


An introductory guide for using AI in education

ARTIFICIAL INTELLIGENCE: TEACHING & LEARNING WITH AI



Alexandros Kofteros, PhD

Artificial Intelligence: Teaching & Learning with AI

An introductory guide for using AI in Education



© 2023 Alexandros Kofteros. All rights reserved.
Created & Distributed under a Creative Commons
Licence.



Nicosia, 2023
ISBN 978-9925-8055-1-8

All images created with Microsoft Bing Create with the exception of screenshots,
unless otherwise stated..

All products mentioned in this book are registered trademarks of their respective
companies.

Special thanks

for their valuable time in assisting in the development of this short guide with their suggestions, proof reading & translations:

Pola Mithou

Vasso Servou

Fotini Tsaliki

Katerina Glezou

Samuel Branco (Portuguese Translation)

Zoi Karagiorgi (Greek Translation)



Contents:

Part A: What's AI?

Chapter 1: A brief history of AI	11
Chapter 2: Introduction to AI	21
Chapter 3: ChatGPT who?	31
Chapter 4: Creating Images	42

Part B: Hitchhiker's Guide to the Galaxy of Education

Chapter 5: "Simple" Maths	59
Chapter 6: Maths for students	79
Chapter 7: ChatGPT for writing	88
Chapter 8: Presenting our work	99
Chapter 9: AI Avatars	117
Chapter 10: AI Music!	127

Part C: Let's Code!

Chapter 11: Machine Learning	143
Chapter 12: Save the oceans!	159
Chapter 13: Machine Learning for Kids.....	170
Chapter 4: RAISE Playground	190



Main topics of this guide:

- A brief history of Artificial Intelligence
- Machine Learning
- AI for image creation
- AI programming




About this guide:

AI has made serious inroads in education, especially with the development of programming tools and technologies that allow even students of younger ages to use them effectively.

This guide is divided in three parts. In the first part, we provide a brief history of the the development of AI, starting with ancient myths and legends and moving to the current era of technological breakthroughs.

In the second part, we look at activities that can be used in the classroom, from primary to secondary education. Some are structured, interactive activities that enable students to learn more about AI and Machine Learning ([EU CodeWeek](#), [code.org](#)), while others are open-ended and can be applied in various subjects of the curriculum (for example, AI image generators). In the third part, we learn more about training models and programming games using Scratch 3.0 variants. Each chapter covers a different tool or set of activities, with lesson plans on how to use them in the classroom. This guide could be extended to a book with more than 400+ pages in size. However, my intention is to give teachers and students (and why not, parents!) access to AI tools and content they can use immediately! And to allow everyone to unleash their imagination!

Welcome!



Welcome to my guide on AI in Education. This has been in and after an amazing EU CodeWeek Leading Teachers/ Ambassadors meeting in Brussels. It was then that the idea of actually creating a short guide on AI came up, since there is rising interest on the subject from educators all over the world. With the development of various resources and technologies, all that was needed was a way of bringing (some of) these under one roof, with suggestions on how to use them effectively in the classroom. Of course, the activities on Machine Learning and the learning resources from the outstanding EU CodeWeek and [Code.org](https://code.org) portals already have fulfill most of our (teachers) needs, but the idea of a coherent, chapter by chapter “old school” ebook was deemed necessary.

I hope you enjoy this guide as much as I did creating it!

Feel free to share suggestions: alexandros@mathisis.org

Alexandros Kofteros, Teacher

“Come writers and critics
Who prophesize with your pen
And keep your eyes wide
The chance won't come again
And don't speak too soon
For the wheel's still in spin
And there's no tellin' who
That it's namin'
For the loser now
Will be later to win
For the times they are a-changin' “

(Bob Dylan)


Part A: What's AI?

What's to learn

In Part A, we learn about the history of AI, starting with the mythos of the first-ever artificially created (at least in Mythology) thinking machine, Talos. We take a tour in the world of science fiction, before we land in modern-day applications of AI that include Machine Learning, virtual assistants, mapping applications, translations, music selections, shopping and more!

In Part B, we use AI tools such as ChatGPT and Bing Image Creator to generate content, while we also use them to develop educational material, including worksheets, lesson plans, presentations.

In Part C, we focus on coding using variants of Scratch 3.0 to create and train models (Machine Learning for Kids) as well as using already trained ones (RAISE Playground).



Chapter 1. A brief history of AI

"Music and mythology,
Einstein and astrology
It all started with the big bang"
(BNL, "Bing Bang Theory")

In this chapter...

In **Chapter 1: "A brief history of AI"**, we will learn about:

- **AI Vs Robotics:** what is Artificial Intelligence, and why Robotics and AI are not the same.
- **Myths and legends about artificial life:** how mythology has given us the first examples of an artificial "life", and its implications in modern fiction.
- **The automatons:** an attempt to create artificial constructs that act as living human beings.
- **AI in Science Fiction:** how AI is depicted in literature and movies / tv series.
- **AI in arts:** how Artificial Intelligence is used to generate graphics and what are its implications.




A brief history of AI

Almost all civilisations have myths and tales of gods or engineers creating sentient machines. To learn more about the history of AI, we will delve into mythology and ancient history, in times where science was not yet prevalent and natural phenomena were attributed to the will of the Gods. From there on, we will move forward through the ages and through various myths and legends, to admire the Automaton, before we finally approach the era of the 1940s and the advent of the first digital computers.

Between mythology and the digital age, we need to examine the impact of Science Fiction in helping us grow an interest to AI and their physical manifestations: that is, (almost) sentient robots. The film industry has introduced deadly robots (Cylons, from "Battlestar Galactica") as well as friendly ones (R2D2 and C3PO from Star Wars).

The digital age saw an emergence of the field of AI with the development of intelligent systems, leading to the second decade of the 21st century and an explosion of new technologies that made AI mainstream.




AI Vs Robotics

Science fiction has done its part in bringing together robots and AI. As a matter of fact, robots in science fiction novels, movies and TV series are almost always “intelligent”. However, robotics and AI are not necessarily the same. Even though the two fields are related, and in many aspects they seem complimentary, in truth they are completely distinctive with different applications and methods. The integration of AI and robotics has been an integral part of cinema since the early days of the science fiction era. For example, in “Forbidden Planet”, a 1956 movie, Robby the Robot has its own character and personality, clear indications of an advanced AI. Robby is a robot and can move around while at the same time exhibiting “intelligence”. Probably the most famous AI in movie history is HAL from 2001: A Space Odyssey. Contrary to the robotic nature of Robby the Robot, HAL is an AI integrated with the space ship, and able to recognise faces, gestures, speech and take its own decisions through

learning. As a conclusion, robots are machines designed to interact with the environment, including objects, and to perform specific tasks (ie lifting objects, soldering circuits etc).

An AI is a software system that can make decisions by analysing information from its environment (ie recognise a red light and stop a car, recognise the face of a person and unlock the computer etc) and even learn.

AI and Robotics are expected to merge as technologies and provide new products.



HAL 9000 from “2001: A Space Odyssey”

Talos, the first(?) robot

Going back in time, the very first known (mythological) example of an artificial intelligence was the legendary bronze warrior Talos, created by Hephaestus to guard the island of Crete. Hephaestus worked in his lab day and night with the help of a cyclops to create the ultimate guardian of the island, which he then offered to king Minos.

Talos was a giant “robot” or automaton, tasked with running around Crete three times a day and repelling invaders. The first depictions of Talos are dated to 400 BCE. According to myth, Jason and the Argonauts reached Crete while returning from their voyage carrying the Golden Fleece.

Talos tried to sink their ship, Argo, by throwing huge boulders at it. Jason in desperation turned to Medea who was a witch. Medea found the weakness of Talos, a huge nail keeping ichor, the liquid that kept him alive, inside its



body. Jason managed to remove the nail, thus releasing ichor and killing Talos. The death of Talos has been depicted in art over the years, helping spread the myth and the legend of the first artificial “life” in mythology.

Ancient Egypt & China

The myth of Ptah (3000 BCE) suggests the creation of Osiris as an artificial being. Similarly, in ancient China there are myths regarding artificial persons. However, the myth of Talos is clearly describing an artificial life made of metal and powered by an “unknown” power source referred to as ichor.



Automatons

Automatons or automata (the term is derived from the Greek word “αὐτόματα” which means “being able to operate on its own” refer to complex constructs that could mimic functions of living beings. For example, a cuckoo clock will open its doors and reveal a small bird that flaps its wings and makes noise to tell time.

There are many examples of automatons throughout the ages. Some of the best known in history were the machines of Heron of Alexandria, who used hydraulics to create automatic doors and other complex machinery, including music organs. Many cultures, including the Chinese and the Japanese, have created their own automatons (or automata). One notable example is the legend of Yan Shi, a creator of artificial beings that presented King Mu of Zhou with a mechanical person. Upon touching its artificial skin, the automaton started singing and moving. However, besides a reference in a text (Lie Zi, 400 BCE), there is no other evidence to support that such a machine was ever



created. Some of the most complex automatons were created during the 18th century. A set of automatons created by Jaquet-Droz are considered to be the most sophisticated ever created. The most complex was “The writer” which could write text on a piece of paper!

The infamous Turk

A very sophisticated automaton was created in 1769. It could play chess and it even managed to defeat some important opponents, including Benjamin Franklin. However, even though it was a very complex mechanism, it proved to be a hoax, since it was found to be operated by a human inside the complex machine!



AI in Science Fiction

Nothing has influenced the depiction of Artificial Intelligence more than (science) fiction. From the writings of Isaac Asimov and his Laws of Robotics, to HAL from "2001: A Space Odyssey" and the modern era of the latest "Terminator" movies, AI and intelligent machines are depicted either as malevolent or benevolent, depending on the writer and the setting. For example, in the dystopian future of "The Matrix", machines have rebelled and enslaved the human population which are actually used as batteries inside a virtual reality. In "Battlestar Galactica" (remake), machines have rebelled against their creators. In the film "AI", Earth has been left devastated and the only intelligent beings are actually artificial ones. AI seems to make a very good background story as the main antagonist (usually), since it exonerates (in most cases) our kind for any fault they have.

On the other hand, we have benevolent machines, with the more prominent ones and most recognisable the droids



from Star Wars (R2D2, C3PO to name two). Interestingly enough, one of the main sci-fi TV shows and movie series -Star Trek- has barely touched the surface of benevolent or malevolent AI, with the exception of the bio-based Borg or Commander Data.

The Laws of Robotics


Isaac Asimov, the famous science fiction writer, has written his three laws of robotics: 1. A robot may not injure a human or allow a human to be harmed, 2. A robot must obey orders except when they conflict with law 1 and 3. A robot must protect its existence as long as it does not conflict with laws 1 and 2.



What we learned so far...

In **chapter 1 "A brief history of AI"**, we travelled back in time to the first mythical representations of an artificial life, the legendary Talos, created by the ancient Greek God Hephaestus to protect the island of Crete. We then learned more about automatons, some of which still exist as exhibits in museums today. And, of course, we looked at the representation of AI in science fiction, in which it is either vilified (HAL 2001, Daleks, Cylons) or turned into sidekicks (R2D2, C3PO to name a few).

The story of AI does not end here and of course it is not covered in detail. There are follow-up activities that we recommend, to learn more about early mentions of artificial life in various legends and cultures.



Further reading:

Almost every civilisation has at least one reference to an artificially intelligent being. We can further research such myths and legends:



- **Mechanical Human:** An engineer in ancient China (300 BCE) created a mechanical device resembling a human that could move like a real person. Even though references are very ambiguous about such a creation, it shows the interest of ancient Chinese engineers in creating complex mechanisms.



- **Brazen Head:** This was a mythical device in medieval legends. One of the first written accounts mentions a talking head that could answer any question with a “yes” or “no”, as far back as the 12th Century CE. It was a very early automaton, resembling a head made of metal. We could see it as a distant concept of an ancestor to ChatGPT. Their origin might go even further back to Hero of Alexandria who created moving statues using steam.



The Jewish Golem: according to a myth, the Golem was an artificial creature created by Rabbi Eliyahu in the 16th Century CE using clay to protect the Jewish community against their enemies.

- **Mechanical Knight:** Leonardo Da Vinci designed an automaton resembling a knight in full armour in late 15th Century CE. According to his design, the mechanical knight could sit, stand, raise its arms and remove its visor. It was expected to wear medieval armour. It is unknown if a prototype was ever created.



- **Turk, the Chess Player:** This automaton was supposed to accurately play chess, and even challenge experienced human players. It was created in the late 18th century. However, it was found to be a hoax, since it was really controlled by a man hiding inside the main body of the machine.

Chapter 2. Introduction to AI

"All those moments will be lost in
time, like tears in the rain"
(Blade Runner, 1982)


In this chapter...

In **chapter 2 “Introduction to AI”**, we learn about AI we already use (but we sometimes don’t even realise we do!).

- Understand what Machine Learning is
- Recognise benefits of Virtual Assistants
- How AI “helps” us decide what to watch, read or buy
- How AI enables us to find places on a map, and do translations in various languages - including Klingon!



Artificial Intelligence (AI)



Link from <https://code.org>


The term "AI," short for Artificial Intelligence, is becoming increasingly commonplace, but what exactly does it entail, and why is it gaining such importance, particularly for educators and the broader public? Artificial Intelligence is not a recent development; it's been in the works for many decades, and, truth be told, we're still quite a way off from achieving what is considered genuine Artificial Intelligence.

When we talk about AI today, we're referring to sophisticated sets of instructions, known as algorithms, that empower computers to perform specific tasks. These tasks

were, until relatively recently, exclusive to human capabilities. To put it simply, AI enables computers to do things that, in the past, only humans could do. For instance, AI has the remarkable ability to comprehend human speech, whether it's spoken or written, and respond in a way that can sometimes be challenging to distinguish from human interaction. It can also identify shapes and objects and, intriguingly, generate content independently based on the descriptions provided to it.

As we delve deeper into the subject, we'll discover that AI has a broad range of capabilities. We can instruct AI to generate extensive written content for school assignments, provide us with valuable insights, create entire presentations complete with original text and images, and much more. The exciting aspect is that AI is continuously advancing, expanding its capabilities and transforming the way we engage with information and produce content. This ongoing evolution is reshaping our approach to work, creativity, and problem-solving.

Machine Learning



Link from <https://code.org>

A popular meme circulating the internet humorously suggests that "people are hooked, while machines are learning." While this meme contains an element of truth, it's essential to clarify the unique nature of machine learning. Unlike human learning, machines acquire knowledge not through traditional methods but by processing vast amounts of data, often concentrated on specific subjects. In many instances, individuals may have inadvertently contributed to this process. Consider, for instance, innocently using a mobile application that ages

your photo. Have you ever wondered about the purpose behind such applications and why they are offered for free?

