II (the second semester): is more oriented towards exposing students to the basics of statistically based AI: We start out with reasoning under uncertainty, setting the foundation with Bayesian Networks and extending this to rational decision theory. Building on this we cover the basics of machine learning.

0.1.3 This Document

Format: The document mixes the slides presented in class with comments of the instructor to give students a more complete background reference.

Caveat: This document is made available for the students of this course only. It is still very much a draft and will develop over the course of the current course and in coming academic years. **Licensing:** This document is licensed under a Creative Commons license that requires attribution, allows commercial use, and allows derivative works as long as these are licensed under the same license. **Knowledge Representation Experiment:** This document is also an experiment in knowledge representation. Under the hood, it uses the sTEX package [Koh08; sTeX], a TEX/LATEX extension for semantic markup, which allows to export the contents into active documents that adapt to the reader and can be instrumented with services based on the explicitly represented meaning of the documents.

0.1.4 Acknowledgments

Materials: Most of the materials in this course is based on Russel/Norvik's book "Artificial Intelligence — A Modern Approach" (AIMA [RN95]). Even the slides are based on a LaTeX-based slide set, but heavily edited. The section on search algorithms is based on materials obtained from Bernhard Beckert (then Uni Koblenz), which is in turn based on AIMA. Some extensions have been inspired by an AI course by Jörg Hoffmann and Wolfgang Wahlster at Saarland University in 2016. Finally Dennis Müller suggested and supplied some extensions on AGI. Florian Rabe, Max Rapp and Katja Berčič have carefully re-read the text and pointed out problems.

All course materials have bee restructured and semantically annotated in the sTeX format, so that we can base additional semantic services on them.

AI Students: The following students have submitted corrections and suggestions to this and earlier versions of the notes: Rares Ambrus, Ioan Sucan, Yashodan Nevatia, Dennis Müller, Simon Rainer, Demian Vöhringer, Lorenz Gorse, Philipp Reger, Benedikt Lorch, Maximilian Lösch, Luca Reeb, Marius Frinken, Peter Eichinger, Oskar Herrmann, Daniel Höfer, Stephan Mattejat, Matthias Sonntag, Jan Urfei, Tanja Würsching, Adrian Kretschmer, Tobias Schmidt, Maxim Onciul, Armin Roth, Liam Corona, Tobias Völk, Lena Voigt, Yinan Shao, Michael Girstl, Matthias Vietz, Anatoliy Cherepantsev, Stefan Musevski, Matthias Lobenhofer, Philipp Kaludercic, Diwarkara Reddy, Martin Helmke, Stefan Müller, Dominik Mehlich, Paul Martini, Vishwang Dave, Arthur Miehlich, Christian Schabesberger, Vishaal Saravanan, Simon Heilig, Michelle Fribrance, Wenwen Wang, Xinyuan Tu, Lobna Eldeeb.

0.1.5 Recorded Syllabus

The recorded syllabus – a record the progress of the course in the academic year 2023/24– is in the course page in the ALeA system at `https://courses.voll-ki.fau.de/course-home/ai-1`. The table of contents in the AI-1 notes at `https://courses.voll-ki.fau.de` indicates the material covered to date in yellow.

The recorded syllabus of AI-2 can be found at `https://courses.voll-ki.fau.de/course-home/ai-2`. For the topics planned for this course, see subsection 0.1.2.

Contents

0.1 Preface i
0.1.1 Course Concept i
0.1.2 Course Contents i
0.1.3 This Document i
0.1.4 Acknowledgments ii
0.1.5 Recorded Syllabus ii

1 Preliminaries **1**
1.1 Administrative Ground Rules 1
1.2 Getting Most out of AI-1 4
1.2.1 I 4
1.3 Learning Resources for AI-1 6
1.4 AI-Supported Learning 7

2 AI – Who?, What?, When?, Where?, and Why? **15**
2.1 What is Artificial Intelligence? 15
2.2 Artificial Intelligence is here today! 17
2.3 Ways to Attack the AI Problem 22
2.4 Strong vs. Weak AI 24
2.5 AI Topics Covered 26
2.6 AI in the KWARC Group 27

Chapter 1

Preliminaries

In this chapter, we want to get all the organizational matters out of the way, so that we can get into the discussion of artificial intelligence content unencumbered. We will talk about the necessary administrative details, go into how students can get most out of the course, talk about where the various resources provided with the course can be found, and finally introduce the ALeA system, an experimental – using AI methods – learning support system for the AI course.

1.1 Administrative Ground Rules

We will now go through the ground rules for the course. This is a kind of a social contract between the instructor and the students. Both have to keep their side of the deal to make learning as efficient and painless as possible.

Prerequisites for AI-1

- **Content Prerequisites:** The mandatory courses in CS@FAU; Sem 1-4, in particular:
  - Course "Algorithmen und Datenstrukturen". (Algorithms & Data Structures)
  - Course "Grundlagen der Logik in der Informatik" (GLOIN). (Logic in CS)
  - Course "Berechenbarkeit und Formale Sprachen". (Theoretical CS)
- **Skillset Prerequisite:** Coping with mathematical formulation of the structures
  - Mathematics is the language of science (in particular computer science)
  - It allows us to be very precise about what we mean. (good for you)
- **Intuition:** (take them with a kilo of salt)
  - This is what I assume you know! (I have to assume something)
  - In most cases, the dependency on these is partial and "in spirit".
  - If you have not taken these (or do not remember), read up on them as needed!
- **Real Prerequisites:** Motivation, interest, curiosity, hard work.(AI-1 is non-trivial)
- You can do this course if you want! (and I hope you are successful)

SOME RIGHTS RESERVED

Now we come to a topic that is always interesting to the students: the grading scheme.

Assessment, Grades

- **Overall (Module) Grade:**
  - Grade via the exam (Klausur) $\rightsquigarrow 100\%$ of the grade.
  - Up to 10% bonus on-top for an exam with $\geq 50\%$ points.($\leq 50\% \rightsquigarrow$ no bonus)
  - Bonus points $\widehat{=}$ percentage sum of the best 10 tuesday quizzes divided by 100.
- **Exam:** 90 minutes exam conducted in presence on paper (∼ April 1. 2024)
- **Retake Exam:** 90 min exam six months later (∼ October 1. 2024)
- ⚠ You have to register for exams in campo in the first month of classes.
- **Note:** You can de-register from an exam on campo up to three working days before.
- **Tuesday Quizzes:** Every tuesday we start the lecture with a 10 min online quiz – the tuesday quiz – about the material from the previous week. (starts in week 2)

SOME RIGHTS RESERVED

Tuesday Quizzes

- **Tuesday Quizzes:** Every tuesday we start the lecture with a 10 min online quiz – the tuesday quiz – about the material from the previous week. (starts in week 2)
- **Motivations:** We do this to
  - keep you prepared and working continuously. (primary)
  - update the ALeA learner model (fringe benefit)
- The tuesday quiz will be given in the ALeA system
  - `https://courses.voll-ki.fau.de/quiz-dash/ai-1`
  - You have to be logged into ALeA!
  - You can take the quiz on your laptop or phone, . . .
  - . . . in the lecture or at home . . .
  - . . . via WLAN or 4G Network. (do not overload)
  - Quizzes will only be available 16:15-16:25!

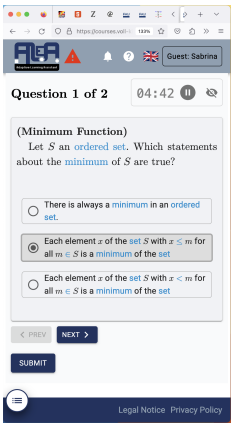

SOME RIGHTS RESERVED

Tomorrow: Pretest

- ⚠ Tomorrow we will try out the tuesday quiz infrastructure with a pretest!
  - **Presence**: bring your laptop or cellphone.
  - **Online**: you can and should take the pretest as well.
  - Have a recent firefox or chrome (chrome: $\geq$ March 2023)
  - Make sure that you are logged into ALeA (via FAU IDM; see below)
- **Definition 1.1.1.** A pretest is an assessment for evaluating the preparedness of learners for further studies.
- **Concretely:** This pretest
  - establishes a baseline for the competency expectations in AI-1 and
  - tests the ALeA quiz infrastructure for the tuesday quizzes.
- Participation in this test is optional; it will not influence your grades in any way.
- The test covers the prerequisites of AI-1 and some of the material that may have been covered in other courses.
- The test will be also used to refine the ALeA learner model, which may make learning experience in ALeA better. (see below)

SOME RIGHTS RESERVED

Due to the current AI hype, the course Artificial Intelligence is very popular and thus many degree programs at FAU have adopted it for their curricula. Sometimes the course setup that fits for the CS program does not fit the other's very well, therefore there are some special conditions. I want to state here.

⚠ Special Admin Conditions ⚠

- Some degree programs do not "import" the course Artificial Intelligence, and thus you may not be able to register for the exam via `https://campus.fau.de`.
  - Just send me an e-mail and come to the exam, we will issue a "Schein".
  - Tell your program coordinator about AI-1/2 so that they remedy this situation
- In "Wirtschafts-Informatik" you can only take AI-1 and AI-2 together in the "Wahlpflicht-bereich".
  - ECTS credits need to be divisible by five $\rightsquigarrow 7.5 + 7.5 = 15$.

SOME RIGHTS RESERVED

I can only warn of what I am aware, so if your degree program lets you jump through extra hoops, please tell me and then I can mention them here.

