Bernhard G. Humm

Applied Artificial Intelligence

An Engineering Approach

Third Edition

Applied Artificial Intelligence

An Engineering Approach

Bernhard G. Humm

This book is available at https://leanpub.com/AAI

This version was published on 2025-10-03

Leanpub

This is a Leanpub book. Leanpub empowers authors and publishers with the Lean Publishing process. Lean Publishing is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

Contents

Preface . . . . **i**

Copyright Notice . . . . **iii**

1. Introduction **1**

1.1 Overview of this Book 2
1.2 What is AI? 2
1.3 A Brief History of AI 5
1.4 Impacts of AI on Society 6
1.5 Prominent AI Projects 7
1.6 Further Reading 11
1.7 Quick Check 11

About the Author **12**

Preface

Why yet another book on Artificial Intelligence?

It is true that hundreds of publications on Artificial Intelligence (AI) have been published within the last decades - scientific papers and text books. Most of them focus on the theory behind AI solutions: mathematical foundations of machine learning, logic, reasoning etc. However, little can be found on engineering AI applications.

Modern, complex IT applications are not built from scratch but by integrating off-the-shelf components: libraries, frameworks, and services. The same applies, of course, for AI applications. Over the last decades, numerous off-the-shelf components for AI base functionality such as machine learning, pre-trained models for natural language processing and computer vision, knowledge graphs and reasoning have been implemented - commercial and open source. Integrating such components into user friendly, high-performance, and maintainable AI applications requires specific engineering skills. This book focuses on those skills.

I have been working in AI since the 1980. After under-/postgraduate studies of Computer Science in Germany and a Ph.D. program in Australia, I worked for more than a decade in a large German software company. In extensive teams, we developed large-scale custom software systems for clients: multinational banks, credit-card issuers, tour operators, telecommunication providers, fashion companies and ATC authorities. My tasks were as diverse as the industry sectors and technologies used. They ranged from development, software architecture, project management to managing a division and running the company's research department.

After re-entering university as a professor 20 years ago, a common theme of my courses has been the professional development of AI applications according to engineering principles and practices. Via R&D projects, funded by industry and the public, I constantly extended expertise in engineering AI applications. In teams with colleagues, postgraduate students and industry partners we developed applications for robotics and manufacturing companies, the cancer department of a hospital, the borderline unit of a psycho-therapeutic clinic, a hotel portal, a library and an creative arts museum. However diverse the industry sectors, many approaches and technologies can be used across the projects in order to build maintainable AI applications with a good user experience that meet functional and non-functional requirements, particularly high performance requirements. Overall, we are combining general software engineering skills with AI expertise. I am constantly learning: from project experience, from my colleagues and partners, from my students, and hopefully also from you, the readers of this book. So please, don't hesitate to contact me under bernhard.humm@h-da.de when you agree or disagree with my findings.

To support you to come to grips with the topics I have added questions and exercises in all chapters. They are indicated by a pencil symbol.

Bernhard G. Humm, Darmstadt, Germany, 2016, 2022 and 2025

Copyright Notice

[1] https://commons.wikimedia.org/wiki/File:Bag_of_words.JPG
[2] https://commons.wikimedia.org/wiki/File:Our_Community_Place_Sandbox.jpg
[3] https://commons.wikimedia.org/wiki/File:Deep_Blue.jpg
[4] https://commons.wikimedia.org/wiki/File:IBMWatson.jpg
[5] https://commons.wikimedia.org/wiki/File:Lee_Se-Dol.jpg
[6] https://commons.wikimedia.org/wiki/File:Lee_Sedol_B_vs_AlphaGo_W_-_Game_4.jpg
[7] https://commons.wikimedia.org/wiki/File:CART_tree_titanic_survivors.png
[8] https://commons.wikimedia.org/wiki/File:Precisionrecall.svg
[9] https://commons.wikimedia.org/wiki/File:The_Creation_of_Adam.jpg
[10] http://lod-cloud.net/
[11] https://commons.wikimedia.org/wiki/File:Lod-cloud_colored_1000px.png
[12] https://commons.wikimedia.org/wiki/File:Face_Capture.PNG
[13] https://commons.wikimedia.org/wiki/File:Restoration.jpg
[14] http://cnx.org/content/col11496/1.6/
[15] https://commons.wikimedia.org/wiki/File:113abcd_Medical_Imaging_Techniques.jpg
[16] https://commons.wikimedia.org/wiki/File:KUKA_robot_for_flat_glas_handling.jpg

