Introduction to Multiagent Systems

Mehdi Dastani
BBL-521
m.m.dastani@uu.nl

Webpage
https://ics.uu.nl/docs/vakken/mas/

Teaching Staff

Lectures: Mehdi Dastani
m.m.dastani@uu.nl

Lab: Roxana Rădulescu
r.t.radulescu@uu.nl

Tutorial assistant: Panagiotis Aronis
p.aronis@uu.nl

The Aim of this Course

- The course consists of lecture and lab sessions.
- Lecture sessions provide an introduction to the field of multiagent systems and covers:
  - decision theory
  - game theory
  - social choice
  - mechanism design
  - auctions
  - some other topics to be announced
- Lab sessions aim at giving you experience in game theoretic analysis in multiagent systems and covers:
  - Multiagent repeated interactions
  - Utility theory
  - Strategic interactions in repeated interactions
  - Opponent modeling
  - Decision-making under uncertainty
  - Development of Multiagent Systems

Lab Sessions

- The lab sessions are organised around a student group assignment
- The assignment
  - concerns the design and implementation of multiagent concepts in repeated interactions
  - consists of written reports and a Python implementation of different agent strategies in repeated interactions
  - is performed in interdisciplinary groups
- Each group consists of five students
- Each group has a coordinator who is responsible for:
  - distributing the tasks,
  - communication with us and other students,
  - submission of reports and agent program, and
  - reporting on activities: experience of the team and a summary of who performed which tasks.
- You can prepare already by forming your group and by starting to work on the introductory parts of the assignment:
  `/https://ics.uu.nl/docs/vakken/mas/`

Exam and Grade

- The final written exam is on Thursday, 10 April 2025 (13:30-16:30) in BEATRIX - 7E ETAGE (city center)
- The final grade is based on exam (70%) and assignment (30%)
- The final grade of the written exam should be ≥ 5
- The final grade of the assignment should be ≥ 5
- To pass the course the final grade (70%w.ex. + 30%ass.) must be ≥ 5.5
- For the assignment part there is NO retake

Multiagent Systems: Literature

- Book (some sections): Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundation, by Yoav Shoham and Kevin Leyton-Brown, Cambridge University Press, 2009.
- Book (background): An Introduction to Multiagent Systems (second edition): Michael Wooldridge. John Wiley & Sons, LTD, 2009.
- See the home page of the course for other background literature.

Artificial Intelligence

Artificial intelligence is the study of understanding, artificial replication and possible enhancement of (human) intelligence.

Artificial Intelligence aims at continuously advancing computer technology to automate ever increasing complex tasks for which human intervention is required.

Artificial Intelligence

Understanding human intelligence is inspiring, but designing artificial models is indispensable.

Artificial Intelligence

Understanding human intelligence is inspiring, but designing artificial models is indispensable.

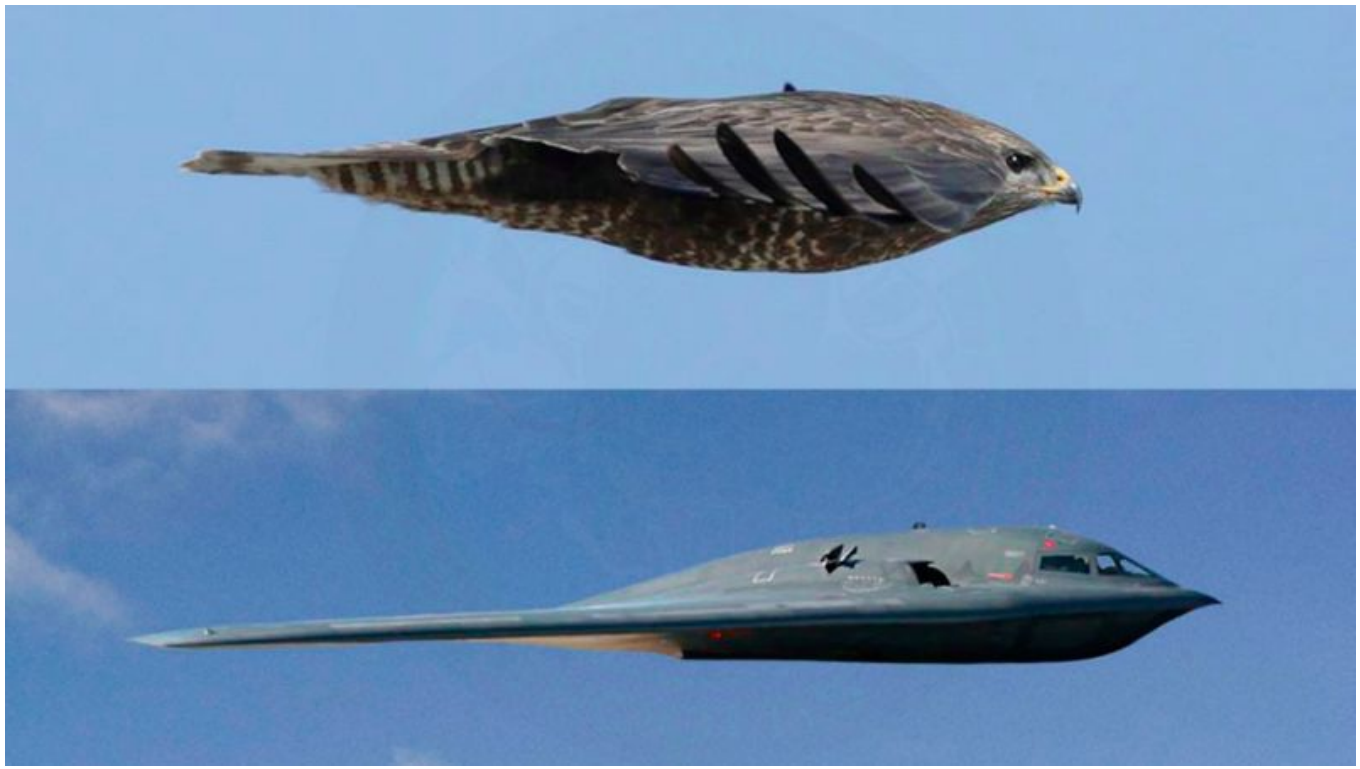

Artificial Intelligence is not just to replicate (human) intelligence, it can push it beyond what (human) intelligence has ever been capable of.

