2025-05-01

21.1 Preliminaries

What you should learn here...

- **What you should learn in AI-2:**
  - In the broadest sense: *A bunch of tools for your toolchest* (i.e. various (quasi-mathematical) models, first and foremost)

What you should learn here...

- **What you should learn in AI-2:**
  - In the broadest sense: *A bunch of tools for your toolchest* (i.e. various (quasi-mathematical) models, first and foremost)
  - the underlying *principles* of these models (assumptions, limitations, the math behind them ...)

What you should learn here...

- **What you should learn in AI-2:**
  - In the broadest sense: *A bunch of tools for your toolchest* (i.e. various (quasi-mathematical) models, first and foremost)
  - the underlying *principles* of these models (assumptions, limitations, the math behind them ...)
  - the ability to describe real-world problems in terms of these models, **where adequate** (...and knowing **when** they are adequate!),

What you should learn here...

- **What you should learn in AI-2:**
  - In the broadest sense: *A bunch of tools for your toolchest* (i.e. various (quasi-mathematical) models, first and foremost)
  - the underlying *principles* of these models (assumptions, limitations, the math behind them ...)
  - the ability to describe real-world problems in terms of these models, **where adequate** (...and knowing **when** they are adequate!), and
  - the ideas behind effective *algorithms* that solve these problems (and to understand them well enough to implement them)
- **Note:** You will likely never get payed to implement an algorithm that e.g. solves Bayesian networks. (They already exist)
  - *But* you might get payed to *recognize* that some given problem *can be* represented as a Bayesian network!
  - **Or:** you can recognize that it is *similar to* a Bayesian network, and reuse the underlying principles to develop new specialized tools.

Compare two employees

- **"We have the following problem and we need a solution: ..."**

Compare two employees

- **"We have the following problem and we need a solution: ..."**
- **Employee 1 – Deep Learning can do everything:** "I just need ≈1.5 million labeled examples of potentially sensitive data, a GPU cluster for training, and a few weeks to train, tweak and finetune the model.
  But *then* I can solve the problem... with a confidence of 95%, within 40 seconds of inference per input. Oh, as long as the input isn't longer than 15`unit`, or I will need to retrain on a bigger input layer..."

Compare two employees

- **"We have the following problem and we need a solution: . . ."**
- **Employee 1 – Deep Learning can do everything:** "I just need ≈1.5 million labeled examples of potentially sensitive data, a GPU cluster for training, and a few weeks to train, tweak and finetune the model.
  But *then* I can solve the problem... with a confidence of 95%, within 40 seconds of inference per input. Oh, as long as the input isn't longer than 15`unit`, or I will need to retrain on a bigger input layer..."
- **Employee 2 – AI-2 Alumna:** "...while you were talking, I quickly built a custom UI for an off-the-shelve `<problem>` solver that runs on a medium-sized potato and returns a *provably correct* result in a few milliseconds. For inputs longer than 1000`unit`, you might need a slightly bigger potato though..."

Compare two employees

- **"We have the following problem and we need a solution: ..."**
- **Employee 1 – Deep Learning can do everything:** "I just need ≈1.5 million labeled examples of potentially sensitive data, a GPU cluster for training, and a few weeks to train, tweak and finetune the model.
  But *then* I can solve the problem... with a confidence of 95%, within 40 seconds of inference per input. Oh, as long as the input isn't longer than 15`unit`, or I will need to retrain on a bigger input layer..."
- **Employee 2 – AI-2 Alumna:** "...while you were talking, I quickly built a custom UI for an off-the-shelve `<problem>` solver that runs on a medium-sized potato and returns a *provably correct* result in a few milliseconds. For inputs longer than 1000`unit`, you might need a slightly bigger potato though..."
- **Moral of the story:** Know your *tools* well enough to select the right one for the job.

21.1.1 Administrative Ground Rules

Prerequisites

- **Remember:** **AI-1** dealt with situations with "complete information" and strictly computable, "perfect" solutions to problems. (i.e. tree search, logical inference, planning, etc.)
- **AI-2** will focus on *probabilistic* scenarios by introducing uncertain situations, and *approximate* solutions to problems. (Bayesian networks, Markov models, machine learning, etc.)

Prerequisites

- **Remember:** **AI-1** dealt with situations with "complete information" and strictly computable, "perfect" solutions to problems. (i.e. tree search, logical inference, planning, etc.)
- **AI-2** will focus on *probabilistic* scenarios by introducing uncertain situations, and *approximate* solutions to problems. (Bayesian networks, Markov models, machine learning, etc.)
- **Weak Prerequisites for AI-2:** (if you do not have them, study up as needed)
  - AI-1 (in particular: PEAS, propositional logic/first-order logic (mostly the syntax), some logic programming)
  - (very) elementary complexity theory. (big Oh and friends)

Prerequisites

- **Remember:** **AI-1** dealt with situations with "complete information" and strictly computable, "perfect" solutions to problems. (i.e. tree search, logical inference, planning, etc.)
- **AI-2** will focus on *probabilistic* scenarios by introducing uncertain situations, and *approximate* solutions to problems. (Bayesian networks, Markov models, machine learning, etc.)
- **Weak Prerequisites for AI-2:** (if you do not have them, study up as needed)
  - AI-1 (in particular: PEAS, propositional logic/first-order logic (mostly the syntax), some logic programming)
  - (very) elementary complexity theory. (big Oh and friends)
  - rudimentary probability theory (e.g. from stochastics)

Prerequisites

- **Remember:** **AI-1** dealt with situations with "complete information" and strictly computable, "perfect" solutions to problems. (i.e. tree search, logical inference, planning, etc.)
- **AI-2** will focus on *probabilistic* scenarios by introducing uncertain situations, and *approximate* solutions to problems. (Bayesian networks, Markov models, machine learning, etc.)
- **Weak Prerequisites for AI-2:** (if you do not have them, study up as needed)
  - AI-1 (in particular: PEAS, propositional logic/first-order logic (mostly the syntax), some logic programming)
  - (very) elementary complexity theory. (big Oh and friends)
  - rudimentary probability theory (e.g. from stochastics)
  - basic linear algebra (vectors, matrices,...)