It's important to note that not every free application engages in this data-driven approach, but some companies strategically choose to release products openly, encouraging thousands, if not millions, to use them regularly. This widespread usage generates valuable data that computers can then leverage to "learn." This is the essence of machine learning – the process of furnishing the computer with substantial amounts of data, empowering it to decipher patterns and trends. In return, the computer utilizes this acquired knowledge to make informed decisions, ultimately enhancing its functionality and problem-solving capabilities.

So, the next time you use a seemingly harmless app that uses photos or music or even text, you might be unknowingly contributing to the vast pool of data that fuels machine learning advancements.

“Hey Siri...”

Did you know that **virtual assistants** such as Siri, Alexa, Cortana (and more) use machine learning to continuously improve how to understand and communicate with you? Every time you use them, you are actually refining and teaching them -without knowing it- how to improve their own service.

This is in fact one of the examples of using machine learning within devices that we use almost daily - and mostly our mobile phones that are now an almost integral part of our life.



Common AI Joke:

“Why are you whispering?”

“Because Siri is listening”

“I laughed, Siri laughed, the TV laughed, the washing machine laughed...”

Virtual assistants are not only tied to our mobile phones but can be found on various types of devices. For example, Alexa, the virtual assistant developed by Amazon, is

featured on a number of different products mainly developed for home automation tasks - that is, to control various parts of your home, turning it into a “smart” house. For example, you can ask Alexa to dim the lights or even control the TV (especially if you have subscribed to Amazon service). With Alexa you can even order things from Amazon (surprise, surprise!).

In general, virtual assistants are built into mobile devices or smart appliances. They respond to voice commands and usually they have a voice of their own (male, female, robot...).



Movies, Music, Marketing...

Have you noticed how even a simple search on the internet can (then) generate constant ads on specific topics? For example, if you do a search for a new mobile phone, you will most definitely start seeing ads about other mobile phones or related products (cases, peripherals) wherever you click - in search engines, in social media etc. This is because search engines use machine learning to analyse our search preferences, or our music and video habits, in order to promote products that appear to be in our liking - be it a tablet that we just researched for purchase or just out of interest, or even an apartment we are interested in renting.

In some cases this might appear to be intrusive and even a nuisance. And it is, in many respects. However, it is hard to dismiss the benefits for both the consumer and the product owner, since machine learning can help push to us products and items that are closer to our liking. It also offers us a glimpse on new products and new items that might otherwise have passed unnoticed.



Movie recommendations

Netflix and other streaming services use machine learning to analyse our preferences and thus promote content that is more to our liking. This way, not only we don't have to search an endless stream of new content, but it helps (the service) keep our interest, and thus remain their subscriber.



Translations

Translation engines use machine learning to improve the quality of translation between one language and another. The most common of these translators are text based, and they usually support all major languages (and then some!). Google has one of the most robust and accurate translators, offering near-instant translation of almost any text we throw at it.

Some services (such as YouTube) use machine learning to automatically create subtitles for the videos uploaded. This is a very important feature as it provides access to many videos for people that would otherwise not understand the language. However, this feature is still not mature, with the exception of generating text in English directly from videos in the same language.

Other services such as Skype already have a beta version of automatic translation (in text) when two people with different languages have a video conference. Even though in its infancy, it shows the tremendous potential that will



allow two or more individuals who do not speak the same or a common language negotiate any type of deal.

As technology keeps churning out more powerful computers with less components.

From Klingon to English

What was once just a part of a fictional universe, and the mechanism with which they explained how everyone speaks the same language (hint: the don't), is now becoming a reality. More and more technologies are released that allow for direct speech-to-text development.



Walking & Driving

Technologies such as Google Maps rely more and more on AI to provide drivers with optimum routes, as well as places of interest along a specific route. What initially was a vaguely detailed map of the world, has become more and more integrated with the Google search engine to provide on-the-spot information on anything the user might require, including search functions that learn from the driver, and thus propose interesting areas to visit.

The AI behind Google Maps is also using thousands of photos uploaded by the users to create virtual reconstructions of places of possible interest to drivers, in an effort to enhance both the use of the specific tool, as to make the driving experience more interesting (and useful). As such, the Maps application will provide the driver with suggestions on interesting or useful places to visit during a trip (ie of historical importance, EV charging stations, coffee shops for rest etc).

As drivers and Google Maps users keep providing large amounts of data to the application, the AI is constantly



improving the information it provides to the user (and drivers).

AI in Google Maps

Google Maps uses literally thousands of photos to create virtual representations of spots on the map that might be of interest to us. Therefore, Google Maps uses an AI that falls under the “Computer Vision” category.



What we learned so far...

In **chapter 2 "Introduction to AI"**, we learned about modern Artificial Intelligence (AI) and Machine Learning.

We examined how we already use AI in our everyday life, from the most common tasks to the most complex. Even services we use on a daily basis, like streaming video and audio, use AI to show content that is of more interest to us. Modern search engines like Google or Bing use AI to promote results, while even our cars employ some kind of Artificial Intelligence, especially with Map applications.



Further reading:

AI is almost everywhere: in our homes, our computers, our smartphones, our cars, even most of the gadgets we use today make at least some use of an AI-related technology.

- **Smartphones:** As mentioned earlier, modern smartphones include AI features for tasks such as auto-completion of sentences when typing text, face recognition of unlocking the device, speech recognition for controlling applications completely hands-free.
- **Online shopping:** most modern online shops feature thousands if not millions of products. To help us find products that interest us, but also to entice us to specific products, online retailers use AI algorithms to analyse and predict our preferences - and show results that might attract our attention and even lead to impulse purchases.
- **Email filters:** Millions of fake emails, with a huge percentage being phishing or other scams, race through the internet every single day. Fortunately, service providers are using many methods, including AI, to filter the vast majority of those before they reach our mailboxes. Still, a lot can reach the end user, but at least a significant percentage does not.

- **Virtual Assistants:** Virtual assistants have existed for decades. One of the most infamous was the Microsoft Office "Clip". Today they can be found in most devices, and provide enormous help by understanding natural language. Solutions such as Siri, Google Assistant and Alexa can be found in common devices that range from smartphones, to washing machines, to speakers!



Chapter 3: ChatGPT who?

"I am the voice of Knight Two
Thousand Micro Processor - K.I.T.T."
(Knight Rider)

In this chapter...


In **chapter 3 "ChatGPT who?"**, we are introduced to ChatGPT. We will:

- Learn more about the capabilities of ChatGPT and how to use it
- How ChatGPT compares to more "traditional" services such as Google Search and Google Translate
- How ChatGPT handles Mathematics
- How to use ChatGPT for creating code




Creating a ChatGPT account

ChatGPT from OpenAI is a jack-of-all-trades but it can be easily understood that it's also a master of several. For generating text, it is super-fast, multilingual and offers (almost) everything we need for writing. We will also examine, later in this chapter, other tools that are more specialised, but ChatGPT is THE tool for instant gratification.



For those unfamiliar with ChatGPT, please refer to Part A of this book. For those starting its use, you must first create an account, starting from the url <https://chat.openai.com>.


You can either create a new account, or link an existing one (ie from Google or Microsoft or Apple).




The creation of a new account is very straight-forward: you click Sign Up and then you enter your email and a password. An email will be sent to the account you provided (ie "alexandros@mathisis.org").

Login to your email account, find the email and just click on the green button "Verify email address". Another window will open, asking further details from you (full name and birth date). That's it! After you complete this step, you are ready to go!

Starting with ChatGPT



The version we use in our examples is ChatGPT 3.5, and it's free for personal use. However, it does not mean it lacks any features we need for text creation. First up, let's familiarise ourselves with the extremely simple environment: on the left side of the window, we can keep a history of previous "chats". In the lower right part of the window, we can see the text box where we write simple sentences or full text. Let's start by introducing ourselves: we will just type "Hello, I am Alexandros" and hit the Enter key. Let's look at the response:



On the left part of the window, ChatGPT has created a title for our first "chat". Also, it has replied to our first sentence. It assumes immediately that we need some assistance and possibly a question will arise OR it prompts us to write a question.


Teaching the machine...

ChatGPT learns constantly by analysing the data and feedback it receives from us. As you can see on the picture above, under each reply we have the option of clicking on the "like", "dislike" icons or on the reload, to issue another response. We can also "vote" on how the conversation is progressing. These details might appear trivial but are essential in the constant improvement of every type of AI.




Multilingual support

One of the main benefits of our chat is the support of multiple languages. As you can see on the screen, it can recognise -and respond in kind- to almost any language we can use. Not surprisingly, it can also tell us how many languages it can understand!



languages. However, it can respond in other languages as well, including Greek, Turkish and Klingon!



Why Klingon?



Even though Klingon is an imaginary language (it exists in the sci-fi universe of "Star Trek"), over the past 6 decades many people, including professional linguists, have created the structure of it, as if it was a real language. Therefore, ChatGPT has also learned to use it, however efficiently compared to real languages spoken by millions!



Because it learns through "crunching" huge amounts of data, it is much more suited to answer in English than other

Working with chats

As we continue writing messages, all the replies and text produced are considered part of the same “chat”. However, during the day it is useful to create new “chats” for different type of replies and content creation. For example, during the same day we might require ChatGPT to create a text on a famous painter, and later on to create activities for our next Maths lesson. By keeping each type of reply in a different “chat”, we can more easily come back later to specific content.



Chats are very effective in organising our engagement with the AI and creating different types of replies and data.


ChatGPT for information search

ChatGPT can answer to most questions we wish to ask. It must be taken into account that not all of this information might be accurate, therefore we need to keep this in mind when using any of the replies generated by ChatGPT. We will see more on search using ChatGPT in the next pages.




ChatGPT Vs Search Engines

We have already established that ChatGPT can be used to return information on specific topics. However, this is far from using a dedicated search engine such as Google or Bing. Whereas in ChatGPT we receive a reply similar to what a knowledgeable person might give us (on a specific subject), search engines will return a list of web pages related to our question. And where ChatGPT gives an instant “answer” to the question, search engines might provide alternative -and equally significant- web pages that might provide us with the information we seek. And, in most cases, identify the source and therefore decide to trust the information or not.




In ChatGPT, we get a precise text with information on Yuri Gagarin. On a search engine search, we get much more than that - we have a summary of his life, images, links to videos and links to multiple pages, including Wikipedia articles.




Generally speaking, if we need to search and receive as much information on a subject, we use a search engine. If we are looking for specific information focused on a search query, we can use ChatGPT. But, we should never take its results at face value.

ChatGPT for translations



ChatGPT can recognise a vast number of languages, and reply (almost) perfectly in most of them. However, should we use it for translations instead of dedicated tools such as Google Translate?

We should keep in mind that Google uses AI technology to constantly improve its translations. Therefore, dedicated tools such as Google Translate already use AI when translating a piece of text (or even a word) from one



language to the other. ChatGPT is very capable in handling translations, but in most cases Google Translate -at least at this point- will do a much better job. On the window (left), we see a translation in Greek of the short bio of Yuri Gagarin. The translation by ChatGPT is full of mistakes that make no sense, whereas the Google Translate results (image top) are very good (with minor adjustments). query, we can use ChatGPT. But, we should never take its results at face value.


We can, of course, "ask" ChatGPT to retry the translation (which brings improved, but not much better results).

ChatGPT for Maths


We can use ChatGPT to solve maths problems. No, we are not going to use it to solve problems that might require significant resources (ie “calculate the Nth Prime number”) but rather use it for simpler problem solving tasks.

Let’s take a typical problem:

“Mary has a number of pens. George has twice as many pens as Mary. Helen has 5 pens less than George. All the pens are 45. How many pens does each person have?”



Maybe the way ChatGPT replied was not that simple and we might have trouble understanding it. We can ask ChatGPT to explain it again in simpler terms:




The picture above shows how ChatGPT broke the explanation in more steps, numbering each step and emphasising it using bold, white(r) headings.

In a next chapter, we will take this approach many steps further and even create charts!

What we learned so far...

In **chapter 3 “ChatGPT who?”**, we were introduced to the amazing capabilities of ChatGPT 3.5. It is an incredibly versatile and powerful tool that replies to our queries almost instantly and with great precision (in most cases). We can ask questions in most languages and it can even translate text from one language to the other. ChatGPT can even handle complex (school) maths and solve equations, or even explain how to solve a (maths) problem. It can even be used for generating programming code. These capabilities will be explored further in chapter 5!




Further Activities:

Here are some activities you can test with ChatGPT v3.5.

You can try these at school with your students as well:

Activity 1: “Broken telephone” (Translation):

Try to translate a sentence from one language to another (it might be a random language - we can ask ChatGPT to show the most common languages it uses). Try translating the translation again and again to other languages. Let’s suppose the 10th language is Croatian. Ask ChatGPT to translate it again to the first language you used. Try to see the end result if it is the same or similar in meaning!



Activity 2: (Mini) Turing test:

Ask students to create a paragraph or a short bio of a person (or an invention, or an item - pick the topic you believe your students would like the most). Ask your students to share the paragraph with each other. One of the paragraphs must be a product of ChatGPT but your students must not know which one. Try to see if they will guess which paragraph was produced by ChatGPT!

Activity 3: Use ChatGPT for assistance with Maths

Prepare maths problems that are deliberately complex - you can even ask ChatGPT to generate them for you. Provide your students (in pairs or in small groups) with at least one mobile device per group with access to ChatGPT. Ask them to use ChatGPT for assistance in how to solve the problems.

Activity 4: Project with references from ChatGPT

Ask students to work on a project (ie “History of music”, “Painters of the 15th century”). Give them access (in small groups) to ChatGPT. They will have to find their information from ChatGPT. Then, they will have to verify that information by comparing the ChatGPT results with at least the Wikipedia related articles.


Chapter 4: Creating Images

"Magic mirror on the wall,
Who's the fairest one of all?"
(Snow White and the Seven Dwarves)

In this chapter...

In **chapter 4 “Creating images”**, we will:

- Generate images from a basic description using Stable Diffusion Web, Deep Dream Generator and DALL-E
- Refine our images
- Use different styles for our images
- Export images to other formats




AI for image creation

Image creation using AI has been a very controversial subject. In fact, at least two AI-generated images have won two awards: the first image won a major photography competition, with a photo created on DALL-E 2, named PSEUDOMNESIA: The Electrician, and the second image, called “Théâtre D’opéra Spatial” won the blue ribbon in the Colorado State Fair Art Competition. Both awards have created controversy, and artists involved in both competitions -as well as many others- were outraged by this, since they consider it to be cheating at best.

In fact, the US Supreme Court has even stated that “art created by an AI cannot be copyrighted”.

Despite these, AI image generation can have a profound effect on how we prepare our teaching material, and how we encourage students to use it for learning purposes, as we will examine in this chapter. When used in good faith, and with proper learning goals, AI image generation can be a very valuable tool both inside and outside the classroom. Therefore, the tool is left to the use for



maximum positive impact - something that we hope will be evident in the following pages!

