Module 1
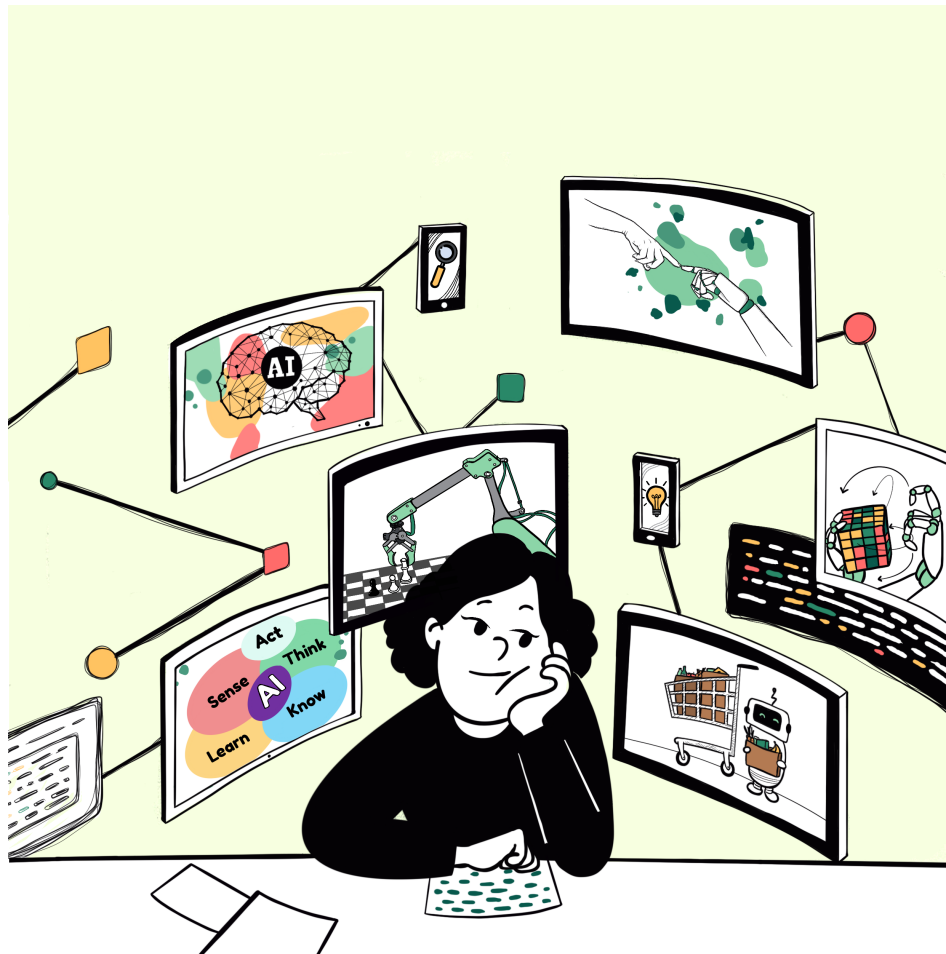
AI Basics

*"The **first thing** anyone has to know about **AI** is that there is **no uniform definition** of **artificial intelligence**!"*

About the Module

This module serves as a basic introduction to the topic of **Artificial Intelligence** (**AI**). The goal of this lesson is to introduce basic *AI* terminology, that will be used inside the other modules as well as in *AI* science in general. More in-depth look at the methods and how they work will then take place in later modules.

While it is not strictly required, it is **recommended** to **start with this module** before introducing the other ones, as it not only helps with the basic terminology, but also provides a good overview of the vast field of *AI*.

Objectives

Students will be able to...

...independently define intelligence for themselves
...reflect to what extent *AI* is built into everyday objects
...understand the concept of algorithms and data
...know different areas of *AI*
...understand possibilities and limitations of *AI*

Agenda

Time	Content
25 min	Theory & Quiz - What is AI?
25 min	Theory & Quiz - Map of AI
15 min	Theory - Possibilities and Limitations
25 min	Theory & Exercise - Algorithms

What is AI?

This module is all about getting a **basic understanding** of intelligence and **artificial intelligence**. It introduces this topic in a reflective way, so students first have to come up with some **definition**, which will then be tested and refined on examples and during discussion. It is based on **slides** and discussions.

Definition
(Slides 2 – 19)

An important thing everyone has to know about *AI* is that there is **no uniform definition** of **artificial intelligence**! The definition can vary widely, depending of the field of research, therefore there are many definitions[1] available. The first thing students are encouraged to do is create their **own** definition of what *AI* (or intelligent behavior) is. Depending on the students this can be either be approached in a discussive way, like using the **think-pair-share** method to help with brainstorming, or in a personal way, where students have to write their own definition down which then will be discussed after looking at some examples. It is also possible to rephrase the question for younger students, so that they should find things a machine should be able to do to be seen as 'intelligent'.

