

Artificial Intelligence 1  
Winter Semester 2024/25  
– Lecture Notes –  
Admin & Overview

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

2025-02-06

This document contains the administrative information and overview chapter of the course notes for the [course](#) “Artificial Intelligence 1” held at FAU Erlangen-Nürnberg in the Winter Semesters 2016/17 ff. Other parts of the [lecture notes](#) can be found at [http://kwarc.info/teaching/AI/notes-\\*.pdf](http://kwarc.info/teaching/AI/notes-*.pdf).

# Contents

<b>1</b>	<b>Preliminaries</b>	<b>5</b>
1.1	Administrative Ground Rules . . . . .	5
1.2	Getting Most out of AI-1 . . . . .	8
1.3	Learning Resources for AI-1 . . . . .	10
1.4	AI-Supported Learning . . . . .	13
<b>2</b>	<b>AI – Who?, What?, When?, Where?, and Why?</b>	<b>23</b>
2.1	What is Artificial Intelligence? . . . . .	23
2.2	Artificial Intelligence is here today! . . . . .	25
2.3	Ways to Attack the AI Problem . . . . .	29
2.4	Strong vs. Weak AI . . . . .	31
2.5	AI Topics Covered . . . . .	33
2.6	AI in the KWARC Group . . . . .	34



# 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](#))

**Note:** I do not literally presuppose the [courses](#) on the slide above – most of you do not have a [bachelor's degree](#) from FAU, so you cannot have taken them. And indeed some of the content of these [courses](#) is irrelevant for AI-1. Stating these [courses](#) is just the easiest way to specifying what content I will be building on – and any [graduate courses](#) has to build on something.

Many of you will have taken the moral equivalent of these [courses](#) in your [undergraduate](#) studies at your home university. If you did not, you will have to somehow catch up on the content as we go along in AI-1. This should be possible with enough motivation.

There are essentially three skillsets that are essential for AI-1:

1. A solid understanding and practical skill in programming (whatever programming language)
2. A good understanding and practice in using [mathematical](#) language to represent complex structures
3. A solid understanding of [formal languages](#) and [grammars](#), as well as applied [complexity theory](#) (basics of [theoretical computer science](#)).

Without (catching up on) these the AI-1 [course](#) will be quite frustrating and hard.

We will briefly go over the most important topics in ?? to synchronize concepts and notation. Note that if you do not have a formal education in [courses](#) like the ones mentioned above you will very probably have to do significant remedial work.

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](#))  $\sim$  100% of the [grade](#).
- ▷ Up to 10% bonus on-top for an [exam](#) with  $\geq 50\%$  points. ( $< 50\% \sim$  no bonus)
- ▷ Bonus points  $\hat{=}$  [percentage sum](#) of the best 10 [prepquizzes](#) divided by 100.

▷ **Exam:** 90 minutes [exam](#) conducted in presence on paper! ( $\sim$  April 1. 2025)

▷ **Retake Exam:** 90 min [exam](#) six months later. ( $\sim$  October 1. 2025)

▷  Register for [exams](#) in <https://campo.fau.de>. (there is a deadline!)

▷ **Note:** You can de-register from an [exam](#) on <https://campo.fau.de> up to three working days before [exam](#). (do not miss that if you are not prepared)



## Preparedness Quizzes


▷ **PrepQuizzes:** Every tuesday 16:15 we start the [lecture](#) with a 10 min online [quiz](#) – the [PrepQuiz](#) – 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 [prepquiz](#) will be given in the [ALEA](#) system

- ▷ <https://courses.voll-ki.fau.de/quiz-dash/ai-1>
- ▷ You have to be **logged into ALEA!** (via FAU IDM)
- ▷ You can take the **prepquiz** on your laptop or phone, ...
- ▷ ... in the **lecture** or at home ...
- ▷ ... via WLAN or 4G Network. (do not overload)
- ▷ **Prequizzes** will only be available 16:15-16:25!



## This Thursday: Pretest



- ▷ This thursday we will try out the **prepquiz** 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: younger than 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 **prequizzes**.
- ▷ Participation in the **pretest** is optional; it will not influence grades in any way.
- ▷ The **pretest** 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)

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 1, and thus you may not be able to register for the exam via <https://campo.fau.de>.
  - ▷ Just send me an e-mail and come to the exam, (we do the necessary admin)
  - ▷ 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 “Wahlpflichtbereich”.
  - ▷ ECTS credits need to be divisible by five  $\leftrightarrow 7.5 + 7.5 = 15$ .


