

Generation AI

IO3 – School Program for Primary Education Students

Tutor Handbook Template



Introduction to the Tutor Handbook

The aim of this short handbook is to support you, as an experienced tutor, to use the Generation AI Lesson Plans with learners in your group. If you are currently working as a Primary Education Teacher, this handbook will help you to introduce the Generation AI Lesson Plans in your workplace. When developing these Lesson Plans, the focus has been to support young learners in developing an evidence-based understanding of the complexities and basic principles of AI, computational thinking, and how they can be integrated in teaching learning to promote creative problem solving, resilience, and design thinking.

This Lesson Plan belongs to the Responsible Use topic and is targeted to children between 6 to 8 years old. This topic aims to provide learners with the basic knowledge of this competence and how AI can be used to solve daily life problems related to this skill.

Introduction to the Competence

Nowadays, we use Artificial Intelligence on a regular basis. Through apps on our smartphone, through Netflix and its recommendation, and our email with a spam filter. However, the above examples seem quite innocent, but AI comes with quite a few risks as well, for example AI bias. Leading to inequalities among people. Therefore, the topic Responsible Use is of great importance when teaching about Artificial Intelligence. With knowledge on the risks of AI, we can become responsible and critical consumers of this modern technology.

Introduction to the lesson

This lesson contains three elements: a video, a learning activity, and a challenge. All elements are designed in a way that your students can work on their own without any additional information needed from you. However, you can help, assist, or take over parts of the lesson if you want to. In this document we explain what you can do for each of the three elements.

Elements of the Lesson Plan

Video

This 1-2 minute video is an explanation of the Responsible use. This video is the first element of the Lesson Plan, and you can find it in the platform of Generation AI project as an embedded video of YouTube. It helps to contextualize the following elements of the lesson plan: Learning activity and Challenge.

What you can do with the video:

As a teacher, you can decide to watch the video as a group, or let students watch the video individually. Afterwards you can ask a few reflecting questions, such as:

- What was the video about?
- Why is Responsible Use important for you?

Learning Activity

The second element that you will find in the Lesson Plan is a learning activity with a problem-based task. This task consists of a guided activity that aims to help students come up with a reflection about a question/problem proposed.

This question/problem aims to give learners the opportunity to know how AI can affect their daily lives in a simple and ludic way.

In this lesson plan, the question/problem is **Can AI do harm?**

This lesson is focused on how frequently we interact with AI and its applications (and how these interactions will become the norm).

Therefore, students will gain a better understanding of the importance of being aware of AI's impact and the potential risks it brings to our lives.

During this lesson, an emphasis will be put on 'how AI sees' (also called computer vision). Students are also introduced to the topic of 'Machine Learning', where they will learn how AI is able to recognize things. Finally,

the students are introduced to possible mistakes that AI can make, for example the AI bias. Where an AI system recognizes something wrong, because we trained it with the wrong data.

The students are introduced to these concepts with help of some PowerPoint slides. The students are given a few reflective questions, which they can try to answer together with classmates after going through the slides.

What you can do with the learning activity:

As this is a guided learning activity, you will need to ensure that students follow each of the steps and links provided in the activity. If you need to use any other resources, feel free to do so and intervene in the course of the activity.

You can also choose to give the presentation yourself. In the Appendix I you can find an example of text that you can use while presenting the slides.

At the end, you can let students discuss the reflective questions in duos, or you can decide to do this with the whole group. In the Appendix II you will find answers and some videos for more information per question.

Challenge

This last element of the Lesson Plan is a Challenge which takes the form of an AI based game. It aims to be the most ludic part of the lesson plan and pretends to conclude with the knowledge acquired during the lesson plan.

The game is related to the Responsible Use topic and is developed using [Machine Learning for Kids](#). The game involves a step-by-step guide about how to create an Machine Learning algorithm that can make a difference between dogs and other animals. This Machine Learning Algorithm can be used for an automated pet door, that is only supposed to allow dogs to enter. The final step of this challenge involves a reflective exercise where the students have to consider the bigger picture when evaluating their newly build pet machine learning algorithm.

What you can do with the challenge:

We would recommend to try this challenge yourself as a teacher and go through all the steps to make sure you know how to help students if they have some problems with one of the steps.