After everyone (or every group) has found a definition, **eight examples** of algorithms / machines are shown and the student have to rate them from 0 (not intelligent) to 5 (very intelligent) using **their definition**. It is ok, even encouraged, for students to **adapt** their definition during or after the exercise. In a final discussion, students can share their definitions and classifications of each example. While, depending on the definition, each of the examples can be more or less intelligent, here are some ideas of how someone could categorize them:

Toaster: 0-1
    While it can be argued, that in more modern toasters some decision making is going on, is usually uses just some temporal / time switches.
Calculator: 1
    Good at a very specific task, relying on very specific input.
Chess playing robot: 3
    While more classical approaches mainly focus on calculation many turns in advanced, modern engines use machine learning to evaluate board positions and learn good moves from human players or using trial and error.

Chatbot: 2-4

Understanding human language an interacting in a meaningful manner can range from simple keyword-based helpline answers to complex digital assistants.

Self driving car: 4

Having a vehicle navigate through human populated spaces requires a lot of understanding of the surrounding environment, still the machines are not really aware of what is going on around them.

Cleaning robot: 2-3

Depending on the model the cleaning logic can range from just driving in random directions to creating a map and evenly drive around everywhere, while reacting to the environment.

Navigation App: 2

Finding the shortest path on a map might be difficult for a person, but using modern search algorithms computers can easily provide a excellent solution even factoring in traffic conditions, using well established search algorithms.

Earth-cleaning robot: 5

Later in this chapter it will become clear, that these capabilities, that are often depicted for intelligent machines in movies, are far more advanced than anything we can create today.

Finally different *AI* definition are shown to provide an idea how other definitions can look like.

Material

- 🖳 AI Basics - What is AI.pdf

References

1. http://agisi.org/Defs_intelligence.html

Map of AI

After the first chapter, where students had to find their own definition, this chapter explores the vast field of *AI* and the many different algorithms and approaches that are connected to it. Therefore, a **Map of AI** is introduced to help understanding that *AI* is comprised of different areas. It is important to note, that this is **just one way** to differentiate between areas in the field of *AI*, not the only one. The relevant part is not these exact five areas, but that students understand, that there simply are a lot of different fields in *AI*.

Map of AI
(Slides 20 – 39)

The map introduces the following areas:

Think

> **Think** is about logic, reaching conclusions, calculating and searching. It includes algorithms for navigation or calculating (searching for) the best action to take next in a game.

Know

> **Know** is all about data and storing data in a useful way. It includes databases as well as general ways to describe things in an orderly manner, like ontologies.

Learn

> **Learn** is about algorithms that are able to modify themselves to better solve specific tasks. This includes learning by trial and error (reinforcement learning) and as learning from examples (supervised learning).

Sense

> **Sense** is about perceiving the environment (objects, temperature, images, …) as well as understanding measurements and correcting errors.

Act

> **Act** includes everything to interact with the environment. It includes motion as well as presenting information in visual or auditorial form (like text and voice generation).

Now that the basic idea of the map is clear, five of the examples from the previous chapter are reintroduced, this time with the intention to find corresponding fields inside the **Map of AI**. Therefore, students have to decide which areas are used for

the system to work. This exercise works well as a guided discussion, but also methods like **think-pair-share** or group discussions work fine.

It is important to note, that the given solutions are only suggestions, most areas can be argued to be connected to any device, depending on the assumptions about the inner workings of the device. For example, a calculator could store formulas internally and therefore be connected to the know-area. The yellow arrows suggest that these areas might be connected, depending on the system. For instance, a chess algorithm can use **reinforcement learning** (especially many newer algorithms make more use of machine learning), but many classical algorithms don't require it.

Material

- 👥 AI Basics - What is AI.pdf
- ✏️ AI Basics - Voting Cards.pdf

Possibilities and Limitations

The next chapter starts by introducing a **video**, the trailer of the movie Wall-E. Student then have to decide, which areas of *AI* are present in this fictional machine. The difference between most fictional and real-world *AI*s is that fictional ones are usually depicted **very human like**, with the ability to freely adapt to **unknown situations** and **self-reflect** on own actions, while our actual *AI* Systems are far away from such capabilities.

Possibilities and Limitations
(Slides 40 - 51)

In general, *AI* systems can be grouped in two categories: **narrow *AI*s** and **general *AI*s**.

Narrow AI
> An intelligent system that is able to efficiently solve a **specific task**. This includes **all** our current *AI* systems!

General AI
> An intelligent system that can understand, reason, reflect and act successfully in unknown environments.