Michael Kohlhas: Artificial Intelligence 1
5
2025-02-06


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



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

None of the things discussed in this section – homeworks, tutorials, study groups, and attendance – are mandatory (we cannot force you to do them; we offer them to you as learning opportunities), but most of them are very clearly correlated with success (i.e. passing the exam and getting a good grade), so taking advantage of them may be in your own interest.

### AI-1 Homework Assignments

---

- ▷ **Goal:** Homework assignments reinforce what was taught in lectures.
- ▷ **Homework Assignments:** Small individual problem/programming/proof task
  - ▷ but take time to solve (at least read them directly  $\leadsto$  questions)
- ▷ **Didactic Intuition:** Homework assignments give you material to test your understanding and show you how to apply it.
- ▷ **⚠ Homeworks** give no points, but without trying you are unlikely to pass the exam.
- ▷ Homeworks will be mainly peer-graded in the ALEA system.
- ▷ **Didactic Motivation:** Through peer grading students are able to see mistakes in their thinking and can correct any problems in future assignments. By grading assignments, students may learn how to complete assignments more accurately and how to improve their future results. (not just us being lazy)


Michael Kohlhas: Artificial Intelligence 1
6
2025-02-06


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!

## AI-1 Homework Assignments – Howto

- ▷ **Homework Workflow:** in [ALEA](#) (see below)
  - ▷ [Homework assignments](#) will be published on thursdays: see <https://courses.voll-ki.fau.de/hw/ai-1>
  - ▷ Submission of solutions via the [ALEA](#) system in the week after
  - ▷ [Peer grading/feedback](#) (and master solutions) via answer classes.
- ▷ **Quality Control:** [TAs](#) and [instructors](#) will monitor and supervise [peer grading](#).
- ▷ **Experiment:** Can we motivate enough of you to make [peer assessment](#) self-sustaining?
  - ▷ I am appealing to your sense of community responsibility here . . .
  - ▷ You should only expect other's to [grade](#) your submission if you [grade](#) their's (cf. Kant's "Moral Imperative")
  - ▷ **Make no mistake:** The [grader](#) usually [learns](#) at least as much as the [gradee](#).
- ▷ **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 groups 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!)



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 [lectures](#), 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 the [lectures](#). (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](mailto:florian.rabe@fau.de)
- ▷ **Tutorials:** One each taught by Florian Rabe (lead); Yasmeen Shawat, Hatem Mousa, Xinyuan Tu, and Florian Guthmann.

- ▷ **Life-saving Advice:** Go to your [tutorial](#), and prepare for it by having looked at the slides and the [homework assignments](#)!

## Collaboration

- ▷ **Definition 1.2.1.** **Collaboration** (or **cooperation**) is the process of groups of **agents acting** together for common, mutual benefit, as opposed to **acting in competition** for selfish benefit. In a **collaboration**, every **agent** contributes to the common goal and benefits from the contributions of others.
- ▷ In **learning** situations, the benefit is “better **learning**”.
- ▷ **Observation:** In **collaborative learning**, the overall result can be significantly better than in **competitive learning**.
- ▷ **Good Practice:** Form **study groups**. (long- or short-term)
  1. ⚠ those **learners** who work most, **learn** most!
  2. ⚠ **freeloaders** – individuals 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)

As we said above, almost all of the components of the AI-1 [course](#) are optional. That even applies to attendance. But make no mistake, attendance is important to most of you. Let me explain, ...

## Do I need to attend the AI-1 Lectures

- ▷ Attendance is not mandatory for the AI-1 [course](#). (official version)
- ▷ **Note:** There are two ways of learning: (both are OK, your mileage may vary)
  - ▷ Approach **B:** Read [a book/papers](#) (here: [lecture notes](#))
  - ▷ Approach **I:** come to the [lectures](#), be [involved](#), interrupt the [instructor](#) whenever you have a question.

The only advantage of **I** over **B** is that books/papers do not answer questions
- ▷ Approach **S:** come to the [lectures](#) and **sleep does not work!**
- ▷ The closer you get to research, the more we need to [discuss](#)!

