AI for Children

Artificial Intelligence Curriculum for Elementary and Secondary Schools

Hoo and Ray Go Shopping

Bias in Computer Vision

07

npi | 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

Bias in Computer Vision – Hoo and Ray Go Shopping

Concept

Hoo and Ray Are off on another adventure! This time, their friend Carl gives them an important task: to bring apples and a pineapple for his big dinner party. Just to be safe, Carl lends them two expert books full of information about orchards and tropical plants. What could possibly go wrong? Well – quite a lot, as it turns out! Hoo and Ray expect to find an entire apple orchard or a bush loaded with pineapples right in the store. In the end, they realize the problem isn't with the store's selection – it's with the mental picture they had of what an apple or a pineapple is.

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.

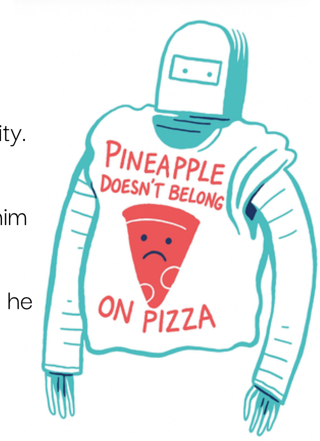

Learning progress map

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.

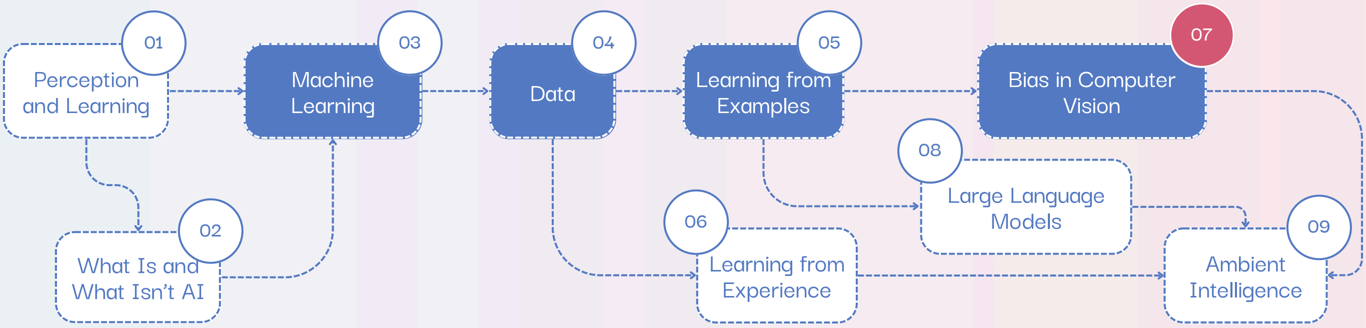

All materials can be found at kurikulum.aidetem.cz/en.

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: 03

Presentation | Editable template in Canva | Feedback form

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.

Computer Vision (CV)
This field began developing even before AI became widely known – mainly for the purpose of recognizing patterns (similarities) in images. The arrival of artificial neural networks brought major progress, and now computers can "see" in ways quite similar to humans. If a system is given enough image data, it can learn to recognize people, objects, or even pets, based on known reference points. Today, we can log into our phones using facial recognition (Face ID), measure distances and identify objects we don't know (Google Lens), or even identify bird species in the forest (BirdNET). Thanks to image recognition, AI can also detect traffic signs, lanes, and obstacles – a key part of self-driving car technology.

Machine Learning (ML)
Just as humans can learn from examples and experiences, so can human-made machines.
Machines use a method called machine learning, which enables AI systems to go beyond simply following pre-programmed instructions and instead come up with new solutions on their own.

Bias
AI systems learn from data that is prepared by humans. Poorly designed, unbalanced, or insufficient data can cause the AI to develop certain biases – or distorted views (in English, this is called "bias"). For instance, if training data reflects long-standing human prejudices (like race or gender), these patterns may show up in the AI's behavior too – for example, it might recommend men more often than women for leadership positions just because that's what it "saw" in past data. Bias can also appear in how well a system recognizes faces: for example, a face recognition system trained mostly on people from Europe might not work well in Africa. If an AI system plays a role in important decisions, we need to be sure it's accurate and fair. That's why it's crucial for AI to be trained on the right data.

Lesson Overview

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

Building Blocks
Bias.

What Are the Students Learning?
If we provide a system with poor data, it will learn to recognize things incorrectly – just like a human would.

Why Are They Learning This?
To help students critically evaluate how AI systems work, based on an understanding of bias.

