

# Modern AI and Education: The Chatbot Revolution

Michael Hahsler, Clinical Associate Professor of CS

Lyle School of Engineering, SMU

*BAES - Chinese Education Symposium, SMU, November 11, 2025*

SMU®



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

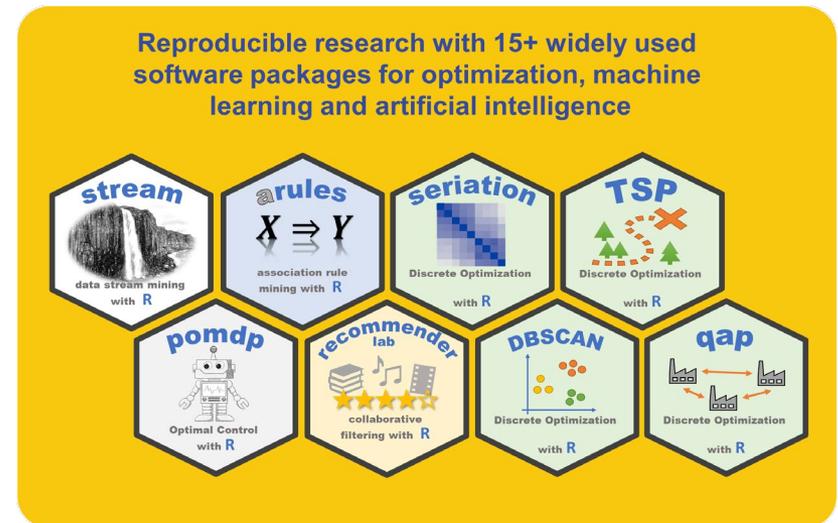


**Dr. Michael Hahsler**  
Department of Computer Science  
Lyle School of Engineering, SMU  
mail: [mhahsler@lyle.smu.edu](mailto:mhahsler@lyle.smu.edu)  
web: [michael.hahsler.net](http://michael.hahsler.net)



Michael Hahsler

- Professor with 25 years of teaching experience.
- Teaching and research areas:
  - Artificial Intelligence
  - Machine Learning
  - Data Mining/Data Science
- 100+ research publications
- Developer and maintainer for 15+ widely used research software packages for R.



Research supported by

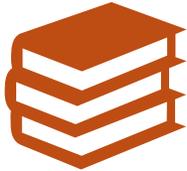


# Agenda

---

This presentation discusses:

- How the technology behind modern AI works.
- How this technology, in the form of chatbots, may revolutionize education.



Applications of AI  
in Education



Modern AI and its  
Goals



Chatbots and  
Education

# myVU: A Next Generation Recommender System Based on Observed Consumer Behavior and Interactive Evolutionary Algorithms

Chapter  
pp 447–457 | [Cite this chapter](#)



Data Analysis

Andreas Geyer-Schulz, Michael Hahsler & Maximillian Jahn

Access th

# Educational and Scientific Recommender Systems: Designing the Information Channels of the Virtual University\*

ANDREAS GEYER-SCHULZ, MICHAEL HAHSLER and MAXIMILLIAN JAHN  
*Department of Information Business, Wirtschaftsuniversität Wien, A-1090 Vienna, Austria. E-mail: Michael.Hahsler@wu-wien.ac.at*

*In this article we investigate the role of recommender systems and their potential in the educational and scientific environment of a Virtual University. The key idea is to use the information aggregation capabilities of a recommender system to improve the tutoring and consulting services of a Virtual University in an automated way and thus scale tutoring and consulting in a personalized way to a mass audience. We describe the recommender services of myVU, the collection of the personalized services of the Virtual University (VU) of the Vienna University of Economics and Business Administration which are based on observed user behavior and self-assignment of experience which are currently field-tested. We show how the usual mechanism design problems inherent to recommender systems are addressed in this prototype.*

# Recommendations for Virtual Universities from Observed User Behavior

Conference paper  
pp 273–280 | [Cite this conference paper](#)



Classification, Automation, and New Media

A. Geyer-Schulz, M. Hahsler & M. Jahn

1997-2004: My first exposure to AI in education was through a project that implemented a prototype of a recommender system for a new concept called a “virtual university.”

# Some Applications of AI in Education

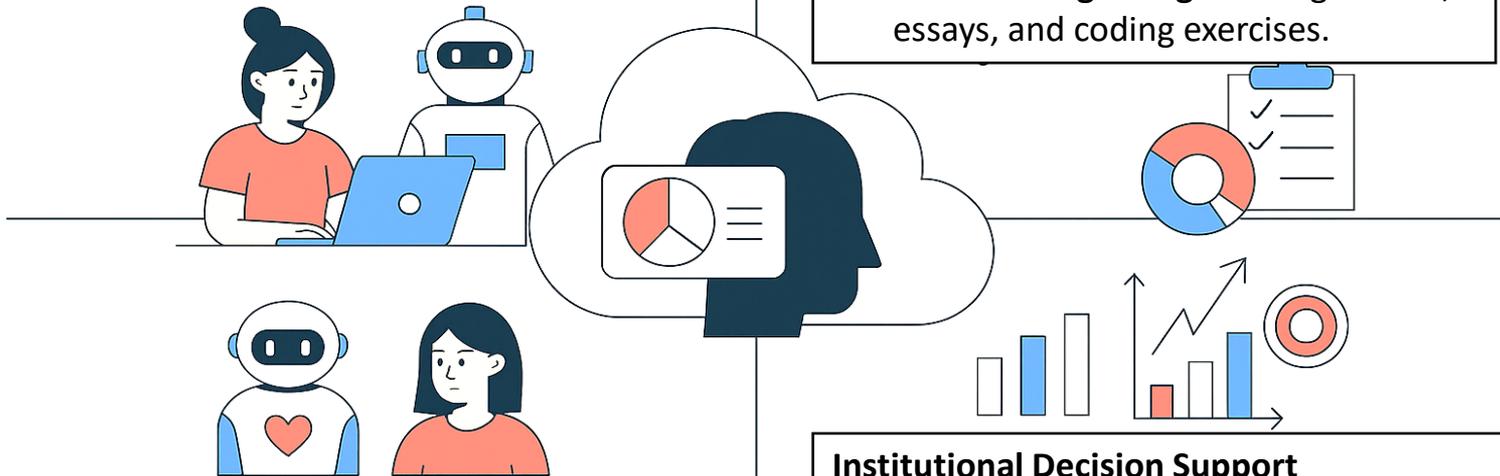
Based on a visualization created with ChatGPT

## Personalized Learning

- **Adaptive learning** platforms/personalized study plans to tailor content to individual students.
- **Intelligent tutoring** systems that provide feedback and hints in real time.

## Assessment and Feedback

- **AI-driven formative assessments** with real-time progress tracking.
- Analytics **identifying at-risk students** or knowledge gaps early.
- **Automated grading** for assignments, essays, and coding exercises.