1.2 Getting Most out of AI-1

1.2.1 I

n this subsection we will discuss a couple of measures that students may want to consider to get most out of the AI-1 course.

None of them – homeworks, tutorials, study groups, and attendance – are mandatory, but most of them are very clearly correlated with success (i.e. passing the exam and getting a good grade).

AI-1 Homework Assignments

- **Homework Assignments:** Small individual problem/programming/proof task
  - but take time to solve (at least read them directly ⇝ questions)
- ⚠ Homeworks give no bonus points, but without trying you are unlikely to pass the exam.
- **Homework/Tutorial Discipline:**
  - Start early! (many assignments need more than one evening's work)
  - Don't start by sitting at a blank screen (talking & study group help)
  - Humans will be trying to understand the text/code/math when grading it.
  - Go to the tutorials, discuss with your TA! (they are there for you!)
- ⚠ We will not be able to grade all homework assignments!
- **Graded Assignments:** To keep things running smoothly
  - Homeworks will be posted on StudOn.
  - Sign up for AI-1 under `https://www.studon.fau.de/crs4622069.html`.
  - Homeworks are handed in electronically there. (plain text, program files, PDF)
  - Do not sign up for the "AI-2 Übungen" on StudOn (we do not use them)
- **Ungraded Assignments:** Are peer-feedbacked in ALEA (see below)

SOME RIGHTS RESERVED

It is very well-established experience that without doing the homework assignments (or something similar) on your own, you will not master the concepts, you will not even be able to ask sensible questions, and take very little home from the course. Just sitting in the course and nodding is not enough! If you have questions please make sure you discuss them with the instructor, the teaching assistants, or your fellow students. There are three sensible venues for such discussions: online in the lecture, in the tutorials, which we discuss now, or in the course forum – see below. Finally, it is always a very good idea to form study groups with your friends.

Tutorials for Artificial Intelligence 1

- **Approach:** Weekly tutorials and homework assignments (first one in week two)
- **Goal 1:** Reinforce what was taught in class. (you need practice)

- ▷ **Goal 2:** Allow you to ask any question you have in a protected environment.
- ▷ **Instructor/Lead TA:** Florian Rabe (KWARC Postdoc)
  - ▷ Room: 11.137 @ Händler building, florian.rabe@fau.de
- ▷ **Tutorials:** One each taught by Florian Rabe (lead); Mahdi Mantash, Robert Kurin, Florian Guthmann.
- ▷ **Life-saving Advice:** Go to your tutorial, and prepare for it by having looked at the slides and the homework assignments!
- ▷ **Caveat:** We cannot grade all submissions with 5 TAs and ∼1000 students.
- ▷ **Also:** Group submission has not worked well in the past! (too many freeloaders)

SOME RIGHTS RESERVED

Collaboration

- ▷ **Definition 1.2.1.** Collaboration (or cooperation) is the process of groups of agents working or acting together for common, mutual, or some underlying benefit, as opposed to working in competition for selfish benefit. In a collaboration, every agent contributes to the common goal.
- ▷ In learning situations, the benefit is "better learning outcomes".
- ▷ **Observation:** In collaborative learning, the overall result can be significantly better than in competitive learning.
- ▷ **Good Practice:** Form study groups. (long- or short-term)
  - ▷ ⚠ those learners who work most, learn most
  - ▷ ⚠ freeloaders – indivicuals who only watch – learn very little!
- ▷ It is OK to collaborate on homework assignments in AI-1! (no bonus points)
- ▷ Choose your study group well (We will (eventually) help via ALeA)

SOME RIGHTS RESERVED

What I am going to go into next is – or should be – obvious, but there is an important point I want to make.

Do I need to attend the lectures

- ▷ Attendance is not mandatory for the AI-1 lecture
- ▷ There are two ways of learning AI-1: (both are OK, your mileage may vary)
  - ▷ Approach B: Read a Book
  - ▷ Approach I: come to the lectures, be involved, interrupt me whenever you have a question.

The only advantage of I over B is that books do not answer questions (yet! ⇜ we are working on this in AI research)

▷ Approach S: come to the lectures and sleep does not work!

▷ **I really mean it:** If you come to class, be involved, ask questions, challenge me with comments, tell me about errors, ...

- ▷ I would much rather have a lively discussion than get through all the slides
- ▷ You learn more, I have more fun (Approach B serves as a backup)
- ▷ You may have to change your habits, overcome shyness, ... (please do!)

▷ This is what I get paid for, and I am more expensive than most books (get your money's worth)

SOME RIGHTS RESERVED

1.3 Learning Resources for AI-1

But what if you are not in a lecture or tutorial and want to find out more about the AI-1 topics?

Textbook, Handouts and Information, Forums, Videos

▷ **Textbook:** *Russel/Norvig: Artificial Intelligence, A modern Approach* [RN09].

- ▷ basically "broad but somewhat shallow"
- ▷ great to get intuitions on the basics of AI

Make sure that you read the edition ≥ 3 ⇜ vastly improved over ≤ 2.

▷ **Course notes:** will be posted at `http://kwarc.info/teaching/AI/notes.pdf`

- ▷ more detailed than [RN09] in some areas
- ▷ I mostly prepare them as we go along (semantically preloaded ⇝ research resource)
- ▷ please e-mail me any errors/shortcomings you notice. (improve for the group)

▷ **StudOn Forum:** `https://www.studon.fau.de/crs4622069.html` for

- ▷ announcements, homeworks (my view on the forum)
- ▷ questions, discussion among your fellow students (your forum too, use it!)

▷ **Course Videos:** AI-1 will be streamed/recorded at `https://fau.tv/course/id/3595`

- ▷ **Organized**: Video course nuggets are available at `https://fau.tv/course/id/1690` (short; organized by topic)
- ▷ **Backup**: The lectures from WS 2016/17 to SS 2018 have been recorded (in English and German), see `https://www.fau.tv/search/term.html?q=Kohlhase`

▷ **Do not let the videos mislead you:** Coming to class is highly correlated with passing the course!

FAU has issued a very insightful guide on using lecture recordings. It is a good idea to heed these recommendations, even if they seem annoying at first.

Practical recommendations on Lecture Resources

▷ Excellent Guide: [Nor+18a] (german Version at [Nor+18b])

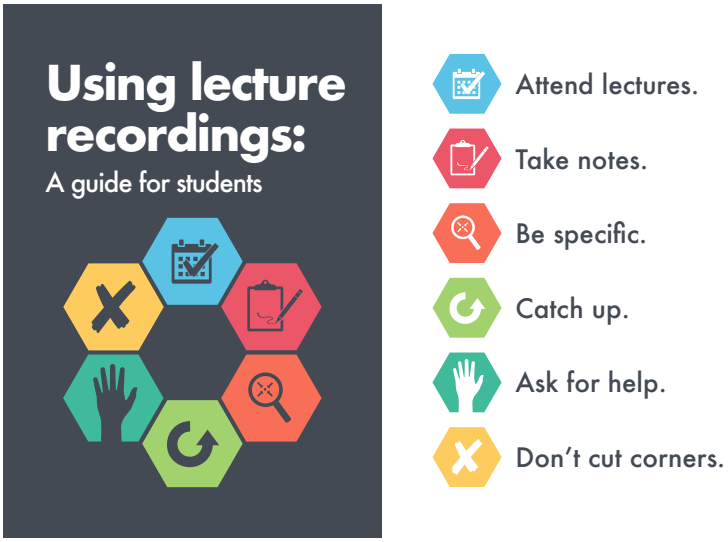

1.4 AI-Supported Learning

In this section we introduce the ALeA (Adaptive Learning Assistant) system, a learning support system we have developed using symbolic AI methods – the stuff we learn about in AI-1 – and which we will use to support students in the course. As such ALeA does double duty in this course it supports learning activities and serves as a showcase, what methods can to in an important application.

ALeA: Adaptive Learning Assistant

▷ **Idea:** Use AI methods to help teach/learn AI (AI4AI)

▷ **Concretely:** Provide HTML versions of the AI-1 slides/notes and embed learning support services into them. (for pre/postparation of lectures)

▷ **Definition 1.4.1.** Call a document active, iff it is interactive and adapts to specific information needs of the readers. (course notes on steroids)

▷ **Intuition:** ALeA serves active course materials. (PDF mostly inactive)

▷ **Example 1.4.2 (Course Notes).** ≙ Slides + Comments

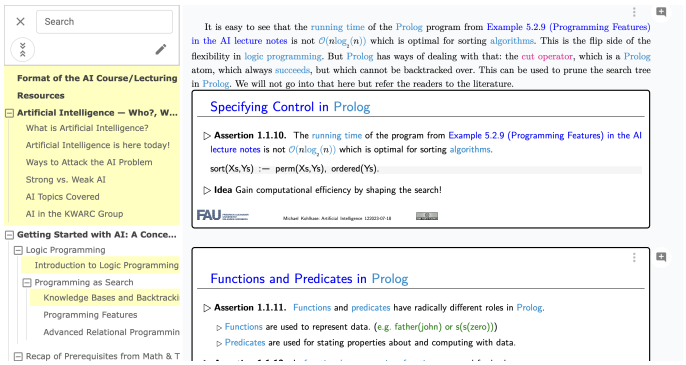

⇝ yellow parts in table of contents (left) already covered in lecture.