AI images that caused controversy


For the Colorado State Fair award, see here: <https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>

For the award of “The Electrician” see here: <https://www.scientificamerican.com/article/how-my-ai-image-won-a-major-photography-competition/>




Stable Diffusion Web


Stable Diffusion is a web based image generator, capable of creating amazing photorealistic images. The web version is free, offering a very easy to use environment. Access is provided via a web browser (on a computer) via the link <https://stablediffusionweb.com> or through a mobile app.



By clicking on the "Get Started Now" link (image above) the web page scrolls directly to the Prompt that allows the creation -through a text description- of our first image. Contrary to other AI image generators, Stable Diffusion Web, at least in basic form, is completely free.




Once we type our description, we can choose one of the available styles (image above). Keep in mind that some of the advanced features, such as higher resolutions, are only available with the Premium Plan. We click on "Generate" to -you guessed it!- generate the image!




Deep Dream Generator

There are many available services that enable image generation with stunning results! How else would a photographer or a graphic artist manage to win international competitions?



Deep Dream Generator is an AI-powered service that enables the creation of amazing images, through simple descriptions (<https://deepdreamgenerator.com>). It provides free "energy" to test the service. We can create a new account or link from another service (ie Gmail).

Like most similar services, it uses a kind of "currency" called Energy - once our "energy" runs out, we have to purchase more to continue generating images.



Before we generate an image, we can place a search in order to identify any similar (and available) pictures other users have generated.


AI images and copyrights

According to a US court, AI-generated art cannot receive copyrights, since only human authors can receive copyrights. This is a very important decision, since (a) all AI-generated images can be used freely and (b) the efforts taken to create an image cannot be protected by law. At the moment of writing, this decision still stands (<http://tinyurl.com/reutersAI2023>)




Generating a “dream”

To generate an image (or “dream” as they are called) with Deep Dream, we click on “Generate” and then type the appropriate description in the Text Prompt (image below). Once we are satisfied with our description, we click on the “Generate” button.



Prior to any generation of image, the exact cost in “energy” is displayed, as shown in the image above. For this example, we will create an image of a Panda playing with a basketball on the surface of the Moon, with Saturn clearly visible (What? A Panda playing with a basketball on the

surface of the Moon is more logical than Saturn being visible?).




Once an image is generated, we can download it or save it. There is a distinctive difference between the two, since saving an image can allow us to share it with the public, or make changes to it. We can also choose to regenerate the image for a better -or alternative- result. However, each regeneration consumes more “energy”.

In the next page, we will edit our created “dream” with some important and powerful capabilities offered by Deep Dream Generator.

AI Editor


All our “dreams” (images created with Deep Dream Generator) can be accessed through our account. We can do some serious editing to our images through the AI Editor, from simple painting on top of our image, to application of various filters.

Image Editor



The tools of the Image Editor allow us to apply filters, draw on top of the generated image, resize it, crop it, apply a frame and much more. In fact, it provides a very basic image editing application to make direct changes to all the pictures we create using the AI service. A very important feature is the capability to upscale the image - that is, add more pixels so that the image can be enlarged without

quality loss. Currently, 1MP (megapixel), 2 and 5MP are supported.



Keep in mind that every time we use AI-powered tools, we consume energy. Upscaling an image can consume up to 15 “energy” credits, therefore we should be very careful how we use this (very important) feature. For most uses, especially in simple worksheets or for internet use, the basic resolution might be adequate.


Image resolution

Images on screen consist of small squares called Pixels (Picture Elements). The quality of an image is based on the amount of pixels used to depict it (its “resolution”). In small sizes (ie an image in a small window), the difference might be negligible between two images with different resolutions. However, once we enlarge the window or the images, the differences start becoming quite apparent.




Advanced generation

An image can be generated by typing a description. However, image generators can take into account many different parameters (or “modifiers”) such as the style (3D, 2D), effects on the image (light, rainbows) or even the style of a famous artist! Of course, we can always include these in our description. Otherwise, we can click on “Add modifiers” to select which ones we would like to include in our image.



Upon selecting the modifiers (image top right) we must take into account that some of these cannot coexist in the same image or are contradicting/redundant. For example, if we select 8K quality, there is no point in selecting 4K as well - 8K is double the resolution.


Also, keep in mind that when we increase the complexity of the image, the end result is not always the best or at least what we might have in mind to create!




For every category, we can select as many modifiers as possible. As mentioned in the previous paragraph, some options contradict the others. On the bottom of the page we can also select the aspect ratio (horizontal, vertical, square) and the quality.

Once we are satisfied with our selections, we can click the “Generate” button on the bottom of the page, to create the image.


Our “dream” has been created, using the modifiers we selected (at least some of them, as evident in the resulting image).




By clicking “Save” on the top of the window, our image will become part of our online collection.



Next, we will remove the background from our generated image. We click on the robot image to view it in the Editor, and then we select “Remove Background” (image below).




By clicking the “Submit” button, the background is removed, at a cost of 10 energy points. However, the end result is quite stunning!




Bing Image Creator

Microsoft has been investing heavily in AI, integrating functionality with its web browser, search engine and other products and services. The Image Creator is based on DALL-E 3, and has been integrated with the Bing engine.



To access and use the service, we can visit the Image Creator website: <https://bing.com/create/>

In order to use the Image Creator, we must first either create a new account or use an existing one (please note: school accounts for Office 365 are not supported at the moment of writing this ebook).



Like most AI generated content, Bing Image Creator uses its own credit system called "Boosts". For every image generated, boosts are used. However, they are regenerated on a monthly basis (at least at the time of writing this ebook).


Explore (and modify!) Ideas!

Before we create our first image, we can explore the available creations of other users. By clicking on one of the images above, we can view not only the picture itself in a larger size, but also the description used to generate it. This is very useful since we can use the same description with minor changes to create variations of the same image!




Generating our first image

To create our first image, we type a description and click the “Create” button. Image Creator will usually create four different images, using one of our available “boosts”.



We can click on each one of our available images to view them in larger size. All our creations have a small “Bing” logo on the lower left side of the image.

As we can see (top right), each created image has a size of 1024x1024 and can be Saved or Downloaded. We can also customize them, as we will see in the following pages.




We can share our image by clicking on the respective button above: a link will be created which we will be able to send to others.

By clicking “Save”, we can save the image in our online “Collection”. We can create collections of our images in order to store them online in an organised fashion.

We can download the image to our computer by clicking on the “Download” link. It will be downloaded to our default location (usually the Downloads folder on our computer).


Customising an Image

An image generated can be further customized using the Microsoft Design Studio. What used to take hours, if not days, in tools like Photoshop, by using AI can be done in mere seconds!



When we click on the Customize link (image above), the Microsoft Designer Studio is launched, offering various tools that allow us to apply filters, effects, and almost endlessly customize our (or at least our AI) creations.

These vary from simple edits, such as resizing or cropping, to more complex ones such as background removal, one of the most tedious tasks in image manipulation.




It should be noted that Microsoft Designer Studio is far more powerful and feature-rich than just a photo - editing application. As a matter of fact, it can be used to create entire brochures, designs or even booklets. However, the features provided by the Designer Studio are outside the scope of this ebook.

We can experiment with the tools provided by the Studio to customize or enhance our image. We must always keep in mind that not everything works perfectly. For example, complex backgrounds are not easily and automatically removed from our image!

What we learned so far. . .

In **chapter 4 “Creating Images”**, we were introduced to the almost magical capabilities of various engines that generate images, before settling to bing.com/Create which is based on DALL-E 3. We generated images from a basic description using Stable Diffusion Web, and also explored the capabilities of Deep Dream Generator and DALL-E. We improved our descriptions to make changes to our generated images, and even used Microsoft Design Studio to make further (manual) changes to an image. We have also saved and exported our results as image files.



Further Activities:

Here are some activities you can test with any image generator. You can try these at school with students

Activity 1: Descriptions!:

In order to create an image, you have to give an accurate description. This might work as an incentive for students to create written text - even a paragraph- and then feed it to an image generator such as ([bing.com](https://www.bing.com)) Create.

A small house on the prairie with children playing with swings in the grass around it, while one child reads a book under the shade of a nearby tree.


A generated image based on the above is this:



Activity 2: Describe an image

We might work a bit differently: as students to describe (on paper or on a word processor) a given image with as much detail as possible.

We then feed the images to the image generator and we compare (can be made into an unofficial “competition” between students) to see which ones are more accurate (closer to the actual image).



Activity 3: Guess the description

This activity can be played as a game: students play in teams. Each team creates a short description (might be one sentence or with a word limit). They then create an image using that sentence. Students in small groups try to write the sentence they used for the image. We try to find the most accurate description, compared to the original one.

Activity 4: Photo stories!:

This is a bit tricky, since image generators are not that good with consistency between generated images. For example, an artist would create characters and use the same ones throughout a story. For example, Mickey Mouse, even with minor changes, is recognisable across decades of use.

We ask students to create a short (photo) story with images and small paragraphs. These short stories can be a simple fairytale, with their own characters or from those they already know (ie Little Red Riding Hood).

Once upon a time, there was a little green grasshopper. He wore blue pants and carried with him an electric guitar. He liked to sit under trees and sing from dusk till dawn.



Suggestion: The actual text might be in the third person, but the description must be written differently in the generator - after all, we do not want the text to appear in the image - not to mention it is rarely written correctly!

Activity 4: Describe yourself!

This activity is a favourite of students, especially in the upper primary schools. We ask students to describe how they look (hair, nose, eyes, ears, mouth). They then generate the image using their descriptions. What is interesting is that the AI generate can mimic many different styles. So we can create our image in photorealistic, cartoony or even LEGO form!

This is also a good activity for students to improve their description so that the end image is closer to how they actually look!

We can also play a game: use all the generated images in a presentation and try to see if the students can guess who is whom!

Part B: Hitchhiker's Guide to the Galaxy of Education

What's to learn

In Part B, we focus on practical uses of AI in education. We focus on tools that enable the integration of AI for preparatory teacher work (Worksheets, presentations, image creation) and will detail the pros and cons of each tool presented. Since there is an abundance of tools, we understand that it is impossible to cover each and every one. Therefore, we focus on specific tools that enable teachers to complete time-consuming preparatory tasks.



Chapter 5: “Simple” Maths!

“What do you get if you multiply six by nine?”

“Six by nine. Forty two.”

“That’s it. That’s all there is”

(The Hitchhiker’s Guide to the Galaxy)

In this chapter...

In **chapter 5 “Simple Maths”**, we take a more in-depth look on how to use ChatGPT to prepare material for teaching Maths. We also use AI-powered Graph tools to create -you guessed it!- graphs! So, without further undo, We will:

- Create variants of problems using ChatGPT
- Generate solutions for maths problems
- Create explanatory instructions
- Create lesson plans
- Solve equations
- Create worksheets
- Create graphs using an AI tool




AI for Maths

Maths... love 'em or hate 'em, they are -in many countries- what is called 'a core subject'. And even the most experienced teacher would like to be able to offer a variety of different material to his or her students, instead of relying on the same old, tried and tested, worksheets, instructions and material created over the years or found off the internet.

ChatGPT can enable experienced -or even less experienced- teachers to create innovative variants of their existing work, and also to cut down on time required to prepare new lessons and new activities.

Over the next few pages we will use ChatGPT to create such material and examine the pros and cons of using it to generate content. We will also focus on how to improve results from our instructions, to enable the delivery of the best possible generated outcome.



ChatGPT & Tutoring


ChatGPT offers amazing capabilities for explaining complex problems. However, If a child exhibits significant learning difficulties, then the teacher must rely on his or her experience to provide the best possible assistance, tailored to the needs of the child.



...simple mathematics

ChatGPT is a real powerhouse when it comes to mathematics. As we have briefly experienced in the previous chapter, it can easily explain how to solve problems. But it can do much more than that, from the simplest task to the most mundane we can think of. For example:


“5+2”




Very politely, it calculates the sum of 5 and 2. Similarly, it can solve much more complex calculations than the above, and even explain us how it did it.

Let us try again, but with prime numbers:

“Calculate the sum of the first and the tenth prime numbers”




Not only it has solved the equation, but it also explains, step by step, the entire procedure. And, of course, we can continue with even more complex instructions.



...just worksheets!

Anyone who teaches mathematics really gets bored with having to create different worksheets for practice (and yes, we need practice...). But, just a simple command in ChatGPT can generate as many exercises as we need:


“Create a student worksheet for first grade, with simple multiplications, additions and subtractions (at least 10 each)”



We can copy and paste the above in a word-processing document, and provide all the necessary changes or customisations. We can even generate images from another AI-powered application (such as DAL-E) to enrich our worksheets with images!

Of course, we are not confined with primary school-level worksheets:

“Generate a student worksheet for high school mathematics with 10 differential equations”




Creating Maths problems

In the previous chapter, we asked ChatGPT to explain (and solve) the following problem:


“Mary has a number of pens. George has twice as many pens as Mary. Helen has 5 pens less than George. All the pens are 45. How many pens each person has?”

ChatGPT can do more than that: it can use the text of a problem to create (almost) infinite variants! So, we (kindly) ask ChatGPT to generate 3 similar problems:



This is an amazing capability, since it significantly reduces the amount of time that is required for a teacher to write problems of a specific type.


Of course, we can change the parameters of the problem and define more elaborate (or more specific) problems. For example, we might ask ChatGPT all the numbers to be between 200 - 500:



Even though Jessica has 6 books fewer than Mike, the number of books for each person is in the 3-digits range.

In the previous page, we used a specific problem to create variants. Let's just give some names and numbers to ChatGPT and see what it can generate:

“Create an algebra problem with Mike, Helen, Rick and Sofia, all holding apples, with three digit numbers”



Not surprisingly, ChatGPT has created a problem with the above simple instruction. Now, let's try to make it a bit more vague:


“Create an algebra problem with 3 different people, using at least one object, with numbers from 100 to 200”

Again, the problem is generated and it allows us to copy & paste it in other programs (ie word processor).

We will give ChatGPT another message:

“Generate a student worksheet with 4 variants of the above problem”

Lo and behold! A worksheet appears in mere seconds!




Adapting problems to student capabilities

A significant, yet extremely time consuming task, is creating problems that are adapted to the needs of all our students. In a typical classroom, we can have as many as 5 different levels of students (mixed-ability classrooms). ChatGPT enables us to generate multiple-level problems in seconds!