- Fig. 7.9: CC Wikimedia Commons, Author Robert Basic, Source Audi A8 2013 Uploaded by AVIA BavARia, 17 October 2013, Audi_A8_2013_(11209949525).jpg[17]
- Fig. 7.10: Wikimedia Commons, Author NASA/JPL/Cornell University, Maas Digital LLC, Source photojournal.jpl.nasa.gov[18], February 2003, NASA_Mars_Rover.jpg[19]
- Fig. 7.11: CC Wikimedia Commons, Author Vazquez88, 5 February 2013, Motion_Capture_-with_Chad_Phantom.png[20]
- Author photo by Markus Schmidt, Darmstadt University of Applied Sciences

MS copilot has been used for text refinement.

[17]https://commons.wikimedia.org/wiki/File:Audi_A8_2013.jpg
[18]http://photojournal.jpl.nasa.gov/catalog/PIA04413
[19]https://commons.wikimedia.org/wiki/File:NASA_Mars_Rover.jpg
[20]https://commons.wikimedia.org/wiki/File:Motion_Capture_with_Chad_Phantom.png

1. Introduction

How relevant is *Artificial Intelligence*?

When I wrote the first edition of this book in 2015, Artificial Intelligence (AI) was hardly noticeably in public. For many people, AI was nothing more than a burst bubble of a 20th century hype. However, even then AI was relevant and ubiquitous in IT applications of business and consumer markets.

Nowadays, nobody questions the relevance of AI. Countries, companies and research organizations specify AI as one of the main strategic areas. AI is hype. Although, as in all hype topics, some expectations and and fears are over-exaggerated, there is no doubt that AI is one of the leading technologies and is here to stay.

A few examples of AI in everyday use are:

- Chatbots and question answering with tools like ChatGPT or Gemini
- Image and video generation with tools like DALL-E and Stable Diffusion
- Speech control for smart-phones, navigation systems and digital assistants, etc.
- Face recognition in cameras
- Learning spam filters in e-mail clients
- AI in computer games
- Semantic Internet search, including machine translation and question answering

Commercial examples are:

- Multi-channel customer service
- Business intelligence
- Robotics and automation
- Supply chain management and logistics
- Self-driving cars, drones, rockets (military and commercial)

I expect AI to continue to change our daily lives and labor markets in an enormous way, just like technology has changed the daily lives of generations since the 19th century with ever increasing speed.

How are AI applications developed?

While most AI publications, such as scientific articles and text books, focus on the (mathematical) foundation of AI, little can be found on engineering AI applications. What kinds of AI libraries, frameworks and services exist? Which ones should be chosen in which situation? How to integrate them into maintainable AI applications with a good user experience? How to meet functional and non-functional requirements, in particular high performance?

The focus of this book is to answer those kinds of questions for software developers and architects.

1.1 Overview of this Book

The remainder of this chapter presents a definition of AI as well as a brief AI history. Chapters 2 and 3 present the main AI approaches: machine learning and knowledge-based AI. Chapter 4 gives guidelines for the architecture of AI applications. The remaining chapters focus on individual AI areas: (5) information retrieval, (6) natural language processing, (7) computer vision, and (8) complex event processing. Chapter 9 concludes the book.

Inspired by one of my AI projects, I use application examples from the domain of art museums throughout the book. Additionally, I present real-world examples from other domains.

At the end of each chapter I repeat the main themes in form of questions that you, the reader, can use as a *quick check*.

The appendix contains product tables as well as source code examples that can be used as starting point for your own developments.

1.2 What is AI?

Encyclopaedia Britannica[1] defines AI as follows.

"Artificial intelligence (AI), the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings."