Prerequisites

- **Remember:** **AI-1** dealt with situations with “complete information” and strictly computable, “perfect” solutions to problems. (i.e. tree search, logical inference, planning, etc.)
- **AI-2** will focus on *probabilistic* scenarios by introducing uncertain situations, and *approximate* solutions to problems. (Bayesian networks, Markov models, machine learning, etc.)
- **Weak Prerequisites for AI-2:** (if you do not have them, study up as needed)
  - AI-1 (in particular: PEAS, propositional logic/first-order logic (mostly the syntax), some logic programming)
  - (very) elementary complexity theory. (big Oh and friends)
  - rudimentary probability theory (e.g. from stochastics)
  - basic linear algebra (vectors, matrices,...)
  - basic real analysis (aka. calculus) (primarily: (partial) derivatives)

Prerequisites

- **Remember:** **AI-1** dealt with situations with "complete information" and strictly computable, "perfect" solutions to problems. (i.e. tree search, logical inference, planning, etc.)
- **AI-2** will focus on *probabilistic* scenarios by introducing uncertain situations, and *approximate* solutions to problems. (Bayesian networks, Markov models, machine learning, etc.)
- **Weak Prerequisites for AI-2:** (if you do not have them, study up as needed)
  - AI-1 (in particular: PEAS, propositional logic/first-order logic (mostly the syntax), some logic programming)
  - (very) elementary complexity theory. (big Oh and friends)
  - rudimentary probability theory (e.g. from stochastics)
  - basic linear algebra (vectors, matrices,...)
  - basic real analysis (aka. calculus) (primarily: (partial) derivatives)
- **Meaning:** I will *assume* you know these things, but some of them we will recap, and what you don't know will make things slightly harder for you, but by no means prohibitively difficult.

“Strict” Prerequisites

- **Most crucially – Mathematical Literacy:** Mathematics is the language that computer scientists express their ideas in! (“*A search problem is a tuple* $(N, S, G, \ldots)$ *such that...*”)
- **Note:** This is a skill that can be *learned*, and more importantly, *practiced!* Not having/honing this skill *will* make things more difficult for you. Be aware of this and, if necessary, work on it – it will pay off, not only in this course.

"Strict" Prerequisites

- **Most crucially – Mathematical Literacy:** Mathematics is the language that computer scientists express their ideas in! ("*A search problem is a tuple* $(N, S, G, \ldots)$ *such that...*")
- **Note:** This is a skill that can be *learned*, and more importantly, *practiced!* Not having/honing this skill *will* make things more difficult for you. Be aware of this and, if necessary, work on it – it will pay off, not only in this course.
- **But also:** Motivation, interest, curiosity, hard work. (AI-2 is non-trivial)

"Strict" Prerequisites

- **Most crucially – Mathematical Literacy:** Mathematics is the language that computer scientists express their ideas in! (*"A search problem is a tuple $(N, S, G, \ldots)$ such that..."*)
- **Note:** This is a skill that can be *learned*, and more importantly, *practiced!* Not having/honing this skill *will* make things more difficult for you. Be aware of this and, if necessary, work on it – it will pay off, not only in this course.
- **But also:** Motivation, interest, curiosity, hard work. (AI-2 is non-trivial)
- **Note:** Grades correlate significantly with invested effort; including, but not limited to:
  - time spent on exercises, (learning is 80% perspiration, only 20% inspiration)
  - being here in presence, (humans are social animals ⇜ mirror neurons)
  - asking questions, (Q/A dialogues activate brains)

"Strict" Prerequisites

- **Most crucially – Mathematical Literacy:** Mathematics is the language that computer scientists express their ideas in! ("*A search problem is a tuple* $(N, S, G, \ldots)$ *such that...*")
- **Note:** This is a skill that can be *learned*, and more importantly, *practiced!* Not having/honing this skill *will* make things more difficult for you. Be aware of this and, if necessary, work on it – it will pay off, not only in this course.
- **But also:** Motivation, interest, curiosity, hard work. (AI-2 is non-trivial)
- **Note:** Grades correlate significantly with invested effort; including, but not limited to:
  - time spent on exercises, (learning is 80% perspiration, only 20% inspiration)
  - being here in presence, (humans are social animals ⇜ mirror neurons)
  - asking questions, (Q/A dialogues activate brains)
  - talking to your peers, (pool your insights, share your triumphs/frustrations)...

  All of these we try to support with the ALeA system. (which also gives us the data to prove this)

Assessment, Grades

- **Overall (Module) Grade:**
  - Grade via the exam (Klausur) $\leadsto$ 100% of the grade.
  - Up to 10% bonus on-top for an exam with $\geq$ 50% points. ($<$ 50% $\leadsto$ no bonus)
  - Bonus points $\widehat{=}$ percentage sum of the best 10 prepquizzes divided by 100.