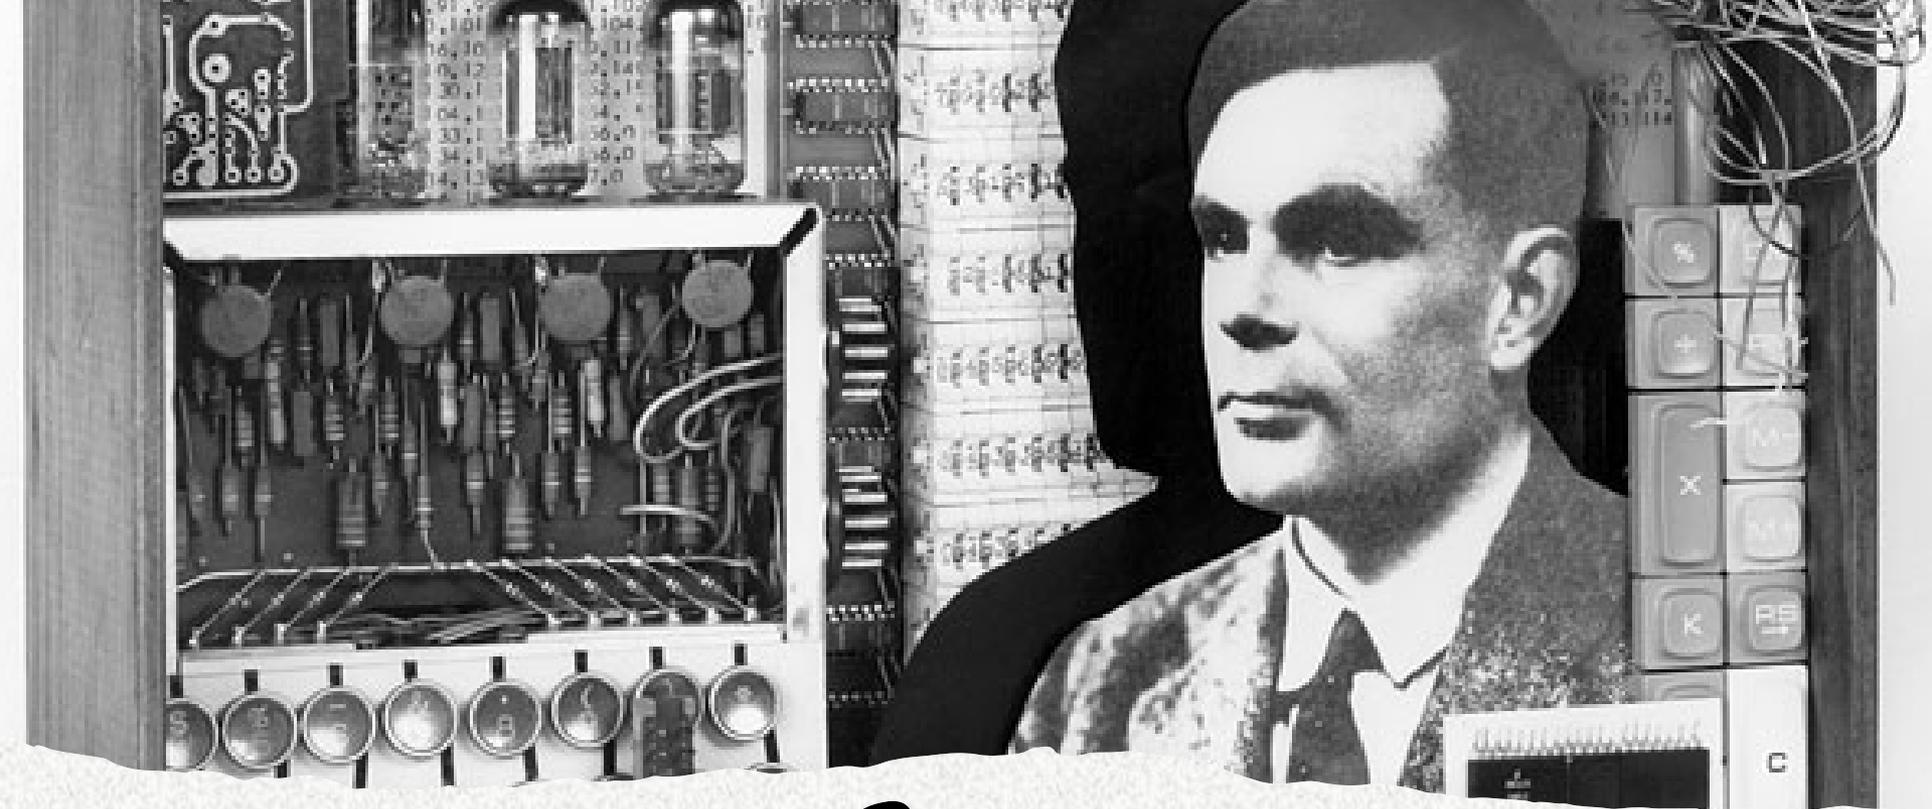


## Student Support and Engagement

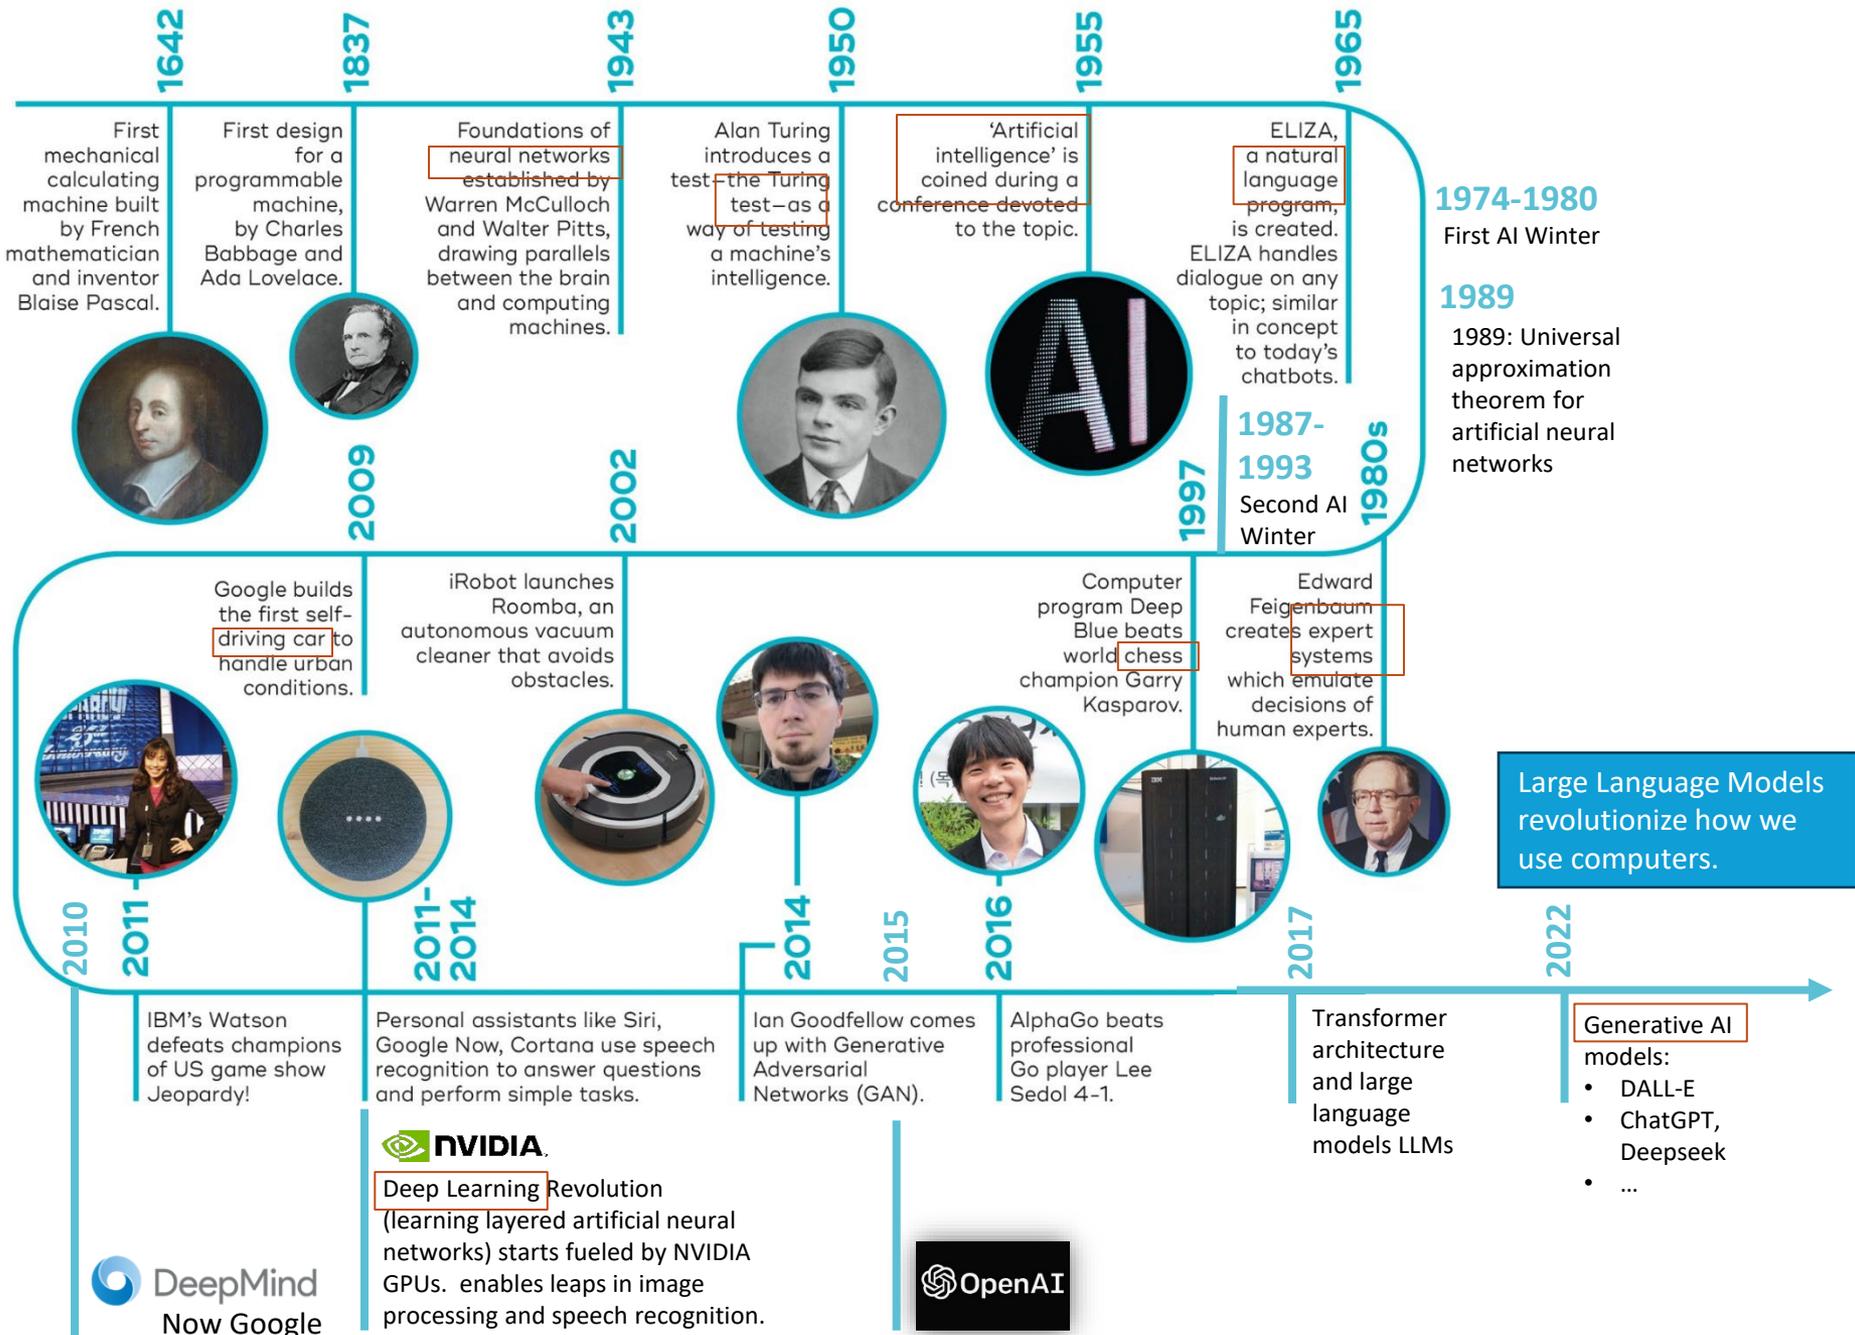
- Chatbots and **virtual assistants** answering academic or administrative queries.
- Sentiment analysis of **student feedback** to improve teaching quality.

## Institutional Decision Support

- Predictive models for **student success**, retention, and enrollment forecasting.
- Data-driven **curriculum optimization** and course demand prediction.
- Recommender systems for **academic advising** and career guidance.



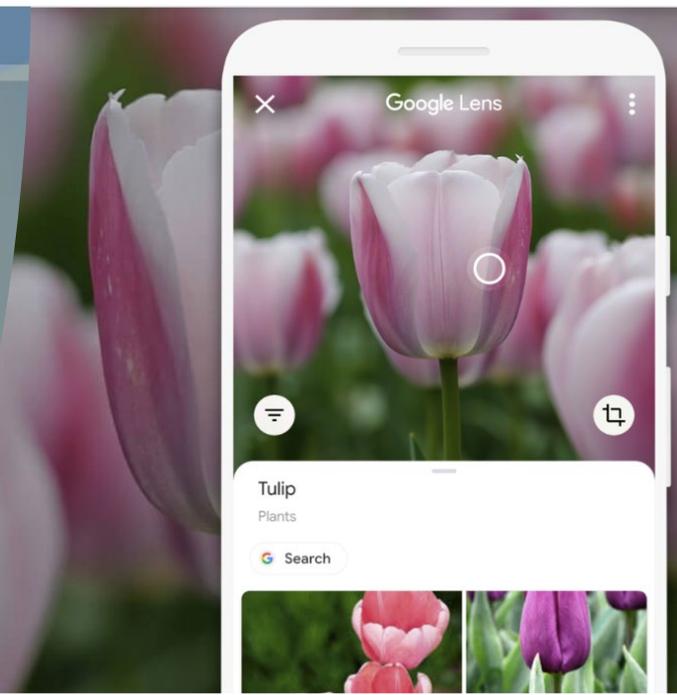
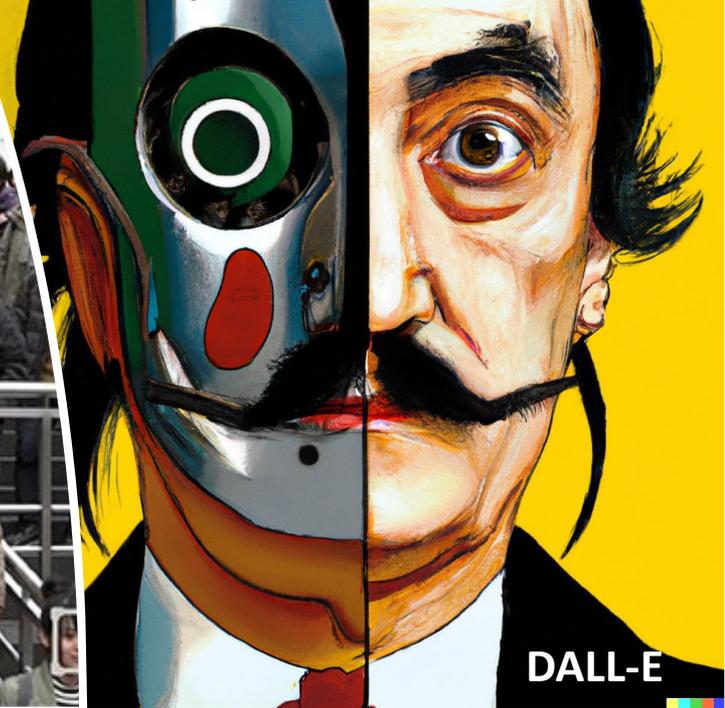
# The History of AI and AI today



# Vision and Image Processing

- **OCR:** read license plates, handwriting recognition (e.g., mail sorting).
- **Face detection:** now standard for smart phone cameras.
- **Vehicle safety systems**
- **Visual search**
- **Image generation**

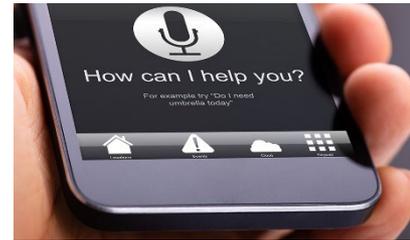
All these technologies operate now at superhuman performance.



# Natural Language Processing



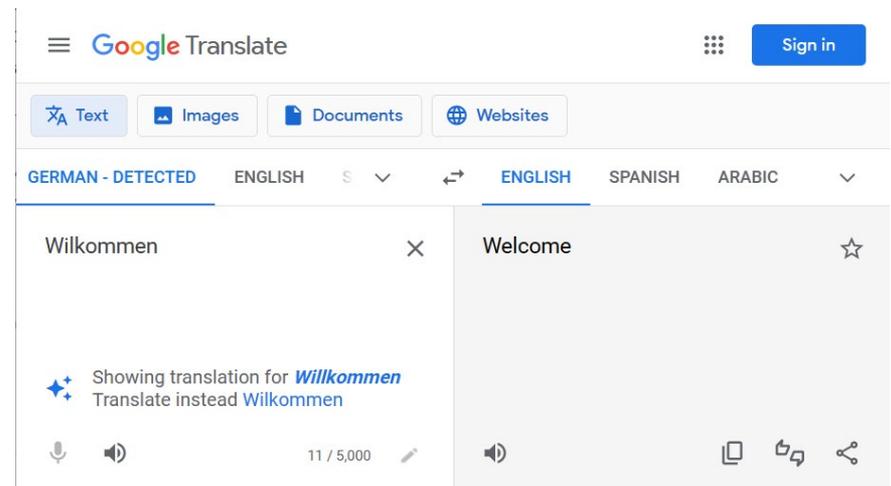
- Text-to-speech
- Speech-to-text to detect voice commands
- Machine translation
- Text generation (Q/A systems) using Large Language Models



**These technologies operate now with close to or even superhuman performance.**

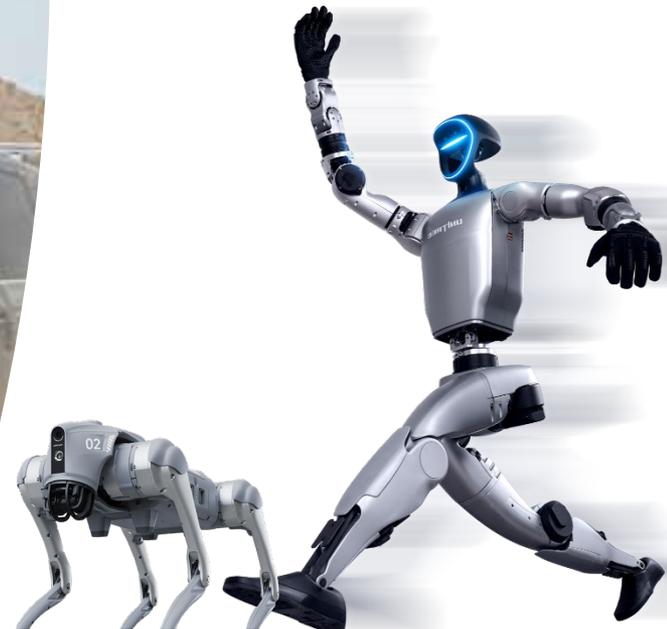
**Humans use language to reason. Does that mean AI that can create high-quality text can reason?**

**Language understanding is still elusive!**



# Robotics

- Mars rovers
- Robot soccer
  - [RoboCup](#)
- Autonomous vehicles
  - [DARPA Grand Challenge](#)
  - Self-driving cars
- Drones
- Personal robotics
  - Humanoid robots
  - [Robotic pets](#)
  - Personal assistants?