Note that this definition does not claim or even postulate that AI applications *are* intelligent nor that AI is comparable ore even equivalent to human intelligence.

What are abilities commonly associated with intelligent beings?

- **Perceiving**: Seeing, listening, sensing, etc.
- **Learning, knowing, reasoning**: Thinking, understanding, planning etc.
- **Communicating**: Speaking, writing, etc.
- **Acting**

See Fig. 1.1.

[1] https://www.britannica.com/technology/artificial-intelligence

Fig. 1.1: Tasks commonly associated with intelligent beings

The different areas of AI can be structured according to those abilities. See the circles in the "landscape of AI" (inspired by an AI Spektrum poster[2]) in Fig. 1.2.

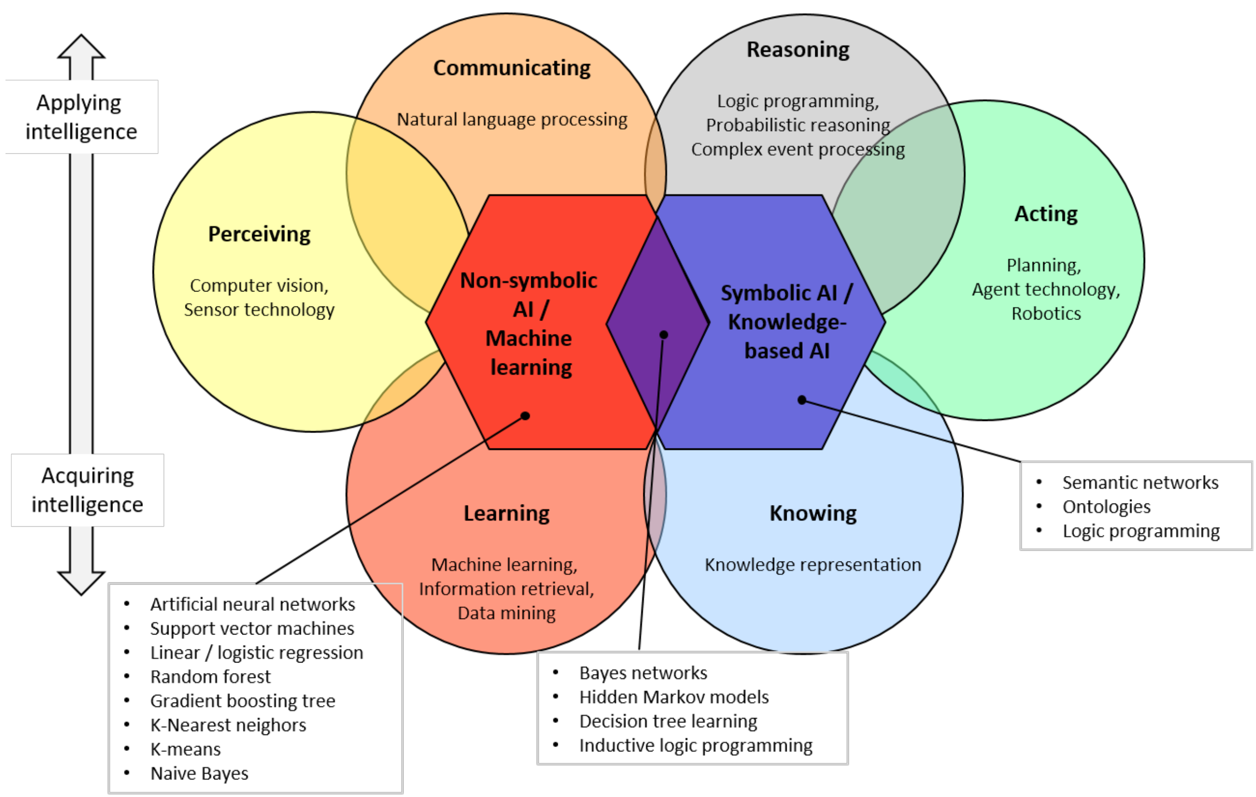

Fig. 1.2: The AI landscape

- **Perceiving** covers the AI areas of *computer vision* and *sensor technology*;

[2]https://www.sigs-datacom.de/order/poster/Landkarte-KI-Poster-AI-2018.php

- **Communicating** covers the area of *natural language processing (NLP)*;
- **Learning** covers *machine learning (ML), information retrieval (IR)* and *data mining*;
- **Knowing** covers *knowledge representation*;
- **Reasoning** covers *logic programming, probabilistic reasoning*, and *complex event processing (CEP)*;
- **Acting** covers planning, agent technology, and robotics.