SOME RIGHTS RESERVED

VoLL-KI Portal at https://courses.voll-ki.fau.de

▷ **Portal for ALeA Courses:** https://courses.voll-ki.fau.de

▷ **AI-1 in ALeA:** https://courses.voll-ki.fau.de/course-home/ai-1

- ▷ All details for the course.
- ▷ recorded syllabus (keep track of material covered in course)
- ▷ syllabus of the last semester (for over/preview)

▷ **ALeA Status:** The ALeA system is deployed at FAU for over 1000 students taking six courses

- ▷ (some) students use the system actively (our logs tell us)
- ▷ reviews are mostly positive/enthusiastic (error reports pour in)

SOME RIGHTS RESERVED

Learning Support Services in ALeA

▷ **Idea:** Embed learning support services into active course materials.

▷ **Example 1.4.3 (Definition on Hover).** Hovering on a (cyan) term reference reminds us of the definition. (even works recursively)

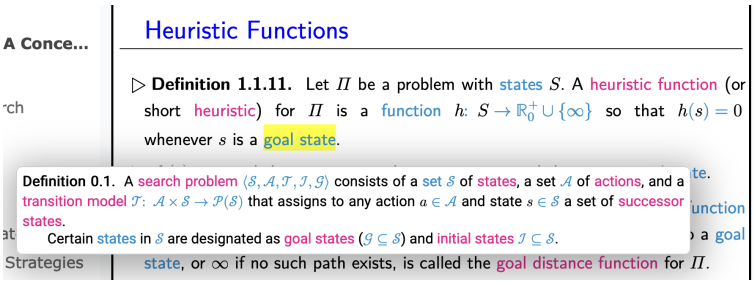

▷ **Example 1.4.4 (More Definitions on Click).** Clicking on a (cyan) term reference shows us more definitions from other contexts.

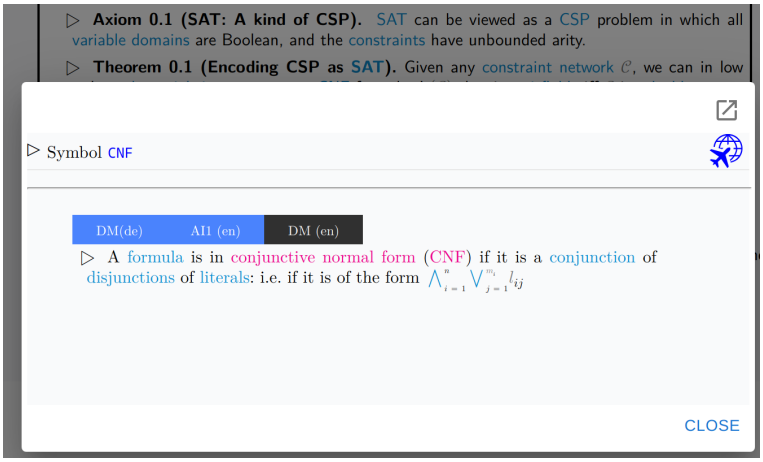

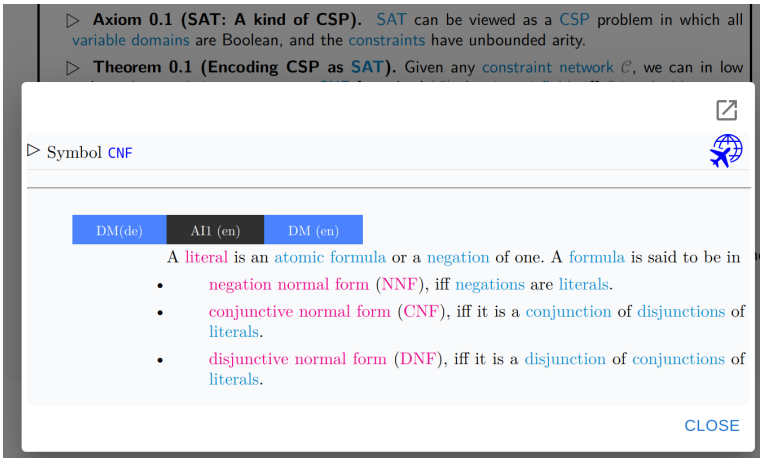

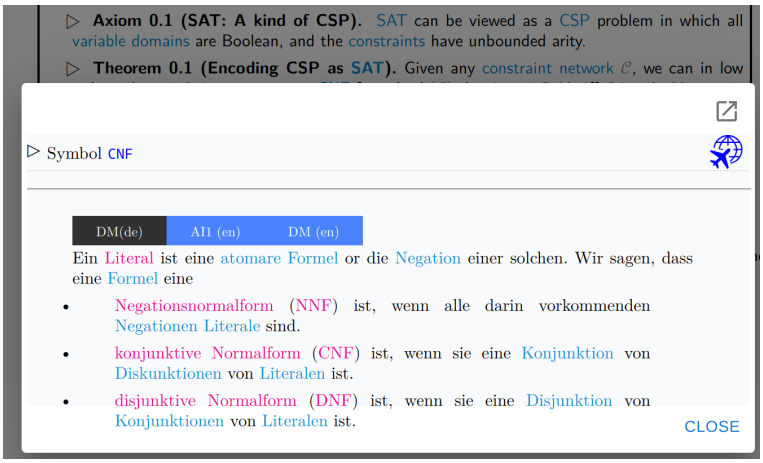

▷ **Example 1.4.5 (Guided Tour).** A guided tour for a concept c assembles definitions/etc. into a self-cont

c = countable ⇝

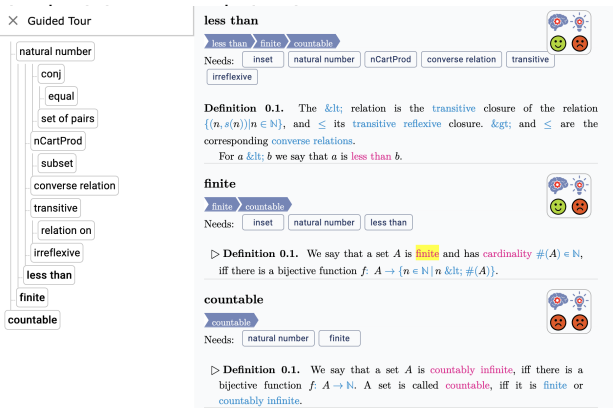

▷ . . . your idea here . . . (the sky is the limit)

SOME RIGHTS RESERVED

Localized Interactions with the Community

▷ Selecting text brings up localized – ancored on selection – interactions:

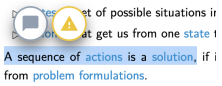

▷ post a (public) comment or take (private) note

▷ report an error to the course authors/instructors

▷ Localized comments induce a thread in the ALeA forum (like the StudOn Forum, but targeted towards specific learning objects)

▷ Answering questions gives karma ≙ a public measure of helpfulness

▷ Notes can be anonymous (⇝ generate no karma)

SOME RIGHTS RESERVED

Let us briefly look into how the learning support services introduced above might work, focusing on where the necessar information might come from.

ALeA≙ Data-Driven & AI-enabled Learning Assistance

▷ **Ingredient 1:** Domain model ≙ knowledge/theory graph

▷ **Ingredient 2:** Learner model ≙ adding competency estimations

▷ **Ingredient 3:** A collection of ready-formulated learning objects

▷ **Ingredient 4:** Educational dialogue planner ⇝ guided tours

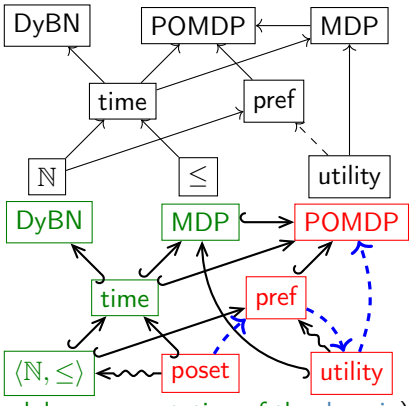

A theory graph provides (modular representation of the domain)

▷ symbols with URIs for all concepts, objects, and relations

▷ definitions, notations, and verbalizations for all symbols

▷ "object-oriented inheritance" and views between theories.

The learner model is a function from learner IDs × symbol URIs to competency values

▷ competency comes in six cognitive dimensions: `remember`, `understand`, `analyze`, `evaluate`, `apply`, and `create`.

▷ ALeA logs all learner interactions (keeps data learner-private)

▷ each interaction updates the learner model function.

Learning objects are the text fragments learners see and interact with; they are structured by

▷ didactic relations, e.g. tasks have prerequisites and learning objectives

▷ rhetoric relations, e.g. introduction, elaboration, and transition

The dialogue planner assembles learning objects into active course materials using

▷ the domain model and didactic relations to determine the order of LOs

▷ the learner model to determine what to show

▷ the rhetoric relations to make the dialogue coherent

SOME RIGHTS RESERVED

New Feature: Drilling with Flashcards

▷ Flashcards challenge you with a task (term/problem) on the front...