# What are the Goals of Artificial Intelligence?



**ASIMO (Advanced Step in Innovative Mobility)** is a humanoid robot created by Honda in 2000

# What is Intelligence?

Definition from Merriam-Webster

## intelligence **noun**

in·tel·li·gence (in-ˈte-lə-jən(t)s)

[Synonyms of intelligence >](#)

**1 a (1)** : the ability to learn or understand or to deal with new or trying situations :

**REASON**

*also* : the skilled use of reason

**(2)** : the ability to apply **knowledge** to manipulate one's environment or to think abstractly as measured by objective criteria (such as tests)

**b** : mental acuteness : **SHREWDNESS**

**c** **Christian Science** : the basic eternal quality of divine Mind

**2 a** : **INFORMATION, NEWS**

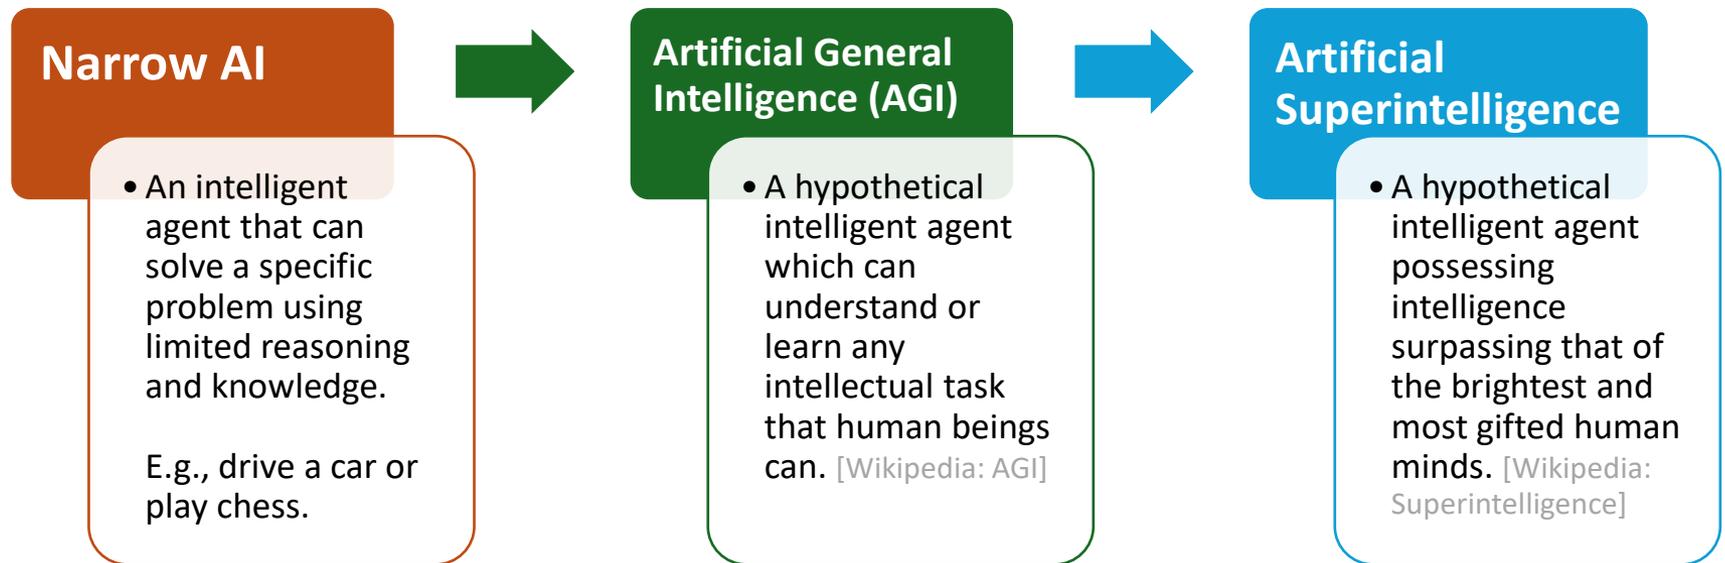
**b** : information concerning an enemy or possible enemy or an area

*also* : an agency engaged in obtaining such information

# The Goal of AI

**“Have machines solve problems that are challenging for humans.”**

We call such a machine an **intelligent agent**.



**How can we achieve this? Create an agent that can:**

Think like a human?

Act like a human?

Think rationally?

Act rationally?

Think like a human?

Act like a human?

Think rationally?

Act rationally?

The brain as an information processing machine.

- Requires scientific theories of **how the brain works**.

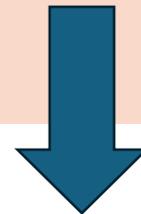
**Note:** The brain does not work like artificial neural networks from machine learning!

How to understand cognition as a computational process?

- Introspection: try to think about how we think.
- **Predict the behavior of human** subjects.
- Image the brain, examine neurological data

AI consciousness

- What does it mean that a machine is conscient/sentient?
- How can we tell?
- What do we do then?



**Cognitive Sciences**

**Moral Philosophy**

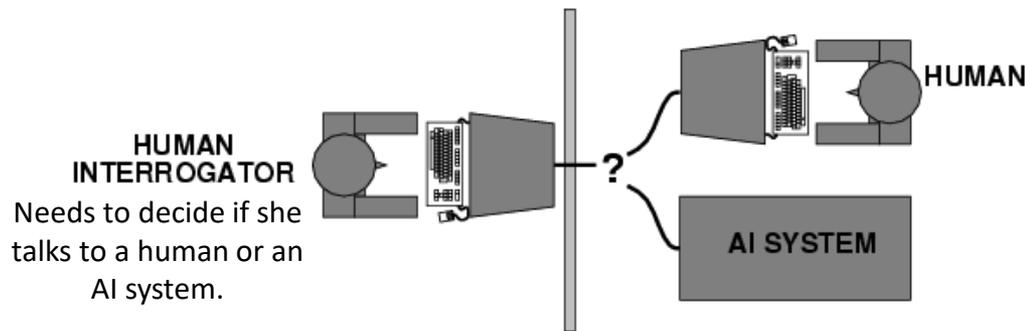
Think like a human?

Act like a human?

Think rationally?

Act rationally?

- Alan Turing rejects the question “Can machines think?”
- The Turing Test tries to define what acting like a human means.



Alan Turing (1950)  
“Computing machinery and intelligence”

- What capabilities would a computer need to have to pass the Turing Test?  
These are still the core AI areas.
  - Natural language processing
  - Knowledge representation
  - Automated reasoning
  - Machine learning
- Turing predicted that by the year 2000, machines would be able to fool 30% of human judges for five minutes.  
Leading Chatbots (2025) much better!

# Turing Test: Criticism

What are some potential problems with the Turing Test?

- Some human behavior is not intelligent.
- Some intelligent behavior may not be human.
- Human observers may be easy to fool.
  - A lot depends on expectations.
  - Anthropomorphic fallacy: humans tend to humanize things.
- Imitate intelligence without intelligence. E.g., the early chatbots ELIZA (1964) simulates a conversation using pattern matching.

Is passing the Turing test a good scientific goal?

- Engineering perspective: Imitating a human is not a good way to solve practical problems.
- We can create useful intelligent agents without trying to imitate humans.

## Chinese Room Argument



Thought experiment by John Searle (1980): Imitate intelligence using rules.

Think like a  
human?

Act like a  
human?

Think  
rationally?

Act  
rationally?

- **Thinking Rationality:** Draw sensible conclusions from facts, logic and data.
- **Logic:** A chain of argument that always yield correct conclusions.  
E.g., "Socrates is a man; all men are mortal; therefore, Socrates is mortal."
- **Logic-based approach to AI:** Describe a problem in formal logic notation and apply general deduction procedures to solve it.  
Issues:
  - Describing real-world problems and knowledge using logic notation is hard.
  - Computational complexity of finding the solution.
  - Much intelligent or "rational" behavior in an uncertain world cannot be defined by simple logic rules.

Example: What about the logical implication

*study hard*  $\Rightarrow$  *A in my AI course*

Should it rather be

*study hard AND be lucky AND ...*  $\Rightarrow$  *A in my AI course*

Think like a human?

Act like a human?

Think rationally?

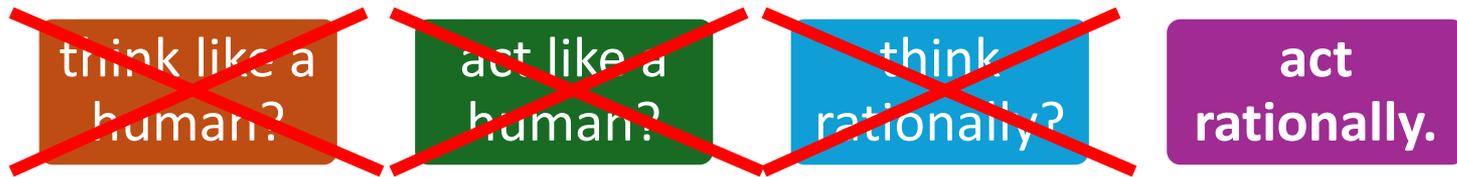
Act rationally?

Acting rational means to try to achieve the “best” outcome.

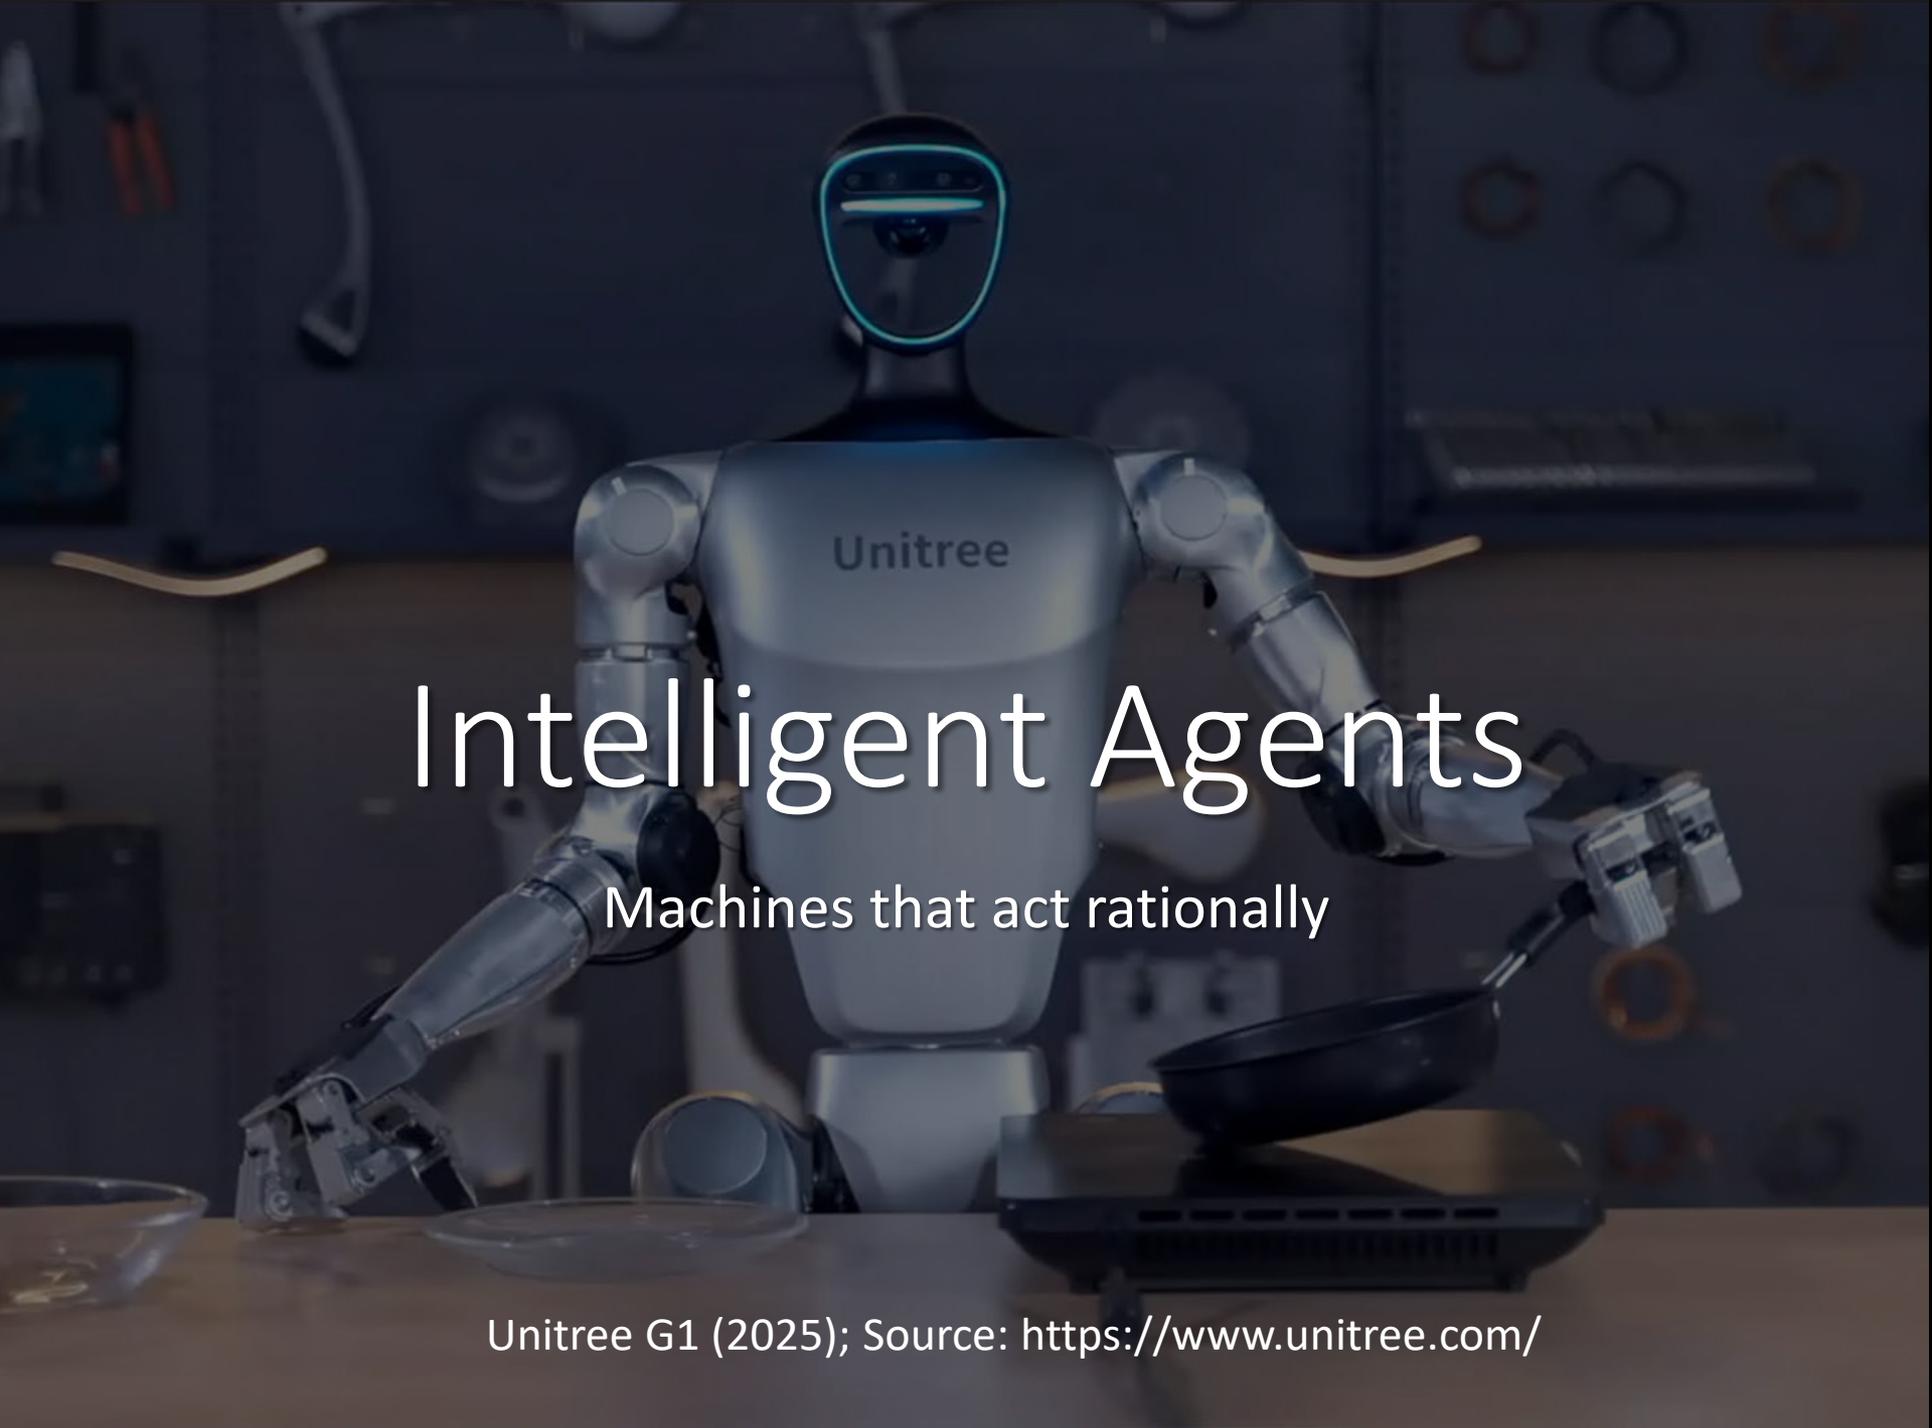
- Best means that we need to do **optimization**.
- The desirability of outcomes can be measured by the economic concept of **utility**. If there is uncertainty about achieving outcomes, then we need to maximizing the **expected utility**.
- Optimization has several advantages:
  - **Generality**: optimization is not limited to logical rules.
  - **Practicality**: can be adapted to many real-world problems.
  - **Well established**: existing solvers and methods for simulation and experimentation.
  - Avoids philosophy and psychology in favor of a **clearly defined objective**.
- **Bounded rationality**: In practice, expected utility optimization is subject to the agent’s knowledge and computational constraints. The agent needs to do the best with its knowledge and resources.