How Do We Know They Have Learned It?
Students prepare data to train a machine learning model and then test whether the model can correctly recognize different things using that data.

Tools
Teacher: Projector and screen or presentation device, camera.
Students: Supplies for creative work, group activity materials.

Digital Competence
Communication and Collaboration.

Bloom's Taxonomy
Applying: Students prepare data to train a machine learning model.
Analyzing: Students analyze the quality and variety of the training data and how it affects model accuracy.
Evaluating: Students evaluate how well the machine learning model performs.

Five Big Ideas
2-A-IV Representation (Feature vectors).
2-C-II Reasoning (Reasoning algorithms).
3-A-VI Nature of Learning (Learning from experience).
3-C-III Datasets (Bias).

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 part of the story to the students.
Hi, kids! This is Hoo and Ray, your masters of searching, learning and… getting lost! This time we went on a super expedition to the supermarket to find apples and pineapples for Carl. We had books full of information, but instead of coming back with fruit, we discovered a lot of things that even Carl had no idea could be bought! Maybe you could help us solve this fruit mystery. We can't wait – because as they say: the more minds, the more we know… and the more we dream!

Imagine you have two friends: one wears glasses and the other doesn't. If you were to imagine one of them holding a book, which one would you picture? Why?
Teacher tip: Ask children why they think the friend with glasses is more connected to books than the one without. Encourage them to consider whether glasses really say anything about what a person likes to do. Encourage them to think about how their idea might have been formed and what would happen if both friends liked reading equally. Children can also think of other examples of bias.

Understand

First activity

min 05

Presentation slide 02–03

Show the students the presentation on slides 02 and 03.
On slide 02 there are drawings of different apples and on slide 03 there are drawings of different forms of pineapple.

Imagine that you want a robot to recognize apples in pictures. First, you have to show it many pictures of apples so that it can find similarities (for example, round shape, color, stem…). Now try to imagine as many types and variations of apples as possible. You would show all these apples to the application.
But wait! Did you think of not only every type of apple, but also apples from all angles, in different colors and with different textures (smooth or speckled)? What about older apples, the kind you wouldn't eat anymore? Or apples with flaws – maybe a bite taken out, or a wormhole? Apples without their peel? Or ones that are sliced in half or chopped into wedges? Artificial intelligence doesn't see apples the way we do. It can't smell or touch them. If you don't show it all the possible variations it might encounter in the real world, it may not recognize an apple when it looks different That's called bias.

min 10

Presentation slide 04–05

Remind students that computers (robots… machines) need lots of examples of something in order to learn what it is. We call those examples training data. When the training data is poorly chosen, the AI ends up recognizing things incorrectly.
Just like in the picture on slides 04–05 – if the computer sees many examples of salmon, but only in sushi and never as a fish, it might generate an image of salmon that's completely wrong. See the **presentation on slide 05** for an example.

Optional warm-up activity
When you hear the word "apple," what ideas, images, or emotions come to mind? Write your thoughts on a small slip of paper and then compare with others: do you all imagine apples the same way? What does an apple look like to you?

min 15

Read the story.

Carl decided to invite the robots over for one of his perfectly imperfect human dinners. He was a lifelong bachelor, and his cooking reflected that. His specialties were apple pie and Hawaiian pizza. Carl was a firm believer in the idea that pineapple definitely belongs on pizza.

But Carl wasn't exactly the best host. When the robots arrived, it turned out he didn't have any apples for the pie – or any pineapple for the pizza. So he asked Hoo and Ray to run to the supermarket and grab the fruit. From previous experiences, however, he knew that robots needed plenty of data (examples) to complete a task. So he sent them off with two books: one called Orchards and Groves, to help them reliably find apples, and the other Growing Tropical Fruit, so they could recognize a pineapple. He hoped the books would help them find everything they needed in the store.

Do you think the robots succeeded? What might have gone wrong?

The robots set off for the supermarket, books in hand, determined to complete their mission. But it didn't take long before they ran into their first problem. Nowhere in the store could they find an orchard section or a pineapple grove, like the ones they'd seen in the books. When they asked the shop assistants, they were met with puzzled looks and shaking heads. So the robots wandered from aisle to aisle, passing right by large crates full of apples and pineapples – because those didn't look anything like the trees and plants pictured in their books. When they couldn't tell by sight or smell, Hoo and Ray simply picked what looked most like the pictures in the books.

What do you think they picked? Remember: one of the key things robots need to learn is that apples grow on trees, and pineapples are surrounded by lots of green leaves.