## 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

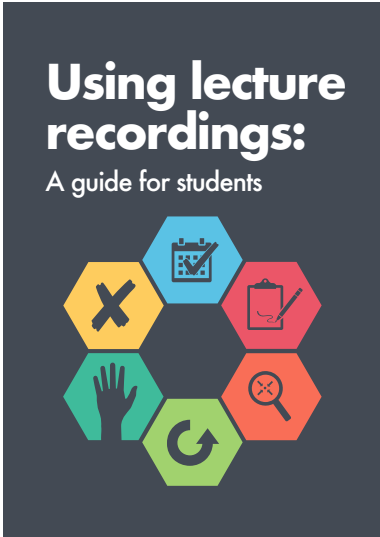
- ▷ **Textbook:** *Russel/Norvig: Artificial Intelligence, A modern Approach* [RN09].
  - ▷ basically “broad but somewhat shallow”
  - ▷ great to get intuitions on the basics of AI







Make sure that you read the **edition  $\geq 3$**   $\leftrightarrow$  vastly improved over  $\leq 2$ .
- ▷ **Lecture notes:** will be posted at <https://kwarc.info/teaching/AI>
  - ▷ more detailed than [RN09] in some areas
  - ▷ I mostly prepare them as we go along (semantically preloaded  $\leadsto$  research resource)
  - ▷ please e-mail me any errors/shortcomings you notice. (improve for the group)
- ▷ **Course Videos:** AI-1 will be streamed/recorded at <https://fau.tv/course/id/4047>
  - ▷ **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 exam!
- ▷ **StudOn Forum:** <https://www.studon.fau.de/crs5832535.html> for
  - ▷ announcements, homeworks (my view on the forum)
  - ▷ questions, discussion among your fellow students (your forum too, use it!)



FAU has issued a very insightful guide on using [lecture videos](#). It is a good idea to heed these recommendations, even if they seem annoying at first.

## Practical recommendations on Lecture Videos

- ▷ **Excellent Guide:** [Nor+18a] (German version at [Nor+18b])

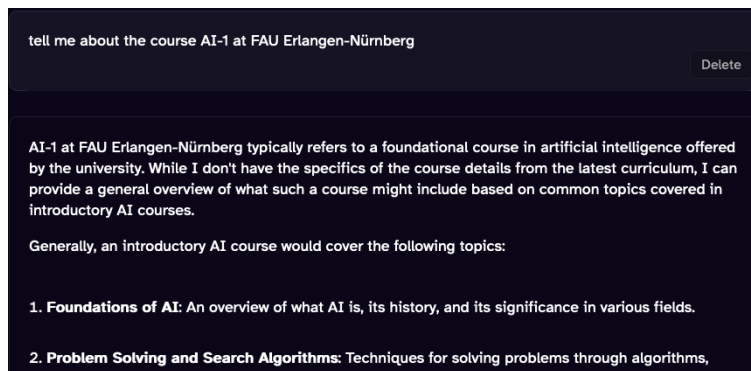


-  Attend lectures.
-  Take notes.
-  Be specific.
-  Catch up.
-  Ask for help.
-  Don't cut corners.


Michael Kohlhase: Artificial Intelligence 1
12
2025-02-06


## NOT a Resource for : LLMs – AI-based tools like ChatGPT

- ▷ **Definition 1.3.1.** A **large language model (LLM)** is a computational model capable of language generation or other natural language processing tasks.
- ▷ **Example 1.3.2.** OpenAI's GPT, Google's Bard, and Meta's Llama.
- ▷ **Definition 1.3.3.** A **chatbot** is a software application or web interface that is designed to mimic human conversation through text or voice interactions. Modern **chatbots** are usually based on **LLMs**.
- ▷ **Example 1.3.4 (ChatGPT talks about AI-1).** (but remains vague)



- ▷ **Note:** LLM-based **chatbots** invent *every word!* (surprisingly often correct)
- ▷ **Example 1.3.5 (In the AI-1 exam).** ChatGPT scores ca. 50% of the points.
  - ▷ ChatGPT can almost pass the exam . . . (We could award it a Master's degree)
  - ▷ But can you? (the AI-1 exams will be in person on paper)

You will only pass the exam, if you can do AI-1 yourself!

- ▷ **Intuition:** AI tools like GhatGPT, CoPilot, etc. (see also [She24])
  - ▷ can help you solve problems, (valuable tools in production situations)
  - ▷ hinders learning if used for homeworks/quizzes, etc. (like driving instead of jogging)
- ▷ **What (not) to do:** (to get most of the brave new AI-supported world)
  - ▷ try out these tools to get a first-hand intuition what they can/cannot do
  - ▷ challenge yourself while learning so that you can also do it (mind over matter!)