Problems generated by ChatGPT can even include geometry. Let's ask ChatGPT to give us two problems that require the pythagorean theorem:

“Generate 2 problems that make use of the Pythagorean Theorem”



ChatGPT has generated (surprise! Surprise!) two problems, and is also offering their solution! But, let's take a look at Problem 1: instead of meters, it uses feet! This is because in some parts of the world (USA, mainly) they still use the Imperial system instead of the Metric. To avoid this, we can specifically ask ChatGPT to use the metric system:

“Generate 2 problems that make use of the Pythagorean Theorem but only use the metric system”

As expected, two new problems are generated, but with the metric system (all measures in meters or centimeters).

We can again refine our instructions with more options. For example, we can specifically state that we need only integers, or we need all the measures to be in cm (centimeters) instead of m (meters). The flexibility allowed is tremendous, and it is up to the teacher to decide how he or she prefers the wording in the problems generated by the AI.

The Imperial (British) system

The Imperial system, also known as the British Imperial system, is a system of units of measurement that evolved from the medieval weights and measures systems of England. It was later refined and standardized during the British Empire. While the Imperial system has been largely replaced by the metric system in many parts of the world, including the United Kingdom, it is still used in some countries, notably the United States.



The Imperial system includes units for measuring length, weight, capacity, and temperature, among other quantities. Some common units in the Imperial system include inches, feet, yards, miles, pounds, ounces, gallons, quarts, and pints.


For example:

- **Length:** 1 foot = 12 inches, 1 yard = 3 feet, 1 mile = 1,760 yards
- **Weight:** 1 pound = 16 ounces
- **Capacity:** 1 gallon = 4 quarts = 8 pints
- **Temperature:** Fahrenheit is the primary unit of temperature measurement in the Imperial system.

Creating notes for students


It is often necessary to create notes for students, with short tutorials on how to solve specific assignments. For example:

“Give simple explanation on how to add two fractions with a different denominator”



ChatGPT is more than happy to comply with precise (and sometimes not so...) instructions on how to add two fractions with a different denominator (or almost anything else we ask it to).

Similarly, we can give more parameters to ChatGPT, for example we can limit the instructions to 200 words, we can ask it to provide a number of examples, and even generate some activities for practice by students:



Double check the instructions

Even though in most cases the instructions will be quite accurate - especially for simpler tasks such as the example shown here, it is advisable that you read them thoroughly and make changes or adapt its content to the needs of your class. This might appear a bit time consuming, but it sure is much better than writing the instructions from scratch!



Creating lesson plans


Lesson plans are a time-consuming task but are very very important in our teaching- they allow us to organise our lesson, find the best available resources for our lesson and our students and help us keep a close watch on time - however, it is understandable that teaching is not a food recipe and we need to make changes on-the-fly.

ChatGPT can create complete lesson plans that include all the suggested material we can use, breakdown of activities, including time allocated to each, as well as follow-up activities and homework.

In fact, we can even make changes to the requirements of the lesson plan to add even more activities, remove or reduce homework, break it down into specific number of lessons etc. Even a simple command can create a very detailed and thorough lesson plan:

“Create a lesson plan on how to teach addition of fractions with the same denominator to primary school students”

As shown in the picture (top right), the lesson plan has been adjusted for primary school students and is very detailed in the breakdown of activities.



The teacher advantage

ChatGPT can create amazing lesson plans – and we have tested it thoroughly with subjects we (a small group of teachers and volunteers contributing to this book) using even trivial subjects for lesson plans. In most cases, lesson plans are almost ready to be taken to the classroom – in other cases, minor changes are required or even adjustments of the material. Combined with the ability to generate worksheets, lesson plan creation can have a profound impact on the time we consume on planning our teaching.

Experienced teachers will always seek to create content that caters to the specific needs of their students. At present, ChatGPT tends to go the other way round (and that's understandable).




No, we can't do graphs


How about we ask ChatGPT to create a graph?

“Create a graph showing the preferences of boys and girls in video game genres”

At the moment, ChatGPT cannot create graphs or other visual content (at least not the free, current version 3.5). What it can do, however, is provide information on how to create graphs in other programs such as spreadsheets (Microsoft Excel, OpenOffice, Google Sheets).



No, we don't do research




ChatGPT will also “refuse” (or very diplomatically indicate that you should refer to other sources) for data on specific subjects. Despite this, it will provide some initial information and estimates, based on the data it has (version 3.5 has access to data up to 2022). It is important to understand that an AI, even though it is capable of suggesting bibliography, will often point out the obvious: read academic journals, official industry reports etc.

Yes, we do sample data!

Ok, ChatGPT 3.5 cannot generate real information from research projects or create graphs. BUT, it will definitely create sample data for use in worksheets, examples, or activities. Let's take check this example:

“Create sample data in a table, showing an imaginary computer company selling computers, printers, and software, with sales in euro currency, in consecutive years from 2018 to 2020”



In the image above, we can clearly see the data created by ChatGPT, based on the information we provided.

Keep in mind that version 3.5 cannot create Word-like tables, but rather uses text to create columns and lines for the data. However, this is still a very flexible way of generating tables with data fast, for use in activities, examples and worksheets.

By clicking on the “Copy code” button (image left), we can very easily copy all its content and then paste it in another application (see table below):

| Product | 2018 Sales (€) | 2019 Sales (€) | 2020 Sales (€) |

|-----|-----|-----|-----|

| Computers | 500,000 | 600,000 | 700,000 |

| Printers | 120,000 | 130,000 | 150,000 |

We can very easily change the appearance of the data. For example, we can select what appears as a column heading, we can limit the amount of sales to 3-digit numbers, or we can even add another column with a total sales in euro!

AI charts!

So, we need to use charts... maybe we need to create a graph that shows data to our students for a project, or maybe we need to include a chart as part of a worksheet. One way to create charts, of course, is to manually use a spreadsheet such as Microsoft Excel, OpenOffice Calc or Google Sheets. Or even a dedicated Charts application. But, again, what we need here is to reduce the time we require to create not just one sample chart but -probably- many!


As we have found out in previous pages, ChatGPT 3.5 can do a LOT of things, but cannot create charts. However, there are many tools available -and some have free features- that allow teachers to create beautiful charts in mere seconds. One of these tools is **Highcharts GPT**.

To access Highcharts GPT, we can use any web browser and type <https://www.highcharts.com/chat/gpt/>.

Chart creation



Highcharts GPT has many advantages as a tool, since it can generate variants of the same chart. It can also create various types of charts using the data we provide, or its own sample data. Data provided by the platform are not accurate and, even though the website -and wisely so- indicates it, we should make sure we identify any information that might lead to misconceptions (ie percentage of girls using computers).




Highcharts GPT is not a free tool - however, we can use most of its functions for free. For educational use, and especially for creating graphs for use in worksheets at school, it is exceptionally easy to use and does not require a licence like other equally powerful alternatives. It must be stated that other similar solutions exist, including ChartGPT and others.

Creating our first chart


We will create a piechart with imaginary data of students playing video games of ages 12-16, for 2022:

“Create a piechart for people of age 12-16 playing video games based on gender, for 2022”



It also gives us notice that the data used is just for demonstration purposes and not accurate! We can change the instruction to reflect more accurate (or more suitable) data:

“Create a piechart for people of age 12-16 playing video games with 45% girls and 55% boys, for 2022 and show the title “Percentage of boys and girls playing video games” and show the percentages of each category”




Highcharts GPT has created a chart with the information we provided. Notice how it also generates computer code!

Chart variants


...what is the use of an advanced AI if it can't produce (automatically) variants of our charts?

“Create 3 different variants of the above chart”



Variants are created automatically, along with their respective code. Of course we can give more specific information and even change the data for each chart or even choose the colors it will use. The variants we can generate are near infinite, based on our choices and the capabilities of the AI, as we can see in the next example (image top right). This will allow us to illustrate our points better, but also to make our charts more appealing.

“Create a piechart with the title "Preferences on movies of high school students", with 4 different categories, "action", "comedy", "drama" and "science fiction". Use only shades of blue for the colors of the chart and create it in 3D”.




For the above chart, we can (again) customise the data. For example, we can specifically state the percentage of each category, add a fifth ("other") or even allocate a specific percentage to one of the categories and let the AI decide on the rest.

As we will see next, we are not confined only on piecharts. We will create more types of charts next.

We need to show the number of sales of a computer company for every year (2018 - 2021). We give the following instructions:


“Create a column chart of computer sales, from 2018 to 2021, with various numbers per year, using 5-digit numbers and the euro as a currency.”




This is a very simple column chart, in which we can clearly see our company going better and better year over year. However, we can give more specific instructions in order to affect the fate of our company (for example, we can

specifically state that the chart must show that the sales of our company are declining:

“Create a column chart of computer sales, from 2018 to 2021, with various numbers per year, using 5-digit numbers and the euro as a currency, with declining numbers.”




We can also get more adventurous and ask for a random 3D column chart, with at least 4 columns:



Similarly, we can create a line chart showing the annual sales of a burger franchise, from 2010 - 2022:

“Create a linechart showing the annual sales of a burger franchise in millions of euros, from 2010 to 2022”



Highcharts supports many other types of charts that we can use. Some of these are:

1. Line Chart: Used to display trends over time or continuous data.
2. Column Chart: Used to compare data across categories or discrete data.
3. Bar Chart: Similar to a column chart but with the x and y-axis flipped.

4. Pie Chart: Used to represent proportions or percentages of a whole.
5. Area Chart: Similar to a line chart, but with the area below the line filled.
6. Scatter Plot: Used to display the relationship between two variables.
7. Bubble Chart: Similar to a scatter plot, but with an additional dimension represented by the size of the bubbles.
8. Polar Chart: Used to display data in a circular plot, often used for radar charts.
9. Box Plot: Used to display the distribution of a dataset using quartiles.
10. Heatmap: Used to represent data using color gradients in a grid-like structure.
11. Treemap: Used to display hierarchical data in nested rectangles.
12. Funnel Chart: Used to visualize stages in a process or sales funnel.


Chart variants

Highcharts supports many more variants of the above chart types. By customising the various types, we can enhance not only their appearance but also their functionality. However, we must always keep in mind that, the simpler the chart, the easier it is for the data to be readable! Don't get overexcited with the color and/or 3D options!



Using charts in worksheets

The creation of a beautiful, functional chart is just the first step- the next is to be able to incorporate our chart in our worksheet! As we have seen in previous pages, we can use ChatGPT to generate the content (text) for a complete worksheet. We can also use Highcharts GPT to create amazing charts with data, in a fraction of the time we would need in a typical spreadsheet (or dedicated chart creation) application.




A simple right-click on the chart will not do; we will not have the capability to copy the chart as an image. However, we can save the chart as an image using the Chart Context

Menu (image bottom left). Interestingly, one of the options is to view the chart in full screen, or to print it right away.

However, the most important option is to save the chart as an image - and here, Highcharts is very generous, as it allows us to save it in various formats, including SVG - a vector graphics format that allows us to enlarge the chart as much as we want without losing detail.

Once we select the file type, it gets downloaded to our computer, and we can import it from there in our favorite application (word processor, presentation etc). Keep in mind that a small "watermark" will appear in the bottom right of the chart (not a problem for educational use).



What we learned so far...

In **chapter 5 "Simple Mathematics"**, we took ChatGPT on a maths tour! We used a problem from a public secondary education textbook and we asked it to explain in simple terms how to solve it. It then generated similar problems for more practice, including their solution. We went a bit further and asked it to give explanatory instructions for each one, and even suggest the prerequisites and how to learn more about them. We also used an amazing AI tool for creating graphs. We could generate highly detailed graphs based on simple descriptions. These graphs can then be used in worksheets or presentations.



Further Activities:

ChatGPT is an immensely helpful tool when organising and preparing for lessons. Try these suggested activities to prepare for your lesson:

Activity 1: Generate Lesson plans

- Organise your weekly or monthly lessons. Generate lesson plans that will help you make the most of your available time. Include activities for students of different levels and abilities (for mixed ability classrooms).

Activity 2: Generate worksheets

Based on your lesson plans, create activities for each day, either to complement the ones found in your textbook or as extra or differentiated work to cater to the needs of all your students.

Worksheets can include activities with solutions and explanations. You can even ask ChatGPT to provide you with links to websites and videos that provide explanations (ie Khan Academy).

Activity 3: Teach maths through graphs

- Generate graphs either with real data or with imaginary data. If you use imaginary data, please make sure you check them carefully to avoid biases.
- Use graphs as an activity with students. They can practice discussing the results of research based on graphs.

Activity 4: Generate graphs

- Try to engage students in research -even on a very small scale (for example, their preferences in movies, music, food, clothing, how often they replace their smartphones etc).
- Use this data to generate graphs using the AI tool (we can even generate graphs with “imaginary” data, just to explain how to “read” graphs
- Use graphs in presentations or posters or on short projects. Students should discuss the findings of the research based on the graphs.


Chapter 6: Maths for students

"Math, Science, History,
Unraveling the mysteries
That all started with the big bang"
(BNL, "Bing Bang Theory")

In this chapter...

In **chapter 6 “Maths for students”**, we use ChatGPT to help students understand maths and utilise it as an “at home” personal tutor! In this chapter, we learn how to:

- Ask ChatGPT to explain how to solve a problem
- Generate worksheets for practice
- Create a “do-it-at-home” evaluation




Maths @ home

Students tend to spend a lot of time at home working on assignments. One of the benefits of AI is that it brings the closest you can get -at this point- to having an actual tutor at home. Taking into consideration that even the most advanced AI is not even near the presence of an experienced human tutor, we will focus on how to make the most of ChatGPT to enable students to get meaningful explanations. Or, at the very least, help them on how to ask the correct questions.

Maybe it's the maths problem that we cannot understand how to solve - maybe it's the equation that we have been trying to solve in vain for the past 30 minutes. ChatGPT can use our information to try and give us some kind of help. At least until we return to the classroom the next day and raise our hands to get help from our teacher!

With full responsibility that no AI (at the moment) can explain as clear and concise as a human teacher, we will offer some tips of using ChatGPT at home and maybe -just



maybe- we might start improving our engagement with our homework!

Homework that has a meaning

Homework can have positive results for learning, as evident by the work of Epstein & Van Voorhis (2001), when given purpose. Therefore, it is absolutely necessary that students understand exactly what they have for homework and why they have it - and practice on maths is not only desirable, but necessary.




ChatGPT as a tutor at home

You are a student. The school gave you some homework. You were paying attention while in class, and you even managed to complete the classroom assignments. However, for some reason, you find it difficult to work with the homework assignments. Maybe they are more difficult. Maybe you forgot some basic information needed. Or maybe in the class, your teacher was providing assistance and you received help while working on your own. One solution is -obviously- to move to the next assignment and ask your teacher for help the next day (and this is what every teacher will suggest you do!). But, you want to give it another try - that's where AI can come in!