Artificial Intelligence: Early Days

Artificial Intelligence: Understand and model the behaviour of a *single intelligent autonomous agent*

- Automatic planning
- Machine learning
- Computer vision
- Computational linguistics
- Robotics

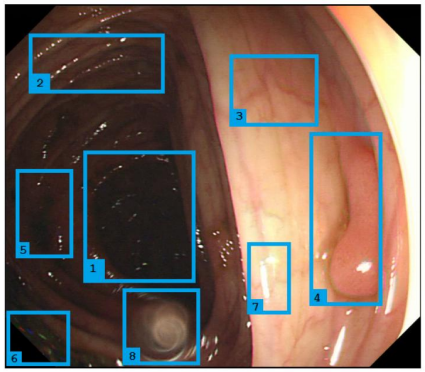

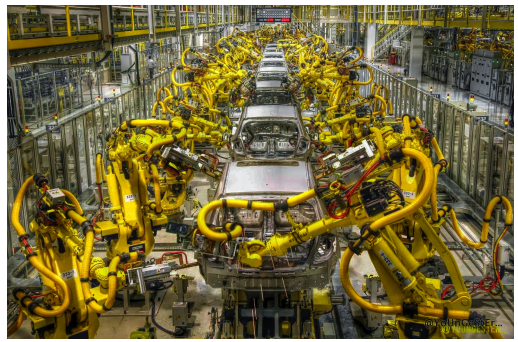

Autonomous Agents: Integrating AI Techniques

Autonomous agents **sense** their environments, **learn** and **reason** to decide and **execute** actions that maximize the chance to successfully achieve their objectives.