## 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 the AI-1 **course** it supports **learning** activities and serves as a showcase, what **symbolic AI** methods can do 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/lecture 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. (lecture notes on steroids)
- ▷ **Intuition:** ALEA serves active course materials. (PDF mostly inactive)
- ▷ **Goal:** Make ALEA more like a **instructor** + **study group** than like a book!
- ▷ **Example 1.4.2 (Course Notes).**  $\hat{=}$  Slides + Comments

The screenshot shows the ALEA interface. On the left is a table of contents with a search bar at the top. The table of contents has yellow highlights on the following items: 'Format of the AI Course/Lecturing Resources', 'Artificial Intelligence – Who?, W...', and 'Getting Started with AI: A Conce...'. On the right is a lecture note titled 'Specifying Control in Prolog'. The note contains text about Prolog's running time and includes code snippets for sorting algorithms. Below the code, there is a section titled 'Functions and Predicates in Prolog' with an assertion about their roles.


↪ yellow parts in table of contents (left) already covered in lectures.



The central idea in the AI4AI approach – using AI to support learning AI – and thus the ALeA system is that we want to make course materials – i.e. what we give to students for preparing and postparing lectures – more like teachers and study groups (only available 24/7) than like static books.

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


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



Artificial Intelligence - I

NOTES SLIDES

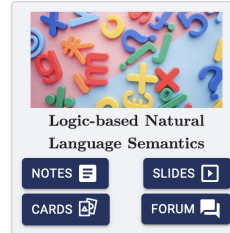
CARDS FORUM



IWGS - I

NOTES SLIDES

CARDS FORUM





Logic-based Natural Language Semantics

NOTES SLIDES

CARDS FORUM

- ▷ **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 semesters (for over/preview)
- ▷ **ALeA Status:** The ALeA system is deployed at FAU for over 1000 students taking eight courses
  - ▷ (some) students use the system actively (our logs tell us)
  - ▷ reviews are mostly positive/enthusiastic (error reports pour in)


Michael Kohlhasse: Artificial Intelligence 1
15
2025-02-06


The ALeA AI-1 page is the central entry point for working with the ALeA system. You can get to all the components of the system, including two presentations of the course contents (notes- and slides-centric ones), the flashcards, the localized forum, and the quiz dashboard.

We now come to the heart of the ALeA system: its learning support services, which we will now briefly introduce. Note that this presentation is not really sufficient to understand what you may be getting out of them, you will have to try them, and interact with them sufficiently that the learner model can get a good estimate of your competencies to adapt the results to you.

### 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 its definition. (even works recursively)

**Heuristic Functions**

▷ **Definition 1.1.11.** Let  $\Pi$  be a problem with states  $S$ . A **heuristic function** (or short **heuristic**) for  $\Pi$  is a function  $h: S \rightarrow \mathbb{R}_0^+ \cup \{\infty\}$  so that  $h(s) = 0$  whenever  $s$  is a **goal state**.

**Definition 0.1.** A **search problem**  $(S, \mathcal{A}, \mathcal{T}, \mathcal{I}, \mathcal{G})$  consists of a set  $S$  of states, a set  $\mathcal{A}$  of actions, and a **transition model**  $\mathcal{T}: \mathcal{A} \times S \rightarrow \mathcal{P}(S)$  that assigns to any action  $a \in \mathcal{A}$  and state  $s \in S$  a set of **successor states**. Certain states in  $S$  are designated as **goal states** ( $\mathcal{G} \subseteq S$ ) and **initial states**  $\mathcal{I} \subseteq S$ .

Strategies state, or  $\infty$  if no such path exists, is called the **goal distance function** for  $\Pi$ .

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

▷ **Axiom 0.1 (SAT: A kind of CSP).** SAT can be viewed as a CSP problem in which all variable domains are Boolean, and the constraints have unbounded arity.

▷ **Theorem 0.1 (Encoding CSP as SAT).** Given any constraint network  $\mathcal{C}$ , we can in low

▷ Symbol CNF

DM(de) AI(en) DM(en)

▷ A **formula** is in **conjunctive normal form (CNF)** if it is a **conjunction** of **disjunctions of literals**: i.e. if it is of the form  $\bigwedge_{i=1}^n \bigvee_{j=1}^{m_i} l_{ij}$

CLOSE

▷ **Axiom 0.1 (SAT: A kind of CSP).** SAT can be viewed as a CSP problem in which all variable domains are Boolean, and the constraints have unbounded arity.

▷ **Theorem 0.1 (Encoding CSP as SAT).** Given any constraint network  $\mathcal{C}$ , we can in low

▷ Symbol CNF

DM(de) AI(en) DM(en)

A **literal** is an **atomic formula** or a **negation** of one. A **formula** is said to be in

- **negation normal form (NNF)**, iff **negations** are **literals**.
- **conjunctive normal form (CNF)**, iff it is a **conjunction** of **disjunctions** of **literals**.
- **disjunctive normal form (DNF)**, iff it is a **disjunction** of **conjunctions** of **literals**.