There is a step-by-step guide that shows students how to do each step. If you want, you can help them a bit by:

- Showing how to create a new project (steps 1 – 11)
- Showing students how to search for a picture on google and how to copy the image address, and upload it in the project either under 'dogs' or 'other' (steps 12 – 18)
- After students have uploaded all their pictures, you can show how they can let their AI model learn and test it (steps 27-36).

After students have tested their Machine Learning Algorithm, you can invite them to test their algorithm even further by searching for:

- Dogs that are NOT let in
- Other animals that are let in (often a wolf)
- Let them reflect on why this happens and what this can mean. Can we trust this Machine Learning algorithm?
- Would this Machine Learning algorithm be appropriate to use for an automatic pet door, that can only let dogs in?
- Can we solve it? Make our Machine learning algorithm better? How?

After this, you can go back to the real life example of Joy from the PowerPoint slides, and discuss how would they feel if it happens to them. Emphasize on the fact that this is called AI bias, and that it can indeed harm us.

If students are very fast, you can give them ideas for other (more challenging) Machine Learning algorithms. Examples:

- Try building a Machine Learning Algorithm that can recognize both dogs and cats (tip: make a label for dogs, a label for cats, and a label for others)
- Try building a Machine Learning Algorithm that can recognize your face and all other faces of your classmates. (tip: make a label for your face, and a label for other faces)
- Try building a Machine Learning Algorithm that can recognize all different kinds of fruit (tip: make a lot of labels for each individual fruit: apple, pear, strawberry, etc.)
- Try building a Machine Learning Algorithm that is never wrong (tip: train the algorithm with A LOT OF pictures).

Appendix I

Slide 1:

Recap on what AI is. They are computer systems that can perform task that require some form of human intelligence. For example, decision making, seeing, recognizing, learning, etc.

Which is great! With these systems we are able to better see who is sick and how we should treat them, we are able to find the fastest routes for when we are visiting family, or we don't have to search for movies, but Netflix knows which movies we would like!

But can it also harm us?

Slide 2:

There are many examples of AI that you use as well. For example, did you know that Snapchat uses AI to detect your face and add these fun filters? And Netflix uses AI to recommend new series or movies to you. So does YouTube, AI tells YouTube which video you may like to watch next. And have you ever asked a question to Siri or Alexa? That's a lot of AI too!

Slide 3:

So let's see if there are other places where we can find AI too.

Slide 4:

Here you see a city. If you click on different places, you will see if there is AI. So where do you want to go first?

Slide 5 - 12:

(descriptions for each place are given on the slides)

Slide 13:

A lot of scientist think that there will be many other applications of AI in the future. So they are really curious of how our lives will look like in fifty years.

Slide 14:

So now that we know that we use AI in many different places, it is important for us to know if AI can harm us. Before I can show you how AI can harm us, I will first explain to you how AI can see.

Slide 15:

Here we have a picture of an apple. This is how we see everything! We see a clear image, colors, a lot of details like the raindrops on the apple.

Slide 16:

This is how AI sees! They see everything is squares.

Slide 17:

We call these squares, **pixels**. Can you still recognize the apple? It is very difficult right?

Slide 18:

That is because there are not many pixels in this picture.

Slide 19:

What about this picture? Can you recognize the apple? Yes, right? But if you look closely, you can still see it is all squares, but very small ones. So this picture has millions of pixels, that's why the picture is a lot sharper than the one we saw before.

So this is how computers see, with pixels.

Slide 20:

Now you can try. How do you think an AI model sees this picture?

Slide 21:

Where should we draw the pixels in this raster on the right?

Slide 22:

It could look a little like this!

Slide 23:

So now we know how AI systems can see things, but how do they understand what they see?

Slide 24:

Well, let's think about how we understand what we see. How do we know these are cats and these are dogs?

Slide 25:

If we look closely, we can see that cats have two eyes, but dogs have too. Cats have a nose, but dogs have too. They both have a tail, hair, four legs... So how do we know the difference?

Slide 26:

Well, because you have seen many examples of cats in your life! And also many examples of dogs. So, after seeing enough examples, you can now recognize a cat and a dog.

Slide 27:

This is also how AI can recognize cats and dogs. We show the AI system many examples of cats and dogs. We start to **train** the AI system. After a while, just like us humans, the AI system will start to recognize cats and dogs. This whole process is called **Machine Learning**.

Slide 28:

Let's do a fun exercise with this! Go to this website on your computers and play the game. Here you can see how the Machine Learning algorithm learns to recognize drawings.

