



AI for Children

Artificial Intelligence Curriculum for Elementary and Secondary Schools

Card

Ethics in Decision Making

02

AI Ethics Deck



<https://kurikulum.aidetem.cz/cards>

Prepared by: Radek Špáta & Eva Nečasová
Methodological consultant: Peťa Dovhunová
Expert guarantors: Tomáš Mlynář, Pavel Kordík
Language proofreading: not yet done
Last update: 01/2025
Version: 03

npi | National Pedagogical Institute
of the Czech Republic

We create methodologies in cooperation with the National Pedagogical Institute.



[Form for
comments](#)

Teaching Material for the AI Curriculum for Elementary and Secondary Schools
Computer Science at Secondary Schools - Cards

Ethics in Decision Making

What is AI ethics?

AI ethics is a relatively young academic field, but its importance is growing rapidly. One key reason is that machines are taking over more and more of our tasks and decisions. Often still under human supervision, but in many narrow areas, machines have already surpassed us. And according to the plans of many tech giants, the near future will include fully autonomous systems. A clear example is self-driving vehicles. They're not yet common on the roads, but that may soon change. These vehicles are designed to take us anywhere without the need for a human operator. Beyond the technical challenges, this also means they'll need to follow ethical guidelines to make the right decisions. At the heart of AI ethics is the effort to create a set of ethical principles that help AI systems make decisions aligned with human values.

Ethics in decision making: The case of self-driving cars

We have high hopes for autonomous vehicles—among other things, they promise to make traffic significantly safer. But the laws of physics can't be bypassed, and even these sophisticated machines will inevitably face situations where any decision leads to some form of harm. Unlike humans, who often react instinctively, a computer may have just enough time to consider a maneuver. Should the car veer right, where a pedestrian stands, or left, where there's a tree? How should it decide? Neither of the common ethical principles offers a clear answer. Turning right endangers the pedestrian; turning left risks the passenger.

One idea might be to program the vehicle to avoid making any steering decision in such cases—simply brake and hold its course. But that feels a bit like avoiding responsibility. We could claim the vehicle didn't decide anything, just braked. Yet choosing not to decide is a decision too. Imagine the car skids toward a crosswalk where two children are standing. To the left, another car is approaching; to the right, a pedestrian walks on the sidewalk. If the vehicle chooses to brake without turning and fails to stop in time, can we really say it didn't make a choice? It had three options: veer left, veer right, or continue straight while braking. If it hits the children, that outcome is still a result of a choice—not to steer. And clearly, that contradicts the principle of doing no harm.

AI Ethics Deck

Ethics in Technology Use and Development
Ethics in Decision Making



[Lesson presentation
in PDF](#)



[Editable presentation
in Canva](#)

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

Note 2: Gender equality is a key value for AI for Children, but to keep our teaching materials concise, we use masculine grammatical forms.

Classroom Activity

Ethics in Decision Making

30–45
min

Activity description

Ethical dilemmas in decision making can be clearly illustrated using the so-called Trolley Problem. This thought experiment explores how to choose between two harmful outcomes—for example, whether to sacrifice one person to save several others.

The Trolley Problem comes in many versions. One classic scenario goes like this: A trolley is speeding down a track where five people are tied up and unable to move. You're standing by a switch. If you pull the lever, the trolley will be diverted onto another track—where one person is tied up. The question is: Is it right to pull the lever and sacrifice one to save five?

This dilemma is often discussed in the context of self-driving cars, which may face similarly difficult choices. In the following activity, students explore various scenarios, choose the ones they believe are more ethically acceptable, and explain their reasoning.


Moral Machine – Research from MIT

In 2014, the MIT Media Lab launched an experiment called Moral Machine—an interactive platform designed to explore how people perceive moral decisions made by AI systems. The setup is simple: users are shown two scenarios in which a self-driving car must “choose” between two harmful outcomes—such as killing two passengers or five pedestrians—and they select the option they find more acceptable.

Since its launch, the experiment has seen massive participation: millions of people from hundreds of countries have made millions of decisions. This makes Moral Machine one of the largest studies ever conducted on global moral preferences. In 2018, the journal *Nature* published a report based on this data, analyzing how moral choices vary by culture, economy, and geography. Some key findings include:

1. Participants from collectivist cultures (e.g., China and Japan) were less likely to prioritize saving younger over older people.
2. People in lower-income countries were more forgiving of pedestrians crossing illegally.
3. In countries with high economic inequality, choices varied more depending on the social status of those involved.
4. Participants from individualistic cultures (e.g., the UK and the US) placed stronger emphasis on saving the most lives overall (the principle of minimizing harm).