The assignment can be as simple as solving a primary school problem or as complex as having to work with differential equations for your last semester at school. You might have tried your luck on YouTube -and yes, there are amazing tutorials on almost any subject there- but, for some reason, you find it hard to understand how to work with your assignment(s). Chances are, ChatGPT can provide simple steps that might help you. And if the first feedback fails, you can always try another. Let's take a look at a simple example used at schools in Cyprus* :

“Calculate the value of $m \in \mathbb{R}$ so that the lines (e1): $(2m+1)x-3my-2=0$ and (e2): $3mx+(m+2)y+1=0$ are perpendicular”.

As stated, this is an example taken from the textbook which already has explanation of the solution. Let's see how ChatGPT might explain the solution:




ChatGPT has provided a very detailed, step-by-step explanation on how to solve the above assignment. These explanations are usually very accurate, and you can always double-check your results the next day in class.


* Example retrieved from https://archeia.moec.gov.cy/sm/268/a_lyk_mathimatika_pros_2.pdf

Let's suppose we still have trouble understanding the explanation given to us. We can ask ChatGPT to provide a simpler explanation:

“Can you provide a simpler explanation of the above problem?”




ChatGPT complies with our request and simplifies the first explanation, taking a more detailed, yet simpler, step-by-step approach. We can also rate this solution, helping improve the feedback we receive even more (photo top right).



What ChatGPT does, is try to change the way it offers instructions every time we ask it to- until it provides instructions that we (might) find simple enough to understand. Let's try again:


“Can you provide an even simpler explanation of the above problem?”



Understanding prerequisites

Sometimes, in order to solve (or understand) how to solve an assignment, we need to fully understand the prerequisites (ie the previous knowledge needed). In this case, we might have a gap in understanding what some of the prerequisites of the previous assignment are. Let's ask ChatGPT about this:


“What are the prerequisites of the above assignment?”



The above reply is extremely helpful- if we cannot recognise what points 1-4 are, then we might not be able to fully understand how to solve the assignment.

Fortunately, ChatGPT can offer help in this. We can either ask ChatGPT to explain our prerequisites (ie “Equation of a Line”) or offer suggestions on where to find online tutorials. And we will do just that:

“Can you recommend an online tutorial on what Equation of a line is?”



ChatGPT is polite enough to provide a link to the tutorial!

Human Vs AI tutor


ChatGPT is incredibly flexible, fast, and often precise when asked questions on assignments. However, at the moment it is far from an experienced -or even a beginner but knowledgeable- human tutor. We suggest you practice using ChatGPT for simple assignments and explanations, but always come back to your school teacher for explanations.



Practice makes perfect


Let's say you have a good understanding of assignments. But, you need a challenge! Or, at least, a bit more practice:

“Can you generate at least 2 similar problems?”



ChatGPT complies and provides not only two similar problems but also suggestions on their solutions. Let's try to make things a bit more challenging for us:

“Generate 3 similar but more complex problems. Do not show the approach or the solution”



Our friendly neighborhood AI has done its best to provide us with 3 more problems but with no suggestions on how to solve them. We can always make changes and request simpler or more complex problems, with or without solutions. And we can always ask ChatGPT to check if we are correct in our assignments.


The trouble with Geometry

ChatGPT can have a profound effect when used correctly and within its limits as a home personalised tutor. However, we must accept and understand that it does have its limits. Especially with subjects that are reliant on images (even simple geometrical shapes), ChatGPT 3.5 cannot provide visual feedback or information. It might be almost perfect for algebra, for example, but not that perfect for geometry.



What we learned so far.

In **chapter 6 “Maths for students”**, we focused on how students themselves can benefit from the use of ChatGPT, especially when we prepare material for them. In this chapter we used ChatGPT to explain (to us) how to solve a problem and then we generated worksheets for further practice. We also used ChatGPT to create worksheets and activities for use at home, including solutions and explanations.



Further Activities:

Here are some activities you can test with any image generator. You can try these at school with students

Activity 1: Problem variants

- Identify problems that might impose a challenge to your students. Ask ChatGPT to create variants to these problems for further practice by your students. Make sure you check them for any inaccuracies before you hand them out to students (for example: “a factory employs 10 men, 12 women and 3 children for carrying heavy weights from point A to point B” - obviously such a problem or scenario should not include children!).
- Generate the solutions for the above problems and print them on a separate paper.
- Generate the explanation for the above problems and include them -if necessary- with the solutions

Activity 2: Students generate problems

Allow students to generate problems and activities for their own practice. Teach them how to use ChatGPT to this purpose, it will be helpful for practice at home!

Activity 4: Identify prerequisites of your students

- Based on student evaluations -either through classwork, homework or assessment- locate the prerequisites required by your students and prepare worksheets for them.
- Offer additional guidance to students based on the prerequisites - ChatGPT can even generate a list of sources or explanations that might be helpful!

Activity 4: Create evaluations for students

- Evaluations take a long time to prepare, especially when we take into account the complexity of the exercises, the knowledge and skill that we need to evaluate, as well as the needs of our students.
- In mixed-ability classrooms (most EU countries in the south) we have to prepare evaluations for 2-3 different levels of students.
- Use ChatGPT to generate evaluations in a fraction of the time it would require otherwise. Based on these evaluations, create preparatory work, as well as exercises after the evaluation, for those students who require further practice.

Chapter 7: ChatGPT for Writing

"Come writers and critics who
prophecies with your pen
And keep your eyes wide
The chance won't come again"
(Bon Dylan "The times are a-changin")

In this chapter...

In **chapter 7 “ChatGPT for Writing”**, we will:


- Generate lesson plans for language teaching
- Create instructional material
- Generate exercises
- Create original text and related questions
- Learn how to create summaries based on text



AI Essay Writing

Let's admit it: even though there are literally thousands of books we can borrow one or two pages for use in our classroom, most often than not the need arises to find a piece of text that is very specific to a specific lesson. And -let's face it- we need to constantly use new material and not regurgitate what we have been using for the past 10 years.

As an example, let's look at the ongoing discussions of using tablets instead of books. That is a solid case for discussion in the classroom. And one approach is to do just that- have a discussion with the students and write down all the arguments. Another approach is to give students a piece of text on the subject, have them read it (on their own or in groups) and then discuss it in the classroom (critical approach to reading). Then how about we find related articles for further study of the subject? Since this book is addressed to teachers, I am not going to "preach to the choir" on the benefits of this approach. However, we all know how time consuming it is to organise the information



required for such a lesson - let alone duplicate this approach even once or twice per week. Fortunately, we have AI tools for that!

Review & Edit

AI has proven to come a long way in the past few years, with incredibly accurate information given on most topics. However, especially when delivering this information to students, we need to review what is generated by the AI tool we use and make changes wherever we feel there is ambiguity or the information might not be correct, or need to adapt for our students.



Creating stories...

ChatGPT can do amazing things. But how about we try to see how it can make our lives easier at school (and during preparatory work).

Let's look at this scenario: you need to create a short story for your students. This story has to do with a specific theme you would like to explore or discuss. Obviously, you could seek information or other written works on the subject. Obviously, you can do that as well (who said that you need to pick one over the other, after all?). However, nothing beats visualising what you already have in mind. Let's suppose you want to offer a short story that ends at a dilemma: that would be an incredible point for discussion in the classroom.

"Create a short story (maximum 300 words) about two siblings who find a treasure. One sibling suggests they give that treasure to the police or local authorities. The other to keep it. End the story with a moral dilemma on what to do and why."

ChatGPT will provide a story based on what we asked it to do. In the next column we can see the entire generated story! We can make changes to it, or further refine it using ChatGPT.

In the attic of their grandparents' creaky house, Mia and Ethan discovered an ancient chest, brimming with gold coins and sparkling gems. Mia's eyes widened with amazement, while Ethan's mind raced with the possibilities of what they could do with their newfound wealth.

"We should give it to the police," Mia insisted, her sense of justice shining through. "Imagine how many people this could belong to. It's the right thing to do."

Ethan, however, had a different perspective. "Think about what we could do with all this treasure! We could buy anything we want, go anywhere we dream. It's a chance of a lifetime!"

The siblings found themselves at odds, torn between morality and the allure of a life transformed by wealth. Mia argued for the greater good, emphasizing the potential impact on the lives of those who may have lost the treasure. Ethan, tempted by the promise of a life beyond their wildest dreams, hesitated to part with the chest.

As the moral dilemma intensified, Mia and Ethan grappled with their values. Mia couldn't shake the thought of someone out there who might be desperately searching for what they had found. Ethan, enchanted by visions of luxury and adventure, felt the weight of the decision pressing on him.

In the end, they agreed to take a night to think it over. The attic echoed with the ticking of the old clock as Mia and Ethan pondered their choices. The next morning, they sat in silence, contemplating the moral crossroads before them.

Mia, with a heavy heart, confessed, "We have to give it to the police. It's the right thing to do." Ethan, torn between dreams and responsibility, nodded reluctantly, realizing that some treasures come with a price that goes beyond gold and jewels. And so, hand in hand, they set out to return the treasure, uncertain of the impact their decision would have on their lives.


dreams and responsibility, nodded reluctantly, realizing that some treasures come with a price that goes beyond gold and jewels. And so, hand in hand, they set out to return the treasure, uncertain of the impact their decision would have on their lives.

Creating questions

In the previous page, we created a short story regarding a moral dilemma. But how about we create some open-ended and closed-ended questions for it?

“Can you create 3 open-ended and 2 closed-ended questions for the above story?”


As you may have guessed, ChatGPT immediately creates some questions! Of course we can always refine them or replace them with our own.




These questions can be a part of written assignment, group work or even for discussion in the classroom.

How about we add some variety in the assignments? How about some True/False questions:

“Create 5 True/False questions for the above story



“Create 5 sentences based on the above story in mixed chronic order”





And here they are! A set of sentences in the wrong chronic order for us to sort!

Creating story variants

Sometimes it is very important to have different texts of similar stories or information. The benefits are numerous, as they allow students, especially when they work in small groups, to identify similarities or differences between different stories or even the characters. We can create as many similar stories as we want, with a simple request:

“Can you generate two similar short stories?”



We have total control of the narrative of each story - for example, we could ask ChatGPT to add just another character in one of the stories, to create alternative endings for each one, to offer “bad” endings etc. Of course, we can always ask it to generate even more variants of the stories, even though it is questionable how much value it might have if we flood students with 10 variations of a text.

What should be important is to focus your attention on what you really want to emphasize with your students - for example, is it comparison between stories to find similar values / characteristics of your “heroes”? Are you trying to compare endings of each story or the moral(s) found in the text? All these must be taken into account when designing your lesson and -of course- prior to generating story descriptions.

In the next pages, we will focus more on developing assignments, lesson plans, grammar activities, images for stories and more!

Consider the environment!


Unless you can provide all students with mobile devices or access to a computer (or at least one device per team), having multiple different stories can require that you print and use large amounts of printer paper. Therefore consider the environment before you get carried away and you print stories upon stories for your students!



Creating worksheets

We have the multiple stories, we have already discussed them with our students, even created some nice activities on our own. How about we (kindly) ask ChatGPT to create worksheets?


“Generate a worksheet based on the three stories”



questions, open-ended reflection, true/false statements and a discussion prompt we can use in the classroom, or as assignment for home.

Since we have 3 similar stories, we can ask ChatGPT to generate questions that compare all three stories together:

“Create questions that compare all 3 stories together”




It almost instantly generates three different worksheets, one for each story -depicted with its own title- with 4 different type of activities per story (Comprehension

ChatGPT has created categories of questions that enable students to compare and use information from all 3 stories. We can modify them as we please!

Summarising books!

Imagine how many times you needed to work with the main content of a book - let's say one of the classics like "Oliver Twist". What you usually have to do is either quote a part of it (useful, and we do it all the time) or have to write a short summary yourself. How about we ask our friend ChatGPT to create, in first person, a summary of said book, in 800 words and less?


"Write Oliver Twist, in the first person, with no more than 800 words"



And yes, here is the summary (actually it's on the bottom left of the page) of "Oliver Twist", written in the first person!

Let us suppose, however, that the text produced is too difficult for younger students -or the opposite! We can ask ChatGPT to re-create it:

"Simplify the above story for primary school students"




A simplified version of the story is generated, more suitable for younger students!

Grammar activities


In the previous page, we generated a short version, in first person, of "Oliver Twist". We can create a worksheet with questions on the text. But, we need to get our hands on some grammar activities!

"Create grammar activities based on this text"



It automatically creates 10 different type of grammar activities, based on the given text (ie "Possessive nouns"). We can focus on one or more of these and create more specific activities:


"Create a worksheet on possessive nouns for the above story"



And here it is! A complete worksheet with some great activities! We can ask ChatGPT to provide us with the solutions as well, or to continue generating more similar activities! The possibilities here are quite remarkably endless; however, it must be stated and emphasised that -at no point- we should just copy and paste (and print) these activities, without reading them first. After all, education is not about using canned food but rather about adapting "food" to the needs and requirements of our students!

What we learned so far...

In **chapter 7 "ChatGPT for writing"**, we worked with activities that enable teachers to more easily prepare for their lessons, including lesson plans and worksheets. We have created exercises and activities, as well as original text, based on our descriptions, and related questions based on it. We also used ChatGPT to practice writing summaries based on specific text.



Further Activities:

ChatGPT can have amazing applications in language learning and writing. We suggest some activities for use with ChatGPT:

Activity 1: Create short stories

- Select a topic you would like your students to work with. Generate multiple short stories based on the topic, and hand them over to students. Ask them to address the common theme of all stories. Students can also find similarities and differences between the stories.

Activity 2: Generate summaries of text

- Give a passage or a text to students and ask them to create a summary. Then ask them to create a summary of their summary if possible. Use ChatGPT to create a summary of the same text and give it to students to compare with their own work (**suggestion**: always check a generated text prior to giving it to students).
- Alternatively, give a summary of a text created by ChatGPT and ask students to enrich it. Prefer to use a story for writing, since it allows for more creative writing.

Activity 4: Generate stories with alternative endings

- Generate short stories using characters from history or from fictional work. Ask ChatGPT to generate multiple alternative endings.
- Ask students to select the ending they prefer and explain their preference.
- Ask students to create their own alternative ending.

Activity 4: Stories by students

- Ask students to generate their own stories using ChatGPT. They have to select a topic, the characters and the setting.
- Ask them to write the description for images that go with their text. They can use an image generator for the images.
- To create both their Stories and their images, they have to give precise instructions to both ChatGPT and the image generator program.
- Ask students to compare their stories, especially if there is a common theme between them.

Chapter 8: Presenting our work

In this chapter...

In **chapter 8 “Presenting our work”**, we will:

- Explore the Tome AI presentation app
- Create AI-generated presentations
- Generate images and text
- Add and modify slides
- Learn Pros and Cons of using Tome
- Use ChatGPT, image generator and Google Slides as free alternatives




Presentations...

Presentations are an important part of teaching (and learning). Presentations, when used properly and only when they have an added value, can elevate a lesson and provide students with visual stimuli.