Slide 29:

So, now we know how AI systems can see and recognize things. So now the question is: Is it dangerous?

Slide 30:

Well, it can be dangerous when AI is making mistakes. This can happen because we train the AI system wrong, for example by showing wrong pictures.

Slide 31:

Imagine, we are teaching AI system what cats are. And we show it these pictures of cats.

Slide 32:

Do you think it will recognize this as a cat?

Slide 33:

Maybe not, because all other cats have hair, and this one doesn't. Or because this one has a lot of wrinkles, but the other cats don't have any wrinkles. So perhaps, that is why the AI system will say: No, this is not a cat, because it doesn't look like the cats we have trained the AI system with.

Slide 34:

What about this one. Will the AI system recognize this as a cat?

Slide 35:

Maybe not, because we cannot see it's face. What if the AI system has started to recognize cats by the fact that they all have two eyes, and a nose?

Slide 36:

If AI makes mistakes, because we have trained it with the wrong pictures, we call it an AI bias.

Slide 37:

These mistakes can happen in real life. In fact, they do. Here is an example. In the picture you see Joy. She is a scientist and an artist.

Slide 38:

In the picture, Joy is trying to get her face recognized by an AI system. But as you can see, the AI system doesn't see her face.

Slide 39:

Then, Joy put on a white mask. And as you can see in the picture, the AI system immediately recognized a face.

Slide 40:

So how is that possible? Do you have any ideas?

Probably because this AI system was trained with pictures of faces of people with a white skin color, and not black a black skin color like Joy has. So the AI system does not recognize a face, because it associates a face with a white color.

Slide 41:

This means that this AI system is biased, because it making mistakes because it was trained with the wrong pictures.

Slide 42:

What do you think about that? Is that fair? How do you think Joy felt? How would you feel?

Slide 43:

Can you think of reasons why this is dangerous? For example, what if this system will be used to let people in buildings. Then Joy could never enter a building anymore. Or what if we are inventing a self-driving car, and this AI system is used to prevent the car from driving people over. That could be very dangerous, if this system only recognizes people with a white skin color.

Slide 44:

Joy shares more on her story in the documentary on Netflix, called: coded bias.

Slide 45:

Now that you know how Machine Learning models can recognize things, by showing them first a lot of examples, and you know how that can go wrong sometimes, it is time for the challenge!

Go to the challenge on the webpage.

Appendix II

Some answers to the questions that students will discuss with each other.

1. Where can we find AI?

Almost everywhere! For example on your smartphone when you use apps like Snapchat. Snapchat uses AI to see your face and make fun filters with it.

Other examples are Netflix, YouTube, Instagram, Spotify, or webshops like Amazon, who all make recommendations on other movies, videos, songs, or products that you will like, based on movies, videos, or songs you watched and listened to before. That is AI!

We can find it in the hospital, where AI is being used to determine who is sick and how we should treat people to get them better.

In the future we will have self driving cars that use a lot of AI.

Bank, the banks use a lot of AI to make sure you are really the person the money belongs too.

Delivery truck, AI can be used to find the fastest our shortest routes in traffic. That is why delivery services like to use AI, so they can do their job a lot faster.

2. Does AI see exactly like we, humans, see?

No, AI and computers in general see things with pixels. Little squares. The more pixels a picture has, the more clearer the picture is.

3. How can AI recognize things?

AI can recognize by training it. We can train AI by showing it many examples of things we want it to recognize. For example, if we want it to recognize cats, we show it lots of examples of cats and other animals. Just like us human, AI will then be able to distinguish cats from other animals. We call this Machine Learning.

4. What is AI bias?

Sometimes, AI can make mistakes. Most of the time it makes mistakes, because we trained it with wrong pictures. The pictures were either wrong (for example we showed a picture of a dog, instead of a cat), or the pictures were not diverse enough (we only showed pictures of one specific breed of a cat and not all kinds of breeds). This way, the AI system can make mistakes and either think something is a cat which isn't, or the other way around and thinks it sees a cat, while it is actually something else.

5. How can an AI bias harm you?

AI bias can harm us, because it can mean that some AI systems who are supposed to recognize us, don't. This can be dangerous. For example, when a self-driving car is supposed to recognize human, to make sure it doesn't ride them over. If the AI system wouldn't see me as a human, the car might not realize it has to stop for me. Therefore, it is important that AI systems are trained with proper pictures, and don't make mistakes.