CLOSE

▷ **Axiom 0.1 (SAT: A kind of CSP).** SAT can be viewed as a CSP problem in which all variable domains are Boolean, and the constraints have unbounded arity.  
 ▷ **Theorem 0.1 (Encoding CSP as SAT).** Given any constraint network  $\mathcal{C}$ , we can in low

▷ Symbol CNF

DM(de)   AII (en)   DM (en)

Ein **Literal** ist eine **atomare Formel** or die **Negation** einer solchen. Wir sagen, dass eine **Formel** eine

- **Negationsnormalform (NNF)** ist, wenn alle darin vorkommenden **Negationen** **Literale** sind.
- **konjunktive Normalform (CNF)** ist, wenn sie eine **Konjunktion** von **Diskjunktionen** von **Literalen** ist.
- **disjunktive Normalform (DNF)** ist, wenn sie eine **Disjunktion** von **Konjunktionen** von **Literalen** ist.

CLOSE

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

$c = \text{countable} \rightsquigarrow$

Guided Tour

- natural number
- conj
- equal
- set of pairs
- nCartProd
- subset
- converse relation
- transitive
- relation on
- irreflexive
- less than
- finite
- countable

**less than**

less than   finite   countable

Needs:   inset   natural number   nCartProd   converse relation   transitive

irreflexive

**Definition 0.1.** The  $\<$  relation is the **transitive closure** of the relation  $\{(n, s(n)) | n \in \mathbb{N}\}$ , and  $\leq$  its **transitive reflexive closure**.  $\&gt;$ ; and  $\leq$  are the corresponding **converse relations**.

For  $a \< b$ ; we say that  $a$  is **less than**  $b$ .

**finite**

finite   countable

Needs:   inset   natural number   less than

▷ **Definition 0.1.** We say that a set  $A$  is **finite** and has **cardinality**  $\#(A) \in \mathbb{N}$ , iff there is a bijective function  $f: A \rightarrow \{n \in \mathbb{N} | n \leq \#(A)\}$ .

**countable**

countable

Needs:   natural number   finite

▷ **Definition 0.1.** We say that a set  $A$  is **countably infinite**, iff there is a bijective function  $f: A \rightarrow \mathbb{N}$ . A set is called **countable**, iff it is **finite** or **countably infinite**.

▷ ... your idea here ...

(the sky is the limit)



Note that this is only an initial collection of **learning support services**, we are constantly working on additional ones. Look out for feature notifications ( ) on the upper right hand of the **ALeA** screen.

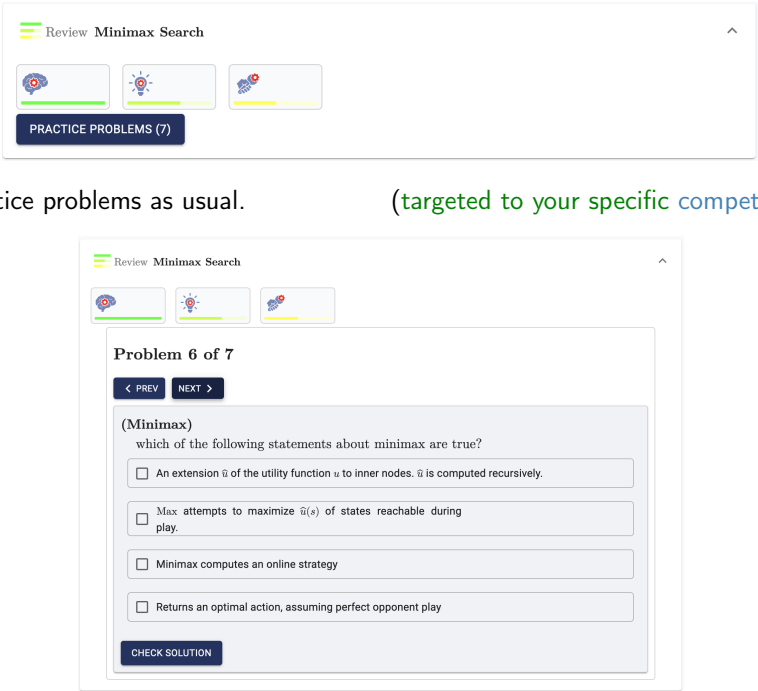
## (Practice/Remedial) Problems Everywhere

- ▷ **Problem:** Learning requires a mix of understanding and test-driven practice.
- ▷ **Idea:** ALeA supplies targeted practice problems everywhere.
- ▷ **Concretely:** Revision markers at the end of sections.
  - ▷ A relatively non-intrusive overview over **competency**



- ▷ Click to extend it for details.