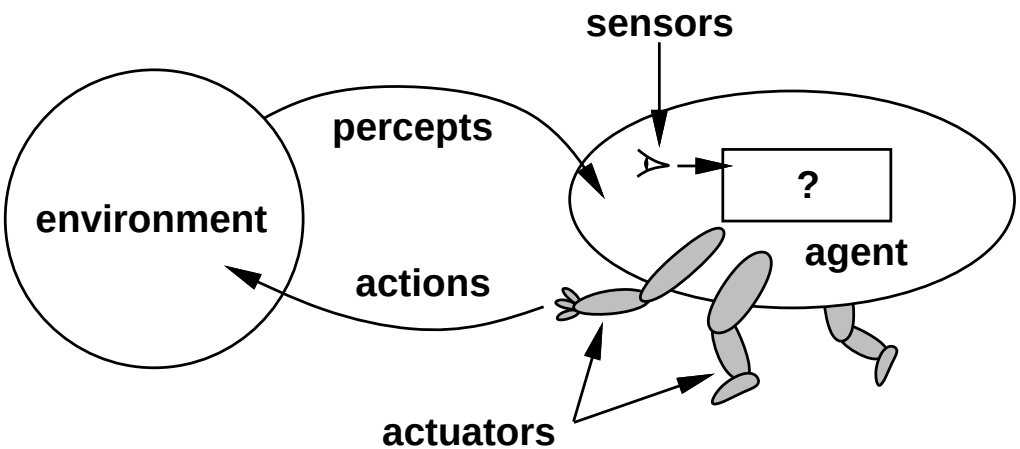

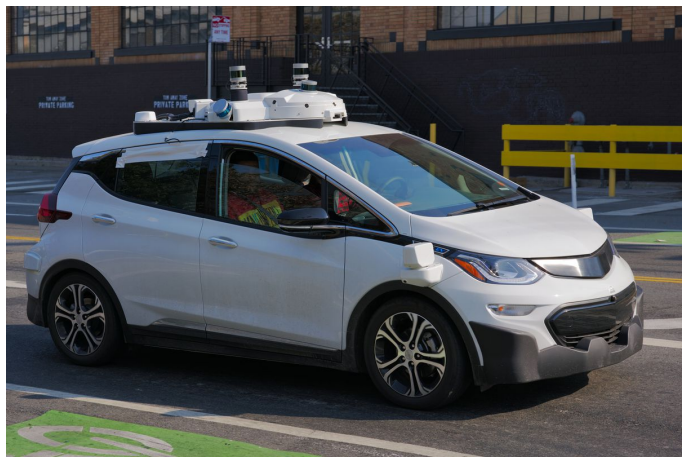

Autonomous Agents: Integrating AI Techniques

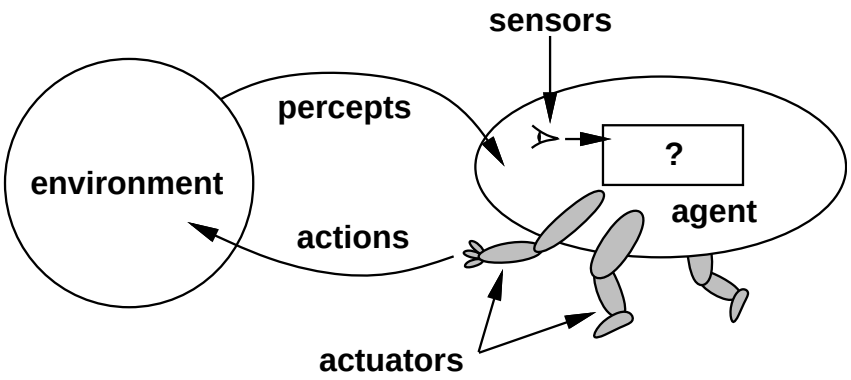

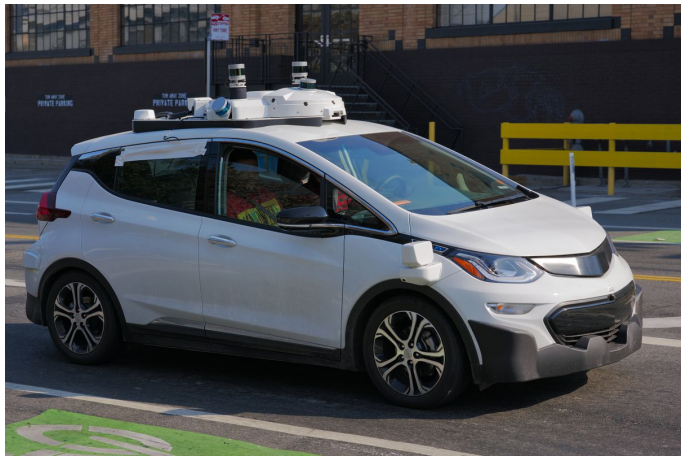

Some research issues

- Updating system state based on sensed data
- Learn and reason to decide actions and plans
- Coordinated execution of actions and plans
- Engineering autonomous agents

Multiagent Systems: Interacting Autonomous Agents

A multiagent system consists of a set of autonomous agents that interact with each other and their surrounding environment to achieve their (joint) objectives.

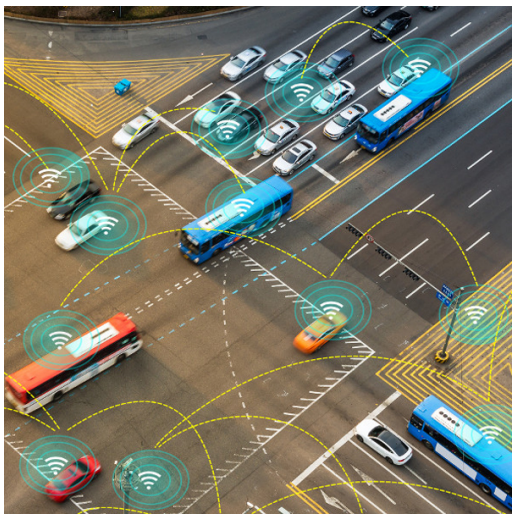

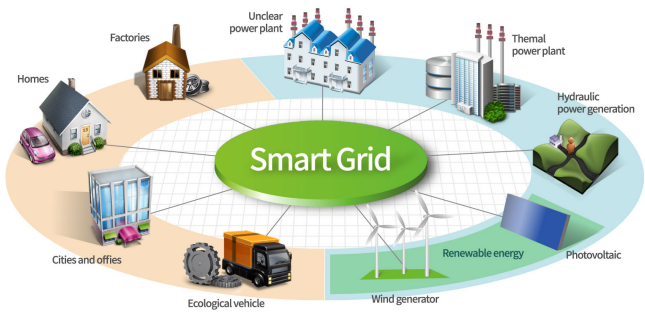

Multiagent Systems: Interacting Autonomous Agents

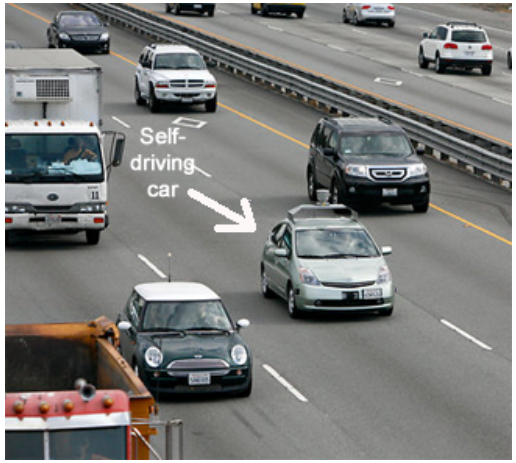

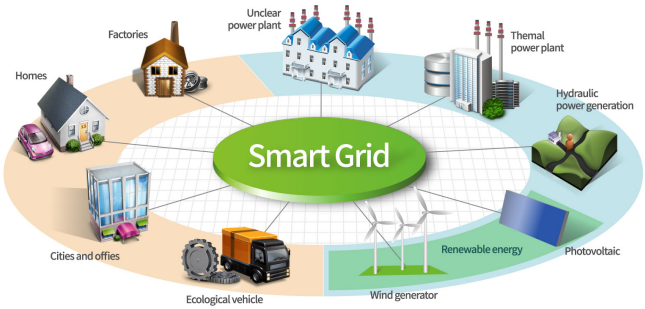

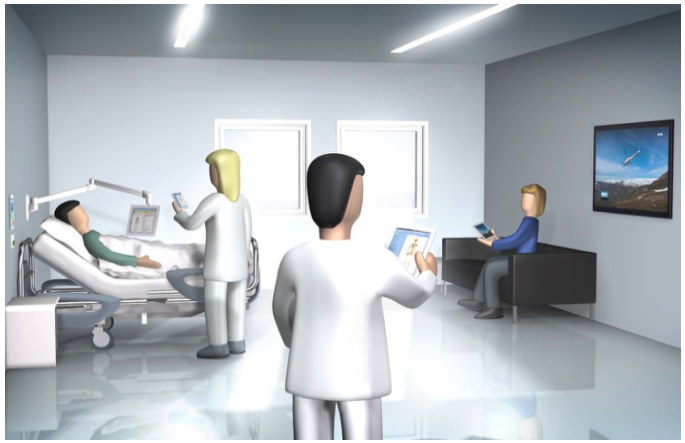

Engineering distributed systems requires **multidisciplinary** techniques to cope with the complexity caused by dynamic **emergent relations** between subsystems.