# Current AI

Current techniques can be used to create **narrow AI agent** that uses limited knowledge and reasoning capabilities and can



That is, create a machines that acts in a way to solve a specific hard problem that traditionally would have been thought to require human intelligence.

A humanoid robot with a white and grey body and a glowing blue visor. The robot is standing in a kitchen, holding a black frying pan over a stove. The background shows a kitchen counter with various items and a wall with circular patterns.

# Intelligent Agents

Machines that act rationally

Unitree G1 (2025); Source: <https://www.unitree.com/>

# Components of an Intelligent Agent

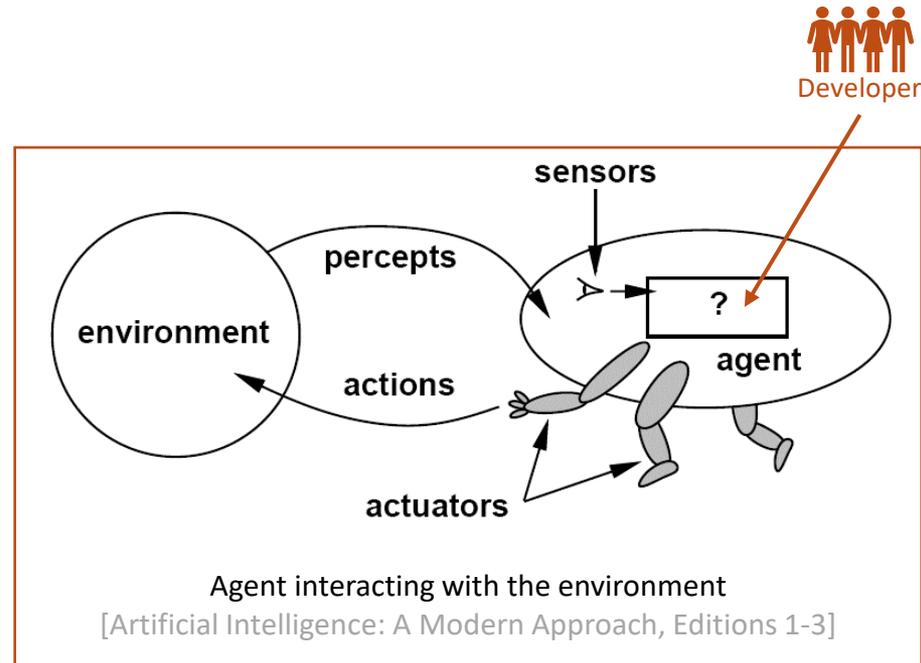
Intelligent agents **act rationally** in their environment.

They need to

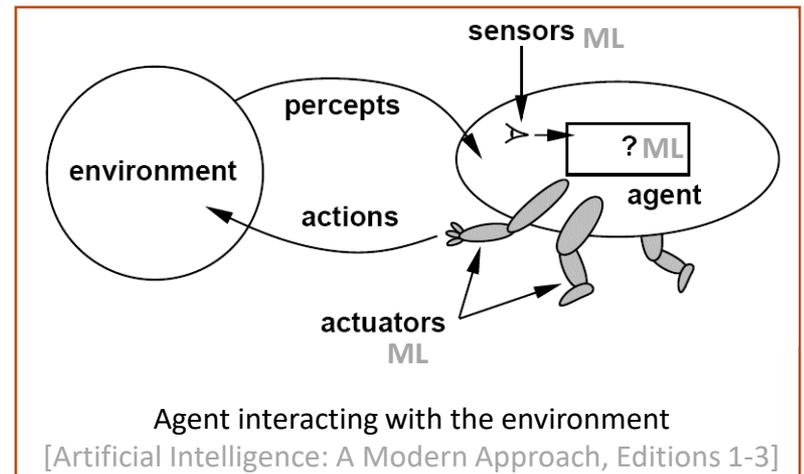
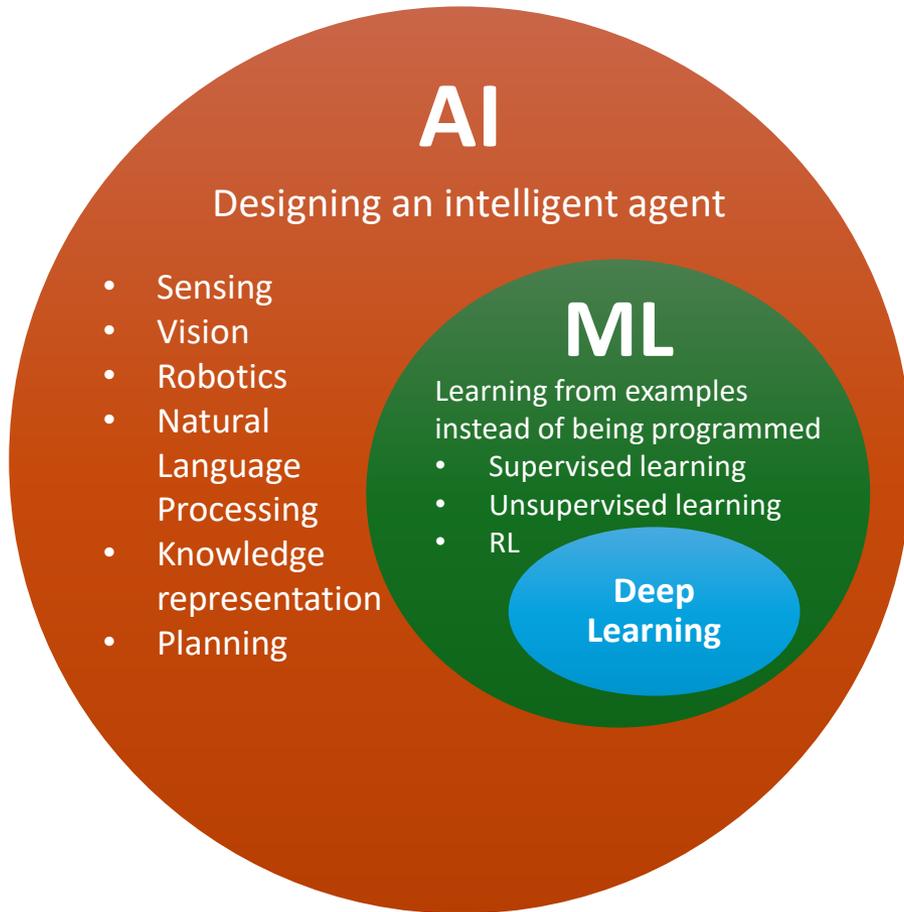
- **Communicate** with the environment using **percepts** and **actions**.
- **Represent knowledge, reason** and **plan** to achieve a desired **outcome**.

Optional

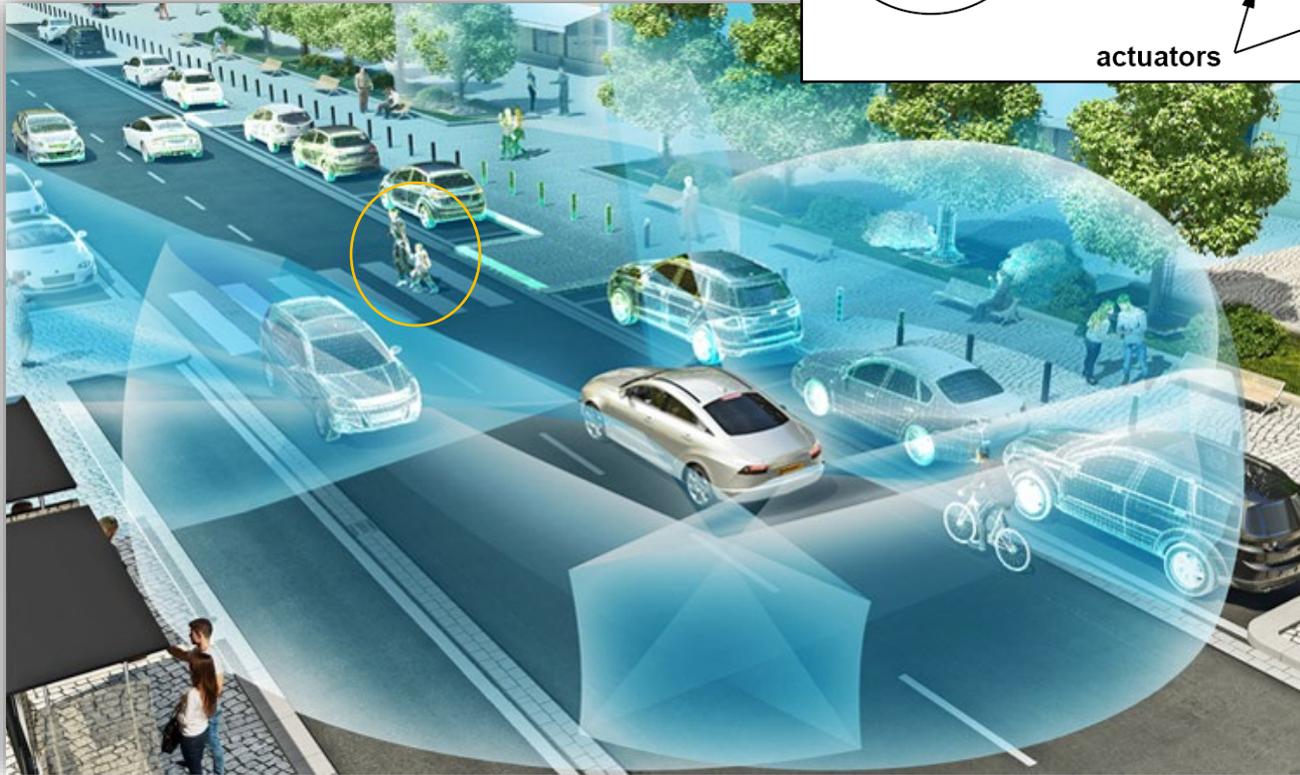
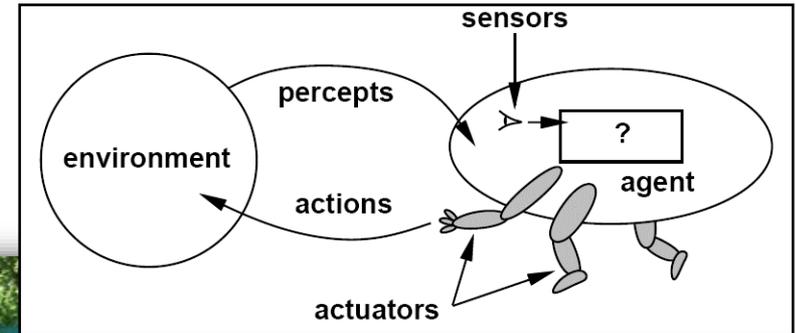
- **Learn from experience** to improve performance.