Review Minimax Search

PRACTICE PROBLEMS (7)

▷ Practice problems as usual. (targeted to your specific competency)

Review Minimax Search

Problem 6 of 7

< PREV NEXT >

(Minimax)  
which of the following statements about minimax are true?

An extension  $\tilde{u}$  of the utility function  $u$  to inner nodes.  $\tilde{u}$  is computed recursively.

Max attempts to maximize  $\tilde{u}(s)$  of states reachable during play.

Minimax computes an online strategy

Returns an optimal action, assuming perfect opponent play

CHECK SOLUTION

FAU Michael Kohlhas: Artificial Intelligence 1 17 2025-02-06

While the [learning support services](#) up to now have been adressed to individual [learners](#), we now turn to services adressed to communities of [learners](#), ranging from [study groups](#) with three [learners](#), to whole [courses](#), and even – eventually – all the alumni of a [course](#), if they have not de-registered from [ALeA](#).

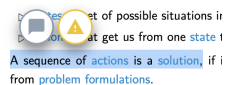
Currently, the community aspect of [ALeA](#) only consists in [localized interactions](#) with the [course materials](#).

The [ALeA](#) system uses the semantic structure of the [course materials](#) to [localize](#) some [interactions](#) that are otherwise often from separate applications. Here we see two:

1. one for reporting content errors – and thus making the material better for all [learners](#) – and“
2. a [localized](#) course forum, where forum threads can be attached to [learning objects](#).

## Localized Interactions with the Community

- ▷ Selecting [text](#) brings up [localized](#) – i.e. anchored on the 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](#).)

problem in the abstract, i.e. make a plan before we actually enter the situation (i.e. *offline*), and then when the problem arises, only execute the plan. If we do not have the knowledge to solve the problem, we can only make partial plans, and have to be in the situation to be able to execute the actions of others). As this is much more difficult w

MY NOTES COMMENTS

1 comments

Michael Kohlhase  Hide Identity

A sequence of actions is a solution

It could equivalently be defined as a sequence of actions: we can compute the state sequence from the action sequence and – given the initial state – the action sequence from the state sequence.

Request response

POST

Michael Kohlhase 4 minutes ago REPLY

A sequence of actions is a solution

I do not understand this, why isn't a solution a sequence of states?

CLOSE

FAU Michael Kohlhase: Artificial Intelligence 1 18 2025-02-06

- ▷ Answering questions gives **karma**  $\hat{=}$  a public measure of **user** helpfulness.
- ▷ Notes can be anonymous ( $\leadsto$  generate no karma)

Let us briefly look into how the **learning support services** introduced above might work, focusing on where the necessary information might come from. Even though some of the concepts in the discussion below may be new to AI-1 **students**, it is worth looking into them. Bear with us as we try to explain the **AI** components of the **ALeA** system.

### ALeA $\hat{=}$ Data-Driven & AI-enabled Learning Assistance

- ▷ **Idea:** Do what a teacher does!  
Use/maintain four models:
- ▷ **Ingredient 1:** **Domain model**  $\hat{=}$  knowledge/theory graph
- ▷ **Ingredient 2:** **Learner model**  $\hat{=}$  adding **competency** estimations
- ▷ **Ingredient 3:** A collection of ready-formulated **learning objects**
- ▷ **Ingredient 4:** Educational dialogue planner  $\leadsto$  **guided tours**

(Good) teachers

- ▷ understand the objects and their properties they are talking about
- ▷ have readimade formulations how to convey them best
- ▷ and understand how these best work together
- ▷ model what the **learners** already **know**/understand and adapts them accordingly

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  $\times$  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

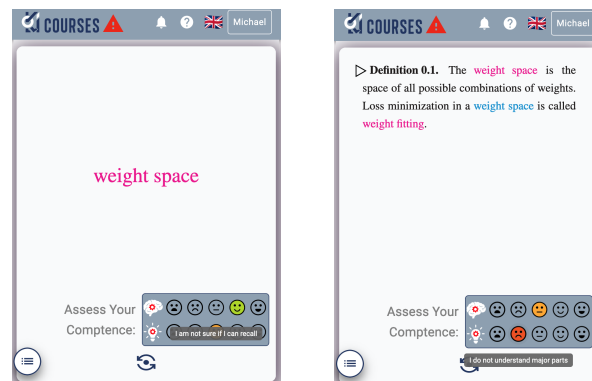
The dialogue planner assembles learning objects into active course material using

- ▷ 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

We can use the same four models discussed in the space of guided tours to deploy additional learning support services, which we now discuss.