Assessment, Grades

- **Overall (Module) Grade:**
  - Grade via the exam (Klausur) $\leadsto$ 100% of the grade.
  - Up to 10% bonus on-top for an exam with $\geq$ 50% points. ($<$ 50% $\leadsto$ no bonus)
  - Bonus points $\widehat{=}$ percentage sum of the best 10 prepquizzes divided by 100.
- **Exam:** exam conducted in presence on paper! ($\sim$ Oct. 10. 2025)
- **Retake Exam:** 90 minutes exam six months later. ($\sim$ April 10. 2026)
- ⚠ You have to register for exams in `https://campo.fau.de` in the first month of classes.
- **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:** Before every lecture we offer a 10 min online quiz – the PrepQuiz – about the material from the previous week. (16:15-16:25; starts in week 2)
- **Motivations:** We do this to
  - keep you prepared and working continuously. (primary)
  - bonus points if the exam has $\geq 50\%$ points (potential part of your grade)
  - update the ALeA learner model. (fringe benefit)
- The prepquizes will be given in the ALeA system

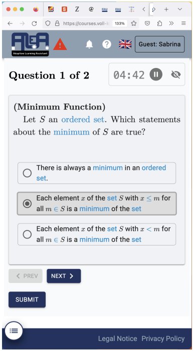

- https://courses.voll-ki.fau.de/quiz-dash/ai-2
- You have to be logged into ALeA! (via FAU IDM)
- You can take the prepquiz on your laptop or phone, . . .
- . . . in the lecture or at home . . .
- . . . via WLAN or 4G Network. (do not overload)
- Prepquizzes will only be available 16:15-16:25!

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 $\leadsto 7.5 + 7.5 = 15$.

21.1.2 Getting Most out of AI-2

AI-2 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 ⇝ 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.
- **Our Experience:** Doing your homework is probably even *more* important (and predictive of exam success) than attending the lecture in person!

AI-2 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 ⇝ 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.
- **Our Experience:** Doing your homework is probably even *more* important (and predictive of exam success) than attending the lecture in person!
- 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)

AI-2 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.

AI-2 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.

AI-2 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!)

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.

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, Privatdozent)
  - Room: 11.137 @ Händler building, `florian.rabe@fau.de`
- **Tutorials:** One each taught by Florian Rabe (lead); Primula Mukherjee, Ilhaam Shaikh, Praveen Kumar Vadlamani, and Shreya Rajesh More.
  - Tutorials will start in week 3. (before there is nothing to do)
  - Details (rooms, times, etc) will be announced in time (i.e. not now) on the forum and matrix channel.
- **Life-saving Advice:** Go to your tutorial, and prepare for it by having looked at the slides and the homework assignments!

Collaboration

- **Definition 1.1.** 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/help most, learn most!
  2. Freeloaders – individuals who only watch – learn very little!
- It is OK to collaborate on homework assignments in AI-2! (no bonus points)
- Choose your study group well! (ALeA helps via the study buddy feature)

Do I need to attend the AI-2 Lectures

- Attendance is not mandatory for the AI-2 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!

21.1.3 Learning Resources for AI-2

Textbooks and supplementary Literature

- **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 $\leftrightsquigarrow$ vastly improved over ≤ 2.

Course Notes, Forum, Matrix

- **Lecture notes** will be posted at `https://kwarc.info/teaching/AI`
  - We mostly prepare/update them as we go along (semantically preloaded ⇝ research resource)
  - Please report any errors/shortcomings you notice. (improve for the group/successors)
- **StudOn Forum:** For announcements – `https://www.studon.fau.de/studon/goto.php?target=lcode_70Bjcaxg`
- **Matrix Channel:** `https://matrix.to/#/#ai-12:fau.de` for questions, discussion with instructors and among your fellow students. (your channel, use it!)
  Login via FAU IDM ⇝ instructions
- **Course Videos** are at at `https://fau.tv/course/id/4225`.
- **Do not let the videos mislead you:** Coming to class is highly correlated with passing the exam!

Practical recommendations on Lecture Videos

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

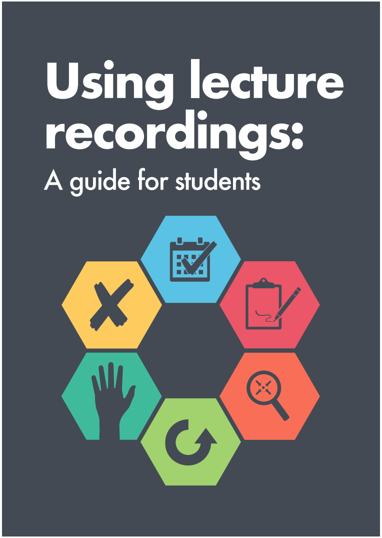

- Attend lectures.
- Take notes.
- Be specific.
- Catch up.
- Ask for help.
- Don't cut corners.

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

- **Definition 1.2.** A large language model (LLM) is a computational model capable of language generation or other natural language processing tasks.
- **Example 1.3.** OpenAI's GPT, Google's Bard, and Meta's Llama.

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

- **Definition 1.7.** A large language model (LLM) is a computational model capable of language generation or other natural language processing tasks.
- **Example 1.8.** OpenAI's GPT, Google's Bard, and Meta's Llama.
- **Definition 1.9.** 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 1.10 (ChatGPT talks about AI-1).** (but remains vague)

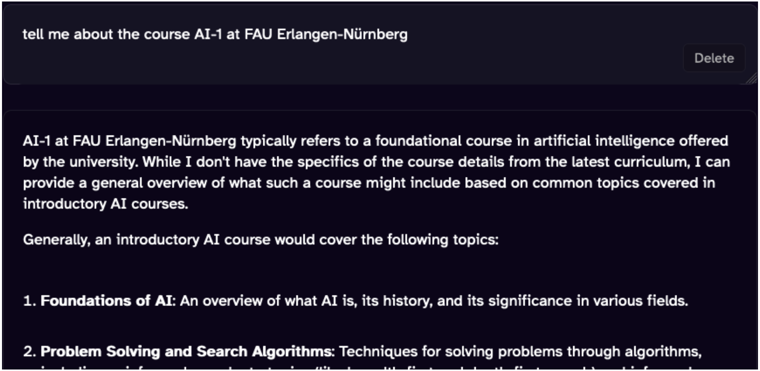

NOT a Resource for : LLMs – AI-based tools like ChatGPT

- **Definition 1.12.** A large language model (LLM) is a computational model capable of language generation or other natural language processing tasks.
- **Example 1.13.** OpenAI's GPT, Google's Bard, and Meta's Llama.
- **Definition 1.14.** 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 1.15 (ChatGPT talks about AI-1).** (but remains vague)
- **Note:** LLM-based chatbots invent *every word*! (surpisingly often correct)
- **Example 1.16 (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!