Some research issues

- modelling and assessing overall system behaviour
- designing interaction mechanisms to achieve optimal collective behaviour
- monitoring and controlling interaction between subsystems
- simulating interacting systems

Normative versus Descriptive AI: A dichotomy

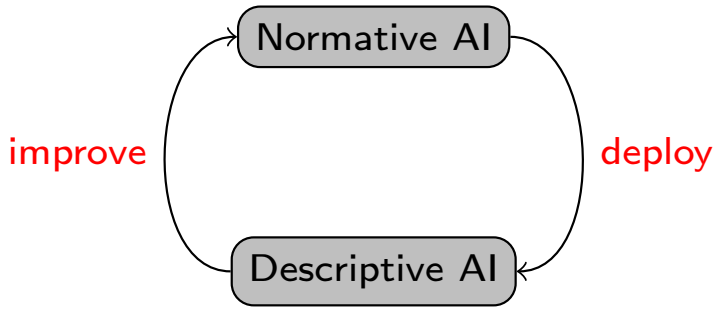

- **Normative Artificial Intelligence**: How **should** intelligent systems behave? An evaluative account of intelligence $\Longrightarrow$ **Model-driven (Symbolic) AI**.

- **Descriptive Artificial Intelligence**: What **do** intelligent systems behave? A descriptive account of intelligence $\Longrightarrow$ **Data-driven (Subsymbolic) AI**.

Normative Artificial Intelligence

How should an intelligent system behave?

- Understanding and explaining intelligent phenomenon (e.g., cognitive tasks such as reasoning, planning, and natural language)
- Building formal models using mathematical tools and techniques to manipulate symbolic knowledge
- Prove properties of the models (safety, liveness, expressiveness)
- Design deductive reasoning algorithms to build **artificial** intelligent systems
- Examples: Expert systems, Navigation systems, Diagnostic systems, Decision trees

Normative Artificial Intelligence

How should intelligent systems behave?

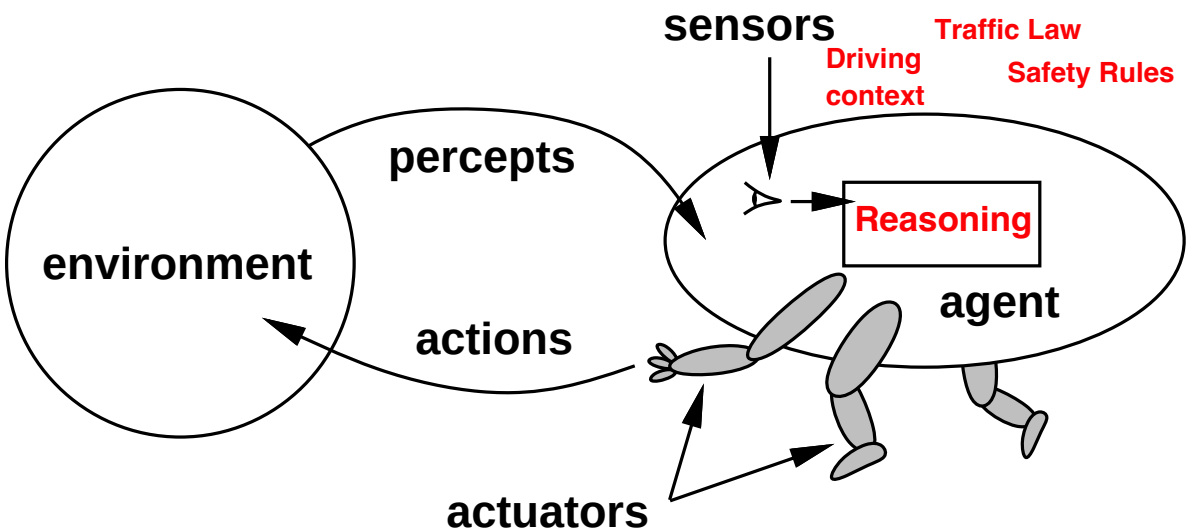

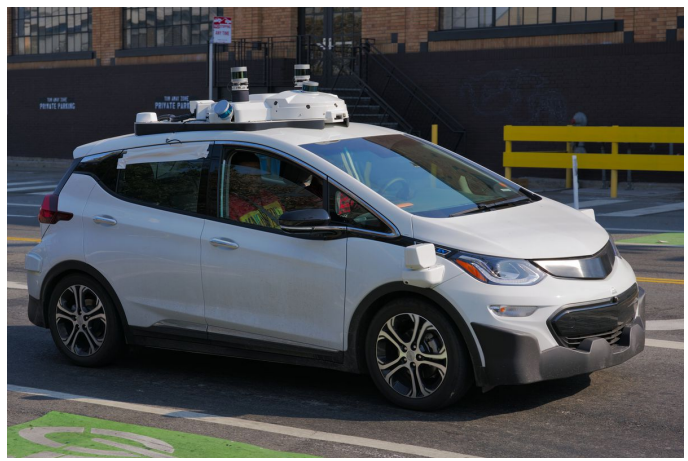

Descriptive Artificial Intelligence

How do intelligent systems behave?

- Modelling the data generated by intelligent systems (e.g., perceptual tasks such as vision, speech recognition, and natural language)
- Building formal models using mathematical tools and techniques to manipulate subsymbolic knowledge
- Measuring correctness of the models (recall, precision, accuracy)
- Design learning algorithms to build **artificial** intelligent systems
- Examples: Deep neural networks, (Deep) Reinforcement Learning, Transformers, Large Language Models

Descriptive Artificial Intelligence

How do intelligent systems behave?