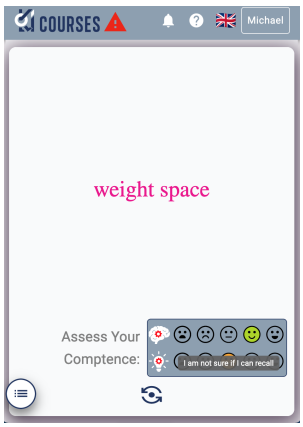

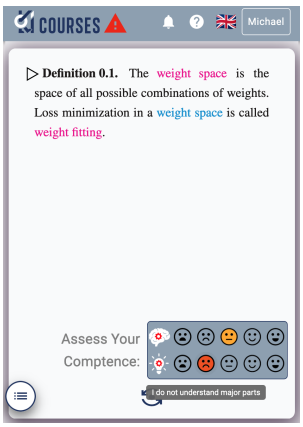

...and the definition/answer is on the back.

▷ Self-assessment updates the learner model (before/after)

▷ **Bonus:** Flashcards can be generated from existing semantic markup (educational equivalent to free beer)

SOME RIGHTS RESERVED

Learner Data and Privacy in ALeA

▷ **Observation:** Most learning support services in ALeA use the learner model; they

▷ need the learner model data to adapt to the invidivual learner!

▷ collect learner interaction data (to update the learner model)

▷ **Consequence:** You need to be logged in (via your FAU IDM credentials) for useful learning support services!

▷ **Problem:** Learner model data is highly sensitive personal data!

▷ **ALeA Promise:** The ALeA team does the utmost to keep your personal data safe. (SSO via FAU IDM/eduGAIN, ALeA trust zone)

▷ **ALeA Privacy Axioms:**

1. ALeA only collects learner models data about logged in users.
2. Personally identifiable learner model data is only accessible to its subject (delegation possible)
3. Learners can always query the learner model about its data.
4. All learner model data can be purged without negative consequences (except usability deterioration)
5. Logging into ALeA is completely optional.

▷ **Observation:** Authentication for bonus quizzes are somewhat less optional, but you can always purge the learner model later.

SOME RIGHTS RESERVED

Concrete Todos for ALeA

▷ **Recall:** You will use ALeA for the tuesday quizzes (or lose bonus points)
All other use is optional (but AI-supported pre/postparation can be helpful)

▷ To use the ALeA system, you will have to log in via SSO (do it now)

- ▷ go to `https://courses.voll-ki.fau.de/course-home/ai-1`
- ▷ in the upper right hand corner you see 

- ▷ log in via your FAU IDM credentials. (you should have them by now)
- ▷ You get access to your personal ALeA profile via 

(plus feature notifications, manual, and language chooser)

▷ **Problem:** Most ALeA services depend on the learner model (to adapt to you)

▷ **Solution:** Initialize your learner model with your educational history!

- ▷ **Concretely**: enter taken CS courses (FAU equivalents) and grades
- ▷ ALeA uses that to estimate your CS/AI competencies (for your benefit)
- ▷ then ALeA knows about you; I don't (ALeA trust zone)

SOME RIGHTS RESERVED

Chapter 2

Artificial Intelligence – Who?, What?, When?, Where?, and Why?

We start the course by giving an overview of (the problems, methods, and issues of) Artificial Intelligence, and what has been achieved so far.

Naturally, this will dwell mostly on philosophical aspects – we will try to understand what the important issues might be and what questions we should even be asking. What the most important avenues of attacks may be and where AI research is being carried out.

In particular the discussion will be very non-technical – we have very little basis to discuss technicalities yet. But stay with me, this will drastically change very soon. **A Video Nugget** covering the introduction of this chapter can be found at `https://fau.tv/clip/id/21467`.

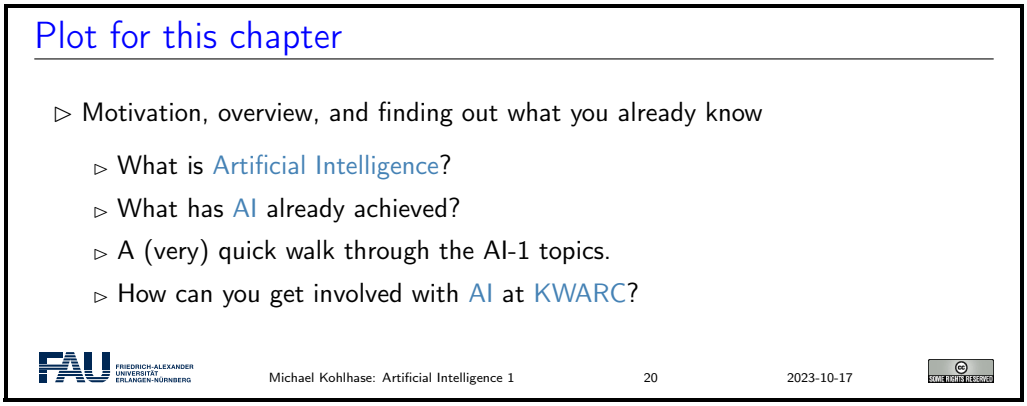

2.1 What is Artificial Intelligence?

A Video Nugget covering this section can be found at `https://fau.tv/clip/id/21701`.

The first question we have to ask ourselves is "What is Artificial Intelligence?", i.e. how can we define it. And already that poses a problem since the natural definition *like human intelligence, but artificially realized* presupposes a definition of Intelligence, which is equally problematic; even Psychologists and Philosophers – the subjects nominally "in charge" of human intelligence – have problems defining it, as witnessed by the plethora of theories e.g. found at [WHI].

▷ **Definition 2.1.1 (According to Wikipedia).** Artificial Intelligence (AI) is intelligence exhibited by machines

▷ **Definition 2.1.2 (also).** Artificial Intelligence (AI) is a sub-field of computer science that is concerned with the automation of intelligent behavior.

▷ **BUT:** it is already difficult to define "Intelligence" precisely

▷ **Definition 2.1.3 (Elaine Rich).** Artificial Intelligence (AI) studies how we can make the computer do things that humans can still do better at the moment.

SOME RIGHTS RESERVED

Maybe we can get around the problems of defining "what Artificial intelligence is", by just describing the necessary components of AI (and how they interact). Let's have a try to see whether that is more informative.

What is Artificial Intelligence? Components

▷ **Elaine Rich:** AI studies how we can make the computer do things that humans can still do better at the moment.

▷ This needs a combination of

- ▷ the ability to learn
- ▷ inference
- ▷ perception
- ▷ language understanding
- ▷ emotion

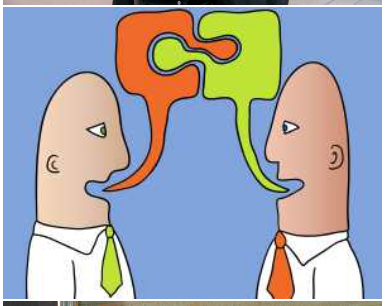

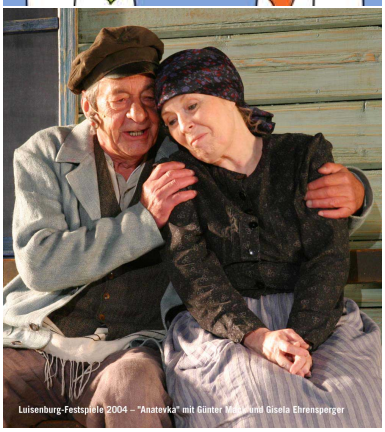

SOME RIGHTS RESERVED

2.2 Artificial Intelligence is here today!

A Video Nugget covering this section can be found at `https://fau.tv/clip/id/21697`.

The components of Artificial Intelligence are quite daunting, and none of them are fully understood, much less achieved artificially. But for some tasks we can get by with much less. And indeed that is what the field of Artificial Intelligence does in practice – but keeps the lofty ideal around. This practice of "trying to achieve AI in selected and restricted domains" (cf. the discussion starting with slide 30) has borne rich fruits: systems that meet or exceed human capabilities in such areas. Such systems are in common use in many domains of application.

> Artificial Intelligence is here today!

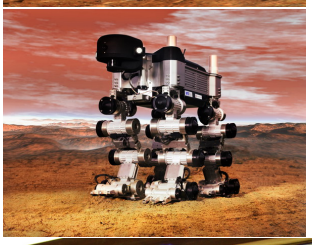

- ▷ in outer space
  - ▹ in outer space systems need autonomous control:
  - ▹ remote control impossible due to time lag

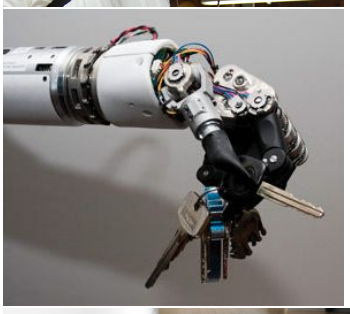

- ▷ in artificial limbs
  - ▹ the user controls the prosthesis via existing nerves, can e.g. grip a sheet of paper.

- ▷ in household appliances
  - ▹ The iRobot Roomba vacuums, mops, and sweeps in corners, . . . , parks, charges, and discharges.
  - ▹ general robotic household help is on the horizon.

- ▷ in hospitals
  - ▹ in the USA 90% of the prostate operations are carried out by RoboDoc
  - ▹ Paro is a cuddly robot that eases solitude in nursing homes.

And here's what you all have been waiting for . . .