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

- **Definition 1.17.** A large language model (LLM) is a computational model capable of language generation or other natural language processing tasks.
- **Example 1.18.** OpenAI's GPT, Google's Bard, and Meta's Llama.
- **Definition 1.19.** 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 1.20 (ChatGPT talks about AI-1).** (but remains vague)
- **Note:** LLM-based chatbots invent *every word*! (surpprisingly often correct)
- **Example 1.21 (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)

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

- **Definition 1.22.** A large language model (LLM) is a computational model capable of language generation or other natural language processing tasks.
- **Example 1.23.** OpenAI's GPT, Google's Bard, and Meta's Llama.
- **Definition 1.24.** 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 1.25 (ChatGPT talks about AI-1).** (but remains vague)
- **Note:** LLM-based chatbots invent *every word*! (surprisingly often correct)
- **Example 1.26 (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!)

ALeA in AI-2

- We assume that you already know the ALeA system from last semester

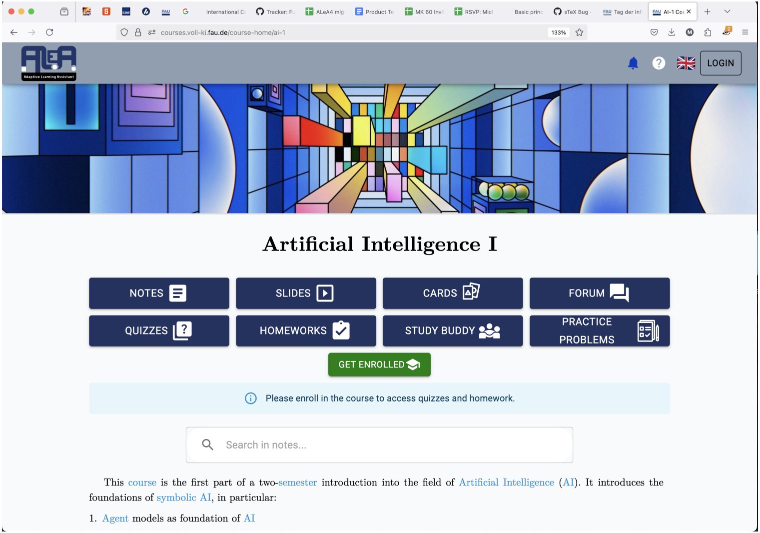

ALeA in AI-2

- We assume that you already know the ALeA system from last semester
- Use it for
  - lecture notes (notes- vs slides-oriented)
  - flashcards (drill yourself on the AI-2 jargon/concepts)
  - course forum (questions, discussions and error reporting)
  - solving and peer-grading homework assignments
  - finding study groups (you need not endure AI-2 alone)
  - practicing with targeted problems (e.g. from old exams)
  - doing the prepquizzes (before each lecture)

21.2 Overview over AI and Topics of AI-II

21.2.1 What is Artificial Intelligence?

What is Artificial Intelligence? Definition

- **Definition 2.1 (According to Wikipedia).** Artificial Intelligence (AI) is intelligence exhibited by machines
- **Definition 2.2 (also).** Artificial Intelligence (AI) is a sub-field of CS that is concerned with the automation of intelligent behavior.
- **BUT:** it is already difficult to define intelligence precisely.
- **Definition 2.3 (Elaine Rich).** artificial intelligence (AI) studies how we can make the computer do things that humans can still do better at the moment.

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

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

Inference

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

Perception

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

Language understanding

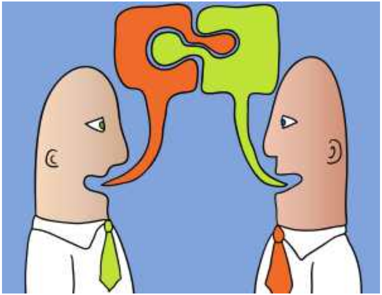

What is Artificial Intelligence? Components

- **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

Emotion

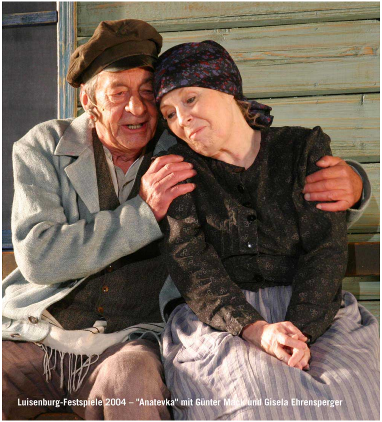

21.2.2 Artificial Intelligence is here today!

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
- in household appliances
- in hospitals
- for safety/security

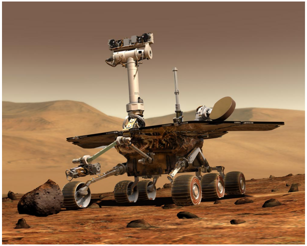

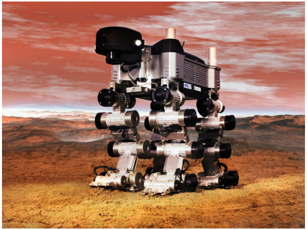

Artificial Intelligence is here today!