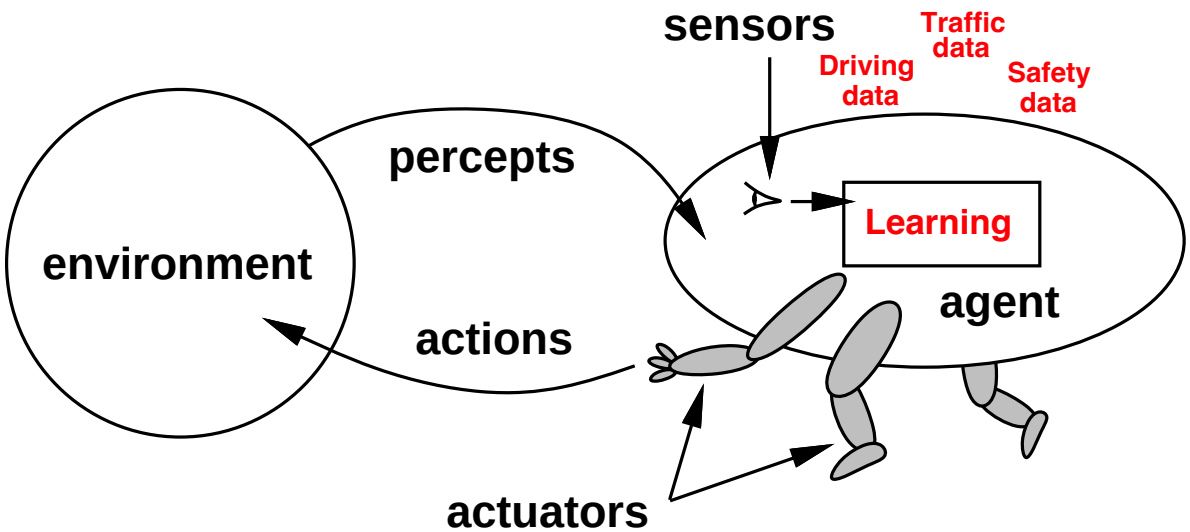

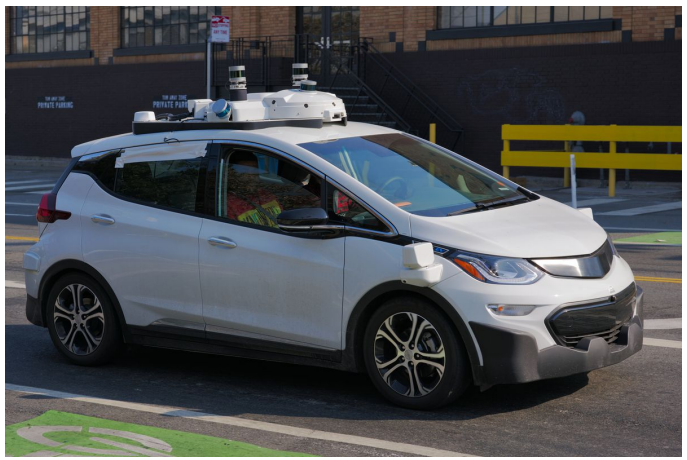

The Rise of Data-driven Artificial Intelligence

Deep Learning

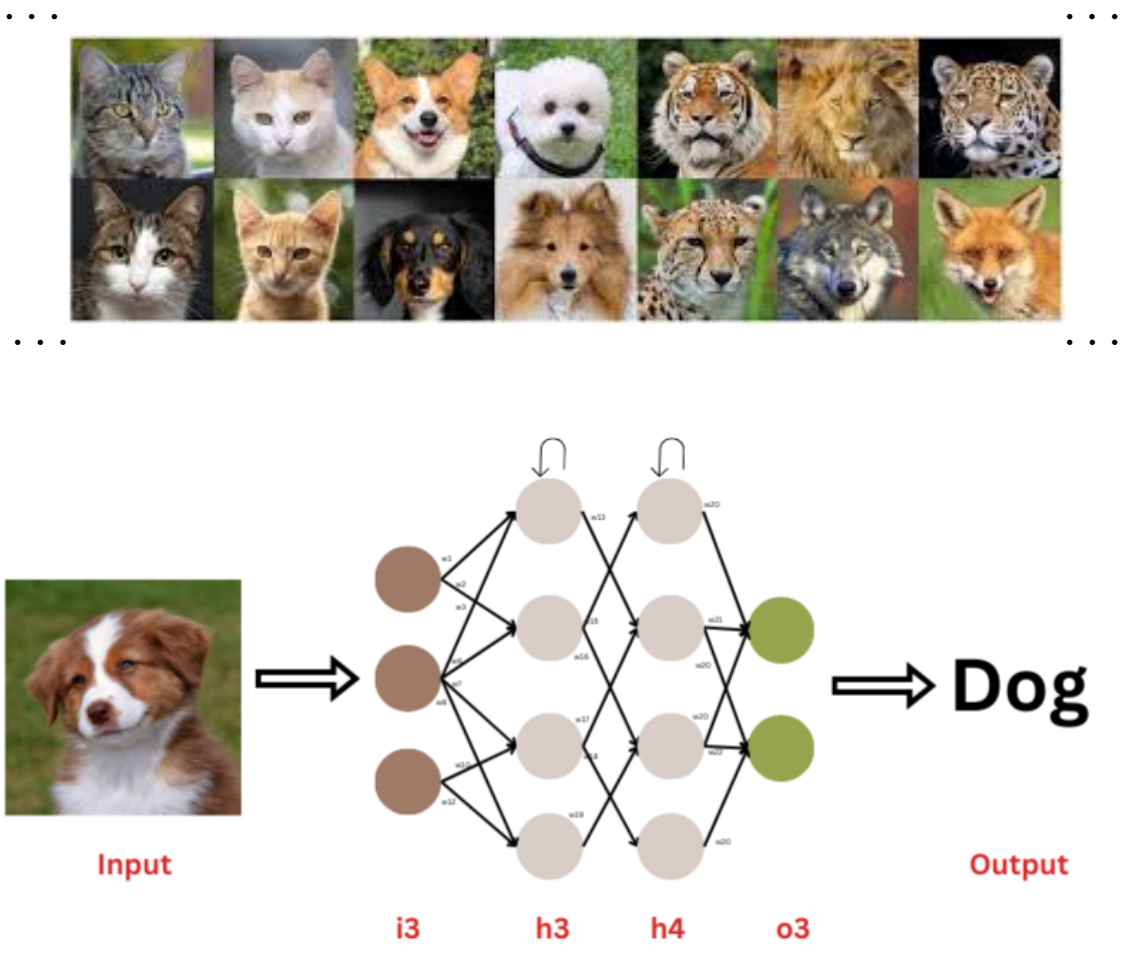

The Rise of Data-driven Artificial Intelligence

(Deep) Reinforcement Learning

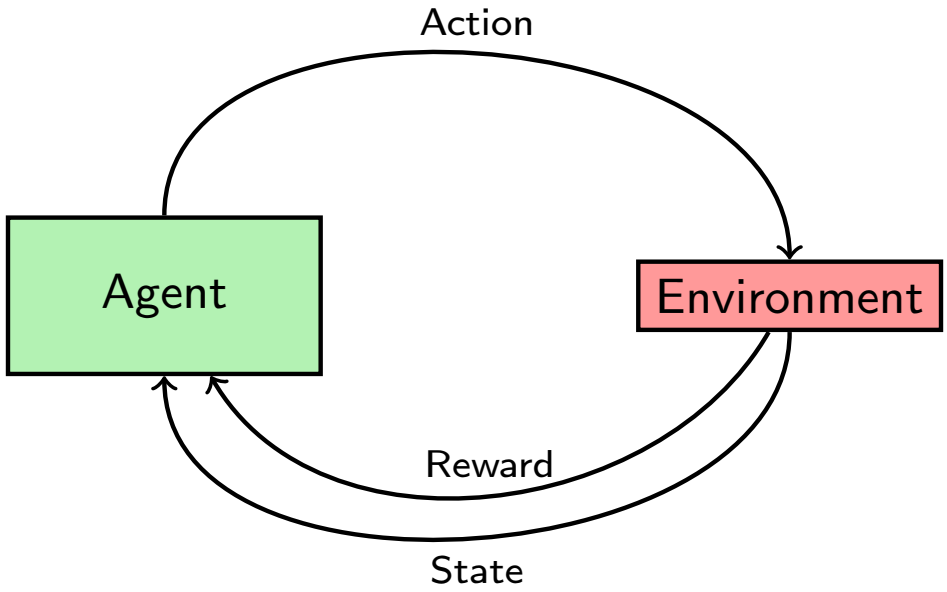

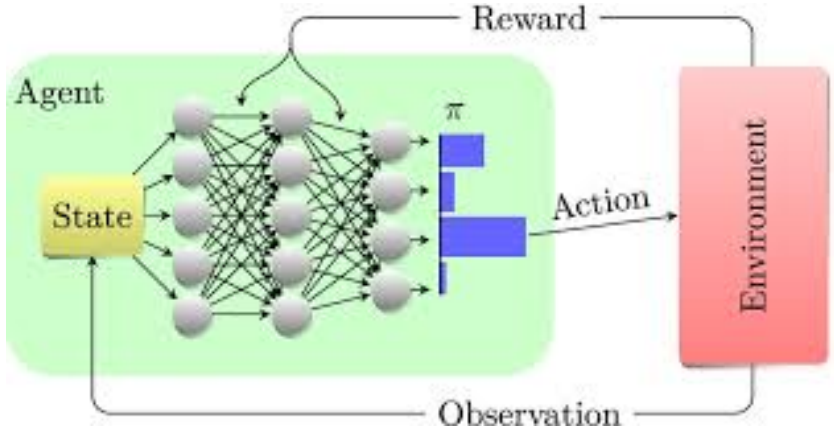

The Rise of Data-driven Artificial Intelligence

Next word Prediction

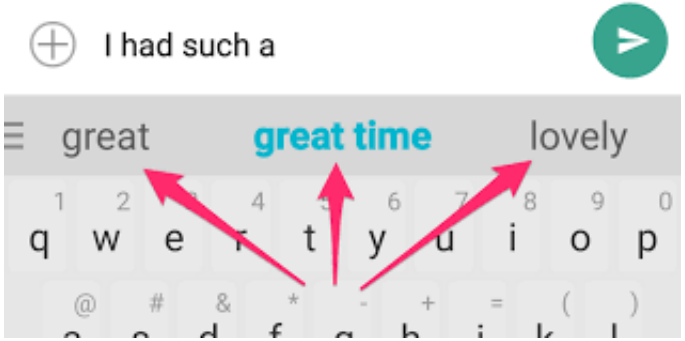