However, even a short presentation can take hours to produce using normal methods - create the actual theme, find the proper images (of good quality), write the text, make necessary edits...

Now imagine that you create short presentations to use at least 3-4 times a week. That would require an additional 5-6 hours every week just for the presentation. When you add to that the workload teachers have for home (correct assignments, prepare for lessons, create worksheets) it becomes unsustainable. Fortunately, AI tools enable the fast development of presentations, as we will explore in this chapter. First, we will use a subscription-based app to develop a presentation. However, for teachers on a very tight budget or for students, we will also explore the use of




"free" tools to mimic as close as possible the functionality -or at least part of it- of Tome.

Presenting to our audience


Presentations can either attract the attention of our audience, or cause our students to lose interest -or even get lost in the information we present. What must be stated here is that presentations must be used sparingly and where they have a clear added value, as mentioned earlier. Otherwise, they can have a detrimental effect on learning.



Tome Vs Canva



Canva (top left) is a versatile, multi-purpose online application that allows even a novice user to create posters, presentations, banners and more. The latest updates have integrated AI tools that make content creation even easier. It has a free version, which allows ample customisation for the novice non-professional user, as well as a flexible subscription that greatly extends its capabilities. It supports hundreds of templates and ready-made art that help users get creative almost immediately.





On the other hand, Tome (top right) is a dedicated online presentation tool, focused exactly on that - creating and delivering presentations. Even though this ebook could feature Canva as an AI-powered presentation tool, Tome represents an excellent example of how a dedicated, focused application can be used to create stunning presentations with AI-generated content almost instantly. Its intuitive user interface and tight integration with AI text and image generation, give Tome the edge as a presentation development tool.

Creating a Tome account


Tome is an online app used to create online presentations. Much like most AI-powered tools that can be found on the internet, Tome is subscription based but allows some functionality for free, in exchange of “credits”. Each time an AI-based feature is used (for example, generation of an image), credits are deduced from the total amount available.

Additionally, each free account can create up to a total of five “tomes” (the equivalent of five PowerPoint presentations).


To use Tome, you must first visit its website and create an account: <https://tome.app>. Alternatively, you can link an account for another service (ie Gmail).



During the creation of our account, Tome will require some information from us.



An important step is deciding on the plan: the Basic is free, with 500 AI credits.



time we use one (ie automatic generation of text or of images) we consume credits.

Contrary to how some of the other AI services work, like the bing.com AI for generating images, our credits on Tome app are not re-generated. If we consume them, the only way to gain more is by subscribing to one of the monthly or yearly services. The price might seem a bit steep, but if you create presentations on a regular basis, Tome can save a lot of time and effort, as we will see in the next pages!

Upon starting the Tome for the first time, we are greeted with the above main screen - from here, we can select a template from one of the provided (and professionally created) ones, or click on the Create button and start with a new presentation or "tome".

As mentioned earlier, the free version of Tome allows each user the creation of five tomes, the equivalent of five powerpoint presentation files. We should always keep in mind, however, that -apart from the number of tomes restriction- there is a limit to the use of AI features: each

For this tutorial, we will focus on the creation of a presentation file and not on the collaborative features of Tome. This might be a subject for another ebook. However, our focus will remain on the generation of a presentation (see following pages) on the history of computers.


Free accounts & personal information

Every service that requires we register with our personal information, has a policy on how it handles our personal data. It is always a good practice to read the policy and decide on giving access to personal information. The Tome app is shown in this book as an AI tool for presentations. Please be careful what kind of information you provide and to which companies!



Presentations using Tome


There are two ways we can create a new presentation (or "tome" as they are called): we can either select one of the templates, and use that as the basis of our presentation, or click on the "Create" button to start from scratch.



In the image above, we see the first slide of our presentation "tome", based on a selected template. Each part of the slide (or page, as they are called in Tome) is an individual object. By clicking inside one of the text boxes,


we can edit its content and use AI tools to improve writing. On the bottom left of the screen we can add a new page, while the text prompt in the bottom (middle) will invoke more options, including upload of media, page generation, image search on the web and more.

On the left side of the page, we can see sample pages of the selected template, while on the right we can find tools to add more content, including tables, charts, images and more.




Creating a new Tome

Templates are tempting, as they provide an immediate design that has a professional appearance. However, in this example, we will focus entirely on creating an AI-enhanced presentation using Tome. From the main (home) screen, we click on “Create” to start using Tome’s versatile and powerful AI tools.




The page is completely blank - we will change this in a few seconds.

From the options (image above), we select “Create presentation about...”




When we click “Generate outline” it will create exactly that:



As we can see on the screen above, we have just created an outline for our tome (presentation).


Creating layouts

If we are not satisfied with the generated outline, we can click on “Regenerate” - a set of new pages will be created. We can then click on “Generate all pages” to immediately have them appear in our tome, or click on “Choose layouts” to customise each and every page.



For every page, four different layouts are generated. We can click on each layout and then click on “Next page”. Alternatively, we can click on “Generate all pages” to have

the entire tome created, and then make changes. Note that images are also generated automatically for each page.




As we create pages, a preview of each one is shown on the left side of the browser window (image above). For each slide, text, titles and images are generated automatically by the AI engine used in the Tome application.

AI-created content

In most cases, the content created by the AI will be reliable. However, it is highly recommended that you proof read the information generated in each slide to ensure its validity and that it complies with the needs of your audience.




If we are satisfied with the creation of the pages by the AI, we can automatically generate the rest of the pages. Keep in mind that Tome uses the AI credits we have (free accounts) whenever content is generated automatically.



The image above shows all pages of our tome, including text and images! The presentations has been created in a fraction of the time it would take us to decide on the content, write the text, find images and build the presentation file using more traditional tools such as PowerPoint or Google Slides!


As we create pages, a preview of each one is shown on the left side of the browser window (image above). For each slide, text, titles and images are generated automatically by the AI engine used in the Tome application.

However, from our initial 500 AI credits, we have expended 200 already. And we still need to make some modifications to our presentation!



Text editing & AI


Once we are satisfied with the layout of our pages, we can start editing text and images. Tome offers some basic editing tools that allow us to change the font, style, color, create bullet or number lists, include links and more.



When we select a paragraph, we can make changes to it using the AI edit option (image above). We can choose to Rewrite the text, fix any spelling and grammar mistakes, reduce it or even extend it. For example, we can select to

extend the text until it fills the entire page or text box -depending on our needs. Alternatively, we can use the old fashioned way of changing the text - that is, to manually input as much information as we require ourselves!


AI can even generate -and include- references to our text. For example, maybe we want to provide more information and a hyperlink to one of the acronyms found in the slide (ARPANET). We select the word and then click on the "Reference" option (image left).



A list of related references appear - we can select which one we prefer and either click "Open" to view it, or "Add" to include it in our page.


Editing pages

Pages can be easily edited. Not only text, but any element on a page can be edited, deleted or inserted. Objects are placed within "tiles" - objects that contain text, images, graphs, videos etc. To create a new text tile, we click on the AI prompt on the bottom of the page (image below) and select "Add text tile".



A new text tile appears in our page. By clicking its header (image top right) we can make changes to its content style.

Any tile can be moved within the page. For example, if we click and hold the header, we can move it to the left or middle of the page.



To delete a tile, we right-click inside it and select "delete". We can always undo any changes we make to a tile.


AI credits and how not to use them

Any standard operation -anything that does not involve AI generated content- does not use AI credits. Therefore, we can do any changes we like to the style, layout and content of our pages without using our AI credits. However, once we start generating slides, text or images using AI, our credits are gone forever - and Tome does not regenerate them!



Generating images

Images can be either uploaded to our page, or generated. We can use the AI prompt on the bottom of the page or the “Add media” shortcut on the right of the page (image below). Image generation expends AI credits.



Once we select the “Generate” option, a prompt will appear in which we can type a description. It must be as detailed as possible, in order to better achieve the desired effect (as seen in a previous chapter of this ebook).

By clicking on the “Generate” button (above), the image will be created and inserted in the tile (note how the image tile is created automatically when we select “Generate”). A total of 4 images will be generated. We can double click on them to add the one we prefer.





Image tile manipulation

As we can see above, the image that was generated is not shown properly in the tile. We can click and drag the sides of the tile to enlarge (or reduce) its width. The other tiles will adjust their size automatically. Tome is a very versatile application in almost every respect.



Page generation


Using the prompt on the bottom (image below), we can generate a new page. We type the description and then click on the “Generate” button.



Four different layouts including text and images will be created. We can either select one of the four and click the button “Keep”, or we can click on “Retry” to generate an additional four layouts. Keep in mind that page generation consumes AI credits!

Each page that is generated can contain images and text. All images are original creations and therefore can be used freely.

Once a page has been created, we can make changes to its content - even upload or generate more images. We can also move a page to another part of our tome by dragging its preview (left part of the screen) to any order we prefer.



To delete a page (or duplicate it) we right-click on its preview and select “Delete”. We can always undo this action.


Content verification

As with any automatically (AI) generated content, we must be very careful to verify that it is correct and can be used by our proposed audience - be it school children or adults during a training/presentation. This also includes the images, since -sometimes- an image might be completely irrelevant to the text it accompanies.



Sharing & Exporting


Sharing a presentation with other people is very easy: by clicking on the Share link (image below), we can add the email of the person(s) we want to provide access to our tome. We can also select the level of access.



We can also generate a link that anyone can use to play (but not edit) our presentation created with Tome.

We can also add people on our Workspace, using the options from the main screen of the Tome website. This enables easier access of our work with our co-workers,

We can also choose to export our tome in Pdf format, but that is only supported with a Pro account for the application.



The basic tools provided to our free account are more than adequate to create stunning presentations in a very short time. However, keep in mind that the AI credits are very limited!


Using Tome for “free”

Tome, on its own, is a very powerful and versatile online presentation tool. Even if we run out of AI credits, we can use ChatGPT and bing.com/ Create (or any other similar service) to generate text and images. It will not be as flexible as being part of the same application, but it will provide us with free to use tools.




Alternative Gamma


Tome is an excellent tool. However, it is not the only kid on its block. Gamma is a very impressive tool with enhanced AI features that can be used to create not only presentations but text documents as well.



The free version offers 400 AI credits, and allows the creation of presentations with up to 10 “cards” (slides). We use text descriptions to automatically generate the content of our presentation, much like the Tome app. A single presentation consumes 40 AI credits. Gamma provides templates that produce professional-looking presentations almost instantly, including images and text!



Unlike Tome, with Gamma we can export our presentations to PDF or PowerPoint files, an invaluable feature for most users!



What we learned so far...

In **chapter 8 “Presenting our work”**, we used the Tome application to generate a complete presentation file (“tome” as its called) with images and text, including professional-looking layout, within seconds, thanks to AI.

We then proceeded to make changes to it, including changing images, text, adding more slides and presenting our work. We also used ChatGPT and an image generator to duplicate similar functionality but without the need of a monthly subscription.

We also took a peek at an alternative application, Gamma, which offers similar features but allows the export of our presentations in pdf and powerpoint formats.



Further Activities:

Creating a presentation is an opportunity to express our ideas to the public. Using AI tools, this process can become easier and faster. We can use AI tools for presentations both for our daily work as teachers and also for use by students, to practice their presenting skills.

Activity 1: Create presentations for various subjects

- A good practice, especially if you teach various and different subjects, is to generate multiple presentations and experience how AI handles them. For example, how would it create a presentation to teach acids and bases in chemistry? How would it show, in 9 slides, how the amoeba multiplies? Or the Punic Wars in Ancient Rome?

Activity 2: Same history, multiple times

- Use the Tome application or ChatGPT (or any other AI-powered tool) to create a short presentation on a specific historic subject (ie history of music).
- Create several different versions of it (based on historical fact, there should be a consistency). Ask students to compare the generated versions and find any inaccuracies or omissions (keep in mind: complex

subjects cannot easily fit in a presentation with 8-10 slides).

- Alternatively, ask students to create 4-5 different variants. Students could present their own variant of the same topic to the rest of the class.
- Students discuss the presentations and make comparisons to find inaccuracies and missing information.

Activity 3: Generate the layout of a presentation


- Generate the layout of a presentation but discard the included information.
- Ask students to populate each slide with appropriate content - this reduces the time and effort required by students to actually decide on the structure of the presentation (which is also important) and focus on the content.
- Ask students to research and write their own text for each slide (not using ChatGPT) but to use image generators to create the images
- Alternatively, ask students to locate images on the web.

Chapter 9: AI Avatars

In this chapter...

In **chapter 9 “AI Avatars”**, we will:

- Learn about avatars
- Create our own AI avatar using templates
- Generate a video avatar using our own photo, with animated lip-syncing
- Create a video presentation using avatars



AI Avatars

In the previous chapter, we discussed about presentations and showed how to use AI to generate -in mere seconds- beautiful “tomes” with pages containing text, images, even charts.

Presentations, when used correctly, can have a positive impact on how we teach, drawing the attention of our students to new information, introducing a new concept or even summarising the lesson of the day.

The use of video presentations takes this many steps further: a short video combines text, audio, animation to provide information. It can be more captivating or even interesting for the students. For decades, companies such as Adobe have created tools to do just that: create video presentations. In recent years, we have seen tools that incorporate animated people (men, women, even robots) to make video presentations.

With AI tools, we can create much more sophisticated representations of digital “people”.



We can even use our own photo as the basis for our video! And this is what we will do in the following pages...

What is an Avatar?


The word comes from Hinduism and Buddhism: when a God, like Vishnu, takes an earthly form, that is called an avatar. In fact, Vishnu had ten avatars himself! In modern day (and technology), an avatar is a digital representation of a human. It can either be something simple as our profile photo, or a 3D representation.




AI Video Generators

There are various tools available that allow the creation of 2D (or even 3D) Avatars. These can range from the absolutely imaginary (ie robots, elves etc) to actual representations -in different styles- of our own photo. For example, we can generate a cartoon image of our uploaded photographs and then set the style -oil painting, cell shaded, pencil drawn...

However, we will focus on video Avatars, using templates and photos, as we look at the various available solutions.




Most of the available platforms that enable the creation of Video Avatars are not free - they have a monthly or yearly subscription fee, and then they still restrict the duration of the video you can create. Synthesia (image bottom left) is one such platform. It features a large number of available avatar templates (140+) and supports more than 120 languages. Even though it is subscription based, it does allow the creation of free AI video.




We have a maximum of 100 characters for our message, and 3 different avatars to choose from. Once we are happy with our text, we can click on "Generate Free AI Video". It will be sent to our email address. Synthesia also has tools for editing our videos.

Polymorf has an extensive library of available avatars to choose from, and even allows us to upload a photo and turn into an avatar - either a personal one, or a photo we have created (ie using an image generator).