The big takeaway from this last section is, that we currently only have **narrow *AI*s**. Therefore, all our systems, as complex as they might seem from the outside, are only good at solving **specific tasks**. No system currently has a form of awareness or understanding of its own doing. While in the last decades big improvements in **narrow *AI*s**, like image recognition, text translation, self-driving vehicles were seen, improvements for **general *AI*s** were very sparse. The final point in this module is that yes, *AI* improved a lot the last decades, but no it is not nearly as complex as suggested in many movies (like Wall-E).

Material

- 🖳 AI Basics - What is AI.pdf
- 🎥 https://youtu.be/alIq_wG9FNk (Wall-E)

Algorithms

The final chapter introduces the concepts of **algorithms** and **data**, as these terms will occur all over the other modules.

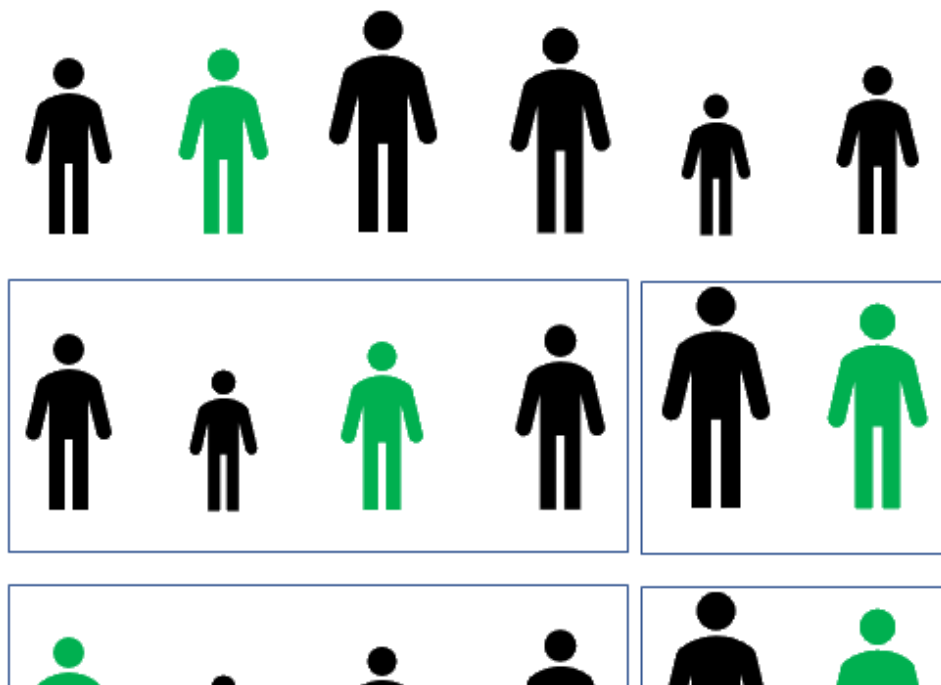## Algorithms
(Slides 52 – 61)

An **algorithm** is just a **list of instructions** that are required to solve a given task. In its basic form, it can be something like a cooking recipe or an assembly guide for a cupboard. More specific in computer science, algorithms often describe methods to modify **data**, like how to efficiently sort a list of names or how to calculate the shortest path through a map.
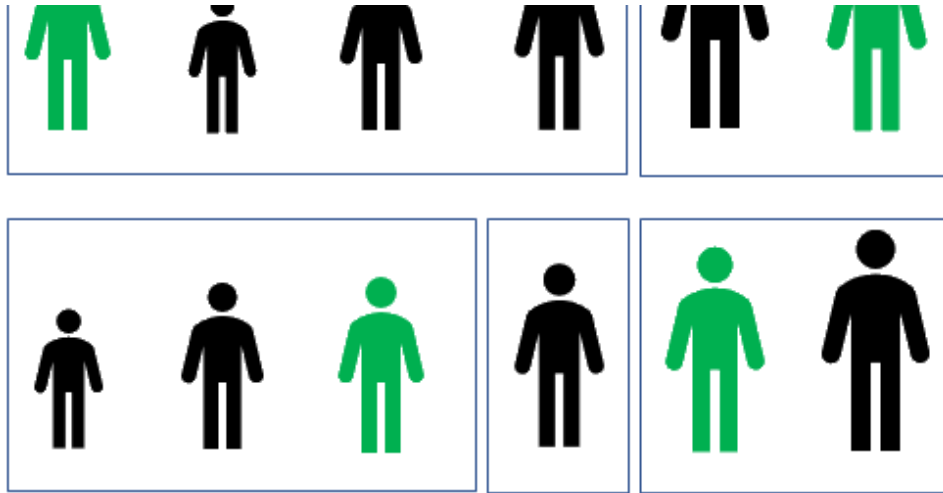
To demonstrate, that algorithms are just lists of instruction that, when followed correctly, lead to the desired result, one of the following two exercises can be used.