- in outer space
- in artificial limbs
  - the user controls the prosthesis via existing nerves, can e.g. grip a sheet of paper.
- in household appliances
- in hospitals
- for safety/security

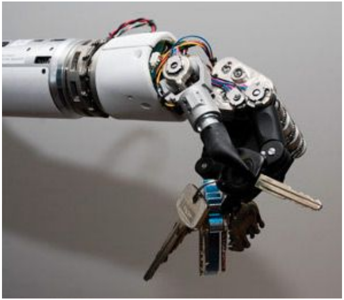

Artificial Intelligence is here today!

- in outer space
- in artificial limbs
- 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
- for safety/security

Artificial Intelligence is here today!

- in outer space
- in artificial limbs
- in household appliances
- 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.
- for safety/security

Artificial Intelligence is here today!

- in outer space
- in artificial limbs
- in household appliances
- in hospitals
- for safety/security
  - e.g. Intel verifies correctness of all chips after the “Pentium 5 disaster”

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 "CS". (becomes a separate subfield of CS)
- **Example 2.4.** Functional/Logic Programming, automated theorem proving, Planning, machine learning, Knowledge Representation, ...
- **Still Consequence:** AI research was alternatingly flooded with money and cut off brutally.

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.

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)

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 2.7.** 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

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 2.8.** 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 summers)

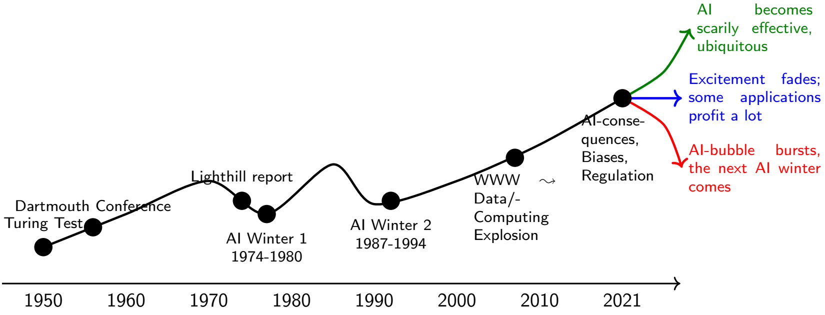

21.2.3 Ways to Attack the AI Problem

Four Main Approaches to Artificial Intelligence

- **Definition 2.9.** 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.

Four Main Approaches to Artificial Intelligence

- **Definition 2.13.** 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 2.14.** 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.

Four Main Approaches to Artificial Intelligence

- **Definition 2.17.** 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 2.18.** 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 2.19.** 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.

Four Main Approaches to Artificial Intelligence

- **Definition 2.21.** 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 2.22.** 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 2.23.** 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 2.24.** 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.

Two ways of reaching Artificial Intelligence?

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

Deep	symbolic AI-1	not there yet cooperation?
Shallow	no-one wants this	statistical/sub symbolic AI-2
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)

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)

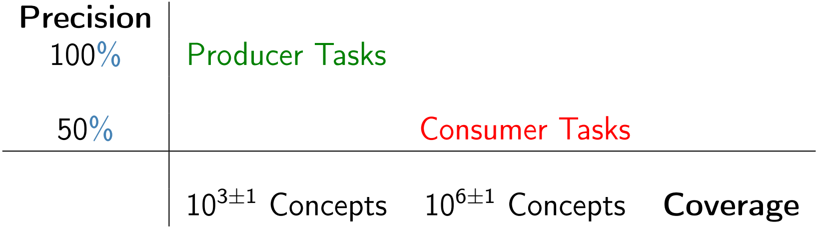

after Aarne Ranta [Ran17].

- **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.

21.2.4 AI in the KWARC Group

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 $\widehat{=}$ 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

Overview: KWARC Research and Projects

<table>
<tr><td colspan="3">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, ...</td></tr>
<tr>
<td>Foundations of Math:
▶ MathML, OpenMath
▶ advanced Type Theories
▶ MMT: Meta Meta Theory
▶ Logic Morphisms/Atlas
▶ Theorem Prover/CAS Interoperability
▶ Mathematical Models/Simulation</td>
<td>KM & Interaction:
▶ Semantic Interpretation (aka. Framing)
▶ math-literate interaction
▶ MathHub: math archives & active docs
▶ Active documents: embedded semantic services
▶ Model-based Education</td>
<td>Semantization:
▶ LaTeXML: LaTeX ⇝ XML
▶ sTeX: Semantic LaTeX
▶ invasive editors
▶ Context-Aware IDEs
▶ Mathematical Corpora
▶ Linguistics of Math
▶ ML for Math Semantics Extraction</td>
</tr>
<tr><td colspan="3">Foundations: Computational Logic, Web Technologies, OMDoc/MMT</td></tr>
</table>

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, ...
  - ...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)

21.2.5 Agents and Environments in AI2

21.2.5.1 Recap: Rational Agents as a Conceptual Framework

Agents and Environments

- **Definition 2.25.** An agent is anything that
  - perceives its environment via sensors (a means of sensing the environment)
  - acts on it with actuators (means of changing the environment).

  Any recognizable, coherent employment of the actuators of an agent is called an action.

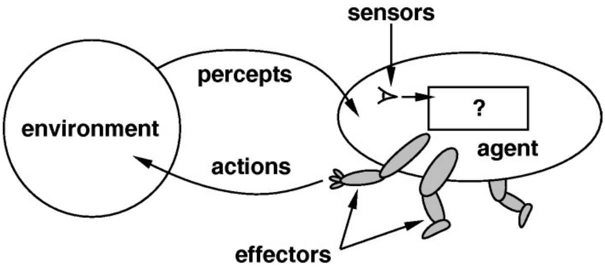

- **Example 2.26.** Agents include humans, robots, softbots, thermostats, etc.
- **Remark:** The notion of an agent and its environment is intentionally designed to be inclusive. We will classify and discuss subclasses of both later.