Please note, that various different namings and groupings of AI areas can be found in literature.

Two fundamentally different approaches to AI can be distinguished (depicted as hexagons in Fig. 1.2):

- **Knowledge based AI** (sometimes called *symbolic AI*): Knowledge is represented explicitly in a human-readable way, e.g., with knowledge graphs, semantic networks, or logic programming languages (see the boxes with technologies attached to the approaches in Fig. 1.2).
- **Machine learning** (sometimes called *non-symbolic AI*): Knowledge is implicit in form of numbers, e.g., as weights in artificial neural networks, support vector machines, in linear / logistic regression etc.

Both approaches, knowledge-based AI and machine learning, have been around from the various beginnings of AI in the 1950s. In the first decades of AI research and practice, knowledge-based AI was most prominent and showed the most convincing results. However, in the last decades this has shifted and today, machine learning approaches are most prominent.

Both approaches have advantages and disadvantages. Machine learning approaches require little upfront knowledge, just (many) samples for training. They exhibit good behavior, e.g., in classifying situations, also with noisy data. However, the reasoning behind decisions cannot be explained to humans. In contrast, the reasoning behind knowledge-based AI is explicit and can be retraced by humans. However, explicit knowledge engineering is required upfront and reasoning under uncertainty is challenging.

Both approaches have been applied in all areas of AI, but knowledge-based AI is more commonly used for the abilities of knowing, reasoning, and acting, whereas machine learning approaches are more commonly used for the abilities of perceiving, communicating, and learning.

I expect that *hybrid AI approaches* will gain importance in future. Hybrid approaches combine the advantages of machine learning and knowledge-based AI. Machine learning may be used for knowledge-based AI, e.g., in the (semi-)automatic extraction of knowledge graphs from texts. Knowledge-based AI may be used for machine learning, e.g., by enriching training data with knowledge-graphs. Explicitly hybrid AI approaches include Bayes networks, hidden Markov models, decision tree learning etc.

1.3 A Brief History of AI

The term "Artificial Intelligence" was coined at the *Dartmouth Workshop* in 1956. Members of this workshop were John McCarthy, Marvin Minsky, Claude Shannon, Allen Newell, Herbert Simon, and others. However, Alan Turing (1912 – 1954) with his fundamental work on computability and the so-called *Turing Test* for assessing intelligent behavior had already laid ground for AI decades before.

The 1960s - 1980s saw unprecedented *AI hype*, triggered by enthusiastic promises of quick AI results. Examples of such promises are:

"Within a generation the problem of creating 'artificial intelligence' will be substantially solved." (Marvin Minsky, 1967)

"Within 10 years computers won't even keep us as pets." (Marvin Minsky, 1970).

or

"Machines will be capable of doing any work that a man can do." (Herbert Simon, 1985) (Citations from (American Heritage, 2001)).

This hype resulted in massive funding of AI projects, particularly in the US.

The effect of those wildly exaggerated promises was the same as in all hype. When people started to notice that the most sophisticated AI applications failed to perform tasks that are easily accomplished by small children they threw the baby out with the bathwater. The disillusionment of the unrealizable expectations resulted in massive cuts in funding and a collapse of the AI market. The 1980s - 1990s have sometimes been called the *AI winter*.

From then on, unnoticed and often not under the term AI, AI methods and technologies have matured enormously driven by major technology companies. For example, Google co-founder Larry Page[3] said in 2006: "We want to create the ultimate search engine that can understand anything. Some people could call that artificial intelligence" .