We also have a selection of languages for our avatar, and also a set of voices (male and female) for each language. We can preview the audio by clicking on the button next to the voice selection field. We can select an avatar from the available ones (or upload or own), then a voice, and then type the text for our digital person to read!

Once we are satisfied with the text we typed, we can click "Render" to generate our animated avatar video.




By clicking "Render", we are presented with a list of choices - we can either use the text we have typed or upload an audio file and use that with our Avatar. We can also view the dimensions of our video, its duration, and the time (credits) we have left. We click on the "Start Render" button to generate the video (note: it can take a few minutes for the video to finish).

Please note: personal information!


Most of the services and tools covered in this ebook require the creation of a new account or linking an existing (email) one. We do not encourage or condone sharing personal information with any platform. Having said that, it is up to you to decide what you link and what you share. At the same time, keep in mind that, when uploading photos to use as avatar, you must make sure you are the owner of that image.




Once the rendering is complete, we will be able to watch our video in our browser. We can also download it as an MP4 video to our computer.



The avatar not only speaks with the voice we selected and the text we typed, but also moves its lips accordingly! Next, we will upload a photo to create a new custom avatar!




To create our avatar from an image, we click on the "Create" button (image bottom left) and we select a picture from our computer. Keep in mind, the picture must have a clear face -real or imaginary- for our avatar to work.




For our example, we will just generate a photo of a friendly elf (male or female) or any other imaginary -or real- person or animal. We save the image to our computer, and then we select it (using the "Create" button).


So, now we have uploaded the image of the elf we have created, and we have selected a voice. Time to type some text and then render the video!



Our custom avatar can speak with good lip-syncing!




All the videos we create can be found in our dashboard. From here we can view them (by clicking on each one) or delete them.




To delete a video (image above), we hover with our mouse on top of its preview. The trash bin appears automatically. By clicking on it, we can delete it. Please note: deleting a video does not increase the available time credits! For example, if you delete a 6-second video, it does not mean that the time available for new videos increases by 6 seconds. Therefore, be careful before you create anything!

Lipsync Video




In the previous pages, we created an avatar that could “speak” the text we typed. We can also use a video we already created (or one of the three available) to lip-sync the character on screen with our text.

If we render the video, the character will say whatever we type - up to the time limit we have available, of course. We will also be able to download the video (like we did before) or even upload an audio so that our character will lip-sync to it. However, we can do more than that - we can actually upload our own video!




We must first upload and select our video. Then we select a voice and we either type some text or we upload some new audio. We click “Rendering” to create the video!



What we learned so far...

In **chapter 9 "AI Avatars"**, we learned how to create digital representations of ourselves for use in images and videos. We used AI tools to generate images based on our photos. We also used AI tools to generate talking videos of Avatars with animation, audio and lip syncing. These Avatars can be used in video files for presentation purposes, to help capture the attention and interest of our students.



Further Activities:

Avatars can be used to great effect in the classroom, by both teachers and students. These activities suggest some ways to use avatars:

Activity 1: Avatar creation by students

- Ask students to use an image generator to create their virtual image. They can create both an imaginary avatar, using any kind of detail (ie green skin, wings) and describe real features.
- Use these avatars in school projects, classroom webpage, presentations etc.

Activity 2: Create a video avatar as a teacher assistant

- Create a video avatar and use it either on its own (video) to explain specific topics of a lesson or integrate the video avatar within a presentation.
- Don't forget to limit the duration of your avatar to less than a minute - most platforms do not support more than this at a time, while students might get bored by watching (regularly) video avatars.

- Use avatars in the classroom or online, as part of homework. The avatar might be used to explain topics they worked with at school.

Activity 3: Create a story for use with avatars

- Students can create a story with a beginning, a middle and an end, and use their avatars as the main actors. This can either be a digital video (video avatars) or a comic book story. For comic book use, generate multiple avatar images with different looks, expression of feelings and poses.
- Remember to test your avatar for consistency - it must resemble each time the same person. Therefore, you must provide very precise instructions or make multiple attempts at recreating your avatar.

Activity 4: Use avatars as sprites in video games

- Avatars can be used in simple video games we can create with Scratch (or even Python). In the following chapters we can find more about creating simple video games. But, if you are already familiar with Scratch 3.0, you can include your avatar right away!

Chapter 10: AI Music!

"Let there be light, sound,
drums, guitar, oh,
Let there be rock!"
(AC/DC)

In this chapter...

In **chapter 10 “AI Music!”**, we will:

- Explore tools that enable the creation of audio clips using AI
- Learn the basics of the Beatoven AI music generator
- Generate our own music based on specific genre
- Edit our audio clips (duration) and change instruments
- Use our music audio with a video



AI Music

Ancient Greeks believed music to be one of the fundamental subjects of studies of an educated person. Hypatia (350 - 415 CE), one of the most renowned mathematicians of all times, regarded music as one of the four most important branches of mathematics, along with arithmetic, geometry and astronomy.

Music, however, is elusive for most of us without a proper education on the subject. At the same time, it is of great importance as music can convey emotions, attract our attention, spice up an otherwise boring or average presentation.

On the internet we can find lots of music we can purchase for listening or for adding to our presentations (or our videos). However, commercial video is restricted by copyright, unless we are willing to pay royalties (for example, we can't upload a video on YouTube with background music from Star Wars unless we paid hefty royalties. And yes, we can find music on royalty-free



websites. But sometimes, we just need something very very specific that suits our needs. AI comes to the rescue!


Royalty-free music

On the internet we can find music that is royalty free and under Public Domain or Creative Commons licence. Usually, websites that support either type, allow us to download and use music in our projects, with some limitations (ie in some cases no commercial use). Some AI Music Generators, especially with free accounts, restrict commercial use and maintain the copyright of our creations!




AI Music Generators

Tools that generate music based on AI have been around for almost a decade. As it is expected, all of the music generators have a licensing fee but allow some kind of free use. Most of these support multiple genres, and even have their own library of (licensed) music you can mix and match to your heart's (and wallets) content!



One of the oldest such tools is [AIVA](#) (image above) that allows the creation of music by generating content that can consist of uploaded (or selected from a library) audio files. It offers a free account with restrictions. Keep in mind that

for the free account, AIVA maintains the copyright of the produced music!




Another AI music generator is [Soundful](#) (image above), a versatile and easy to use online tool, with a flexible pricing policy that also includes a free (with restrictions, again) account. Soundful can combine music clips you either own and uploaded or you can buy through its own service. It has a very easy to use interface that allows the almost instant creation of music.

There are many more AI music generators that offer similar features. For this ebook, we will focus on [Beatoven.ai](#)!


Beatoven, our friendly AI!

Like most other AI music generators, Beatoven offers a very flexible monthly (or even minute) scheme, as well as a free account with restrictions.




Beatoven offers a very friendly to use environment that can even accept a description via text and produce music accordingly. It can be compared in simplicity to image generators. It also allows the creation of music by combining audio from other clips as well, a feature we will explore in later pages. The free account offers all the features, but it does not allow for downloads.

Fortunately, Beatoven offers a access to all its features in the free version, except downloading music. For that, we have to pay per minute, but it at least allows the use of all its tools. To sign-up for the tool, we can use either our Gmail or create a new account.



Once we sign into our account, we can start creating music! There are multiple ways of generating our music, but we can also use the prompt (image above) to describe what we want to create. We will start with a humble rock song, Iron Maiden style and see what will happen:

"A short intro of rock music with guitar, bass and drums, Iron Maiden style"



Our 1-minute track has been created. Not only that, but four versions of the track have been created

We can click on each of the four tracks (top right of the screen) to listen to the end result. We can click on "GET NEW TRACKS" to regenerate another 4 (we will get a notice in case we have not listened to all 4 of them). Keep in mind that, contrary to how other AI tools work, there are no credits or tokens or "energy" here, so you can experiment almost indefinitely until you are satisfied.

On the same screen, we can select to make changes to the volume of the song, including the fade-in and fade-out parts of it, edit the instruments used in each track, as well as change the genre.


Music genres

There is no single "style" of music. For example, you can immediately tell the difference between a heavy metal song and an RNB song. At the same time, there are different variations of each style. The main category of each kind of music is called "genre". Sometimes we combine elements from various genres to create a new kind of music. Some people like it, some don't!




Recompose track


If we are not satisfied with any of our tracks, we can regenerate them (see previous page). However, most often than not, we might prefer to make changes to part of a track. For that purpose, we can click the "Select & Recompose" button (image below).



Once we click the button "Select & Recompose", we can drag the part of the song we want to change. When we are satisfied with the part of the song we wish to change, we click on the "Recompose". We will have the option, not only to change the specific (selected) part of the song but also to choose the instruments we wish to change. At least one instrument will have to be selected (see image top right). We click "Recompose" to continue.




Once the procedure is complete (it might require some time, based on the duration selected), we can listen to the track again to make sure we are satisfied with the result (image below).




Creating “emotions”


Beatoven has been designed to be used even by people with limited to no music studies. The “emotions” option (image below) is an easy way to make additions to our selected track, to better convey just that: an emotion.



Once we click on the “Create Emotion” button, we will be presented with a matrix of 16 different “emotions” ranging from Sad to Triumphant to Fearful. We can click on the preview button (image top right) of each one to preview how our (selected) track will sound. Once we are satisfied with the result, we can click on the Apply button on the bottom of the screen, to use the selected emotion. It will be applied at the selected point of our track.




The “emotion” we created, appears on the timeline of the track. It will be created at the selected point of the tracks’ timeline (image below). If we are satisfied, we can click on the “Recompose track” button to finalise the change.





Change Instruments


On the right of the screen, under the tracks, we can find the names of the instruments used in the generated song. We can mute some of them, by clicking on the "Edit" button (image below). Keep in mind, at least one must remain!



We can click on any instrument to mute or unmute it. We click "Apply Changes" to finish editing the track.




As we can see in the image above, we only left one instrument in our track! The change, of course, was only made in the selected track - in this case, the first one in our list. If we click on another track, its instruments remain in place (image below).




Change Genre


As strange as it might sound, changing the genre of a song is not something unfamiliar - actually, it happens all the time. We have many examples, including jazz versions of heavy metal songs from Metallica, or rock versions of famous pop songs like “Dancing with tears in my eyes” by Avantasia. Beatoven allows us to do a simplified genre change, using its AI capabilities.



The genre of a song is displayed in the right side of the main screen (image above). We click on “edit” to change it to one of the available options. Sorry, no sub-genres here, therefore the options are rather limited for hardcore fans!




By clicking on a genre, we can listen to its preview. Some genres (like Rnb here) are disabled based on the mood and tempo of the selected track. Click “Apply” to proceed.



Once we apply the new genre and Recompose the track, we are clear to listen to the changes, including instruments used! The tempo, however, remains the same!


Add a video!

So... we have created an amazing 1-minute song. How about we add it as the introduction (or background music) to one of our videos? We just click on the "Add Video or Podcast" link to select our video.




After we locate the video we need to use, it begins uploading to the Beatoven web platform (remember this is an online tool!). Please make sure the video you are using is yours and there are no issues with personal information (people appearing in the video) or other royalty issues!

While the track is uploaded, you will be able to preview the video and the audio by clicking on the "play" button. Beatoven will then proceed to recompose the track.




Once the video has been uploaded and has been recomposed, we will be able to view it and make further edits to the audio. We can even click on the "Replace Media" (image above) to choose another video.

Keep in mind that, when using a video file, it takes a lot longer to compose the track again.




Creative Commons Audio

Before you start using a subscription based service, including one that uses AI for generating music, it is a good idea to see what is already available in public domain or Creative Commons.



The Free Music Archive is an organisation that specialises in storing and distributing royalty-free music for your projects. This means that you can re-use the music you download from this platform in all your projects (some restrictions apply for some content) without having to worry about copyright issues, since the creators have already given you this right.

A simple search can bring thousands of results. You can preview each song online and choose to download it. A lot of information is given for each song, including genre, album, artist and duration.



Music downloaded from sites such as this can be used with the AI generated tracks, to enhance your projects!

Royalty-free Vs free

Royalty-free does not necessarily mean that it's "free". It means that, for a (usually small) fee, you can use a digital work in your projects. Some Creative Commons projects require fees for commercial use, and they usually come with some restrictions. Please read carefully the licence before you use Creative Commons files.



What we learned so far...

In **chapter 10 "AI Music!"**, we learned how to use AI tools to generate music. We used Beatoven.ai, a very friendly online tool that allows the creation of short music clips of various genres. The music is generated through text descriptions, with Beatoven using its available instruments to create 4 tracks each time. As most AI tools, Beatoven is not free, and it has a monthly subscription cost and/or the option to purchase "minutes" for downloads. After we create our tracks, based on a specific genre, we have the option to change (mute) instruments used or to change the entire genre of a (selected) track. We can also use our tracks with videos we import to Beatoven.



Further Activities:

Teachers often have to find suitable music for school projects, performances, presentations, and more. In many cases, teachers tend to use readily-available music from other sources (ie known commercial songs). These might sound good in the short run, but when used in school projects that might circulate on the internet, copyright issues ensue. Let's take a look at what we can (further) do with music generators:

Activity 1: Ask students to generate music for a project

- Assign a project to your students. They can work in small groups. Give them access to a music generator like Beatoven. The project might be a presentation (in which they can import their sound or play it directly from Beatoven), or a video (in which they will have to import the music).
- Students will have to be creative, not only in the description they give for generating music but they must be able to explain why they picked the specific genre. This might be a good starting point to teach genres in Music (see Activity 2).

Activity 2: Make a project on music genres

- Assign students to groups and give each group a specific genre supported by the AI tool of preference (either Beatoven or any other).
- Students should be able to use multiple sources of information (including ChatGPT) to learn more about the genre they selected.
- Create at least 1 track based on their assigned music genre.

Activity 3: Compare AI Music Generators


- Create a short matrix with features (ie "Text to music generation", "Video clip editing", "Cost per minute/song", "Instrument selection", "genre support", etc).
- Ask students to use and compare at least 3 different AI music generators apart from Beatoven.
- Ask them to use at least one of the tools to generate a short track (1 minute average).
- Students should be able to select on one of the tools as their choice of AI music generator, and to explain the reason(s) of their decision.

Part C: Let's Code!

What's to learn

In Part C, we will focus on Machine Learning and AI programming, with emphasis on younger students. We will train our AI to recognise images. Then, we will use ready-made material to focus more on AI & Machine Learning, and we will create an innovative lesson plan and activities using AI.

We will also use online tools with AI-enabled features to “teach” our AI to recognise shapes and colors. We will also create a simple game using AI extensions for Scratch.




Chapter 11: Machine Learning

"I'll teach you to laugh and to cry
They're really the same you'll see
All the why's in your life
Are under my control"
(Queensryche)

In this chapter...