# Artificial Intelligence vs. Machine Learning



# Example: Self-Driving Car

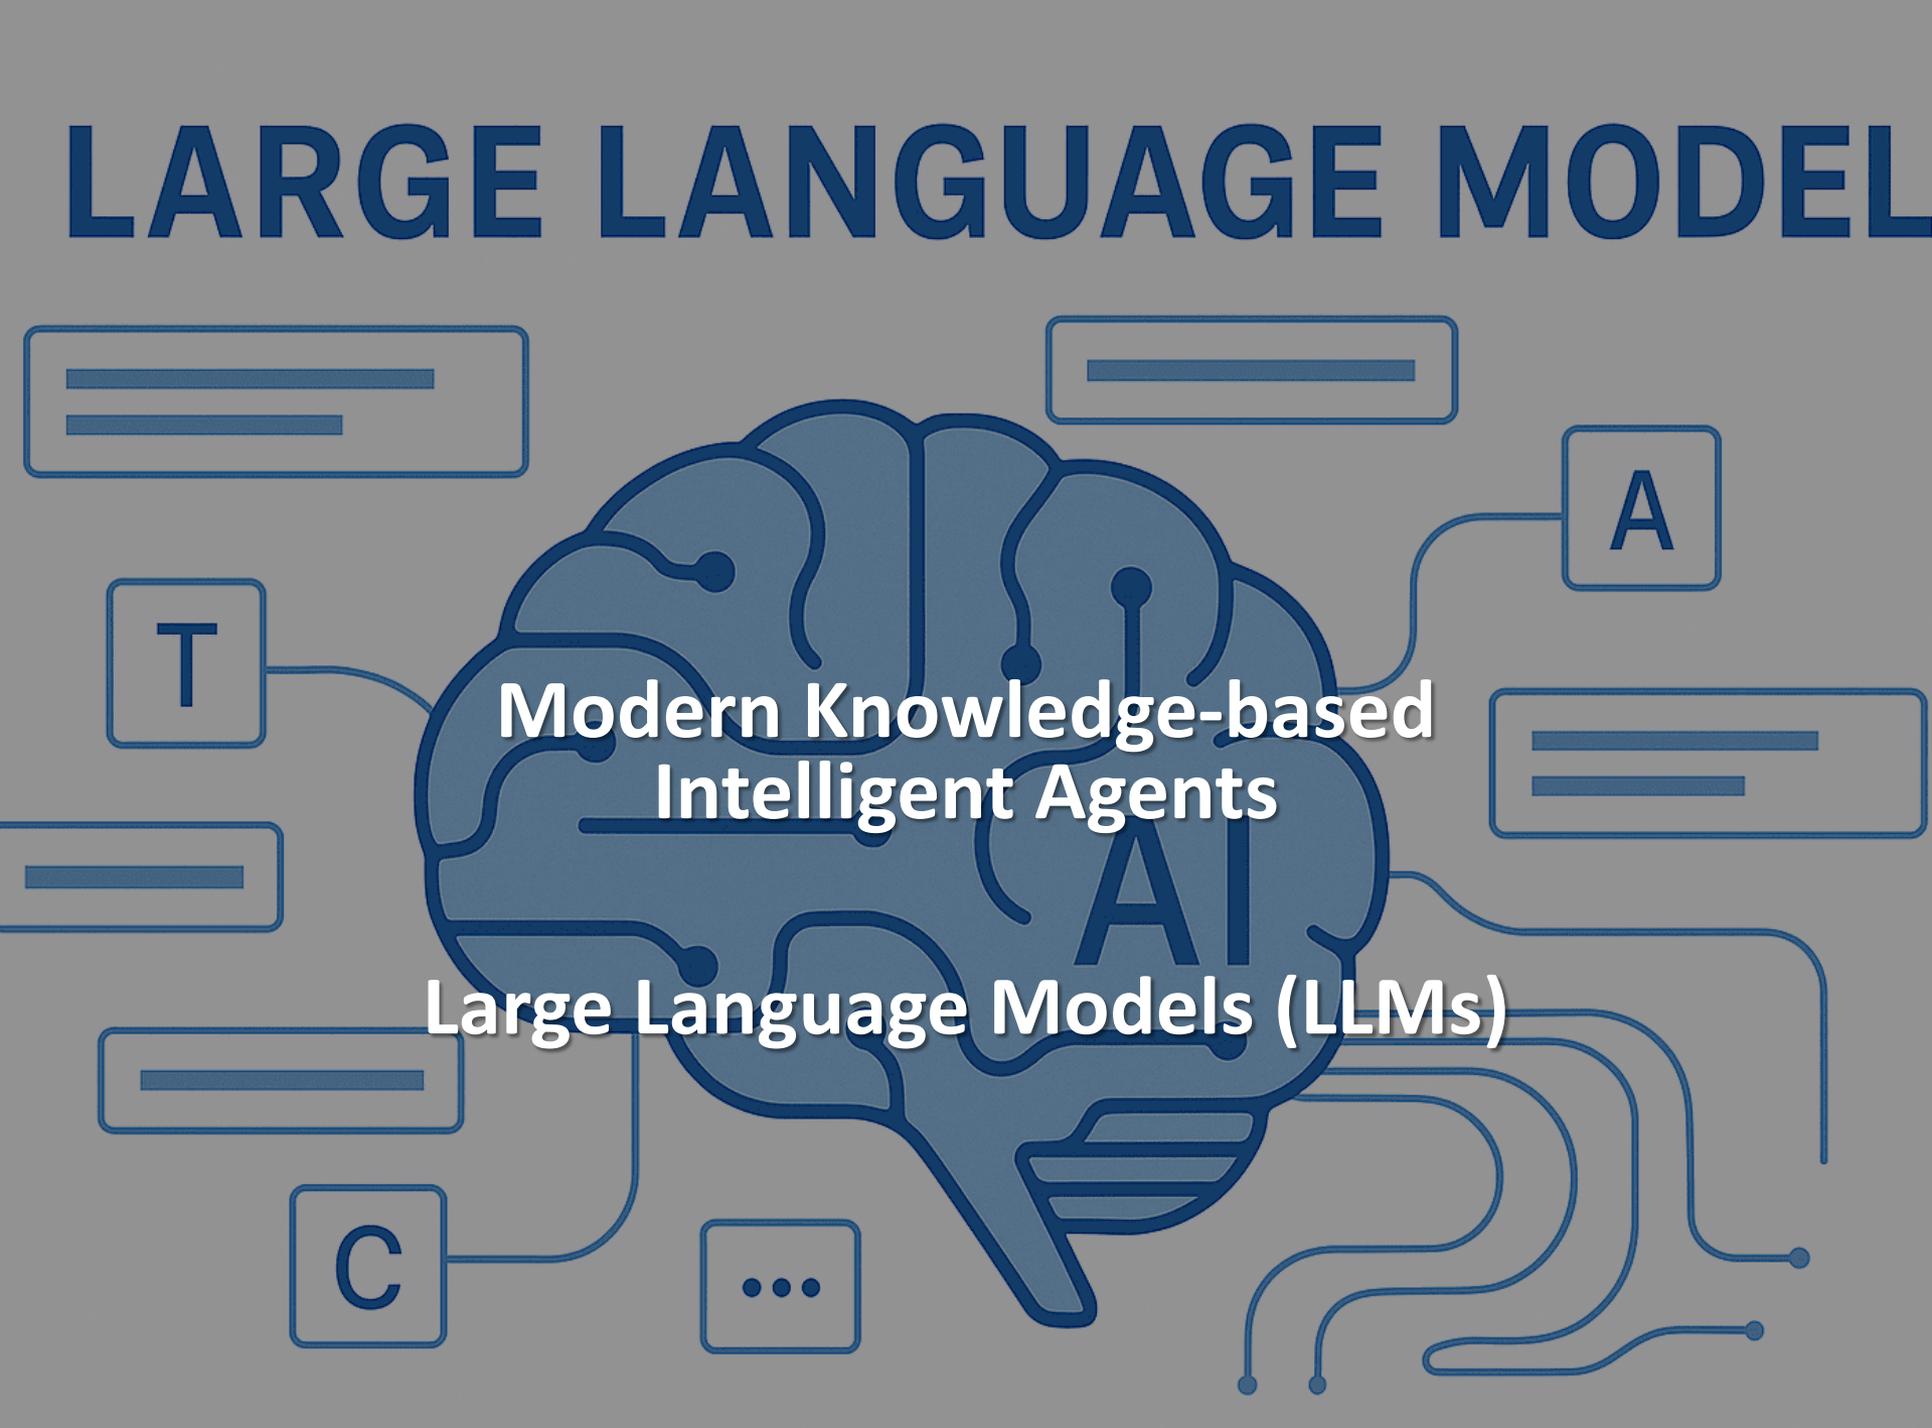


Percepts: Other objects including people crossing the street.

Action: Slow down or stop the car.

Objective: Reach the destination safely

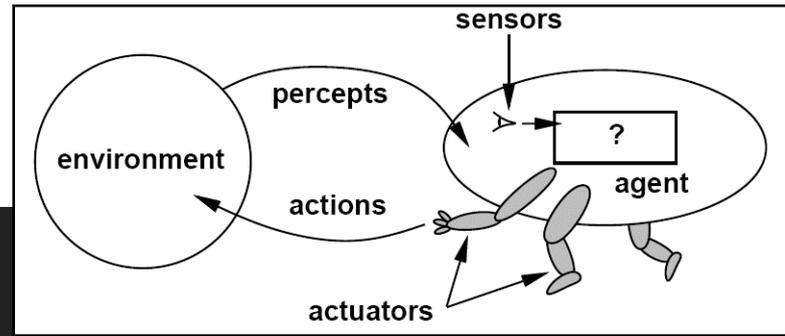
# LARGE LANGUAGE MODEL



Modern Knowledge-based  
Intelligent Agents

Large Language Models (LLMs)

# Example: Homework and LLMs



ChatGPT ▾

Please write a 500-word essay about Sokrates' life.



Here is a 500-word essay on Sokrates' life:

## **The** Life and Legacy of Sokrates

Sokrates, one of history's most influential philosophers, lived in Athens, Greece, from 470/469 BCE to 399 BCE. His contributions to Western philosophy laid the groundwork for ethics, epistemology, and logic. Despite never writing down his teachings, Sokrates' ideas have endured through the works of his students, most notably Plato.

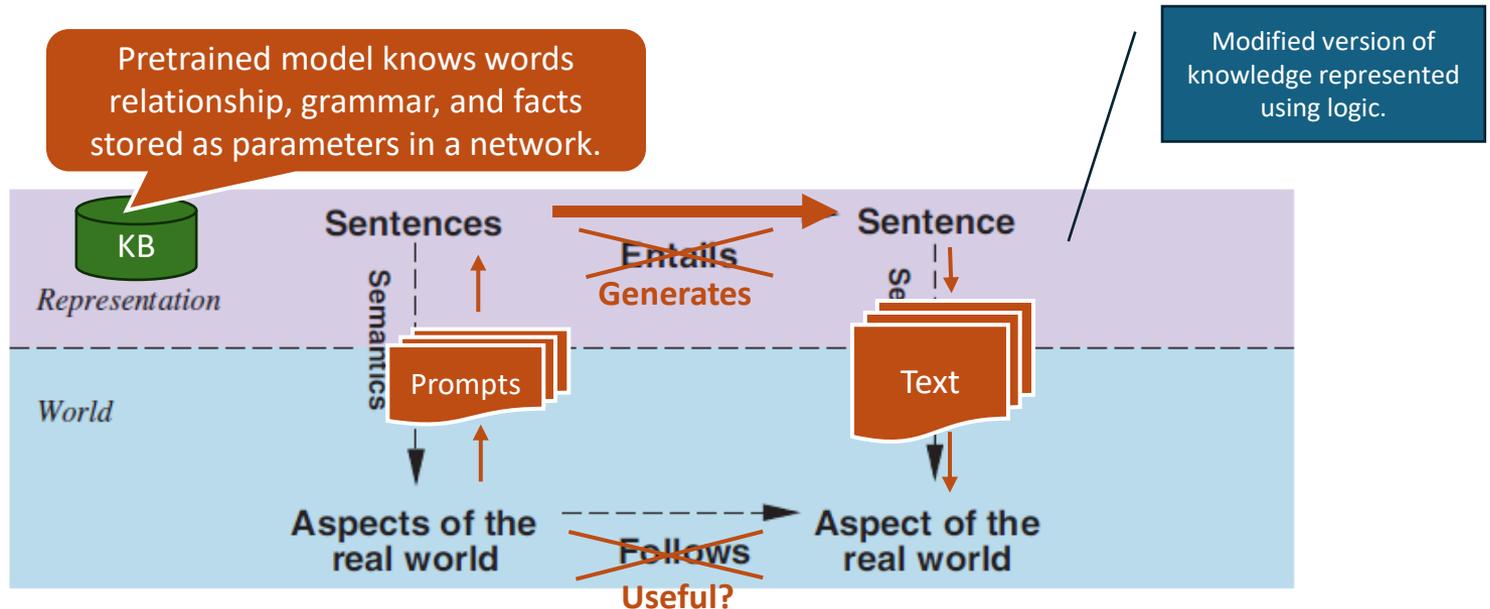
This knowledge has to come from somewhere!

Percept: Your prompt

Action: Next most likely word... More words are created word-by-word.

Objective: You may like a useful answer, but what is ChatGPT's objective?