## 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)

- ▷ **Idea:** Challenge yourself to a **card stack**, keep drilling/assessing flashcards until the **learner model** eliminates all.
- ▷ **Bonus:** Flashcards can be generated from existing semantic markup (**educational equivalent to free beer**)

We have already seen above how the **learner model** can drive the **drilling with flashcards**. It can also be used for the configuration of **card stacks** by configuring a **domain** e.g. a section in the **course materials** and a **competency threshold**. We now come to a very important issue that we always face when we do **AI systems** that **interface** with humans. Most web technology companies that take one the approach “the user pays for the services with their **personal data**, which is sold on” or integrate advertising for remuneration. Both are not acceptable in university setting.



But abstaining from monetizing **personal data** still leaves the problem how to protect it from intentional or accidental misuse. Even though the **GDPR** has quite extensive exceptions for research, the **ALeA** system – a research prototype – adheres to the principles and mandates of the **GDPR**. In particular it makes sure that **personal data** of the **learners** is only used in **learning support services** directly or indirectly initiated by the **learners** themselves.


## Learner Data and Privacy in ALeA

- ▷ **Observation:** Learning support services in **ALeA** use the **learner model**; they
  - ▷ need the **learner model** data to adapt to the individual **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.

So, now that you have an overview over what the **ALeA** system can do for you, let us see what you have to concretely do to be able to use it.

### Concrete Todos for ALeA

- ▷ **Recall:** You will use ALeA for the **prepquizzes** (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)

FAU Michael Kohlhas: Artificial Intelligence 1 22 2025-02-06 

Even if you did not understand some of the **AI jargon** or the underlying methods (yet), you should be good to go for using the **ALeA** system in your day-to-day work.



## 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>.

### Plot for this chapter

- ▷ Motivation, overview, and finding out what you already know
  - ▷ What is [Artificial Intelligence](#)?
  - ▷ What has [AI](#) already achieved?
  - ▷ A (very) quick walk through the AI-1 topics.
  - ▷ How can you get involved with [AI](#) at [KWARC](#)?

 Michael Kohlhase: Artificial Intelligence 1 23 2025-02-06 

## 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 [natural intelligence](#) – have problems defining it, as witnessed by the plethora of theories e.g. found at [WHI].

[What is Artificial Intelligence? Definition](#)

- ▷ **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.



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



**Note** that list of components is controversial as well. Some say that it lumps together cognitive capacities that should be distinguished or forgets others, . . . We state it here much more to get AI-1 **students** to think about the issues than to make it normative.




## 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 32) 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.



FAU Michael Kohlhas: Artificial Intelligence 1 26 2025-02-06

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. (and still promises answers soon)
- ▷ **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.

FAU Michael Kohlhas: Artificial Intelligence 1 27 2025-02-06

All of these phenomena can be seen in the growth of AI as an academic discipline over the course of its now over 70 year long history.

### The current AI Hype — Part of a longer Story

- ▷ The history of AI as a discipline has been very much tied to the amount of funding – that allows us to do research and development.
- ▷ Funding levels are tied to public perception of success (especially for AI)
- ▷ **Definition 2.2.2.** An AI winter is a time period of low public perception and funding for AI, mostly because AI has failed to deliver on its – sometimes overblown – promises  
An AI summer is a time period of high public perception and funding for AI
- ▷ A potted history of AI (AI summers and winters)



Of course, the future of AI is still unclear, we are currently in a massive hype caused by the advent of deep neural networks being trained on all the data of the Internet, using the computational power of huge compute farms owned by an oligopoly of massive technology companies – we are definitely in an AI summer.

But AI as a academic community and the tech industry also make outrageous promises, and the media pick it up and distort it out of proportion, . . . So public opinion could flip again, sending AI into the next winter.

## 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.

### Four Main Approaches to Artificial Intelligence

- ▷ **Definition 2.3.1.** **Symbolic AI** is a subfield of AI based on the assumption that many aspects of intelligence can be achieved by the manipulation of symbols, combining them into meaning-carrying 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** (also called connectionism or neural AI) is a subfield of AI that posits that intelligence is inherently tied to brains, where information is represented by a simple sequence pulses that are processed in parallel via simple calculations realized by neurons, and thus concentrates on neural computing.
- ▷ **Definition 2.3.4.** **Embodied AI** posits that intelligence cannot be achieved by reasoning about the state of the world (symbolically, statistically, or connectivist), but must be embodied i.e. situated in the world, equipped with a “body” that can

interact with it via **sensors** and **actuators**. Here, the main method for realizing **intelligent behavior** is by **learning** from the world.

