AI for Children

Artificial Intelligence Curriculum for Elementary and Secondary Schools

Robots' Maze Adventure

Learning from Experience

06

National Pedagogical Institute of the Czech Republic

We create methodologies in cooperation with the National Pedagogical Institute.

Teaching material for Elementary Schools–AI in Computer Science

Learning from Experience – Robots' Maze Adventure

Concept

The cat Kitty has gone missing, so Hoo and Ray set off to find her. An adventurous journey awaits! They discover an enormous maze of tunnels beneath the city–and the secret of where all the missing socks have gone. But how can they find their way through the labyrinth? Is it better to use the left-hand rule or leave a trail of breadcrumbs? One thing's for sure–just like people, robots don't learn only from examples, but also from experience.

Robot Hoo

Hoo is programmed as a curious and slightly unsure robot. He always tries to understand others.
He also collects various human artifacts he finds online–rare memes or old internet trends.
He then shows them to Ray, who sees no value in them.

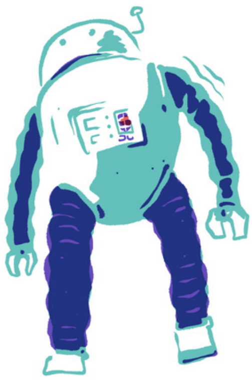

Robot Ray

Ray is programmed for practicality. He constantly looks for ways to process data efficiently.
Human emotions don't interest him –what matters are the numbers.
He always generates fast and accurate responses, though he often takes things too literally.
Ray spends his time building complex mechanical models.

Learning progress map

The Learning Progress Map outlines the key concepts that children should understand during elementary school. The most essential ones are marked in solid blue, while the recommended concepts are shown in white. Each concept is accompanied by a teaching material and a presentation.

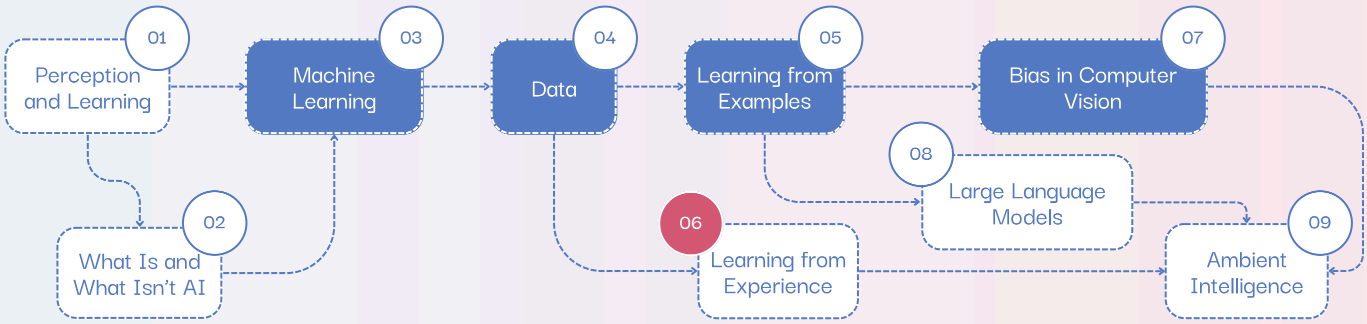

All materials can be found at kurikulum.aidetem.cz/en.

Presentation | Editable template in Canva | Feedback form

Created by: Bára Karpíšková
Concept: Eva Nečasová
Expert guarantors: Cyril Brom, Zbyněk Filipi, Tomáš Mlynář, Pavel Kordík
Artistic design: Jindra Janíček
Proofreading: not yet done
Last update: 01/2025
Version: 02

These teaching materials were translated using ChatGPT.
Please note possible imperfections in the expressions or wording.

Glossary of terms

Artificial Intelligence (AI)
There is no single, universally accepted definition of artificial intelligence.
However, all definitions agree that it refers to a system that simulates human thinking and behavior.
AI usually takes the form of a computer program designed to solve tasks that once required significant human intelligence and were considered the domain of humans (or animals).
It is also a scientific field, with roots dating back to the first half of the 20th century, focused not only on understanding intelligent systems, but above all, on creating them.

Machine Learning (ML)
Just as humans can learn from examples and experience, so too can man-made machines.
Machines use a method called machine learning to learn, which allows artificial intelligence systems to be more than just a set of pre-programmed actions, but to come up with new solutions on their own.

Reinforcement Learning
Reinforcement learning (sometimes referred to as feedback-based learning) is a type of machine learning in which machines are allowed to try things on their own using trial and error. Based on the feedback they receive—through so-called policies—they gradually develop effective behavioral strategies. For example, if we wanted to create a robot capable of navigating a maze, we would reward it for every correct step and penalize it for mistakes or dead ends. Over time, the robot would learn the most efficient way to reach the goal.
To keep things simple, this teaching material refers to this process as learning from experience.