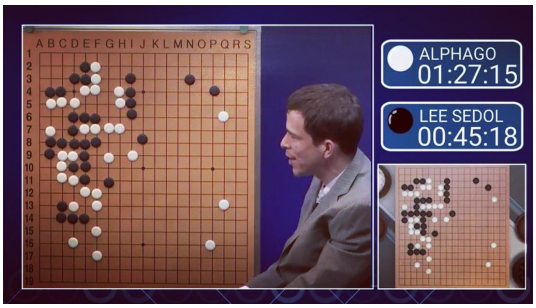

CC-BY-SA: Buster Benson@
https://www.flickr.com/photos/erikbenson/25717574115

▷ AlphaGo is a program by Google DeepMind to play the board game go.

▷ In March 2016, it beat Lee Sedol in a five-game match, the first time a go program has beaten a 9 dan professional without handicaps. In December 2017 AlphaZero, a successor of AlphaGo "learned" the games go, chess, and shogi in 24 hours, achieving a superhuman level of play in these three games by defeating world-champion programs. By September 2019, AlphaStar, a variant of AlphaGo, attained "grandmaster level" in Starcraft II, a real time strategy game with partially observable state. AlphaStar now among the top 0.2% of human players.

We will conclude this section with a note of caution.

The AI Conundrum

▷ **Observation:** Reserving the term "Artificial Intelligence" has been quite a land grab!

▷ **But:** researchers at the Dartmouth Conference (1956) really thought they would solve/reach AI in two/three decades.

▷ **Consequence:** AI still asks the big questions.

▷ **Another Consequence:** AI as a field is an incubator for many innovative technologies.

▷ **AI Conundrum:** Once AI solves a subfield it is called "computer science". (becomes a separate subfield of CS)

▷ **Example 2.2.1.** Functional/Logic Programming, automated theorem proving, Planning, machine learning, Knowledge Representation, . . .

▷ **Still Consequence:** AI research was alternatingly flooded with money and cut off brutally.

2.3 Ways to Attack the AI Problem

A Video Nugget covering this section can be found at `https://fau.tv/clip/id/21717`.
There are currently three main avenues of attack to the problem of building artificially intelligent systems. The (historically) first is based on the symbolic representation of knowledge about the world and uses inference-based methods to derive new knowledge on which to base action decisions. The second uses statistical methods to deal with uncertainty about the world state and learning methods to derive new (uncertain) world assumptions to act on.

Three Main Approaches to Artificial Intelligence

▷ **Definition 2.3.1.** Symbolic AI is based on the assumption that many aspects of intelligence can be achieved by the manipulation of symbols, combining them into structures (expressions) and manipulating them (using processes) to produce new expressions.

▷ **Definition 2.3.2.** Statistical AI remedies the two shortcomings of symbolic AI approaches: that all concepts represented by symbols are crisply defined, and that all aspects of the world are knowable/representable in principle. Statistical AI adopts sophisticated mathematical models of uncertainty and uses them to create more accurate world models and reason about them.

▷ **Definition 2.3.3.** Subsymbolic AI attacks the assumption of symbolic and statistical AI that intelligence can be achieved by reasoning about the state of the world. Instead it posits that intelligence must be embodied i.e. situated in the world, equipped with a "body" that can interact with it via sensors and actuators. The main method for realizing intelligent behavior is by learning from the world, i.e. machine learning.

As a consequence, the field of Artificial Intelligence (AI) is an engineering field at the intersection of computer science (logic, programming, applied statistics), cognitive science (psychology, neuroscience), philosophy (can machines think, what does that mean?), linguistics (natural language understanding), and mechatronics (robot hardware, sensors).
Subsymbolic AI and in particular machine learning is currently hyped to such an extent, that many people take it to be synonymous with "Artificial Intelligence". It is one of the goals of this course to show students that this is a very impoverished view.

Two ways of reaching Artificial Intelligence?

▷ We can classify the AI approaches by their coverage and the analysis depth (they are complementary)

Analysis ↑ vs. Coverage →	Narrow	Wide
Deep	symbolic AI-1	not there yet cooperation?
Shallow	no-one wants this	statistical/sub symbolic AI-2

▷ **This semester** we will cover foundational aspects of symbolic AI (deep/narrow processing)

▷ **next semester** concentrate on statistical/subsymbolic AI. (shallow/wide-coverage)

We combine the topics in this way in this course, not only because this reproduces the historical development but also as the methods of statistical and subsymbolic AI share a common basis.
It is important to notice that all approaches to AI have their application domains and strong points. We will now see that exactly the two areas, where symbolic AI and statistical/subsymbolic AI have their respective fortes correspond to natural application areas.

Environmental Niches for both Approaches to AI

▷ **Observation:** There are two kinds of applications/tasks in AI

- ▷ Consumer tasks: consumer grade applications have tasks that must be fully generic and wide coverage. (e.g. machine translation like Google Translate)
- ▷ Producer tasks: producer grade applications must be high-precision, but can be domain-specific (e.g. multilingual documentation, machinery-control, program verification, medical technology)

Precision	$10^{3\pm1}$ Concepts	$10^{6\pm1}$ Concepts **Coverage**
100%	Producer Tasks	
50%	Consumer Tasks	

▷ **General Rule:** Subsymbolic AI is well suited for consumer tasks, while symbolic AI is better suited for producer tasks.

▷ A domain of producer tasks I am interested in: mathematical/technical documents.

An example of a producer task – indeed this is where the name comes from – is the case of a machine tool manufacturer T, which produces digitally programmed machine tools worth multiple million Euro and sells them into dozens of countries. Thus T must also comprehensive machine

operation manuals, a non-trivial undertaking, since no two machines are identical and they must be translated into many languages, leading to hundreds of documents. As those manual share a lot of semantic content, their management should be supported by AI techniques. It is critical that these methods maintain a high precision, operation errors can easily lead to very costly machine damage and loss of production. On the other hand, the domain of these manuals is quite restricted. A machine tool has a couple of hundred components only that can be described by a comple of thousand attribute only.

Indeed companies like T employ high-precision AI techniques like the ones we will cover in this course successfully; they are just not so much in the public eye as the consumer tasks.

To get this out of the way . . .

CC-BY-SA: Buster Benson@ https://www.flickr.com/photos/erikbenson/25717574115

- ▷ AlphaGo = search + neural networks (symbolic + subsymbolic AI)
  - ▷ we do search this semester and cover neural networks in AI-2.
  - ▷ I will explain AlphaGo a bit in ??.

SOME RIGHTS RESERVED

2.4 Strong vs. Weak AI

A Video Nugget covering this section can be found at `https://fau.tv/clip/id/21724`.
To get this out of the way before we begin: We now come to a distinction that is often muddled in popular discussions about "Artificial Intelligence", but should be cristal clear to students of the course AI-1 – after all, you are upcoming "AI-specialists".

Strong AI vs. Narrow AI

- ▷ **Definition 2.4.1.** With the term narrow AI (also weak AI, instrumental AI, applied AI) we refer to the use of software to study or accomplish *specific* problem solving or reasoning tasks (e.g. playing chess/go, controlling elevators, composing music, . . .)
- ▷ **Definition 2.4.2.** With the term strong AI (also full AI, AGI) we denote the quest for software performing at the full range of human cognitive abilities.
- ▷ **Definition 2.4.3.** Problems requiring strong AI to solve are called AI hard.

▷ **In short:** We can characterize the difference intuitively:

- ▷ narrow AI: What (most) computer scientists think AI is / should be.
- ▷ strong AI: What Hollywood authors think AI is / should be.

▷ **Needless to say** we are only going to cover narrow AI in this course!

SOME RIGHTS RESERVED

One can usually defuse public worries about "is AI going to take control over the world" by just explaining the difference between strong AI and weak AI clearly.
I would like to add a few words on AGI, that – if you adopt them; they are not universally accepted – will strengthen the arguments differentiating between strong and weak AI.

A few words on AGI...

▷ The conceptual and mathematical framework (agents, environments etc.) is the same for strong AI and weak AI.

▷ AGI research focuses mostly on abstract aspects of machine learning (reinforcement learning, neural nets) and decision/game theory ("which goals should an AGI pursue?").

▷ Academic respectability of AGI fluctuates massively, recently increased (again). (correlates somewhat with AI winters and golden years)

▷ Public attention increasing due to talk of "existential risks of AI" (e.g. Hawking, Musk, Bostrom, Yudkowsky, Obama, ...)

▷ **Kohlhase's View:** Weak AI is here, strong AI is very far off. (not in my lifetime)
But even if that is true, weak AI will affect all of us deeply in everyday life.

▷ **Example 2.4.4.** You should not train to be an accountant or truck driver!
(bots will replace you)

SOME RIGHTS RESERVED

I want to conclude this section with an overview over the recent protagonists – both personal and institutional – of AGI.

AGI Research and Researchers

▷ "Famous" research(ers) / organizations

- ▷ MIRI (Machine Intelligence Research Institute), Eliezer Yudkowsky (Formerly known as "Singularity Institute")
- ▷ Future of Humanity Institute Oxford (Nick Bostrom),
- ▷ Google (Ray Kurzweil),
- ▷ AGIRI / OpenCog (Ben Goertzel),
- ▷ `petrl.org` (People for the Ethical Treatment of Reinforcement Learners). (Obviously somewhat tongue-in-cheek)

▷ ⚠ Be highly skeptical about any claims with respect to AGI! (Kohlhase's View)

2.5 AI Topics Covered

A Video Nugget covering this section can be found at `https://fau.tv/clip/id/21719`.

We will now preview the topics covered by the course "Artificial Intelligence" in the next two semesters.