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**)

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

- ▷ **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	
100%	Producer Tasks
50%	Consumer Tasks
	10 <sup>3±1</sup> Concepts    10 <sup>6±1</sup> Concepts    Coverage

after Aarne Ranta [Ran17].

- ▷ **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.

FAU Michael Kohlhasse: Artificial Intelligence 1 31 2025-02-06 

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 provide 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 couple of thousand attributes 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**.

## 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**, and **AI complete**, iff **AGI** should be able to solve them all.

- ▷ **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!**

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 soon**)

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),
  - ▷ petr1.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

FAU Michael Kohlhas: Artificial Intelligence 1 36 2025-02-06 

## 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, there 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  $\hat{=}$  250-300 hours involvement by definition (1/3 of your time/ECTS)
- ▷ **BTW:** There will also be an AI2SysProj next semester! (another chance)

FAU Michael Kohlhas: Artificial Intelligence 1 37 2025-02-06 



## 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 ( $\triangle$  **Mathematics**  $\hat{=}$  **Anything Formal**  $\triangle$ )
  - ▷ 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


Michael Kohlhase: Artificial Intelligence 1
38
2025-02-06


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

<b>Applications:</b> eMath 3.0, Active Documents, Active Learning, Semantic Spreadsheets/CAD/CAM, Change Management, Global Digital Math Library, Math Search Systems, <b>SMGloM:</b> Semantic Multilingual Math Glossary, Serious Games, ...		
<b>Foundations of Math:</b> <ul style="list-style-type: none"> <li>▷ <b>MathML</b>, <i>OpenMath</i></li> <li>▷ advanced Type Theories</li> <li>▷ <b>MMT:</b> Meta Meta Theory</li> <li>▷ Logic Morphisms/Atlas</li> <li>▷ Theorem Prover/CAS Interoperability</li> <li>▷ Mathematical Models/Simulation</li> </ul>	<b>KM &amp; Interaction:</b> <ul style="list-style-type: none"> <li>▷ Semantic Interpretation (aka. Framing)</li> <li>▷ math-literate interaction</li> <li>▷ <b>MathHub:</b> math archives &amp; active docs</li> <li>▷ Active documents: embedded semantic services</li> <li>▷ Model-based Education</li> </ul>	<b>Semantization:</b> <ul style="list-style-type: none"> <li>▷ <b>LaTeXML:</b> <math>\text{LaTeX} \rightsquigarrow \text{XML}</math></li> <li>▷ <b>STeX:</b> Semantic <b>LaTeX</b></li> <li>▷ invasive editors</li> <li>▷ Context-Aware IDEs</li> <li>▷ Mathematical Corpora</li> <li>▷ Linguistics of Math</li> <li>▷ ML for Math Semantics Extraction</li> </ul>
<b>Foundations:</b> Computational Logic, Web Technologies, <b>OMDoc/MMT</b>		


Michael Kohlhase: Artificial Intelligence 1
39
2025-02-06


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 **students** 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, (Meta)<sup>n</sup>-Frameworks
  - ▷ Math Corpus Linguistics: Semantics Extraction
  - ▷ Serious Games, Cognitive Engineering, Math Information Retrieval, Legal Reasoning, ...
  - ▷ ... last but not least: KWARC is the home of [ALeA!](#)
- ▷ We always try to find a topic at the intersection of your and our interests.
- ▷ We also sometimes have positions!. (HiWi, Ph.D.:  $\frac{1}{2}$  E-13, PostDoc: full E-13)



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

- [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>.
- [Ran17] Aarne Ranta. *Automatic Translation for Consumers and Producers*. Presentation given at the Chalmers Initiative Seminar. 2017. URL: <https://www.grammaticalframework.org/~aarne/mt-digitalization-2017.pdf>.
- [RN09] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 3rd. Prentice Hall Press, 2009. ISBN: 0136042597, 9780136042594.
- [She24] Esther Shein. 2024. URL: <https://cacm.acm.org/news/the-impact-of-ai-on-computer-science-education/>.
- [WHI] *Human intelligence* — *Wikipedia The Free Encyclopedia*. URL: [https://en.wikipedia.org/w/index.php?title=Human\\_intelligence](https://en.wikipedia.org/w/index.php?title=Human_intelligence) (visited on 04/09/2018).