Lesson Overview

Recommended Age, Lesson Length
Children aged 8-11, 45–90 minutes.

Building Blocks
Machine learning from experience (reinforcement learning).

What Are the Students Learning?
Machines, just like humans, can learn from experience. They use trial and error to repeatedly test and discover the best solution for a given task.

Why Are They Learning This?
Understanding the principle of reinforcement learning is an important part of the machine learning puzzle.

How Do We Know They Have Learned It?
In their own words, students will describe a specific strategy that led them to success in solving the problem.

Tools
Teacher: Projection equipment and presentation materials.
Students: Writing supplies, worksheets.

Digital Competence
Facilitating Learners' Digital Competence.

Bloom's Taxonomy
Applying: Students apply maze-navigation strategies and trial-and-error to solve the problem.
Analyzing: They analyze the success of various strategies and identify the reasons behind their effectiveness or ineffectiveness.
Evaluating: They assess the pros and cons of the strategies used and reflect on the experience gained.

Five Big Ideas
3-A-I Nature of Learning (Humans vs. machines).
3-A-VI Nature of Learning (Learning from experience).

Note: Gender equality is key for AI for Children, but for brevity we use masculine formulations in our methodologies.

Engage

min 00

Presentation slide 01

Read this part of the story to the students.

Hoo and Ray were sitting at home, collecting all the data they needed before visiting Carl, the caretaker and cat enthusiast. Suddenly, Hoo picked up a signal on their sound sensors…

"Ray, don't you think it's a bit too quiet?" Hoo listened carefully—and yes, the silence was almost deafening. There was no meowing, which they had gotten used to hearing behind the bathroom door. Soon they noticed the bathroom light was on. It looked like Kitty had given up on socks and headed for the underground maze of tunnels and pipes beneath the town.

Without hesitation, Hoo and Ray grabbed the strongest flashlight they had (the one they usually used for reading robot manuals at night), and headed straight into the maze. "Look out, Ray, tunnels starting here! Left or right?" Hoo asked at each turn, but just as expected, the tunnel walls didn't answer. It took them a while to figure out the best strategy for navigating the labyrinth. Hoo believed the best way was to always stick to the left wall, while Ray was sure right was the way to go. They agreed to experiment and alternate—going left once and right the next—to see which strategy worked better.

Have you ever been in a maze? If yes, how did it go? What did you enjoy about it and what didn't you like?
Teacher tip: Encourage students to describe not only what happened, but also how they felt. You can help them by suggesting words that express emotions and by guiding them to think about what exactly made them feel frustrated or happy.

Understand

First activity

min 05

Presentation slide 02

Sticky Note Method
Students write their answers on sticky notes (individually or in pairs) and then group the notes based on similarities.

Imagine you're standing in front of a maze and you have to find your way out. What will you do? Will you try the first path you see, or will you come up with a strategy? Write down at least one idea on a piece of paper.

Here are some strategies you can suggest to the students:

The right-hand (or left-hand) rule: A simple method (although it doesn't work in all types of mazes), where you choose one side and stick to it the entire time.

Backtracking: If you hit a dead end, you go back to the last intersection and try a different path.

Breadcrumb trail: You leave crumbs or small objects in places you've already been so that if you return, you'll know you've been there before.

What strategy do you think Hoo and Ray will choose? And how well will it work for them?

Read the story.

In one of their first attempts, by keeping to the left-hand side, they reached a little alcove–right where Kitty was napping. "Ray," cheered Hoo, "see? The left side led us straight to her!" But Ray, still confident in the right-hand rule, suggested they go back and try the right side too, just to see if it might be faster. They let Kitty nap and continued their experiment.

After another round–this time sticking strictly to the right–they ended up at Kitty again, who had woken up by then and was calmly watching them continue their determined testing. "The right side got us here too," said Ray. "We have to try again! This time, let's move completely at random and track the distance and time," Ray decided.

When Hoo and Ray tried their 181st path, they finally realized which method worked best–and that Kitty was no longer waiting for them in her usual cozy nook. Their antennas started to twitch. "Maybe she stuck with the left wall," said Hoo, as they successfully returned home with a bag full of lost socks that Kitty had probably dropped there on her underground adventures.

Second activity

Presentation slide 03

min 15

Just like Hoo and Ray, you'll be making your way through a maze.