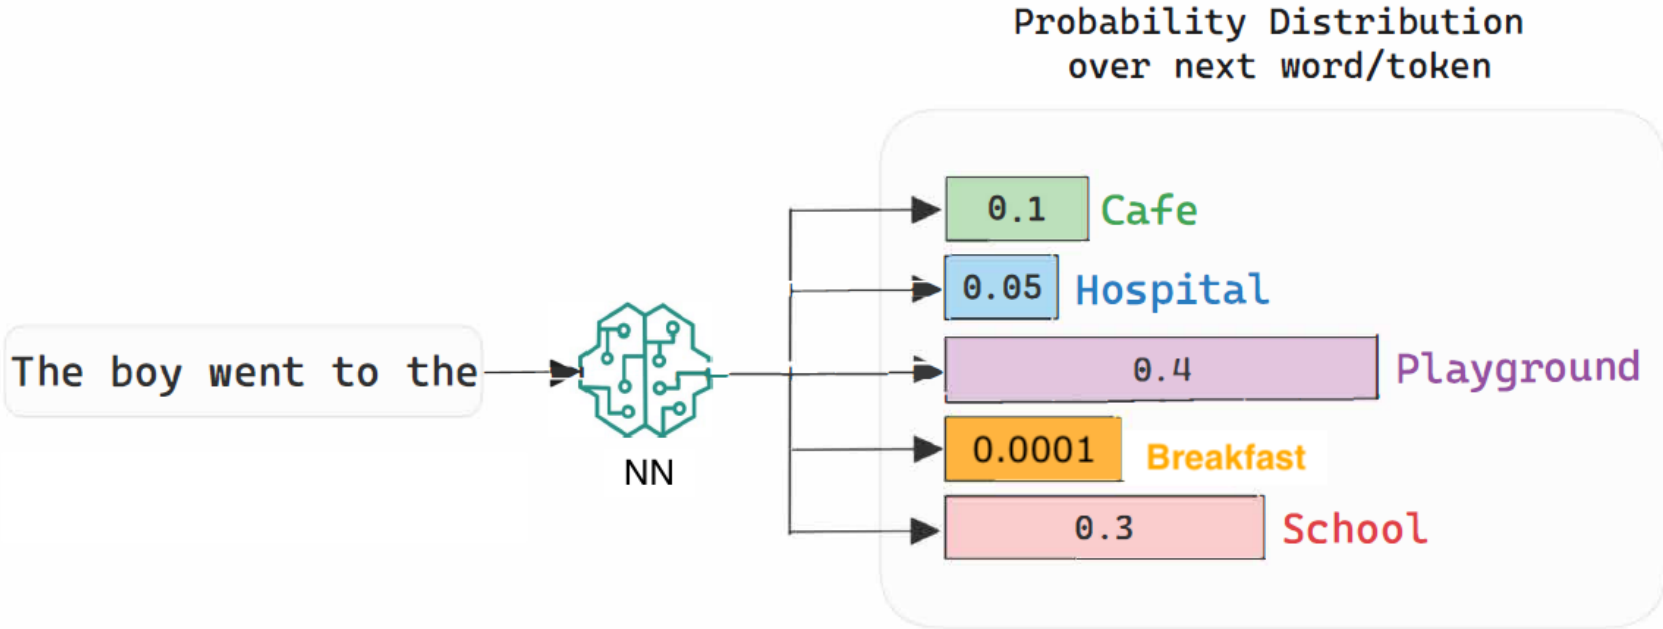

The Rise of Data-driven Artificial Intelligence

Large Language Models

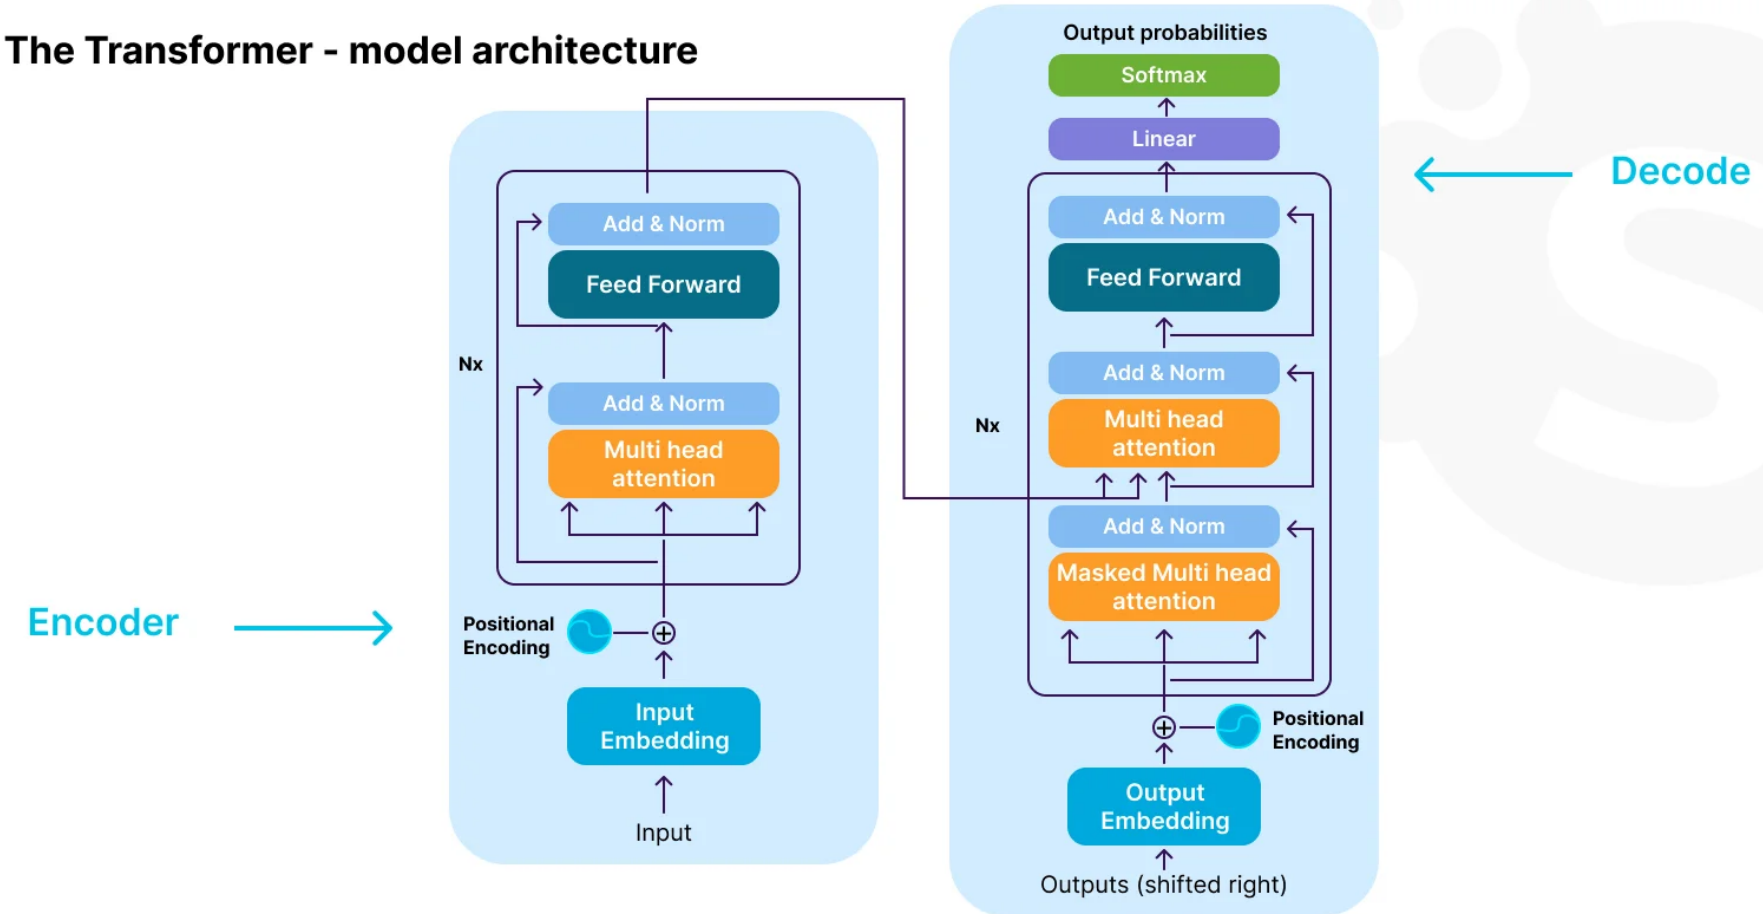

Responsible Artificial Intelligence

The design, development, and deployment of AI systems have far-reaching legal, social, and ethical implications that must be considered closely and carefully.

Virtues that AI systems should possess:

- **Responsibility & Accountability**: decisions/choices are traceable.