Agent Schema: Visualizing the Internal Agent Structure

- **Agent Schema:** We will use the following kind of agent schema to visualize the internal structure of an agent:

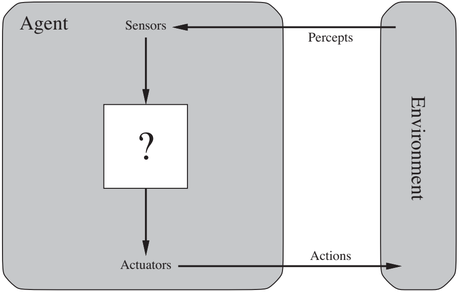

Different agents differ on the contents of the white box in the center.

Rationality

- **Idea:** Try to design agents that are successful! (aka. "do the right thing")
- **Problem:** What do we mean by "successful", how do we measure "success"?
- **Definition 2.27.** A performance measure is a function that evaluates a sequence of environments.
- **Example 2.28.** A performance measure for a vacuum cleaner could
  - award one point per "square" cleaned up in time T?
  - award one point per clean "square" per time step, minus one per move?
  - penalize for $> k$ dirty squares?
- **Definition 2.29.** An agent is called rational, if it chooses whichever action maximizes the expected value of the performance measure given the percept sequence to date.
- **Critical Observation:** We only need to maximize the expected value, not the actual value of the performance measure!
- **Question:** Why is rationality a good quality to aim for?

Consequences of Rationality: Exploration, Learning, Autonomy

- **Note:** A rational agent need not be perfect:
  - It only needs to maximize expected value (rational $\neq$ omniscient)
    - need not predict e.g. very unlikely but catastrophic events in the future
  - Percepts may not supply all relevant information (rational $\neq$ clairvoyant)
    - if we cannot perceive things we do not need to react to them.
    - but we may need to try to find out about hidden dangers (exploration)
  - Action outcomes may not be as expected (rational $\neq$ successful)
    - but we may need to take action to ensure that they do (more often) (learning)
- **Note:** Rationality may entail exploration, learning, autonomy (depending on the environment / task)
- **Definition 2.30.** An agent is called autonomous, if it does not rely on the prior knowledge about the environment of the designer.
- Autonomy avoids fixed behaviors that can become unsuccessful in a changing environment. (anything else would be irrational)
- The agent may have to learn all relevant traits, invariants, properties of the environment and actions.

PEAS: Describing the Task Environment

- **Observation:** To design a rational agent, we must specify the task environment in terms of performance measure, environment, actuators, and sensors, together called the PEAS components.
- **Example 2.31.** When designing an automated taxi:
  - **Performance measure:** safety, destination, profits, legality, comfort, ...
  - **Environment:** US streets/freeways, traffic, pedestrians, weather, ...
  - **Actuators:** steering, accelerator, brake, horn, speaker/display, ...
  - **Sensors:** video, accelerometers, gauges, engine sensors, keyboard, GPS, ...
- **Example 2.32 (Internet Shopping Agent).** The task environment:
  - Performance measure: price, quality, appropriateness, efficiency
  - Environment: current and future WWW sites, vendors, shippers
  - Actuators: display to user, follow URL, fill in form
  - Sensors: HTML pages (text, graphics, scripts)

Environment types

- **Observation 2.33.** *Agent design is largely determined by the type of environment it is intended for.*
- **Problem:** There is a vast number of possible kinds of environments in AI.
- **Solution:** Classify along a few “dimensions”. (independent characteristics)
- **Definition 2.34.** For an agent a we classify the environment e of a by its type, which is one of the following. We call e
  1. fully observable, iff the a’s sensors give it access to the complete state of the environment at any point in time, else partially observable.
  2. deterministic, iff the next state of the environment is completely determined by the current state and a’s action, else stochastic.
  3. episodic, iff a’s experience is divided into atomic episodes, where it perceives and then performs a single action. Crucially, the next episode does not depend on previous ones. Non-episodic environments are called sequential.
  4. dynamic, iff the environment can change without an action performed by a, else static. If the environment does not change but a’s performance measure does, we call e semidynamic.
  5. discrete, iff the sets of e’s state and a’s actions are countable, else continuous.
  6. single-agent, iff only a acts on e; else multi-agent(when must we count parts of e as agents?)

Reflex Agents

- **Definition 2.35.** An agent $\langle\mathcal{P},\mathcal{A},f\rangle$ is called a reflex agent, iff it only takes the last percept into account when choosing an action, i .e. $f(p_1,\ldots,p_k)=f(p_k)$ for all $p_1,\ldots,p_k\in\mathcal{P}$.
- **Agent Schema:**

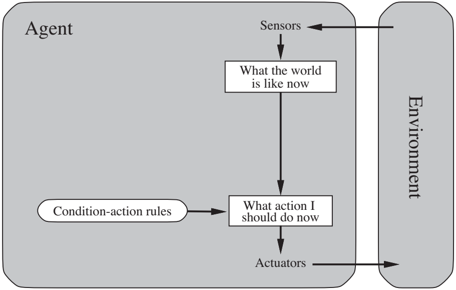

- **Example 2.36 (Agent Program).**

procedure Reflex−Vacuum−Agent [location,status] **returns** an action
  if status = Dirty **then** . . .