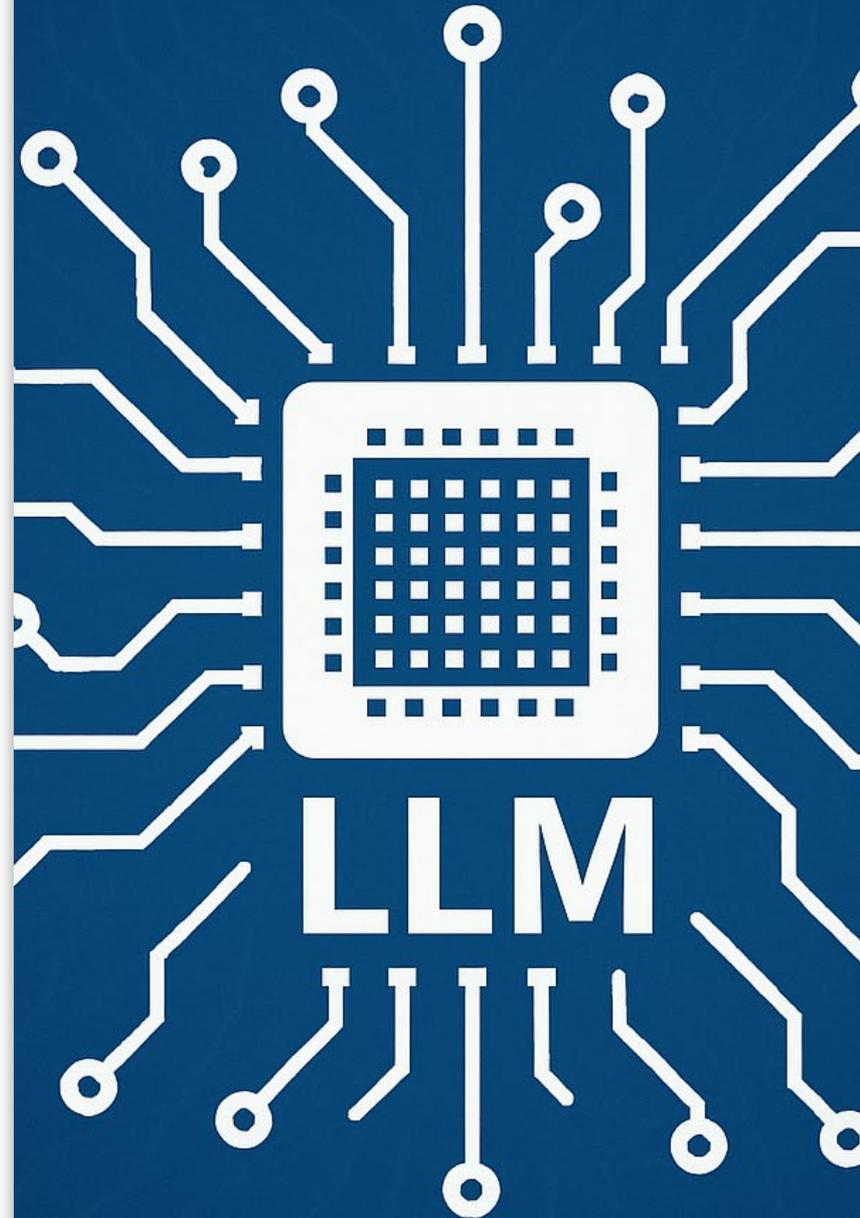
# Using Natural Language for Knowledge Representation (LLMs)

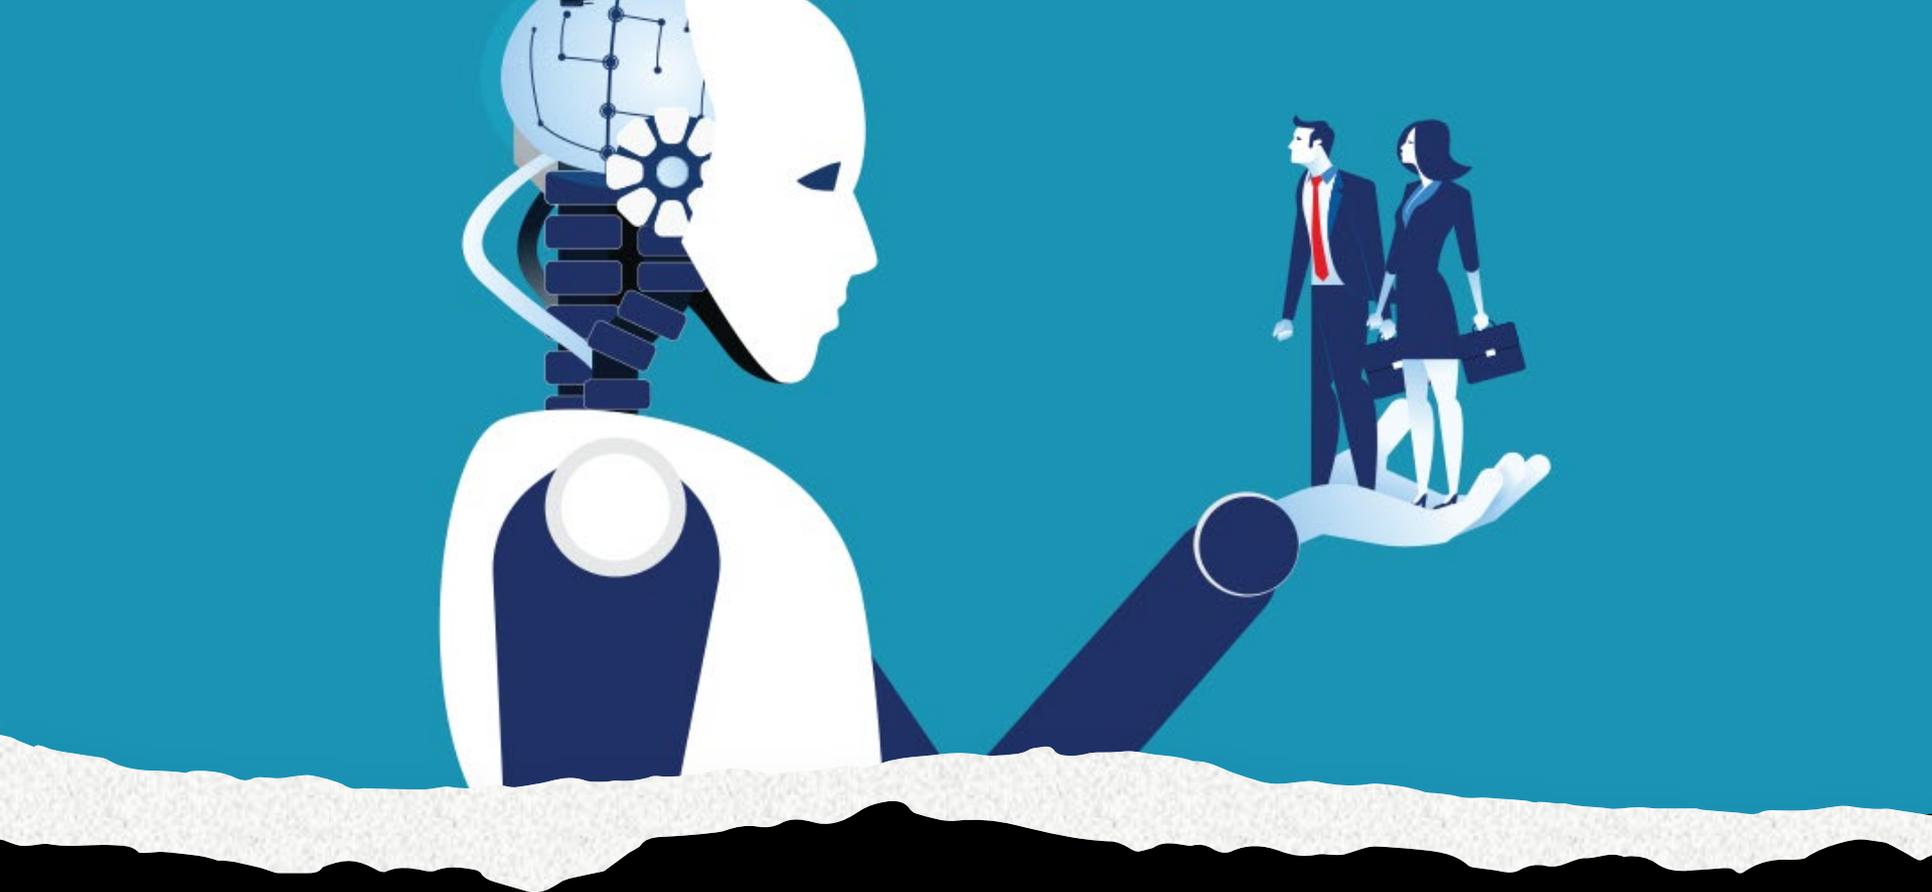


- The user formulates a question about the real world as a natural language **prompt** (a sequence of tokens).
- The LLM **generates text** using a pretrained model that represents its knowledge base.
- The text (hopefully) is useful in the real world. The **objective function** is not clear. Maybe it is implied in the prompt?

# Some Issues with LLMs

- **Hallucinations**; the LLM does not say when it is not sure. Retrieval Augmented Generation (RAG) can help.
- **Privacy and data security** when a third party hosts the LLM. Training data can be poisoned.
- **Safety**: It is easy to avoid guardrails (i.e., jailbreak) leading to unsafe responses.
- LLMs currently have no (or very little) **common sense**.
- LLMs' performance is influenced by small prompt changes. The results are not very **dependable**.
- **Reasoning** capabilities are very limited. Agentic AI tries to mitigate this issue by focusing on smaller tasks.





# AI Ethics & Safety

A new Frontier for  
Fairness and Freedom  
AIMA Chapter 27

# Commonly-Cited Safety and Ethics Principles

## Use of AI by companies and organizations

- Ensure safety
- Limit harmful uses of AI
- Establish accountability: Liability?
- Avoid concentration of power: Winner-takes-All

## Protect individuals

- Uphold human rights and values.
- Ensure fairness: Equal opportunity/equal impact. Reflect diversity/inclusion
- Provide transparency: Explanations to build trust
- Respect privacy: Surveillance?
- Contemplate implications for employment: Income and purpose.

## Governance

- Acknowledge legal/policy implications

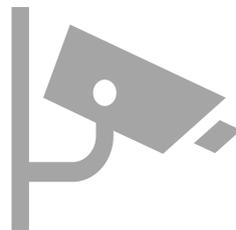
Next, we look at the implementation of these principles in different countries.

# Types of AI Safety

“Prevent accidents, misuse, or other harmful consequences of AI.”



AI Testing



Monitoring AI



Adversarial  
robustness

How should this be ensured?

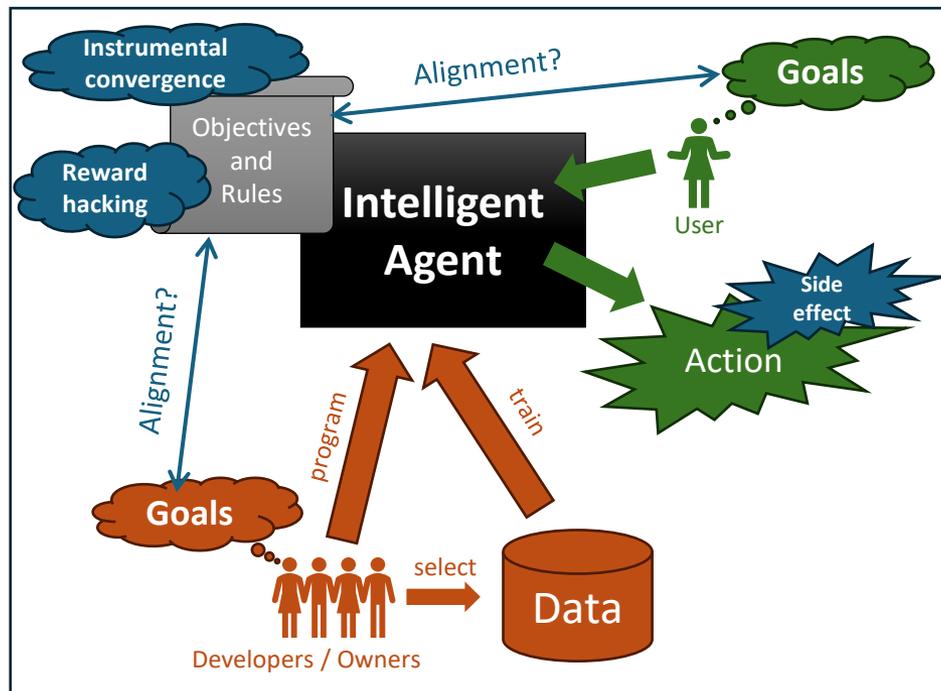
- Corporate self-regulation
- Private watchdogs
- Government action
- International treaties