This development of AI applications by major technology drivers lead to the situation today where AI is relevant and ubiquitous in everyday use. Moreover, AI has become a hype topic, again, particularly boosted by generative AI. This is promoted by continuous media attention, science fiction movies, and bold statements by individuals in the AI community that are remarkably similar to the ones in the 1970s:

"Artificial intelligence will reach human levels by around 2029. Follow that out further to, say, 2045, we will have multiplied the intelligence, the human biological machine intelligence of our civilization a billion-fold." (Ray Kurzweil, 2005[4])

I personally cannot see the slightest evidence for such bold claims and regard them as pure science fiction. Such exaggerated claims create much media attention but may eventually lead to another AI winter. A most profound response to such claims has been written by Rodney Brooks in 2017[5]

[3] http://bigdata-madesimple.com/12-famous-quotes-on-artificial-intelligence-by-google-founders
[4] https://www.brainyquote.com/quotes/ray_kurzweil_591137
[5] http://rodneybrooks.com/the-seven-deadly-sins-of-predicting-the-future-of-ai/

which I recommend reading.

1.4 Impacts of AI on Society

Software applications and in particular AI applications may have massive impacts on society. I believe that software engineers must be aware of those implications and act responsibly.

Automation technology has always had enormous impact on the labor market. We can see this development particularly since the industrial revolution in the 19th century. With the advent of information technology in the 20th century, the speed of innovation has accelerated. AI technology is a continuation of this highly dynamic process, with ever increasing innovation speed. Like in earlier technology shifts, groups of jobs will become obsolete and new jobs will emerge. With agentic coding, also the software engineering industry is undergoing massive changes.

There are numerous predictions about the implications on AI on the labor market. Some predict a replacement rate of human labor by robots for certain business sectors of up to 99%[6]. Some predict that fears of massive displacement of workers are unfounded in the longer term[7].

To quote Mark Twain (or Niels Bohr maybe others), "It's Difficult to Make Predictions, Especially About the Future"[8] ;-). I personally assume that the ever increasing automation technology will, in fact, decrease the need for human labor since we cannot (and should not) increase consumption forever. In this case it will be necessary to come to a societal agreement of how to distribute the wealth generated by machines. Some discuss a basic income[9] as one model.

Also, as a society, we must come to an agreement, which decisions we may leave to machines and which ones we must not. Alan Bundy pointedly and correctly states in an CACM viewpoint[10]: "Smart machines are not a threat to humanity. Worrying about machines that are too smart distracts us from the real and present threat from machines that are too dumb".

Bias in machine learning is an issue which is in public attention. Since machine learning depends on data, the quality of data severely affects the quality of decision being made. But data can be biased due to the creators or origins of data - and thus, decisions based on machine learning can be biased.

More and more decisions are delegated to machines although, of course, no piece of software is perfect. This may be fine for self-driving cars, at least as soon as they drive safer than average humans (which I expect in the foreseeable future). But it may be dangerous in areas where decisions by machines may have far-reaching impacts that no human can foresee. One of those areas may be high-frequency trading where machines autonomously make selling and buying decisions, but the impact on the global financial market is unclear. Even more critical are autonomous weapons that are programmed to detect alleged attacks automatically and to start a counter-strike without human

[6] https://futurism.com/will-artificial-intelligence-take-job-experts-weigh
[7] https://ec.europa.eu/epsc/sites/epsc/files/ai-report_online-version.pdf
[8] https://quoteinvestigator.com/2013/10/20/no-predict/
[9] https://en.wikipedia.org/wiki/Basic_income
[10] https://cacm.acm.org/magazines/2017/2/212436-smart-machines-are-not-a-threat-to-humanity/

intervention. An open letter from AI and robotics researchers on autonomous weapons[11] in which this issue is discussed was signed by more than 30,000 people.

Data-intensive applications like machine learning may use an enormous amount of computing power and, thus energy. Thus, positive societal effects of AI may be alleviated by a large carbon footprint. *Green AI* is the field which tries to reduce negative effects of AI on the environment.