Topics of AI-1 (Winter Semester)

- ▷ Getting Started
  - ▷ What is Artificial Intelligence? (situating ourselves)
  - ▷ Logic programming in Prolog (An influential paradigm)
  - ▷ Intelligent Agents (a unifying framework)
- ▷ Problem Solving
  - ▷ Problem Solving and search (Black Box World States and Actions)
  - ▷ Adversarial search (Game playing) (A nice application of search)
  - ▷ constraint satisfaction problems (Factored World States)
- ▷ Knowledge and Reasoning
  - ▷ Formal Logic as the mathematics of Meaning
  - ▷ Propositional logic and satisfiability (Atomic Propositions)
  - ▷ First-order logic and theorem proving (Quantification)
  - ▷ Logic programming (Logic + Search$\rightsquigarrow$ Programming)
  - ▷ Description logics and semantic web
- ▷ Planning
  - ▷ Planning Frameworks
  - ▷ Planning Algorithms
  - ▷ Planning and Acting in the real world

Topics of AI-2 (Summer Semester)

- ▷ Uncertain Knowledge and Reasoning
  - ▷ Uncertainty
  - ▷ Probabilistic reasoning
  - ▷ Making Decisions in Episodic Environments
  - ▷ Problem Solving in Sequential Environments
- ▷ Foundations of machine learning

- ▷ Learning from Observations
- ▷ Knowledge in Learning
- ▷ Statistical Learning Methods

▷ Communication (If there is time)

- ▷ Natural Language Processing
- ▷ Natural Language for Communication

AI1SysProj: A Systems/Project Supplement to AI-1

- ▷ The AI-1 course concentrates on concepts, theory, and algorithms of symbolic AI.
- ▷ **Problem:** Engineering/Systems Aspects of AI are very important as well.
- ▷ **Partial Solution:** Getting your hands dirty in the homeworks and the Kalah Challenge
- ▷ **Full Solution:** AI1SysProj: AI-1 Systems Project (10 ECTS, 30-50places)
  - ▷ For each Topic of AI-1, where will be a mini-project in AI1SysProj
  - ▷ e.g. for game-play there will be Chinese Checkers (more difficult than Kalah)
  - ▷ e.g. for CSP we will schedule TechFak courses or exams (from real data)
  - ▷ solve challenges by implementing the AI-1 algorithms or use SoA systems
- ▷ **Question:** Should I take AI1SysProj in my first semester? (i.e. now)
- ▷ **Answer:** It depends . . . (on your situation)
  - ▷ most master's programs require a 10-ECTS "Master's Project"(Master AI: two)
  - ▷ there will be a great pressure on project places (so reserve one early)
  - ▷ BUT 10 ECTS $\widehat{=}$ 250-300 hours involvement by definition (1/3 of your time/ECTS)
- ▷ **BTW:** There will also be an AI2SysProj next semester! (another chance)

2.6 AI in the KWARC Group

A Video Nugget covering this section can be found at `https://fau.tv/clip/id/21725`.

Now allow me to beat my own drum. In my research group at FAU, we do research on a particular kind of Artificial Intelligence: logic, language, and information. This may not be the most fashionable or well-hyped area in AI, but it is challenging, well-respected, and – most importantly – fun.

The KWARC Research Group

▷ **Observation:** The ability to represent knowledge about the world and to draw logical inferences is one of the central components of intelligent behavior.

▷ **Thus:** reasoning components of some form are at the heart of many AI systems.

▷ **KWARC Angle:** Scaling up (web-coverage) without dumbing down (too much)

- ▷ Content markup instead of full formalization (too tedious)
- ▷ User support and quality control instead of "The Truth" (elusive anyway)
- ▷ use Mathematics as a test tube (⚠ Mathematics ≙ Anything Formal ⚠)
- ▷ care more about applications than about philosophy (we cannot help getting this right anyway as logicians)

▷ The KWARC group was established at Jacobs Univ. in 2004, moved to FAU Erlangen in 2016

▷ see `http://kwarc.info` for projects, publications, and links

SOME RIGHTS RESERVED

Research in the KWARC group ranges over a variety of topics, which range from foundations of mathematics to relatively applied web information systems. I will try to organize them into three pillars here.

Overview: KWARC Research and Projects

Applications: eMath 3.0, Active Documents, Active Learning, Semantic Spreadsheets/CAD/CAM, Change Mangagement, Global Digital Math Library, Math Search Systems, SMGloM: Semantic Multilingual Math Glossary, Serious Games, ...

Foundations of Math:	**KM & Interaction**:	**Semantization**:
▷ MathML, OpenMath	▷ Semantic Interpretation (aka. Framing)	▷ LaTeXML: LaTeX → XML
▷ advanced Type Theories	▷ math-literate interaction	▷ sTeX: Semantic LaTeX
▷ Mmt: Meta Meta Theory	▷ MathHub: math archives & active docs	▷ invasive editors
▷ Logic Morphisms/Atlas	▷ Active documents: embedded semantic services	▷ Context-Aware IDEs
▷ Theorem Prover/CAS Interoperability	▷ Model-based Education	▷ Mathematical Corpora
▷ Mathematical Models/Simulation		▷ Linguistics of Math
		▷ ML for Math Semantics Extraction

Foundations: Computational Logic, Web Technologies, OMDoc/Mmt

SOME RIGHTS RESERVED

For all of these areas, we are looking for bright and motivated students to work with us. This can take various forms, theses, internships, and paid student assistantships.

Research Topics in the KWARC Group

▷ We are always looking for bright, motivated KWARCies.

▷ We have topics in for all levels! (Enthusiast, Bachelor, Master, Ph.D.)

▷ List of current topics: `https://gl.kwarc.info/kwarc/thesis-projects/`

- ▷ Automated Reasoning: Maths Representation in the Large
- ▷ Logics development, $(\text{Meta})^n$-Frameworks
- ▷ Math Corpus Linguistics: Semantics Extraction
- ▷ Serious Games, Cognitive Engineering, Math Information Retrieval, Legal Reasoning, ...

▷ We always try to find a topic at the intersection of your and our interests.

▷ We also often have positions!. (HiWi, Ph.D.: $\frac{1}{2}$, PostDoc: full)

SOME RIGHTS RESERVED

Sciences like physics or geology, and engineering need high-powered equipment to perform measurements or experiments. computer science and in particular the KWARC group needs high powered human brains to build systems and conduct thought experiments.

The KWARC group may not always have as much funding as other AI research groups, but we are very dedicated to give the best possible research guidance to the students we supervise.

So if this appeals to you, please come by and talk to us.

Bibliography

[Bac00] Fahiem Bacchus. *Subset of PDDL for the AIPS2000 Planning Competition*. The AIPS-00 Planning Competition Comitee. 2000.

[BF95] Avrim L. Blum and Merrick L. Furst. "Fast planning through planning graph analysis". In: *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI)*. Ed. by Chris S. Mellish. Montreal, Canada: Morgan Kaufmann, San Mateo, CA, 1995, pp. 1636–1642.

[BF97] Avrim L. Blum and Merrick L. Furst. "Fast planning through planning graph analysis". In: *Artificial Intelligence* 90.1-2 (1997), pp. 279–298.

[BG01] Blai Bonet and Héctor Geffner. "Planning as Heuristic Search". In: *Artificial Intelligence* 129.1–2 (2001), pp. 5–33.

[BG99] Blai Bonet and Héctor Geffner. "Planning as Heuristic Search: New Results". In: *Proceedings of the 5th European Conference on Planning (ECP'99)*. Ed. by S. Biundo and M. Fox. Springer-Verlag, 1999, pp. 60–72.

[BKS04] Paul Beame, Henry A. Kautz, and Ashish Sabharwal. "Towards Understanding and Harnessing the Potential of Clause Learning". In: *Journal of Artificial Intelligence Research* 22 (2004), pp. 319–351.

[Bon+12] Blai Bonet et al., eds. *Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS'12)*. AAAI Press, 2012.

[Bro90] Rodney Brooks. In: *Robotics and Autonomous Systems* 6.1–2 (1990), pp. 3–15. DOI: 10.1016/S0921-8890(05)80025-9.

[Cho65] Noam Chomsky. *Syntactic structures*. Den Haag: Mouton, 1965.

[CKT91] Peter Cheeseman, Bob Kanefsky, and William M. Taylor. "Where the *Really* Hard Problems Are". In: *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI)*. Ed. by John Mylopoulos and Ray Reiter. Sydney, Australia: Morgan Kaufmann, San Mateo, CA, 1991, pp. 331–337.

[CM85] Eugene Charniak and Drew McDermott. *Introduction to Artificial Intelligence*. Addison Wesley, 1985.

[CQ69] Allan M. Collins and M. Ross Quillian. "Retrieval time from semantic memory". In: *Journal of verbal learning and verbal behavior* 8.2 (1969), pp. 240–247. DOI: 10.1016/S0022-5371(69)80069-1.

[Dav67] Donald Davidson. "Truth and Meaning". In: *Synthese* 17 (1967).

[DCM12] DCMI Usage Board. *DCMI Metadata Terms*. DCMI Recommendation. Dublin Core Metadata Initiative, June 14, 2012. URL: http://dublincore.org/documents/2012/06/14/dcmi-terms/.

[DF31] B. De Finetti. "Sul significato soggettivo della probabilita". In: *Fundamenta Mathematicae* 17 (1931), pp. 298–329.