Presentation slide 06

min 20

When Carl saw Hoo and Ray standing in the doorway holding a plastic Christmas tree with red plastic ornaments, he realized they had misunderstood the apple situation yet again. And the pineapple? That didn't go any better. But Carl was a flexible man. He knew the robots wouldn't starve, so he decided to make the best of it – instead of his fancy dinner party, he threw a holiday-themed celebration.

Second activity

Presentation slides 07 and 08

Just like Hoo and Ray, students will now explore how AI learns to recognize apples and pineapples.

In this task, students will train a so-called machine learning model. A **machine learning model** is, among other things, a program that can be taught to recognize different things by showing it examples of them.

To begin, you can demonstrate how the model works using the Teachable Machine app. To test the model, you'll need test images, which you can find here.

Divide the students into two groups.
The students' task now is to create as many examples as possible for two different categories (classes) to help train a machine learning model in Teachable Machine. They can draw, use modeling clay, take photos – anything they like – to show different examples. These don't have to be apples or pineapples; they can be anything the students choose. Each group will choose one of the two categories, and each student will create one example. Most of these examples will be used for the training set – these are the ones you'll upload into Teachable Machine as part of the classes. A small number of examples should be selected as the test set – these are the ones you'll later use to test the model in the preview window on the right.

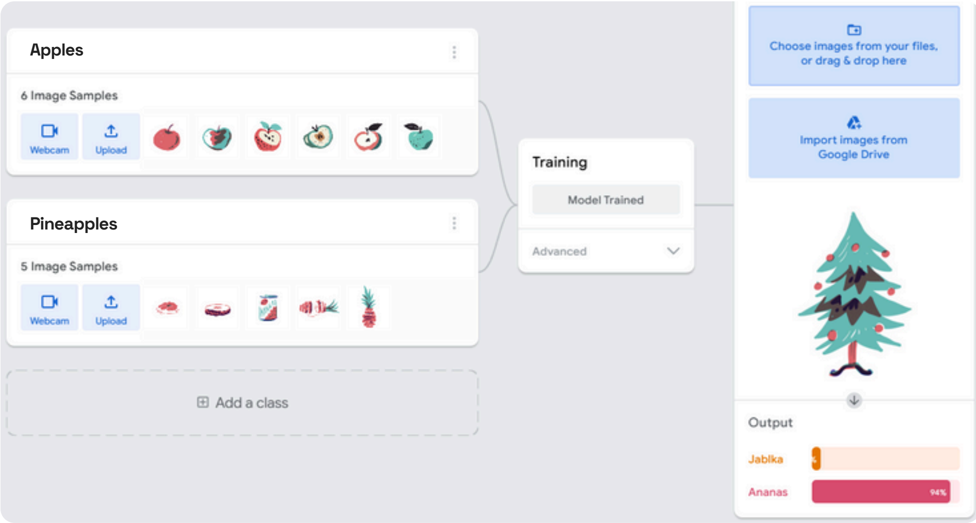

Reflect

min 40

Read the story.

After a long day of wandering through the shelves, looking for trees and bushes, and comparing fruits to pictures in our books, we ended up with something that really doesn't belong on a pizza. We figured out that even if we have information from books, sometimes it's not enough. We need a lot of it, and we also need to see, feel, and experience things with our own sensors – that is, on our own metal! So next time we'll be smarter and maybe bring someone with us who can recognize fruits better than us. Thanks for helping us today, and we're looking forward to another adventure – and this time, hopefully without Christmas decorations... Although Carl liked them quite a bit.

Why do you think robots couldn't find apples and pineapples in the supermarket even though they had books full of pictures?
The books only provided examples of plants in their natural environment, which may not correspond to how we see fruit in the supermarket. Robots learn from what they have available to them, and if they saw apple trees and pineapple trees in a book, it is not surprising that they did not recognize apples in a crate or sliced pineapple. You can lead children to think that robots do not use creativity or imagination, and therefore have difficulty comparing reality with images that do not correspond to what they have learned from books.

What data would help robots recognize things better? What should it look like?
Students can suggest what better data might be: a variety of photos from different angles, close-ups, videos, or info about object size and surface texture.

What if we gave the robots only drawings of apples and pineapples – would they still recognize them in photos?
Probably not – because robots (computers, machines) don't recognize things like humans do. They need very specific examples, which is why it's so hard to teach them to recognize things well.

More to explore

Test the model with real photographs.
Children can also take photos of real objects and use them to train a model in Teachable Machine.