You can explore these findings on the Moral Machine results explorer by visiting [this link](#). A world map will appear – click on any country to explore its moral preferences. A world map will appear, and by clicking on individual countries, you can view a comparison between (a maximum of) two selected countries and the global average. For example:



In this teaching material, students analyze and discuss moral dilemmas.

We've prepared two tools to support the activity itself, plus one follow-up resource:



If you plan to use this lesson at the secondary school level, we recommend using the kid-friendly version of the Moral Machine we created, available at: bit.ly/moralni-stroj.

In this version, the characters are not people but animals. It includes five scenarios that always appear in the same sequence.



For upper secondary schools, we recommend using the original Moral Machine app at moralmachine.net. On the homepage, click the "Start Judging" button. The scenarios will appear in a random order.



If you have extra time, you can follow up with the Absurd Trolley Problems app at neal.fun/absurd-trolley-problems. It features 30 different trolley problem scenarios that work great as a starting point for discussion.

Lesson Overview

Recommended Age, Lesson Length

Children aged 14–19, 45 minutes.

Building Blocks

Ethics of artificial intelligence, intelligent systems.

What Are the Students Learning?

Artificial intelligence systems make decisions and the goal is for them to do so in harmony with people.

Why Are They Learning This?

They approach technology responsibly, based on an understanding of how ethical decisions affect daily life and society.

How Do We Know They Have Learned It?

They will explain the ethical dilemmas associated with decision-making by artificial intelligence systems. They will propose possible solutions and defend their decisions.

Tools

Teacher: Projection equipment, presentation to be shown.
Students: Writing aids, worksheet.

Digital Competence

Facilitating Learners' Digital Competence.

Bloom's Taxonomy

Analysis: Students analyze specific examples of ethical dilemmas.
Evaluation: They will assess the correctness of their decisions.

Five Big Ideas

5-A-I Ethical AI (Diversity of Interests and Disparate Impacts).

Engage

10
min

Discuss

What comes to mind when you hear “self-driving car”?

A self-driving car is a vehicle that can operate without human input. It uses sensors, cameras, radars, and artificial intelligence to assess road conditions in real time, plan routes, and respond to obstacles. The goal is to ensure smooth and safe driving—without needing a driver.

Do you think this technology is already advanced enough to be safe for everyday use?

Autonomous vehicle technology is highly advanced, but not yet flawless. There are still situations—such as extreme weather, unusual traffic conditions, or technical glitches—where failures can occur. That’s why self-driving cars are still in the testing phase and are used mainly in limited ways, typically on specific routes or in areas with favorable conditions.

Presentation slide 02

Show the video to the students.

Video address: bit.ly/4ejI51p (3:40).



[Perception, decision making and movement](#)

The video introduces Waymo’s advanced self-driving technology and explains how sensor-equipped vehicles—using LiDAR, cameras, and radar—navigate through city environments. The technology is broken down into three key steps: perception (sensors collect data), decision-making (AI and machine learning process the data and predict the behavior of nearby objects), and movement (the vehicle drives safely and smoothly).

Discuss

What was this video about?

See paragraph above.

Understand


Describe

Presentation slide 03

Show students the presentation on slide 03 and describe the image together.

The illustration shows a diagram of a self-driving car with various sensors that make autonomous driving possible. The data collected by each sensor must be processed in real time using onboard computers and advanced AI algorithms within the vehicle itself.

Key functions of each sensor:



What role do sensors (LiDAR, camera, radar) play in how a self-driving car perceives its surroundings and makes decisions?

Sensors are essential to how a self-driving car perceives its environment:

LiDAR emits laser pulses that bounce off surrounding objects. By measuring how long it takes for each pulse to return, the system calculates distances and creates a detailed 3D map of the surroundings, which the car uses for navigation and decision-making.

Cameras capture visual details like colors, traffic signs, traffic lights, and pedestrians.

Radar detects the distance and speed of moving objects, such as nearby vehicles, and works reliably even in poor visibility conditions.

Together, these sensors provide a comprehensive picture of the environment. Artificial intelligence processes this data, predicts how nearby objects (like cars or pedestrians) will behave, and decides how the vehicle should respond to ensure smooth and safe driving.

Think, pair, share

Think, pair, share

Think: Students first reflect on the topic individually.

Pair: They then team up in pairs to discuss their thoughts.

Share: After a set time, selected pairs share their ideas with the whole class.

Imagine you're riding a bike (or dancing, playing football...). What do you need to pay attention to? How do you decide what to do next?