Model-based Reflex Agents: Idea

- **Idea:** Keep track of the state of the world we cannot see in an internal model.
- **Agent Schema:**

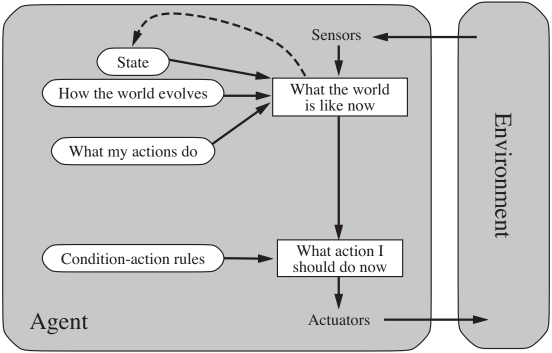

Model-based Reflex Agents: Definition

- **Definition 2.37.** A model-based agent $\langle \mathcal{P}, \mathcal{A}, \mathcal{S}, \mathcal{T}, s_0, S, a \rangle$ is an agent $\langle \mathcal{P}, \mathcal{A}, f \rangle$ whose actions depend on
  1. a world model: a set $\mathcal{S}$ of possible states, and a start state $s_0 \in \mathcal{S}$.
  2. a transition model $\mathcal{T}$, that predicts a new state $\mathcal{T}(s, a)$ from a state s and an action a.
  3. a sensor model S that given a state s and a percept p determine a new state $S(s, p)$.
  4. an action function $a \colon \mathcal{S} \to \mathcal{A}$ that given a state selects the next action.

  If the world model of a model-based agent A is in state s and A has last taken action a, and now perceives p, then A will transition to state $s' = S(p, \mathcal{T}(s, a))$ and take action $a' = a(s')$.
  So, given a sequence $p_1, \ldots, p_n$ of percepts, we recursively define states $s_n = S(\mathcal{T}(s_{n-1}, a(s_{n-1})), p_n)$ with $s_1 = S(s_0, p_1)$. Then $f(p_1, \ldots, p_n) = a(s_n)$.
- **Note:** As different percept sequences lead to different states, so the agent function $f() \colon \mathcal{P}^* \to \mathcal{A}$ no longer depends only on the last percept.
- **Example 2.38 (Tail Lights Again).** Model-based agents can do the ??? if the states include a concept of tail light brightness.

21.2.5.2 Sources of Uncertainty

Sources of Uncertainty in Decision-Making

Where's that d. . . Wumpus?
And where am I, anyway??

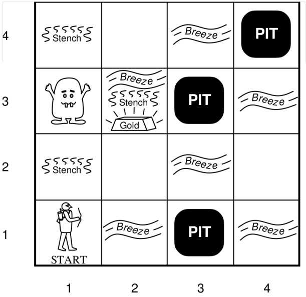

- **Non-deterministic actions:**
  - "When I try to go forward in this dark cave, I might actually go forward-left or forward-right."

Sources of Uncertainty in Decision-Making

Where's that d. . . Wumpus?
And where am I, anyway??

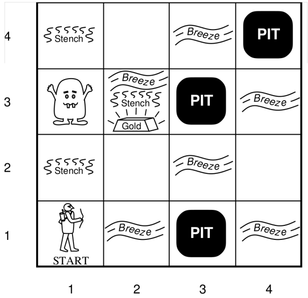

- **Non-deterministic actions:**
  - "When I try to go forward in this dark cave, I might actually go forward-left or forward-right."
- **Partial observability with unreliable sensors:**
  - "Did I feel a breeze right now?";
  - "I think I might smell a Wumpus here, but I got a cold and my nose is blocked."
  - "According to the heat scanner, the Wumpus is probably in cell [2,3]."

Sources of Uncertainty in Decision-Making

Where's that d. . . Wumpus?
And where am I, anyway??

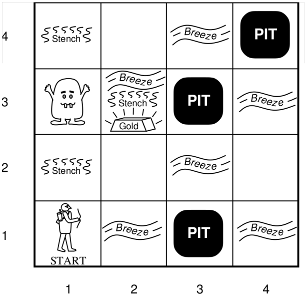

- **Non-deterministic actions:**
  - "When I try to go forward in this dark cave, I might actually go forward-left or forward-right."
- **Partial observability with unreliable sensors:**
  - "Did I feel a breeze right now?";
  - "I think I might smell a Wumpus here, but I got a cold and my nose is blocked."
  - "According to the heat scanner, the Wumpus is probably in cell [2,3]."
- **Uncertainty about the domain behavior:**
  - "Are you *sure* the Wumpus never moves?"

Unreliable Sensors

- **Robot Localization:** Suppose we want to support localization using landmarks to narrow down the area.
- **Example 2.39.** *If you see the Eiffel tower, then you're in Paris.*

Unreliable Sensors

- **Robot Localization:** Suppose we want to support localization using landmarks to narrow down the area.
- **Example 2.40.** *If you see the Eiffel tower, then you're in Paris.*
- **Difficulty:** Sensors can be imprecise.
  - Even if a landmark is perceived, we cannot conclude with certainty that the robot is at that location.
  - *This is the half-scale Las Vegas copy, you dummy.*
  - Even if a landmark is *not* perceived, we cannot conclude with certainty that the robot is *not* at that location.
  - *Top of Eiffel tower hidden in the clouds.*
- Only the probability of being at a location increases or decreases.

21.2.5.3 Agent Architectures based on Belief States

World Models for Uncertainty

- **Problem:** We do not know with certainty what state the world is in!

World Models for Uncertainty

- **Problem:** We do not know with certainty what state the world is in!
- **Idea:** Just keep track of all the possible states it could be in.
- **Definition 2.42.** A model-based agent has a world model consisting of
  - a belief state that has information about the possible states the world may be in,
  - a sensor model that updates the belief state based on sensor information, and
  - a transition model that updates the belief state based on actions.

World Models for Uncertainty

- **Problem:** We do not know with certainty what state the world is in!
- **Idea:** Just keep track of all the possible states it could be in.
- **Definition 2.43.** A model-based agent has a world model consisting of
  - a belief state that has information about the possible states the world may be in,
  - a sensor model that updates the belief state based on sensor information, and
  - a transition model that updates the belief state based on actions.
- **Idea:** The agent environment determines what the world model can be.