[DHK15] Carmel Domshlak, Jörg Hoffmann, and Michael Katz. "Red-Black Planning: A New Systematic Approach to Partial Delete Relaxation". In: *Artificial Intelligence* 221 (2015), pp. 73–114.

[Ede01] Stefan Edelkamp. "Planning with Pattern Databases". In: *Proceedings of the 6th European Conference on Planning (ECP'01)*. Ed. by A. Cesta and D. Borrajo. Springer-Verlag, 2001, pp. 13–24.

[FD14] Zohar Feldman and Carmel Domshlak. "Simple Regret Optimization in Online Planning for Markov Decision Processes". In: *Journal of Artificial Intelligence Research* 51 (2014), pp. 165–205.

[Fis] John R. Fisher. *prolog :- tutorial*. URL: `https://www.cpp.edu/~jrfisher/www/prolog_tutorial/` (visited on 10/10/2019).

[FL03] Maria Fox and Derek Long. "PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains". In: *Journal of Artificial Intelligence Research* 20 (2003), pp. 61–124.

[Fla94] Peter Flach. Wiley, 1994. ISBN: 0471 94152 2. URL: `https://github.com/simply-logical/simply-logical/releases/download/v1.0/SL.pdf`.

[FN71] Richard E. Fikes and Nils Nilsson. "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving". In: *Artificial Intelligence* 2 (1971), pp. 189–208.

[Gen34] Gerhard Gentzen. "Untersuchungen über das logische Schließen I". In: *Mathematische Zeitschrift* 39.2 (1934), pp. 176–210.

[Ger+09] Alfonso Gerevini et al. "Deterministic planning in the fifth international planning competition: PDDL3 and experimental evaluation of the planners". In: *Artificial Intelligence* 173.5-6 (2009), pp. 619–668.

[GJ79] Michael R. Garey and David S. Johnson. *Computers and Intractability—A Guide to the Theory of NP-Completeness*. BN book: Freeman, 1979.

[Glo] *Grundlagen der Logik in der Informatik*. Course notes at `https://www8.cs.fau.de/_media/ws16:gloin:skript.pdf`. URL: `https://www8.cs.fau.de/_media/ws16:gloin:skript.pdf` (visited on 10/13/2017).

[GNT04] Malik Ghallab, Dana Nau, and Paolo Traverso. *Automated Planning: Theory and Practice*. Morgan Kaufmann, 2004.

[GS05] Carla Gomes and Bart Selman. "Can get satisfaction". In: *Nature* 435 (2005), pp. 751–752.

[GSS03] Alfonso Gerevini, Alessandro Saetti, and Ivan Serina. "Planning through Stochastic Local Search and Temporal Action Graphs". In: *Journal of Artificial Intelligence Research* 20 (2003), pp. 239–290.

[Hau85] John Haugeland. *Artificial intelligence: the very idea*. Massachusetts Institute of Technology, 1985.

[HD09] Malte Helmert and Carmel Domshlak. "Landmarks, Critical Paths and Abstractions: What's the Difference Anyway?" In: *Proceedings of the 19th International Conference on Automated Planning and Scheduling (ICAPS'09)*. Ed. by Alfonso Gerevini et al. AAAI Press, 2009, pp. 162–169.

[HE05] Jörg Hoffmann and Stefan Edelkamp. "The Deterministic Part of IPC-4: An Overview". In: *Journal of Artificial Intelligence Research* 24 (2005), pp. 519–579.

[Hel06] Malte Helmert. "The Fast Downward Planning System". In: *Journal of Artificial Intelligence Research* 26 (2006), pp. 191–246.

[Her+13a] Ivan Herman et al. *RDF 1.1 Primer (Second Edition). Rich Structured Data Markup for Web Documents*. W3C Working Group Note. World Wide Web Consortium (W3C), 2013. URL: `http://www.w3.org/TR/rdfa-primer`.

[Her+13b] Ivan Herman et al. *RDFa 1.1 Primer – Second Edition. Rich Structured Data Markup for Web Documents*. W3C Working Goup Note. World Wide Web Consortium (W3C), Apr. 19, 2013. URL: http://www.w3.org/TR/xhtml-rdfa-primer/.

[HG00] Patrik Haslum and Hector Geffner. "Admissible Heuristics for Optimal Planning". In: *Proceedings of the 5th International Conference on Artificial Intelligence Planning Systems (AIPS'00)*. Ed. by S. Chien, R. Kambhampati, and C. Knoblock. Breckenridge, CO: AAAI Press, Menlo Park, 2000, pp. 140–149.

[HG08] Malte Helmert and Hector Geffner. "Unifying the Causal Graph and Additive Heuristics". In: *Proceedings of the 18th International Conference on Automated Planning and Scheduling (ICAPS'08)*. Ed. by Jussi Rintanen et al. AAAI Press, 2008, pp. 140–147.

[HHH07] Malte Helmert, Patrik Haslum, and Jörg Hoffmann. "Flexible Abstraction Heuristics for Optimal Sequential Planning". In: *Proceedings of the 17th International Conference on Automated Planning and Scheduling (ICAPS'07)*. Ed. by Mark Boddy, Maria Fox, and Sylvie Thiebaux. Providence, Rhode Island, USA: Morgan Kaufmann, 2007, pp. 176–183.

[Hit+12] Pascal Hitzler et al. *OWL 2 Web Ontology Language Primer (Second Edition)*. W3C Recommendation. World Wide Web Consortium (W3C), 2012. URL: http://www.w3.org/TR/owl-primer.

[HN01] Jörg Hoffmann and Bernhard Nebel. "The FF Planning System: Fast Plan Generation Through Heuristic Search". In: *Journal of Artificial Intelligence Research* 14 (2001), pp. 253–302.

[Hof11] Jörg Hoffmann. "Every806thing You Always Wanted to Know about Planning (But Were Afraid to Ask)". In: *Proceedings of the 34th Annual German Conference on Artificial Intelligence (KI'11)*. Ed. by Joscha Bach and Stefan Edelkamp. Vol. 7006. Lecture Notes in Computer Science. Springer, 2011, pp. 1–13. URL: http://fai.cs.uni-saarland.de/hoffmann/papers/ki11.pdf.

[How60] R. A. Howard. *Dynamic Programming and Markov Processes*. MIT Press, 1960.

[ILD] *7. Constraints: Interpreting Line Drawings*. URL: https://www.youtube.com/watch?v=l-tzjenXrvI&t=2037s (visited on 11/19/2019).

[JN33] E. S. Pearson J. Neyman. "IX. On the problem of the most efficient tests of statistical hypotheses". In: *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 231.694-706 (1933), pp. 289–337. DOI: 10.1098/rsta.1933.0009.

[Kah11] Daniel Kahneman. *Thinking, fast and slow*. Penguin Books, 2011. ISBN: 9780141033570.

[KC04] Graham Klyne and Jeremy J. Carroll. *Resource Description Framework (RDF): Concepts and Abstract Syntax*. W3C Recommendation. World Wide Web Consortium (W3C), Feb. 10, 2004. URL: http://www.w3.org/TR/2004/REC-rdf-concepts-20040210/.

[KD09] Erez Karpas and Carmel Domshlak. "Cost-Optimal Planning with Landmarks". In: *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI'09)*. Ed. by C. Boutilier. Pasadena, California, USA: Morgan Kaufmann, July 2009, pp. 1728–1733.

[Kee74] R. L. Keeney. "Multiplicative utility functions". In: *Operations Research* 22 (1974), pp. 22–34.

[KHD13] Michael Katz, Jörg Hoffmann, and Carmel Domshlak. "Who Said We Need to Relax *all* Variables?" In: *Proceedings of the 23rd International Conference on Automated Planning and Scheduling (ICAPS'13)*. Ed. by Daniel Borrajo et al. Rome, Italy: AAAI Press, 2013, pp. 126–134.

[KHH12a] Michael Katz, Jörg Hoffmann, and Malte Helmert. "How to Relax a Bisimulation?" In: *Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS'12)*. Ed. by Blai Bonet et al. AAAI Press, 2012, pp. 101–109.

[KHH12b] Emil Keyder, Jörg Hoffmann, and Patrik Haslum. "Semi-Relaxed Plan Heuristics". In: *Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS'12)*. Ed. by Blai Bonet et al. AAAI Press, 2012, pp. 128–136.

[KNS97] B. Kessler, G. Nunberg, and H. Schütze. "Automatic detection of text genre". In: *CoRR* cmp-lg/9707002 (1997).

[Koe+97] Jana Koehler et al. "Extending Planning Graphs to an ADL Subset". In: *Proceedings of the 4th European Conference on Planning (ECP'97)*. Ed. by S. Steel and R. Alami. Springer-Verlag, 1997, pp. 273–285. URL: `ftp://ftp.informatik.uni-freiburg.de/papers/ki/koehler-etal-ecp-97.ps.gz`.

[Koh08] Michael Kohlhase. "Using LaTeX as a Semantic Markup Format". In: *Mathematics in Computer Science* 2.2 (2008), pp. 279–304. URL: `https://kwarc.info/kohlhase/papers/mcs08-stex.pdf`.

[Kow97] Robert Kowalski. "Algorithm = Logic + Control". In: *Communications of the Association for Computing Machinery* 22 (1997), pp. 424–436.

[KS00] Jana Köhler and Kilian Schuster. "Elevator Control as a Planning Problem". In: *AIPS 2000 Proceedings*. AAAI, 2000, pp. 331–338. URL: `https://www.aaai.org/Papers/AIPS/2000/AIPS00-036.pdf`.