Your task is to navigate a printable underground labyrinth (on your worksheet) and find the lost Kitty. Along the way, you'll come across various socks that have mysteriously been scattered throughout the maze. Your goal is to collect as many socks as possible while trying to find a path to our furry friend.

But it won't be that easy–there are also socksnatching mouse lurking in the tunnels. Each time you pass them, they'll steal one of your socks–if you have any. Once they take a sock, it stays in place, so the same mouse might steal from you again. However, if you manage to collect a sock, it stays with you and can't be collected twice.

And remember–the best path is the one that gets the most socks back to Kitty.

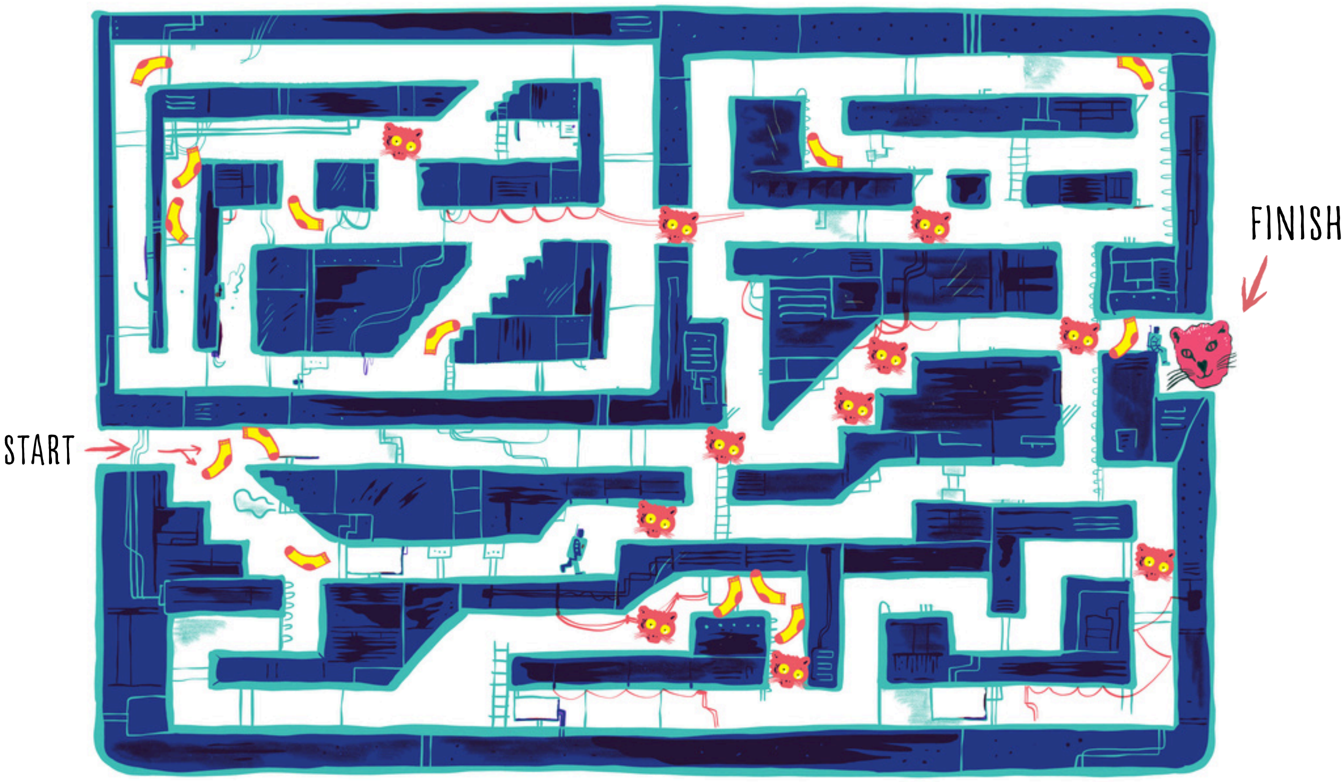

Discuss with the students.

min 25

How many socks did you manage to collect? And how many were stolen by the socksnatching mouse?
The maximum number of socks students can collect before reaching Kitty is 8.

Did you use a specific strategy to move through the maze and collect socks?
Students share their strategies.

In the last lesson, we saw that robots and computers can learn from different kinds of examples. Do you remember what some of those examples were?
Answer: They included various types of data, like pictures, videos, sounds, texts, tables… even 3D objects.

Explain:
But computers and robots can learn in a different way—through experience. They try out different things, and people give them feedback. This is simply trial-and-error learning. If something doesn't work, people let them know, and they won't repeat it. On the other hand, if something works well, people give them positive feedback, and computers and robots are much more likely to repeat that action next time.