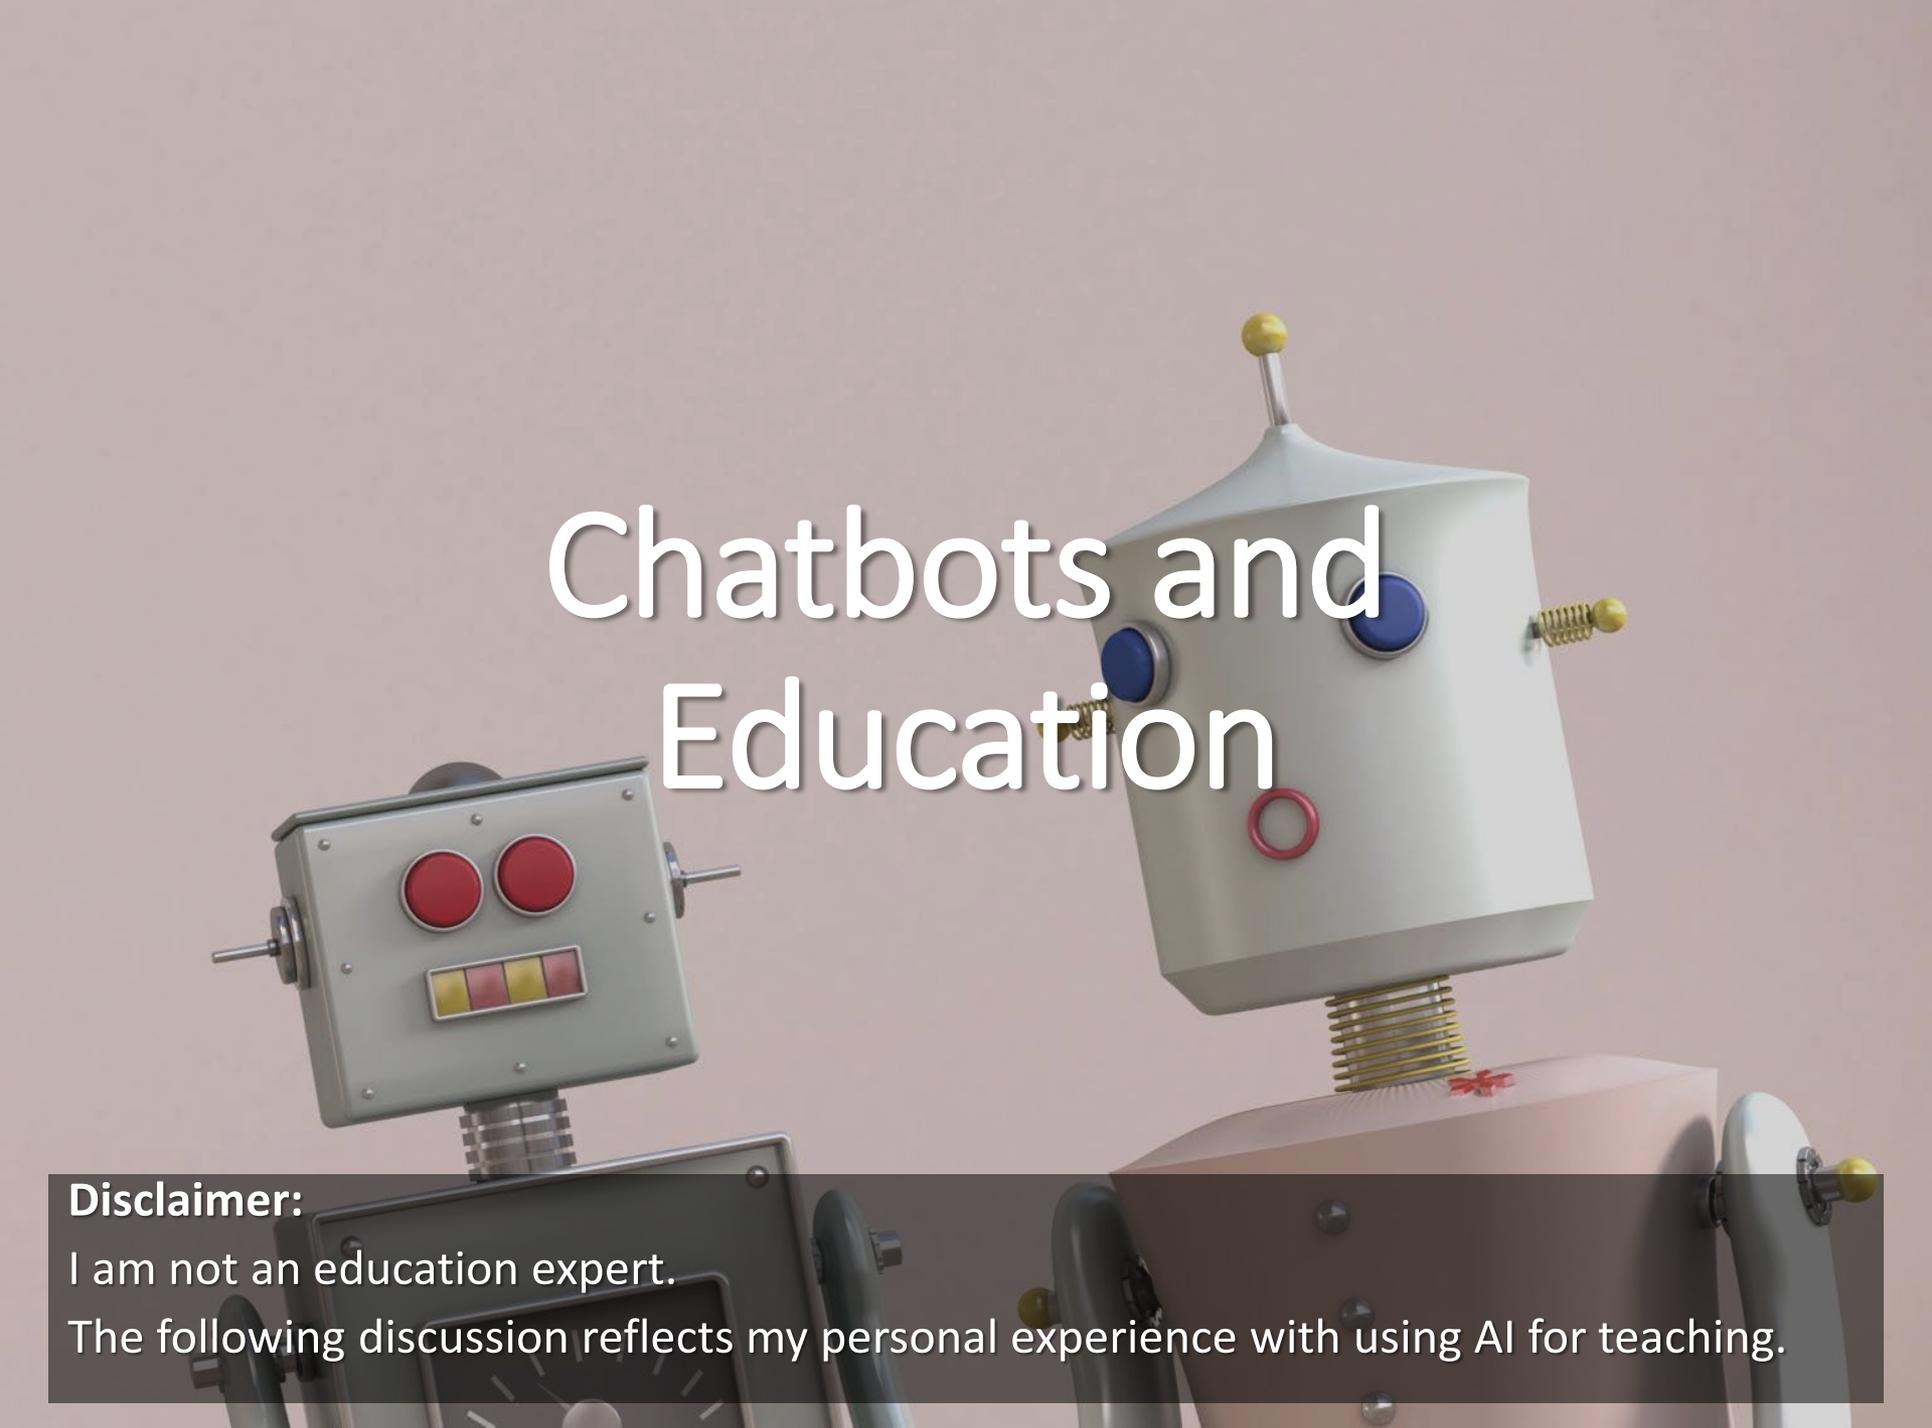
Is it reasonable to assume that a superintelligent AI can get around being tested and monitored?

# AI Safety and Optimizers

Intelligent Agents are “optimizers!”

- **Goal/reward alignment:** How do we specify a robust objective function? Whose objectives are used?
- **Reward hacking:** The AI learns to exploit unintended side effects to get a high “score” without solving the objective. AI needs to follow social norms.
- **Instrumental convergence:** All intelligent agents will pursue common subgoals like the need for more power to get better at reaching its objectives. How will this need be balanced with human’s needs?



The background features two stylized, retro-style robots. On the left is a smaller, grey robot with two red circular eyes and a yellow and red rectangular mouth. On the right is a larger, white robot with two blue circular eyes, a red ring-shaped mouth, and a yellow antenna on top. Both robots have a spring-like neck. The title text is centered over the white robot.

# Chatbots and Education

**Disclaimer:**

I am not an education expert.

The following discussion reflects my personal experience with using AI for teaching.

# Some Applications of AI in Education

## Personalized Learning

- **Adaptive learning** platforms/personalized study plans to tailor content to individual students.
- **Intelligent tutoring** systems that provide feedback and hints in real time.

## Assessment and Feedback

- **AI-driven formative assessments** with real-time progress tracking.
- Analytics **identifying at-risk students** or knowledge gaps early.
- **Automated grading** for assignments, essays, and coding exercises.



## Student Support and Engagement

- Chatbots and **virtual assistants** answering academic or administrative queries.
- **Sentiment analysis of student feedback** to improve teaching quality.

## Institutional Decision Support

- Predictive models for **student success**, retention, and enrollment forecasting.
- Data-driven **curriculum optimization** and course demand prediction.
- **Recommender systems** for **academic advising** and career guidance.

Large language models and chatbots are very capable for many applications.

# Intelligent Tutoring in my AI Course

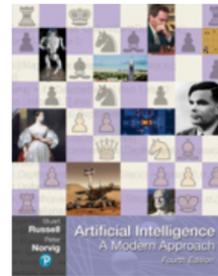
README CC-BY-SA-4.0 license

<https://github.com/mhahsler/CS7320-AI>

## Artificial Intelligence - Lecture Material, Simple Python Code Examples and Assignments

This repository contains lecture material, simple Python code examples, and assignments for the course CS 5/7320 Artificial Intelligence taught by [Michael Hahsler](#) at the [Department of Computer Science at SMU](#).

The code examples cover several chapters of the textbook [Artificial Intelligence: A Modern Approach \(AIMA\)](#) by Russell and Norvig. The code in this repository is intended to be simple to focus more on the basic AI concepts and less on the use of advanced implementation techniques (e.g., object-oriented design and flexibility). More complex code examples accompanying the textbook can be found at the [GitHub repository aimacode](#).



[Ask the AIMA Scholar \(GPT\)](#) a question about the content of the textbook.

### Use cases:

- Ask questions
- Discuss topics
- Help with coding
- Give feedback for assignments
- Quiz the student for studying

A screenshot of the AIMA Scholar chatbot interface. At the top, it says 'AIMA Scholar 4.0' and 'Share'. Below that, a user prompt says 'please quiz me on concepts of AI'. The chatbot responds with a message: 'Great — here's a quiz to test your understanding of core AI concepts, based on *Artificial Intelligence: A Modern Approach*. Questions range from basic to advanced. Let me know if you'd like a specific topic (e.g., search, logic, learning, planning, etc.)'. Below the message, there is a section titled 'AI Concepts Quiz (Mixed Topics)' with a list of topics. The first topic is '1. Agent Architectures' and the first question is 'Q: What is the key difference between a model-based reflex agent and a simple'.

AIMA Scholar 4.0 Share

please quiz me on concepts of AI

Great — here's a quiz to test your understanding of core AI concepts, based on *Artificial Intelligence: A Modern Approach*. Questions range from basic to advanced. Let me know if you'd like a specific topic (e.g., search, logic, learning, planning, etc.).

**AI Concepts Quiz (Mixed Topics)**

1. Agent Architectures

Q: What is the key difference between a model-based reflex agent and a simple

# Using Intelligent Tutoring Systems to Cheat

[Find a journal](#)

[Publish with us](#)

[Track your research](#)

[Search](#)

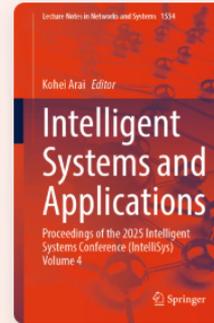
[Cart](#)

[Home](#) > [Intelligent Systems and Applications](#) > Conference paper

## AI Content Self-Detection for Transformer-Based Large Language Models

Conference paper | First Online: 20 August 2025

pp 147–168 | [Cite this conference paper](#)



[Intelligent Systems and Applications](#)  
(IntelliSys 2025)

[Antônio Junior Alves Caiado](#)  & [Michael Hahsler](#)

[Access this chapter](#)

**Abstract:** The usage of generative artificial intelligence (AI) tools based on large language models, including ChatGPT, Bard, and Claude, for text generation has many exciting applications with the potential for phenomenal productivity gains. One issue is authorship attribution when using AI tools. This is especially important in an academic setting where the inappropriate use **of generative AI tools may hinder student learning or stifle research by creating a large amount of automatically generated derivative work. Existing plagiarism detection systems can trace the source of submitted text but are not yet equipped with methods to accurately detect AI-generated text...**

# Other Teaching Related Uses We Experiment with

- **Analysis of student feedback for large lectures**
  - Using a Chatbot works very well.
- **Automated grading and feedback generation for project reports.**
  - My teaching assistant experiments with different language models.
  - The grading performance is currently not up to standard and depends heavily on the model and used prompts.
  - Automatically generated feedback is often too general.
  - Chatbots are improving.

# A Recommender System Architecture for University Curriculum Advising

**Zerui Ma**

Southern Methodist University

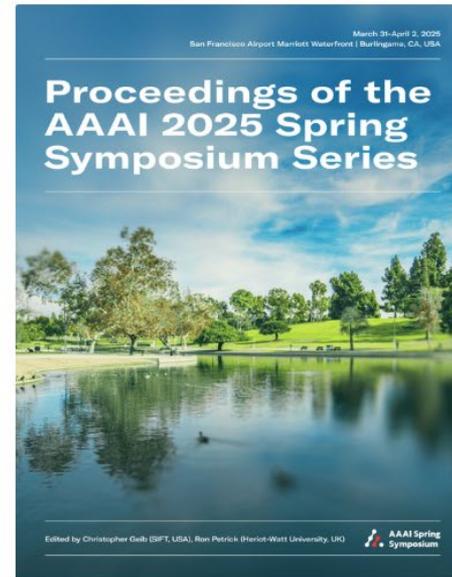
**Michael Hahsler**

Southern Methodist University

**Peter Moore**

Southern Methodist University

DOI: <https://doi.org/10.1609/aaais.v5i1.35593>



**Abstract:** Effective academic advising plays a crucial role in student success, yet universities face challenges in optimizing advising processes and course enrollment. This task is complicated by the fact that several graduation requirements have to be met while also taking the students' interests into account. Academic advising has historically been performed by a skilled human adviser. **Universities can optimize course planning and help students make informed decisions about their academic path with recommender systems.** This case study develops a goal-based agent recommender system based on a large language model tailored to undergraduate students, depending on curriculum requirements, prerequisite dependencies, and student preferences...