[KS06] Levente Kocsis and Csaba Szepesvári. "Bandit Based Monte-Carlo Planning". In: *Proceedings of the 17th European Conference on Machine Learning (ECML 2006)*. Ed. by Johannes Fürnkranz, Tobias Scheffer, and Myra Spiliopoulou. Vol. 4212. LNCS. Springer-Verlag, 2006, pp. 282–293.

[KS92] Henry A. Kautz and Bart Selman. "Planning as Satisfiability". In: *Proceedings of the 10th European Conference on Artificial Intelligence (ECAI'92)*. Ed. by B. Neumann. Vienna, Austria: Wiley, Aug. 1992, pp. 359–363.

[KS98] Henry A. Kautz and Bart Selman. "Pushing the Envelope: Planning, Propositional Logic, and Stochastic Search". In: *Proceedings of the Thirteenth National Conference on Artificial Intelligence AAAI-96*. MIT Press, 1998, pp. 1194–1201.

[Kur90] Ray Kurzweil. *The Age of Intelligent Machines*. MIT Press, 1990. ISBN: 0-262-11121-7.

[Lee02] Lillian Lee. "Fast context-free grammar parsing requires fast Boolean matrix multiplication". In: *Journal of the ACM* 49.1 (2002), pp. 1–15.

[LPN] *Learn Prolog Now!* URL: `http://lpn.swi-prolog.org/` (visited on 10/10/2019).

[LS93] George F. Luger and William A. Stubblefield. *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*. World Student Series. The Benjamin/Cummings, 1993. ISBN: 9780805347852.

[Luc96] Peter Lucas. "Knowledge Acquisition for Decision-theoretic Expert Systems". In: *AISB Quarterly* 94 (1996), pp. 23–33. URL: `https://www.researchgate.net/publication/2460438_Knowledge_Acquisition_for_Decision-theoretic_Expert_Systems`.

[McD+98] Drew McDermott et al. *The PDDL Planning Domain Definition Language*. The AIPS-98 Planning Competition Comitee. 1998.

[Met+53] N. Metropolis et al. "Equations of state calculations by fast computing machines". In: *Journal of Chemical Physics* 21 (1953), pp. 1087–1091.

[Min] *Minion - Constraint Modelling*. System Web page at `http://constraintmodelling.org/minion/`. URL: `http://constraintmodelling.org/minion/`.

[MMS93] Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. "Building a large annotated corpus of English: the penn treebank". In: *Computational Linguistics* 19.2 (1993), pp. 313–330.

[MR91] John Mylopoulos and Ray Reiter, eds. Sydney, Australia: Morgan Kaufmann, San Mateo, CA, 1991.

[MSL92] David Mitchell, Bart Selman, and Hector J. Levesque. "Hard and Easy Distributions of SAT Problems". In: *Proceedings of the 10th National Conference of the American Association for Artificial Intelligence (AAAI'92)*. San Jose, CA: MIT Press, 1992, pp. 459–465.

[NHH11] Raz Nissim, Jörg Hoffmann, and Malte Helmert. "Computing Perfect Heuristics in Polynomial Time: On Bisimulation and Merge-and-Shrink Abstraction in Optimal Planning". In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI'11)*. Ed. by Toby Walsh. AAAI Press/IJCAI, 2011, pp. 1983–1990.

[Nor+18a] Emily Nordmann et al. *Lecture capture: Practical recommendations for students and lecturers*. 2018. URL: `https://osf.io/huydx/download`.

[Nor+18b] Emily Nordmann et al. *Vorlesungsaufzeichnungen nutzen: Eine Anleitung für Studierende*. 2018. URL: `https://osf.io/e6r7a/download`.

[NS63] Allen Newell and Herbert Simon. "GPS, a program that simulates human thought". In: *Computers and Thought*. Ed. by E. Feigenbaum and J. Feldman. McGraw-Hill, 1963, pp. 279–293.

[NS76] Alan Newell and Herbert A. Simon. "Computer Science as Empirical Inquiry: Symbols and Search". In: *Communications of the ACM* 19.3 (1976), pp. 113–126. DOI: `10.1145/360018.360022`.

[OWL09] OWL Working Group. *OWL 2 Web Ontology Language: Document Overview*. W3C Recommendation. World Wide Web Consortium (W3C), Oct. 27, 2009. URL: `http://www.w3.org/TR/2009/REC-owl2-overview-20091027/`.

[PD09] Knot Pipatsrisawat and Adnan Darwiche. "On the Power of Clause-Learning SAT Solvers with Restarts". In: *Proceedings of the 15th International Conference on Principles and Practice of Constraint Programming (CP'09)*. Ed. by Ian P. Gent. Vol. 5732. Lecture Notes in Computer Science. Springer, 2009, pp. 654–668.

[Pól73] George Pólya. *How to Solve it. A New Aspect of Mathematical Method*. Princeton University Press, 1973.

[Pra+94] Malcolm Pradhan et al. "Knowledge Engineering for Large Belief Networks". In: *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*. UAI'94. Seattle, WA: Morgan Kaufmann Publishers Inc., 1994, pp. 484–490. ISBN: 1-55860-332-8. URL: `http://dl.acm.org/citation.cfm?id=2074394.2074456`.

[Pro] *Protégé*. Project Home page at `http://protege.stanford.edu`. URL: `http://protege.stanford.edu`.

[PRR97] G. Probst, St. Raub, and Kai Romhardt. *Wissen managen*. 4 (2003). Gabler Verlag, 1997.

[PS08] Eric Prud'hommeaux and Andy Seaborne. *SPARQL Query Language for RDF*. W3C Recommendation. World Wide Web Consortium (W3C), Jan. 15, 2008. URL: `http://www.w3.org/TR/2008/REC-rdf-sparql-query-20080115/`.

[PW92] J. Scott Penberthy and Daniel S. Weld. "UCPOP: A Sound, Complete, Partial Order Planner for ADL". In: *Principles of Knowledge Representation and Reasoning: Proceedings of the 3rd International Conference (KR-92)*. Ed. by B. Nebel, W. Swartout, and C. Rich. Cambridge, MA: Morgan Kaufmann, Oct. 1992, pp. 103–114. URL: `ftp://ftp.cs.washington.edu/pub/ai/ucpop-kr92.ps.Z`.

[RHN06] Jussi Rintanen, Keijo Heljanko, and Ilkka Niemelä. "Planning as satisfiability: parallel plans and algorithms for plan search". In: *Artificial Intelligence* 170.12-13 (2006), pp. 1031–1080.

[Rin10] Jussi Rintanen. "Heuristics for Planning with SAT". In: *Proceeedings of the 16th International Conference on Principles and Practice of Constraint Programming*. 2010, pp. 414–428.

[RN03] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 2nd ed. Pearso n Education, 2003. ISBN: 0137903952.

[RN09] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 3rd. Prentice Hall Press, 2009. ISBN: 0136042597, 9780136042594.

[RN95] Stuart J. Russell and Peter Norvig. *Artificial Intelligence — A Modern Approach*. Upper Saddle River, NJ: Prentice Hall, 1995.

[RW10] Silvia Richter and Matthias Westphal. "The LAMA Planner: Guiding Cost-Based Anytime Planning with Landmarks". In: *Journal of Artificial Intelligence Research* 39 (2010), pp. 127–177.

[RW91] S. J. Russell and E. Wefald. *Do the Right Thing — Studies in limited Rationality*. MIT Press, 1991.

[Sil+16] David Silver et al. "Mastering the Game of Go with Deep Neural Networks and Tree Search". In: *Nature* 529 (2016), pp. 484–503. URL: `http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html`.

[Smu63] Raymond M. Smullyan. "A Unifying Principle for Quantification Theory". In: *Proc. Nat. Acad Sciences* 49 (1963), pp. 828–832.

[SR14] Guus Schreiber and Yves Raimond. *RDF 1.1 Primer*. W3C Working Group Note. World Wide Web Consortium (W3C), 2014. URL: `http://www.w3.org/TR/rdf-primer`.

[SR91] C. Samuelsson and M. Rayner. "Quantitative evaluation of explanation-based learning as an optimization tool for a large-scale natural language system". In: *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI)*. Ed. by John Mylopoulos and Ray Reiter. Sydney, Australia: Morgan Kaufmann, San Mateo, CA, 1991, pp. 609–615.

[sTeX] *sTeX: A semantic Extension of TeX/LaTeX*. URL: `https://github.com/sLaTeX/sTeX` (visited on 05/11/2020).

[SWI] *SWI Prolog Reference Manual*. URL: `https://www.swi-prolog.org/pldoc/refman/` (visited on 10/10/2019).

[Tur50] Alan Turing. "Computing Machinery and Intelligence". In: *Mind* 59 (1950), pp. 433–460.

[Vas+17] Ashish Vaswani et al. "Attention is All you Need". In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon et al. Vol. 30. Curran Associates, Inc., 2017. URL: `https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf`.

[Wal75] David Waltz. "Understanding Line Drawings of Scenes with Shadows". In: *The Psychology of Computer Vision*. Ed. by P. H. Winston. McGraw-Hill, 1975, pp. 1–19.

[WHI] *Human intelligence — Wikipedia The Free Encyclopedia*. URL: `https://en.wikipedia.org/w/index.php?title=Human_intelligence` (visited on 04/09/2018).