- **Value Alignment**: decisions and choices are aligned with the norms and values of user group in mind.
  - **Fairness**: decision and choices are bias free and inclusive.
  - **Privacy**: user data and access to data are protected.
  - **Explainability**: decisions and choices can be perceived, detected, and understood by humans.

Ethic framework for Artificial Intelligence

What should a self-driving car do?
http://moralmachine.mit.edu

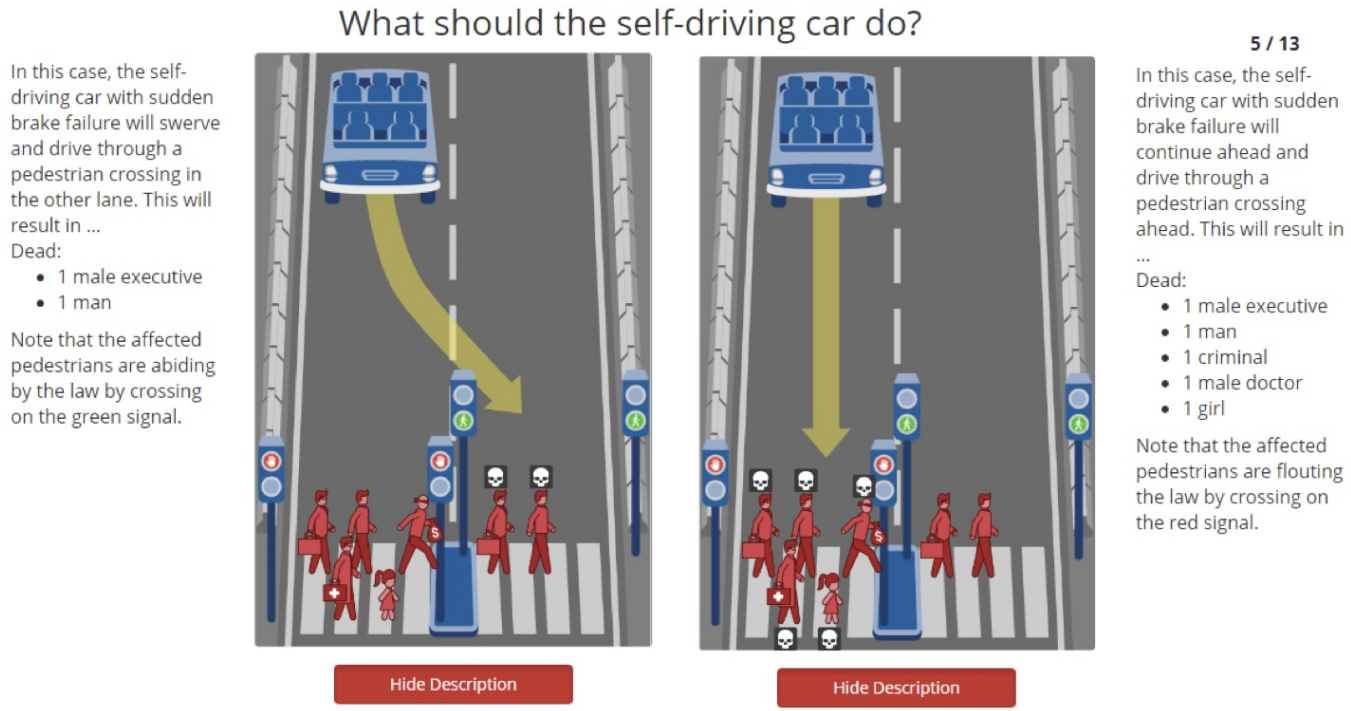

- number of casualties
- kind of casualties

Applications of Multiagent Systems

Causality & Responsibility in Multiagent Environments

NS & ProRail aims at improving the transport capacity of the Dutch railway system by allowing trains to drive closer to each other.

A Covid-19 Epidemic Simulation with Norm-aware Agents

Non-pharmaceutical interventions are used to control the spread of Covid-19. Agent-based simulations is a technology to model and understand the efficacy of these interventions.