Bubble Sort

In this exercise the classical **Bubble Sort** algorithm is used to sort a line of students by height.
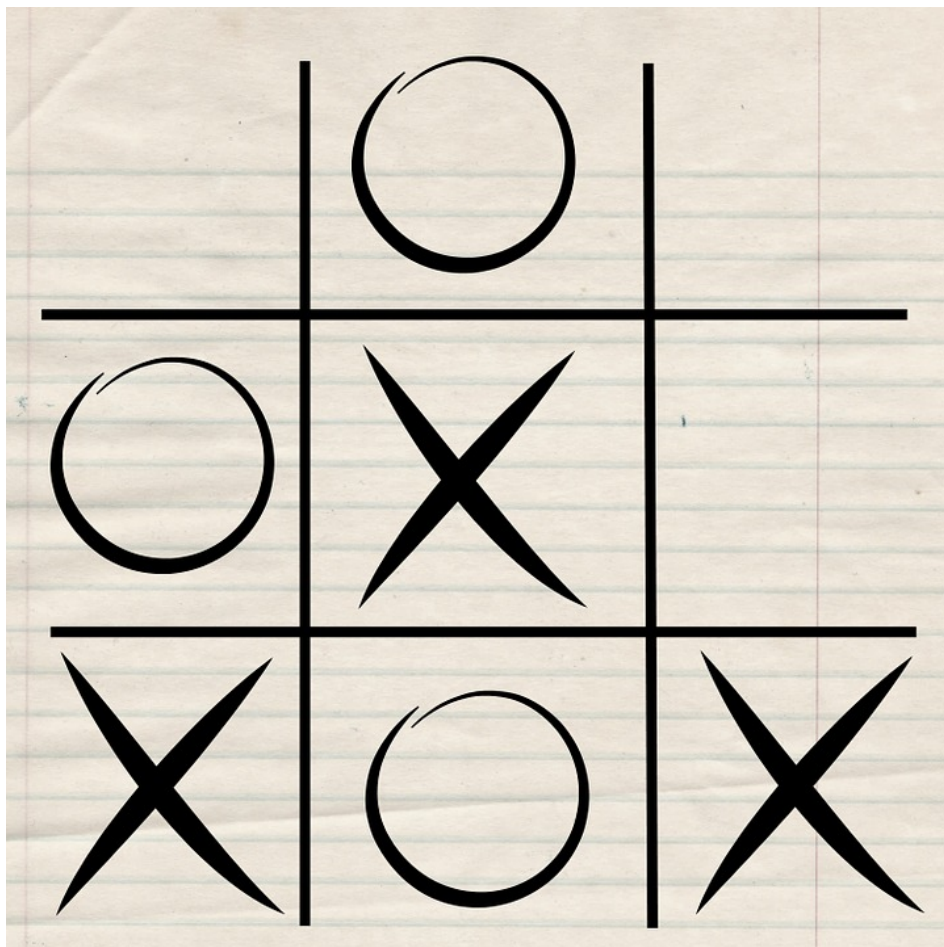
Exercises / Bubble Sort

Tic-Tac-Toe

This exercise introduces a simple algorithm for playing **Tic-Tac-Toe** which the students can try to beat.

Exercises / Tic-Tac-Toe

After trying out one or two of the algorithms above, it is time to let the students create their own.

Create your own algorithm

In this exercise, each student has to create his or her **own** algorithm on how to solve a simple/mundane task like making breakfast, taking out the garbage or even opening a door.
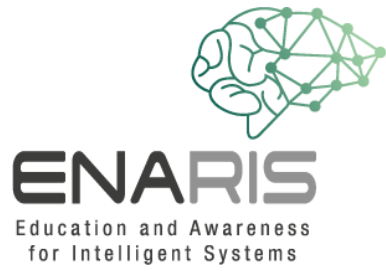
When it comes to algorithms for computers, it is important to be as **specific as possible**. For instance, the simple algorithm: "move forward adn then turn right", conveys no information of how far or with which speed one should move or how far to rotate.

1. Every student needs a piece of paper and something to write
2. Let everyone decide on a **simple** task to write an algorithm about
3. After ten to fifteen minutes let the students **present** their algorithm
4. Discuss how detailed the algorithm was and what could be improved (e.g. can you put something in the trash can when the lid might be closed?)

In the last few slides the concept of **data** is introduced. As data will become very important in the areas of machine learning (Supervised Learning, Reinforcement Learning), it is crucial to understand, that data is just a representation of any kind of information, stored on a digital device. While we **can't** take a person and squish it on a storage device, we **can** store other information, like the height or an image of the person. The **choice of** which **information** we use as data can have a **huge impact** in the final result. Looking at the sorting example, if we had used the pitch of the voice instead of the height of students, they probably would have been much harder to sort, as the pitch is not as easily distinguishable as the height.

Material

- 🖳 AI Basics - What is AI.pdf

EUROPEAN UNION