Computer scientists developing AI applications must be aware of their impact and act responsibly. Current keywords in the debate are responsible AI, ethical AI, human-centered AI or explainable AI. We will cover aspects of these in the book.

1.5 Prominent AI Projects

Some AI projects have become milestones in AI history due to their public visibility. I shortly mention just three examples.

In 1997, *IBM Deep Blue* beat the chess world champion Garry Kasparov. This was a milestone since chess is considered one of the most complex board games.

See Fig. 1.3.

Fig. 1.3: IBM Deep Blue beats the chess champion Garry Kasparov

However, when people examined the computational algorithms used in the Deep Blue Application, people quickly realized the difference between the IBM approach and the way human chess players act intelligently. For people who believed in machines being intelligent like humans, this was a disillusionment. For practitioners this was no disillusionment at all but an important milestone in the development of applications which exhibit behavior of human intelligence.

[11]https://futureoflife.org/open-letter-autonomous-weapons/

In 2011, IBM succeeded with another important AI project: *IBM Watson*. While Deep Blue targeted a most specific ability, namely playing chess, IBM Watson was able to answer natural language questions about general knowledge. The media highlight of this project was beating the human champions at the popular US quiz show Jeopardy! This was remarkable since not only questions about history and current events, the sciences, the arts, popular culture, literature, and languages needed to be answered but also play on words as well as execution speed and strategy needed to be considered.

See Fig. 1.4.

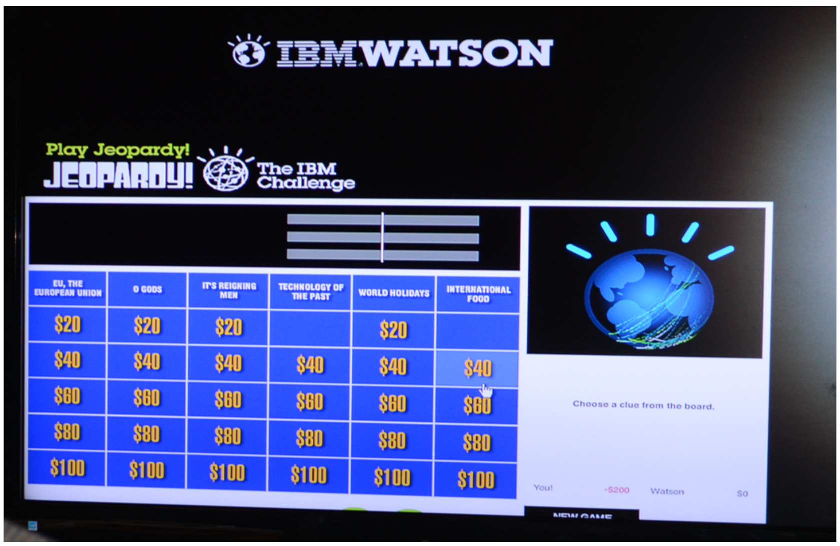

Fig. 1.4: IBM Watson for the Jeopardy! quiz show

In 2012, the ImageNet Large Scale Visual Recognition Challenge marked a turning point in AI history. A deep convolutional neural network called *AlexNet*, developed by Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, dramatically outperformed all previous models in image classification. It reduced the top-5 error rate by more than 10 percentage points, showcasing the power of deep learning on large datasets. This breakthrough demonstrated that neural networks, when trained with sufficient data and computational power, could rival human-level perception in visual tasks. The success of AlexNet ignited a wave of innovation across computer vision, speech recognition, and natural language processing. It also accelerated GPU adoption for AI workloads and laid the foundation for modern architectures like ResNet and Transformers. The ImageNet moment is widely credited with launching the deep learning revolution.

Go is a complex board game that requires intuition, creative and strategic thinking. It has long been considered a difficult challenge in AI and is considerably more difficult to solve than chess. In 2015, AlphaGo, a computer Go program developed by Google DeepMind, played a five-game Go match with the 18-time world champion Lee Sedol[12]. The received high media attention. AlphaGo won all but the fourth game. This was considered another breakthrough in AI.

See Fig. 1.5.