World Models for Uncertainty

- **Problem:** We do not know with certainty what state the world is in!
- **Idea:** Just keep track of all the possible states it could be in.
- **Definition 2.44.** A model-based agent has a world model consisting of
  - a belief state that has information about the possible states the world may be in,
  - a sensor model that updates the belief state based on sensor information, and
  - a transition model that updates the belief state based on actions.
- **Idea:** The agent environment determines what the world model can be.
- In a fully observable, deterministic environment,
  - we can observe the initial state and subsequent states are given by the actions alone.
  - Thus the belief state is a singleton (we call its sole member the world state) and the transition model is a function from states and actions to states: a transition function.

World Models by Agent Type in AI-1

- **Search-based Agents:** In a fully observable, deterministic environment
  - goal-based agent with world state $\widehat{=}$ "current state"
  - no inference. (goal $\widehat{=}$ goal state from search problem)
- **CSP-based Agents:** In a fully observable, deterministic environment
  - goal-based agent withworld state $\widehat{=}$ constraint network,
  - inference $\widehat{=}$ constraint propagation. (goal $\widehat{=}$ satisfying assignment)
- **Logic-based Agents:** In a fully observable, deterministic environment
  - model-based agent with world state $\widehat{=}$ logical formula
  - inference $\widehat{=}$ e.g. DPLL or resolution.
- **Planning Agents:** In a fully observable, deterministic, environment
  - goal-based agent with world state $\widehat{=}$ PL0, transition model $\widehat{=}$ STRIPS,
  - inference $\widehat{=}$ state/plan space search. (goal: complete plan/execution)

World Models for Complex Environments

- In a fully observable, but stochastic environment,
  - the belief state must deal with a set of possible states.
  - $\rightsquigarrow$ generalize the transition function to a transition relation.

World Models for Complex Environments

- In a fully observable, but stochastic environment,
  - the belief state must deal with a set of possible states.
  - $\rightsquigarrow$ generalize the transition function to a transition relation.
- **Note:** This even applies to online problem solving, where we can just perceive the state. (e.g. when we want to optimize utility)

World Models for Complex Environments

- In a fully observable, but stochastic environment,
  - the belief state must deal with a set of possible states.
  - $\leadsto$ generalize the transition function to a transition relation.
- **Note:** This even applies to online problem solving, where we can just perceive the state. (e.g. when we want to optimize utility)
- In a deterministic, but partially observable environment,
  - the belief state must deal with a set of possible states.
  - we can use transition functions.
  - We need a sensor model, which predicts the influence of percepts on the belief state – during update.

World Models for Complex Environments

- In a fully observable, but stochastic environment,
  - the belief state must deal with a set of possible states.
  - $\leadsto$ generalize the transition function to a transition relation.
- **Note:** This even applies to online problem solving, where we can just perceive the state. (e.g. when we want to optimize utility)
- In a deterministic, but partially observable environment,
  - the belief state must deal with a set of possible states.
  - we can use transition functions.
  - We need a sensor model, which predicts the influence of percepts on the belief state – during update.
- In a stochastic, partially observable environment,
  - mix the ideas from the last two. (sensor model + transition relation)

Preview: New World Models (Belief) $\rightsquigarrow$ new Agent Types

- **Probabilistic Agents:** In a partially observable environment
  - belief state $\widehat{=}$ Bayesian networks,
  - inference $\widehat{=}$ probabilistic inference.

Preview: New World Models (Belief) $\leadsto$ new Agent Types

- **Probabilistic Agents:** In a partially observable environment
  - belief state $\widehat{=}$ Bayesian networks,
  - inference $\widehat{=}$ probabilistic inference.
- **Decision-Theoretic Agents:** In a partially observable, stochastic environment
  - belief state + transition model $\widehat{=}$ decision networks,
  - inference $\widehat{=}$ maximizing expected utility.
- We will study them in detail this semester.

Overview: AI2

- Basics of probability theory (probability spaces, random variables, conditional probabilities, independence,...)

Overview: AI2

- Basics of probability theory (probability spaces, random variables, conditional probabilities, independence,...)
- Probabilistic reasoning: Computing the *a posteriori* probabilities of events given evidence, causal reasoning (Representing distributions efficiently, Bayesian networks,...)

Overview: AI2

- Basics of probability theory (probability spaces, random variables, conditional probabilities, independence,...)
- Probabilistic reasoning: Computing the *a posteriori* probabilities of events given evidence, causal reasoning (Representing distributions efficiently, Bayesian networks,...)
- Probabilistic Reasoning over time (Markov chains, Hidden Markov models,...)

$\Rightarrow$ We can update our world model episodically based on observations (i.e. sensor data)

Overview: AI2

- Basics of probability theory (probability spaces, random variables, conditional probabilities, independence,...)
- Probabilistic reasoning: Computing the *a posteriori* probabilities of events given evidence, causal reasoning (Representing distributions efficiently, Bayesian networks,...)
- Probabilistic Reasoning over time (Markov chains, Hidden Markov models,...)

⇒ We can update our world model episodically based on observations (i.e. sensor data)

- Decision theory: Making decisions under uncertainty (Preferences, Utilities, Decision networks, Markov Decision Procedures,...)

⇒ We can choose the right action based on our world model and the likely outcomes of our actions

Overview: AI2

- Basics of probability theory (probability spaces, random variables, conditional probabilities, independence,...)
- Probabilistic reasoning: Computing the *a posteriori* probabilities of events given evidence, causal reasoning (Representing distributions efficiently, Bayesian networks,...)
- Probabilistic Reasoning over time (Markov chains, Hidden Markov models,...)

$\Rightarrow$ We can update our world model episodically based on observations (i.e. sensor data)

- Decision theory: Making decisions under uncertainty (Preferences, Utilities, Decision networks, Markov Decision Procedures,...)

$\Rightarrow$ We can choose the right action based on our world model and the likely outcomes of our actions

- Machine learning: Learning from data (Decision Trees, Classifiers, Neural Networks,...)

References I

[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/.