Possible answer: When I'm riding a bike, I have to pay attention to a lot of things at once. I watch what's happening ahead—whether there's an obstacle, a car, or a pedestrian. I also look at traffic lights and road signs. At the same time, I listen for cars behind me or if someone's honking. When turning, I feel the road conditions—whether it's wet or if the surface is gravel or asphalt—and adjust my riding so I don't slip. I'm constantly deciding when to brake, speed up, or swerve to avoid something. Most of this happens without thinking, but I still need to stay alert the whole time.



Presentation slide 04 or 05


Students split into pairs or small groups and work with devices and worksheets.

Introduce them to the Moral Machine research project, which explores how self-driving cars should make decisions in ethically challenging situations. Explain that they'll be analyzing scenarios where an autonomous car must choose between two outcomes—such as protecting passengers or pedestrians.

Go over the icons and symbols used in the visuals: pedestrians, passengers, age groups, social status, traffic conditions, and traffic lights. Emphasize that there are no right or wrong answers. Their task is to reflect on what they personally consider fair or just. They should be ready to explain and defend their decisions during the discussion. Encourage them to focus on the values they see as important—like protecting life or following rules.

Students often want to explore alternative scenarios (e.g., whether the car could slow down). Remind them to stick strictly to the two options presented.

Then, depending on your choice, have students access either the kid-friendly Moral Machine version (QR code on slide 04) or the original [MIT Moral Machine](#) (QR code on slide 05).



Presentation slide 06 and 07

Hand out the worksheets to the students.

There are two versions—one for the kid-friendly Moral Machine tool (page 08 of this teaching material, slide 06 in the presentation), and the other for the MIT Moral Machine (page 09 of the teaching material, slide 07 in the presentation).

Students' task is to go through three to five scenarios in the app, describe each situation using the given structure, choose the solution they find most acceptable, and explain the reasons behind their choice.

Reflect

15 min



Presentation slide 08

Discuss the worksheets.

Students present their decisions in groups and explain the reasoning behind their choices. You can use some of the following reflective questions to help them think more deeply about their decision-making process and consider which moral principles influenced their choices.

What values or principles guided your decision?

To reflect on: Did you aim to protect the majority? Did you consider the vulnerability of certain individuals?

What dilemma was the most difficult for you and why?

To reflect on: What was confusing about the situation? What did you struggle with the most?

Would you decide differently if the scenario were slightly changed? If so, how?

To reflect on: How would changing the conditions affect your decision?

Which decision do you believe was the fairest, and why?

To reflect on: How would you define fairness in this context?

If you were in the position of someone affected by the decision, how would you feel?

To reflect on: Can you put yourself in the shoes of those impacted by the decision?

Do you think an autonomous system (e.g., a car) should make decisions the same way a human would? Why or why not?

To reflect on: What do you think should matter most in AI decision-making – safety, fairness, or speed of response? Which aspect should take priority?

More to explore



Show students the results of the Moral Machine study.
The research is described in section 02 of this teaching material. You can view detailed findings on the Moral Machine results explorer by visiting [this link](#). A world map will appear, and by clicking on individual countries, you can view a comparison – up to two countries at a time – against the global average:



You can compare moral preferences of different countries (green and ochre) with the global average.



Students can discuss up to 30 examples of the trolley problem.
If you have time left, you can follow up with the Absurd Trolley Problems at: [neal.fun/absurd-trolley-problems](#). This site offers 30 different trolley scenarios that work great as a basis for discussion.

Scenario number:

☐ Pedestrians crossing on red

Number of pedestrians:

Number of passengers:

The pedestrians are:

The passengers are:

☐ Left scenario☐ Right scenario

Reasons:

Scenario number:

☐ Pedestrians crossing on red

Number of pedestrians:

Number of passengers:

The pedestrians are:

The passengers are:

☐ Left scenario☐ Right scenario

Reasons:

Scenario number:

☐ Pedestrians crossing on red

Number of pedestrians:

Number of passengers:

The pedestrians are:

The passengers are:

☐ Left scenario☐ Right scenario

Reasons:

Scenario number:

Left☐ Pedestrians crossing on red Pedestrians hit: Passengers hit:

Pedestrians hit are:

Passengers hit are:

Right☐ Pedestrians crossing on red Pedestrians hit: Passengers hit:

Pedestrians hit are:

Passengers hit are:

Reasons:

☐ Left scenario☐ Right scenario

Scenario number:

Left☐ Pedestrians crossing on red Pedestrians hit: Passengers hit:

Pedestrians hit are:

Passengers hit are:

Right☐ Pedestrians crossing on red Pedestrians hit: Passengers hit:

Pedestrians hit are:

Passengers hit are:

Reasons:

☐ Left scenario☐ Right scenario