[12] https://en.wikipedia.org/wiki/AlphaGo_versus_Lee_Sedol

Fig. 1.5: Lee Sedol, Go match with AlphaGo

In late 2022, OpenAI's ChatGPT captured global attention as the first large-scale conversational AI to reach mainstream adoption. Built on GPT-3.5 and later GPT-4, ChatGPT demonstrated remarkable fluency, contextual understanding, and versatility across domains—from coding and writing to tutoring and entertainment. Within weeks, it amassed over 100 million users, becoming the fastest-growing consumer application in history. Its intuitive chat interface made advanced AI accessible to the general public, sparking widespread fascination and debate. ChatGPT's viral success marked a cultural shift, bringing generative AI into classrooms, workplaces, and daily life. It also catalyzed a wave of innovation in agentic systems, prompting tech giants to accelerate their own AI offerings. The launch of ChatGPT redefined expectations for human–machine interaction and set the stage for the agentic AI era.

In 2023, AI-generated art and music surged into mainstream culture, transforming creative industries and public perception of AI. Tools like *DALL·E 2*, *Midjourney*, and *Stable Diffusion* enabled users to generate stunning visuals from text prompts, while platforms such as *AIVA* and *Soundraw* composed original music with minimal human input. Artists, designers, and hobbyists embraced these models for rapid prototyping and inspiration, while debates around authorship, copyright, and artistic authenticity intensified. Social media was flooded with AI-generated portraits, album covers, and surreal compositions, blurring the line between human and machine creativity. The boom democratized access to high-quality content creation and sparked new genres of digital expression. AI art exhibitions, music collaborations, and viral remix culture became global phenomena. This moment marked a shift from AI as a tool of analysis to AI as a co-creator, redefining the boundaries of imagination. See Fig. 1.6. for a screenshot from DALL-E 3.