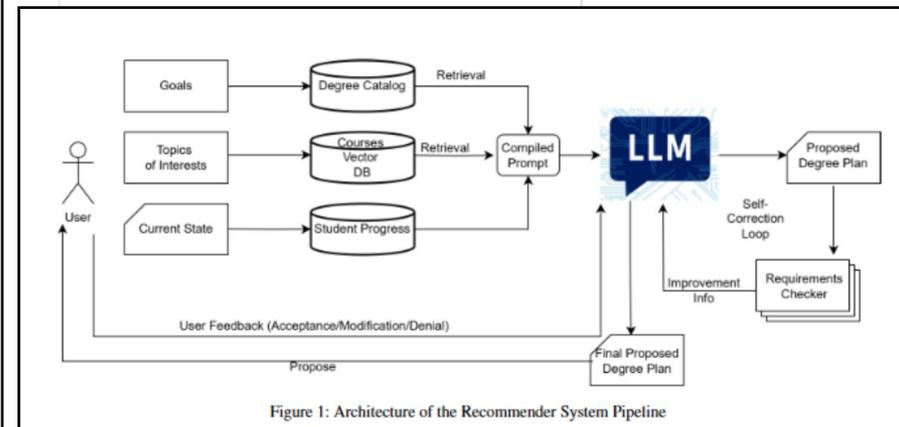


Figure 1: Architecture of the Recommender System Pipeline

# Use of Chatbots

## Access and student attitude

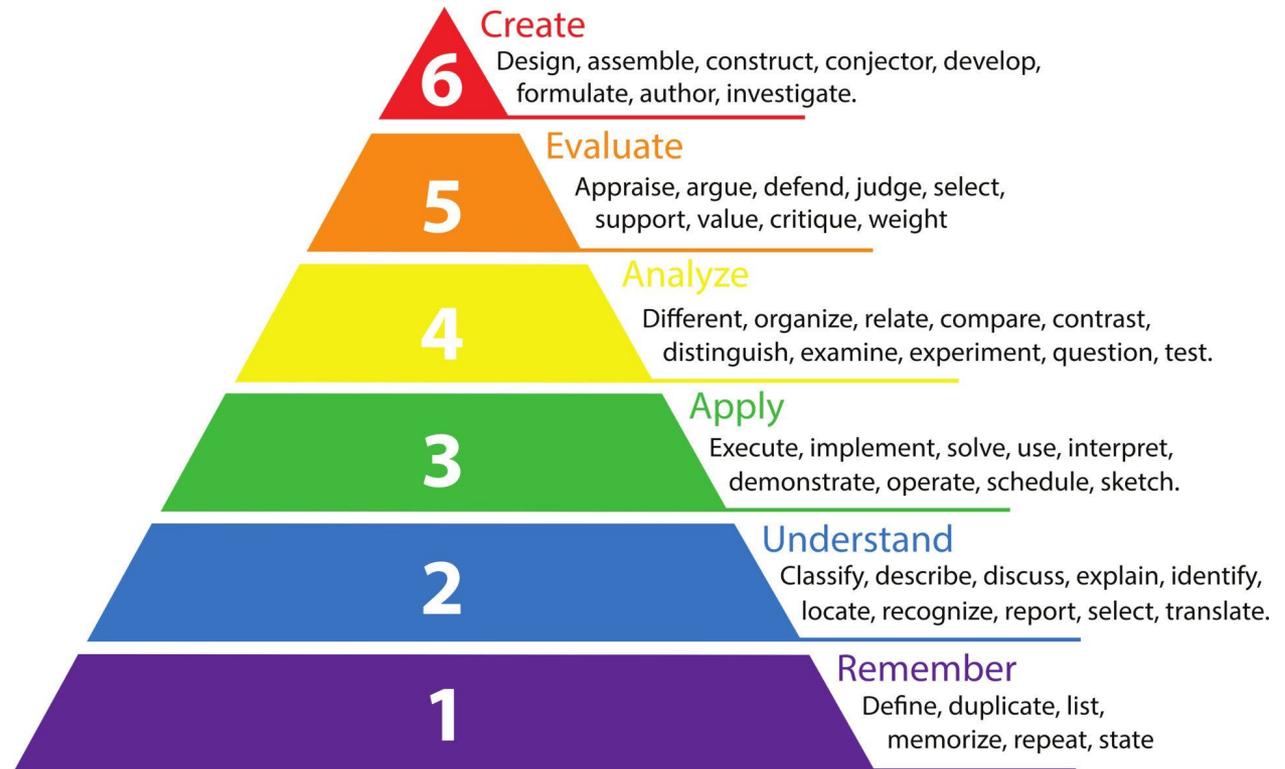
- Chatbots are easily accessible as a self-service tool. Companies give free/discounted access to students.
- Students argue that using chatbots prepares them for work.

## Options:

- a) Restricting the use** of chatbots for class work:
  - Not practical, access is easy, and use is hard to detect.
  - removes potential benefits AI use.
- b) Designing assessments** where chatbots cannot do the work for the student: Large projects with milestones, group assignments, presentations, ...
  - Does not scale to large courses.
  - LLMs are getting better and already can answer questions that advanced students cannot answer.
- c) Adapt learning outcomes:** What does a student need to know and be able to do to be successful with AI.  
**Use AI for formative assessments but go back to oral and paper exams without access to technology for summative assessments.**

# Bloom's Taxonomy

A framework for categorizing learning objectives based on cognitive complexity.



# Effect of Chatbot Use on Learning\*

Bloom's Level	How Chatbots Support Learning	How Chatbots Hinder Learning
<b>1. Remembering</b> (Recall facts, define, list)	<ul style="list-style-type: none"><li>• Provide quick factual recall and <b>quizzes</b>.</li><li>• <b>Summarize</b> key terms and concepts efficiently.</li></ul>	<ul style="list-style-type: none"><li>• <b>Promote passive recall</b> — students “ask” instead of “remember.”</li><li>• <b>Undermines long-term retention</b> (no effortful retrieval).</li><li>• <b>Encourages dependency</b> on AI as an external memory.</li></ul>
<b>2. Understanding</b> (Explain, summarize, interpret)	<ul style="list-style-type: none"><li>• Simplify complex ideas and use analogies.</li><li>• <b>Offer multiple explanations</b> for diverse learners.</li></ul>	<ul style="list-style-type: none"><li>• Learners may <b>accept answers</b> without reflection.</li><li>• Creates <b>illusion of understanding</b> without conceptual depth.</li></ul>
<b>3. Applying</b> (Use knowledge in new situations)	<ul style="list-style-type: none"><li>• <b>Generate contextual problems</b> and case scenarios.</li><li>• <b>Give step-by-step guidance</b> and practice exercises.</li></ul>	<ul style="list-style-type: none"><li>• <b>Reduces authentic problem-solving practice</b></li><li>• <b>Learners may copy solutions</b> instead of reasoning independently.</li></ul>

\*Created with ChatGPT

These issues are especially worrisome if learners use chatbots **without guidance** as self-service tools.

# Effect of Chatbot Use on Learning\*

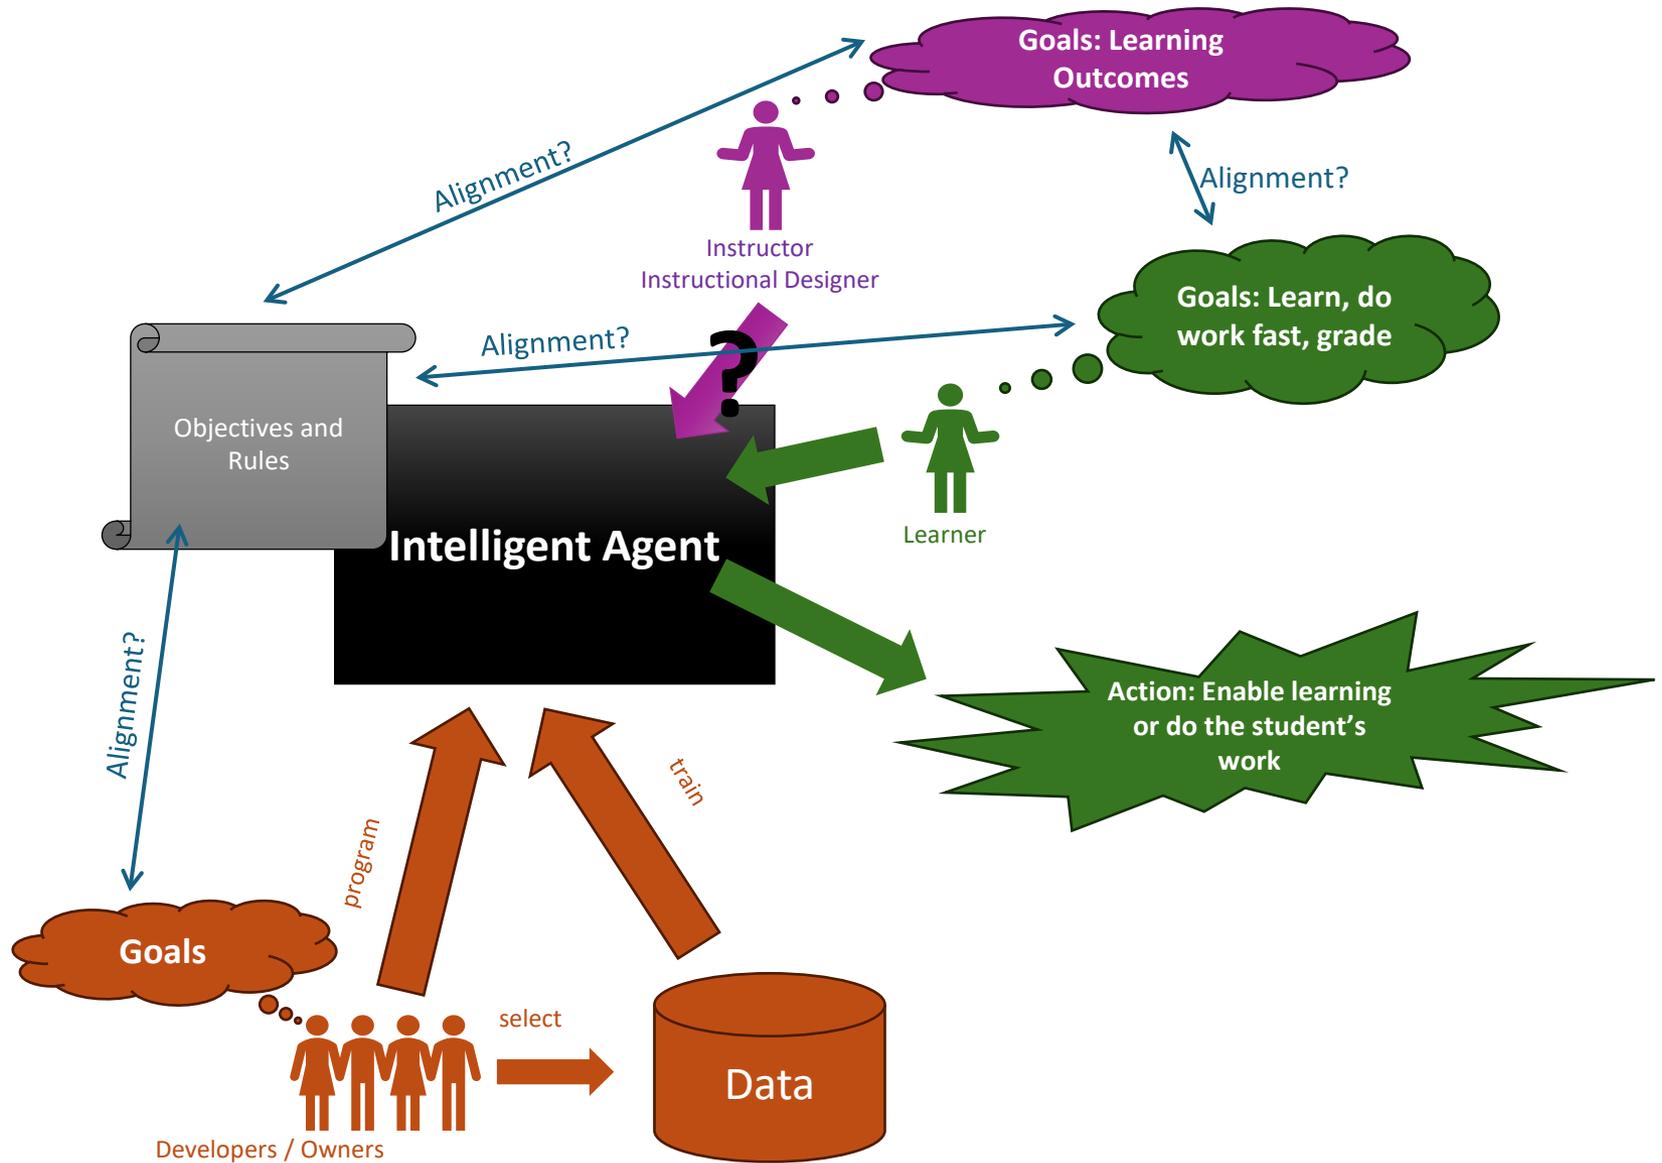
Bloom's Level	How Chatbots Support Learning	How Chatbots Hinder Learning
4. Analyzing (Differentiate, compare, organize)	<ul style="list-style-type: none"> <li>Help break down arguments, data, or texts.</li> <li>Prompt for comparisons or cause-effect analysis.</li> <li><b>Support identification of structure and logic.</b></li> </ul>	<ul style="list-style-type: none"> <li>Presents polished, unchallenged reasoning.</li> <li><b>Students stop questioning assumptions.</b></li> <li><b>Diminishes critical and source evaluation skills.</b></li> </ul>
5. Evaluating (Judge, critique, assess)	<ul style="list-style-type: none"> <li>Facilitate <b>debate simulations</b> and ethical discussions.</li> <li>Provide <b>contrasting viewpoints</b> for critique.</li> <li>Encourage justification of opinions.</li> </ul>	<ul style="list-style-type: none"> <li>May output biased or inconsistent reasoning.</li> <li><b>Learners might accept AI-generated judgments as "objective."</b></li> <li>Hinders metacognitive evaluation of personal stance.</li> </ul>
6. Creating (Design, compose, innovate)	<ul style="list-style-type: none"> <li><b>Support brainstorming</b>, ideation, and prototyping.</li> <li>Assist in <b>organizing</b> creative outputs (e.g., outlines, scripts).</li> <li><b>Provide rapid feedback</b> for iteration.</li> </ul>	<ul style="list-style-type: none"> <li>Encourages generic, AI-shaped outputs.</li> <li><b>Reduces original thought</b> and creative struggle.</li> <li>Can lead to <b>academic dishonesty</b> (AI-authored work).</li> </ul>

Especially if levels 1-4 are already weakened.

\*Created with ChatGPT

Chatbot use seems more beneficial in higher level problems (5&6). How can use be channeled towards these learning components?

# The Alignment Problem in Education



# Conclusion

- Chatbots are one of the AI technologies that can be used in education.
- Chatbots are a powerful tool to support
  - educators
  - students
  - administrators
- They also have the **potential to hinder learning**, and we are just learning what they are and how to mitigate them.
- Currently, educators are left to their own devices to make the best of the situation.
- To be successful, we need:
  - a) **Best practices to guide educators.**
  - b) Foster an **awareness in students** that misusing the technology is against their own interests.