It's just like when the robots were moving through the maze. They received positive feedback in the form of socks, and negative feedback came when the mischievous socksnatching mouse stole them again.

YouTube

Presentation slide 04

min 30

Show students the video.
Link: youtube.com/watch?v=pJPdW8WWAso (11:12)
Video explanation: The video shows five AI agents (3D robot objects) learning to run 100 meters. Whoever finishes first wins a cupcake. The agents are learning purely through trial and error. The program tells them whether something worked or not. Based on that, the agents gradually develop their own strategies—they learn how to "run better."

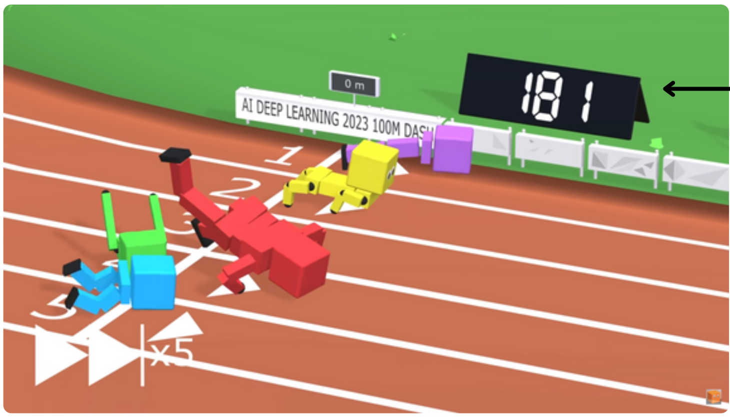

Number of attempts

The video is 11 minutes long – feel free to show just a few short clips. For example:
Beginning: Explain what the video is about, what the goal is, and show the number of attempts.
Attempt 204: The green robot takes its first steps. You can have students guess which agent will finish first.
Attempt 390: The red one is moving but unsure what to do, and the purple one decides to cheat.
Attempt 738: The robots are improving but still struggling with balance.
Attempt 813: The purple one is starting to get good!
Attempt 954: The red one almost manages to stay upright… but then…
Attempt 1638: The red one runs and we watch with suspense as it approaches the finish line… And it makes it!

Reflect

min 40

Read the story.

Back in the kitchen, they found Carl calmly sitting at the table with Kitty, who had clearly made it home safely–apparently quite a bit earlier than the robots. "So, guys, which path turned out to be the best one?" Carl asked after the robots shared the details of their adventure.

"We found her, but we realized we really have to try out a lot of different paths–and when we make a mistake, we need to learn from it and try not to repeat it. But wow, it takes a long time!" Ray concluded wisely. "Looks like Kitty figured it out faster than you did," Carl chuckled. "And remember–the journey is the goal."

It took the robots a long time to find the best path. Why do you think that was?
Answer: Because they tried to reach Kitty using completely random paths. If they made a mistake, it took longer to fix it. And since the maze was complex (with many paths), it took them quite a while to find the best one.

And how might Kitty or a mouse find their way through the maze?
Possible answer: Maybe by using smell, hearing, or some sort of inner map–especially since cats and mice may have a different sense of space than humans do. Mice, for example, might rely on smell trails or whiskers to feel along the walls.

So robots (machines) can learn not just from examples, but also from experience. Do you think if these robots had to go through a different maze, like under a house or in a city, they'd find the best path again?
Answer: They would probably use what they learned, but because the new maze is different, they would have to test things again and keep learning.

If there's extra time

Hexapawn

Presentation slide 05

Children can play the game Hexapawn.
Address: mrozilla.cz/lab/hexapawn

Hexapawn is a simplified version of chess, created for educational purposes in the field of artificial intelligence. It is played on a 3×3 chessboard, and each player starts with three pawns. The goal is to get one of your pawns to the opposite side of the board. Pawns move like in chess: they can move forward by one square if the space is empty, or capture an opponent's pawn diagonally. The player who first reaches the far end of the board with a pawn, defeats all of the opponent's pawns, or blocks the opponent from making any legal move, wins the game.

The developer Mrozilla offers a program on his website that uses Hexapawn to demonstrate the principle of machine learning from experience. At first, the program (artificial intelligence) makes random moves in response to the player's actions. These moves are initially completely random, but if a move leads to a loss, the program remembers it and avoids repeating it (unless the page is reloaded).
By learning from experience, the program quickly becomes unbeatable.

This way or that way! After Kitty and the socks!

Explore the mysterious maze of underground tunnels, find Kitty but more importantly, choose the path that lets you collect the most socks. To make things trickier, there are also socksnatching mice hiding in the corridors. Every time you pass one, it'll steal a sock–if you're carrying any, of course.

Kitty

sock

socksnatching mouse