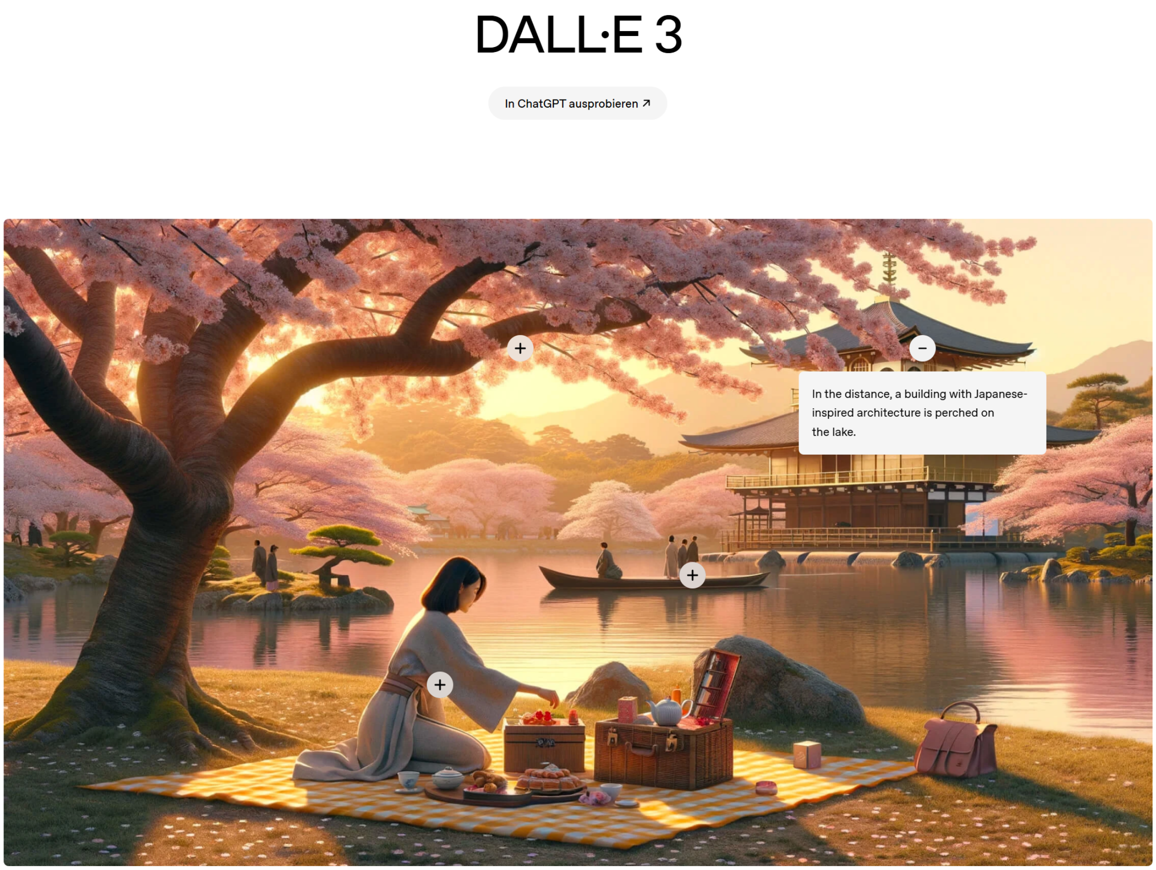

Fig. 1.6: DALL-E screenshot

In 2025, OpenAI introduced *GPT-4o*, a groundbreaking multimodal model capable of processing text, images, and audio in real time. This leap enabled the creation of real-time agentic systems that could hold spoken conversations, interpret visual input, and act autonomously across digital environments. GPT-4o's responsiveness and contextual depth redefined human–AI interaction, making agents feel more intuitive, adaptive, and lifelike. Developers rapidly integrated GPT-4o into virtual assistants, customer service bots, and educational tutors. Its release marked a shift from static chatbots to dynamic, multimodal agents capable of reasoning and reacting in real time. GPT-4o became the foundation for next-generation AI applications across industries.

To conclude, AI applications have made enormous progress in the last two decades - applications that were unthinkable in the late 20th century. All those applications exhibit behavior of human intelligence in certain situations (answering questions, performing tasks, playing games, driving cars, etc.) without necessarily imitating human intelligence as such.

The following chapters give insights into engineering such AI applications.

1.6 Further Reading

There are numerous publications on AI, scientific papers, text books and online resources. Many text books an scientific papers focus on the mathematical foundations and theories behind AI mechanisms. Many tutorials by vendors like OpenAI provide excellent entry points in using specific AI frameworks. Many online articles on platforms like "Towards Data Science (Medium)" introduce AI concepts and methods on a practical level. Online courses by organizations like Udacity, Coursera or edX provide theoretical background as well as hands-on experience with AI technologies. In the end, also chats with chatbots like ChatGPT or Gemini can deepen knowledge on engineering AI applications.

1.7 Quick Check

The quick check shall help you assessing whether you have understood the main topics of this chapter. Answer the following questions.

1. What does the term "Artificial Intelligence" mean?
2. What are main areas of AI?
3. Sketch the history of AI
4. What is the role of AI today?
5. What are potential impacts of AI applications on society?
6. Name a few prominent AI projects

About the Author

Bernhard G. Humm has been working in AI since the 1980. He is a professor for software engineering at the Computer Science Department of Darmstadt University of Applied Sciences, Germany. He received a Ph.D. from the University of Wollongong, Australia and the degree Dipl.-Inform. from Kaiserslautern University, Germany. His research focus is on applied AI and software architecture. He is the managing director of the institute of applied informatics, Darmstadt (aIDa) and Ph.D. coordinator. He is running several national and international research projects in co-operation with industry and research organizations and is publishing his results regularly. Before re-entering university as a professor, he worked in the IT industry for 11 years as a software architect, chief consultant, IT manager and head of the research department of a large German software company. His clients were in the industry sectors finance, tourism, trade, and aviation.

Contact:
Prof. Bernhard G. Humm, Dipl.-Inform., Ph.D. Computer Science (AUS)
Darmstadt University of Applied Sciences, Computer Science Department
Schoeffertstr. 3, 64295 Darmstadt, Germany
bernhard.humm@h-da.de, https://fbi.h-da.de/personen/bernhard-humm