In **chapter 11 “Machine Learning”**, we will train our computer to recognise images, sounds and poses, using “Teachable Machine”. We will:


- Learn how to create a machine learning model
- Use data to teach the computer to recognise images
- Use data to teach the computer to recognise sounds
- Use data to teach the computer to recognise poses



Teach the machine

In this chapter, we will use an incredible online tool developed by Google, “Teachable Machine”, to do just that: teach our computer! But, how does this occur? How do machines learn? There are two aspects in this question: the first is the technological one, how the “engine” really works. The other aspect -which is within the scope of this ebook- is how this learning occurs. And the answer to this is quite simple: similarly to how people learn! For example: let’s imagine you want to teach someone who hasn’t seen a computer for the first time, to recognise tablets from laptops. What do you do? You show them many laptops and many different tablets (probably with different colors, add-ons, cases, even shapes) until they recognise a laptop from a tablet.

In the above example, we used images (or actual items) to teach someone how to recognise two distinctive objects. What our friend has learned through the use of this information, is called a learning model. Similarly, we can create other learning models, for example to help



someone learn the difference between a car and a motorbike, an aircraft and a rocket (and the list goes on...).

The science of machine learning


Feeding data to a computer is the easy part, and one that is within the scope of this book. How this data is handled, this is the tricky part and this is where scientists have been investing countless hours for the past few decades to reach the point we are today – to use simple AI tools that would have been unimaginable just a decade ago.




Teachable Machine

One of the best places to start and to understand through examples how to train a machine, is this. It is an amazingly simple and intuitive tool that allows the easy and fast creation of learning models.


It is completely free and has been developed by Google.
<https://teachablemachine.withgoogle.com>



On the first page we can see animations featuring what we can achieve in mere seconds with Teachable Machine. It even includes a short video (which we suggest you watch) on how Teachable Machine works..




This is probably one of the most well-designed websites and tools available. The home (first) page is self-explanatory - the visitor is quickly shown all the information that is required to understand how machine learning models work, how to train the computer, and how to test or export your model. You can also find simple tutorials as well as examples made with Teachable Machine!




Getting started...

To start a new project, we click on the “Get started” button on the top right of the home page (see previous page).



We are immediately shown three types of projects (Image, Audio, Pose) and the options to Open an existing one from our (Google) Drive or from a file. Through the menu (top left), we can create a new project, open one from our Drive or from a file, and watch videos on how to gather samples, train our model and then export it.

It is a good idea to watch these three videos at least once, regardless of the examples we show in this ebook.



As we will see in the next pages, all our models can be exported and used on websites and on devices such as BBC Micro:bit and Arduino. However, we will focus more on testing our models on our own computer or mobile device. There are many tutorials that show how to use Arduino and Micro:bit with Teachable Machine to create incredible AI-powered solutions.


Extending the application of Teachable Machines

A very easy and interesting project is to train a model to recognise your own face or even voice and follow you around. All you really need is an inexpensive Micro:bit and some other easy to find components. For a more detailed tutorial, check the page <http://tinyurl.com/buildairobot>




Starting an Image Project

For our first project we will create a model based on images. We will create a standard model for use on our computer or website.



It is possible -and supported- to create a model for IoT devices. As such, we could create, for example, a robot that recognises objects based on the image we provide, or an unlocking door mechanism that recognises our face! However, in this ebook we will focus solely on the standard image model and not the embedded for devices.




The main interface allows us to insert images using a webcam (especially when we want the AI to recognise faces or our face) or by uploading files. Each class must consist of similar objects (ie Class 1 might be faces, or coffee mugs, or shoes or whatever you want!).

The latest version of Teachable Machine allows the creation of as many classes as we want. The more the classes, the more complex our model becomes but it can recognise more objects!

For our example, we will need our model to recognise smiling or sad faces, therefore we will use our webcam and upload images to feed as much data as we can to our model!

Populating Classes

Before we start uploading images, it is a good idea to define the number of classes we need for our project, and then name this.




For this example, we will rely on two classes, one named “Smiling faces” and the other “Sad faces”. Once we are satisfied with the creation of our classes, we can start “feeding” them with images. Since we are focusing on facial expressions, we will use both uploaded images and images captured through our web camera.


The more (images) the merrier...




Obviously, the more images we “feed” our class, the better the recognition! But, we must set a proper limit: what is the maximum or minimum? For Teachable Machine, it is a good idea to have at least 15 images for every class but you can always



To upload images, we click on the “Upload” button (image left). We can drag and drop images from our computer or import from our Google Drive (remember, this is a Google product!). All images will be cropped to square shape.




Click on the “X” button (above) to close the class.




Before we capture an image, we can select to crop it or flip it. Once we are satisfied, we click on the “x” button in the blue webcam capture frame (image bottom left).

Our first class has been populated with images (14 samples) of smiling people (and a smiley image!). We will add another image (or more) using our webcam.




To remove a single image, we hover with the mouse above it. A trash can icon appears. We click on it to delete the image. By clicking on the class menu (see image above) we can select to delete an entire class, to remove all samples or selected, to save the samples to Drive (useful for recurring models) or disable class.



We can very easily grab multiple images using the webcam and by clicking (once) on “Hold to Record”. These will be added to our selected class.

When we disable a class, it remains in our model, including all its images. However, it will not be used when training the model. This feature is extremely useful if we want to test training with various classes or create different models with a mixture of classes.




For this example, we will use images from our Google Drive. Teachable Machines will request permissions to use our Google Drive. We can either grant it full or just read-only permissions.

We continue with populating the second class that we named “Sad faces”. We can use images from our computer or from the webcam. Alternatively, if we have a library of images on our Google Drive (useful, if we regularly present how to train models) we can choose to select those instead.

Include the audience



If you regularly -or even sometimes- present to an audience or train others, try to engage them during your training session or presentations. In the above case, you can ask members of the audience/trainees to walk to your computer and webcam so you can capture their facial expressions (please do not distribute/show their images without their consent).




We will use the content of the folder “Sad faces”. We can hold down shift and click all images to select them. Then we click the “Select” button to insert all images to our class.



By uploading the image files, we will have two classes with at least 15 images in each! Next, we will train our model!

Training our model!


Our classes have their images. Next, we train our model!



We can just click on "Train model" (image above) to start the training. The time required is based on the complexity of the classes and the parameters we set (more on that later).




We have two ways to check our model: we can use the webcam and test if it can recognise a smiling from a sad face OR we can use a file (image below).



To use an image file for testing, we can either drag and drop it on the blue box shown in the image above, or select it from Google Drive. On the "Output" we will see to what class it belongs, based on a percentage.

Adding a class


In some cases, having just two classes makes little to no sense, not to mention that results are ambiguous. For example, we never have either a sad OR a happy face (actually we have many expressions, including angry, surprised etc. but let's focus on Normal, Sad, Happy..




Once we upload our images, we can start the training of our model again - this time with all three of them!

Increasing the complexity

We could create a model that recognises all our emotions that are shown with our facial expressions, by adding more classes (anger, surprise, frustration, disgust etc). By adding more classes (and images in the classes) we increase the complexity of the model and -subsequently- the time required to train it. However, the results will be much better (theoretically), since all emotions will be recognised and seen as such.




With all three classes ready, we click "Train Model" to do just that! Notice that it takes a bit more time for training, since we have increased the number of classes (and number of images). We can test our model again:



Advanced Training

If we want to ensure a more advanced training that (might) ensure better results, we can click on the “Advanced” option and change the parameters available.




Epochs refers to the number of times all our samples are fed into the training model. As we increase the number of epochs, the accuracy of the model is expected to increase. However, it is recommended to first check your model with the standard training parameters and then make changes accordingly.

Batch size refers to the number of items (images here) used every time. For example, if the number is set to 16 and we have 48 images, they will be split into 3 different batches. When all 3 are complete, an epoch will be completed.


The **Learning Rate** measures how fast the model learns from the training data. Usually we do not change this, as it greatly increases the training time without -at least for simple classes with small numbers of items- offering significant gains.

Under the hood displays in graph form how our model works.




Exporting a model

We can always use the preview screen of our model to show how machine learning works using classes. However, in most cases, we will need to export our model and use on a website or with an IoT device.




After we are happy with the results using Preview, we can click the “Export Model”. The easiest way to use your model is by uploading it (image top right). It will be immediately hosted as a link by Teachable Machine, and a URL will be created instantly. This is the easiest way to use your model publicly. Alternatively, you can download the model and use it in your own website.



For the above example, a shareable link has been created:


<https://teachablemachine.withgoogle.com/models/z49iWNdlj/>

We can preview the above link in a web browser:




Audio & Pose Projects

Similarly, we can create an Audio or a Pose Project. From the menu (top left of the window) we select “New project”. Alternatively, we can click on the project boxes (image below).



Similar to how we worked with an image project, we will create classes made of items. For example, for an audio project, we will need to record or upload at least a set of 8 audio samples for each class.

If, for example, we want to create a model that distinguishes between laughter, and cry, we need at least three classes, one of which is Background noise. Then we can upload samples into each one (image top right).



To add audio files to our classes, we can either use our microphone to record short 20 second audio files (at least 20 files for Background Noise) or upload from our computer or Google Drive. As soon as we have uploaded the necessary files, we can start training our model.

We will follow a similar approach for a Pose Project.

Background noise

On Audio projects, this is a compulsory class. It cannot be deleted or disabled, and it requires at least 20 audio samples. We can use our microphone to record the samples, or upload audio files. This is a necessity, in order to differentiate the actual audio from the classes we create, from any and all background noises from the surroundings.




What we learned so far...

In **chapter 11 “Machine Learning”**, we used “Teachable Machine”, an online tool, to help train models using images (and audio and poses).

The use of Teachable Machine enabled us to understand what machine learning is and how to teach a computer to recognise images, sounds or even body poses.

The models we train can also be used to control robots or to create games using tools such as Scratch (as we will see in a later chapter).



Further Activities:

Machine Learning can enable teachers to create projects for students, with which they (the learners) can understand or perform even complex activities.

Activity 1: Train a model to recognise facial expressions

- Ask students to identify all possible emotions expressed through facial expressions.
- Train a model using multiple classes to identify at least 5 different emotions conveyed by facial expressions.
- Ask students to perform these facial expressions, but also to use (a) photos from the web, (b) drawings of their own, (c) images generated using AI.
- Test the model to ensure its accuracy. If needed, retrain the model and add more images to each class.

Activity 2: Teach sign language!

- Teach the basics of sign language by creating multiple classes. For this interesting activity, you might be required to ask the help of a professional, or refer to one of the tutorials on YouTube or the web.

- Test the model with photos and hand gestures and make improvements.
- If possible, present the trained models (in person, or via web conference) to professionals or even other students with hearing disability. Gather their feedback and make improvements.

Activity 3: Create a sorting model for recycling

- More often than not, people get confused on what kind of trash goes where - create and train a model that provides such feedback (ie in which bin to place the empty soda can).
- If possible, place a laptop near the recycling bins of your school (if you do regular recycling).
- Ask students to test their items before they place them in the recycling bin.
- Make changes and improvements.
- Keep in mind that usually schools sell the same type of bottled water or other items, therefore it is much easier to train such a model in a school than for the general public.

Chapter 12: Save the Oceans!

"Water and air, the two essential
fluids on which all life depends,
Have become global garbage cans"
(Jacques-Yves Cousteau)

In this chapter...

In **chapter 12 "Save the oceans!"**, we learn more about training an AI using "massive" amounts of data. We take a look at an interactive activity created for learning coding (<https://code.org>) and how we can train an AI to recognise trash from sea animals. We also use online Machine Learning tools to train our computer to recognise basic shapes. We will:

- Learn more about Machine Learning through an interactive activity
- Use Machine Learning to train our computer to recognise shapes



Let's clean the ocean!

Our oceans are full of waste we throw every year in the sea. As a matter of fact, in the Pacific there are millions of tons of floating plastic. This gathering of waste is so huge, it has been named "The Great Pacific Garbage Patch", and it is wreaking havoc to marine life. Furthermore, this is just a percentage of the huge amount of trash under the sea.

As a good introductory to students on Machine Learning, we can use the outstanding material created for Code.org. The code.org website (<https://code.org>) has an incredible section with amazing content on AI, including videos, lesson plans and interactive activities.

Through well-structured lesson plans and activities, including offline and online ones, we get to experience how machines learn and how they improve their learning through crunching of large pieces of data. During every step of the interactive activity, we get more information on the problem of plastic waste in the seas.



It is not within the scope of this ebook to present all the material of code.org on AI. However, we will focus -as a good example- on the "AI for Oceans" activity.

Shoes out of plastic


In an effort to help clean the plastic out of the oceans and beaches, many companies have announced and already started, ambitious programs that gather this material and use it to create new products. One of the products created using ocean plastic are running shoes from a high profile international company.



Code.org & AI Content


The code.org portal is one of the most amazing places on the internet to learn about coding - and much more!

It is a multilingual portal with fun coding activities that include lesson plans, worksheets, videos and even resources for teachers to keep track on the progress of their students.



Recognising the importance of AI in all aspects of our lives, the code.org team have create a special section.

Through the code.org/ai page, we can find a plethora of material focusing on what AI and Machine Learning is. From simple explanatory videos, to complete lesson plans, worksheets and online activities, this is the place to go if you are an educator needing to make a start on the subject with your students!




Not all the material for AI has been translated to every language. However, with proper teacher assistance, it is extremely easy to work with the activities - and gain a proper understanding on how AI changes the world around is in amazing ways!

Getting started...

The first thing we need to do is decide our course of action: are we (as educators) going to use the material to teach Machine Learning as the primary learning goal OR are we going to use the activity as part of an Environmental Education lesson? Or, maybe, we can do both!


For starters, we can use the amazing material already available: <https://studio.code.org/oceans/lessons/1>



As you can see, there is everything here to get started in a complete setting that even includes warm up activities. Of


course, we can use this lesson plan as a blueprint and modify it to better serve our students (age, interests, abilities). This page is full of interesting resources and videos we can use, and, based on my experience, there is much more material here than we need, almost eliminating any need to seek for additional content.

The interactive activity is part of the lesson plan: to teach our AI to recognise garbage from sea animals, and get rid of only those that do not belong anywhere but the recycle bin (or the trash bin, for that matter!).



The first thing we need to do is decide our course of action: are we (as educators) going to use the material to teach Machine Learning as the primary learning goal OR are we going to use the activity as part of an Environmental Education lesson? Or, maybe, we can do both!


For starters, we can use the amazing material already available: <https://studio.code.org/oceans/lessons/1>



As you can see, there is everything here to get started in a complete setting that even includes warm up activities. Of course, we can use this lesson plan as a blueprint and modify it to better serve our students (age, interests,

abilities). This page is full of interesting resources and videos we can use, and, based on my experience, there is much more material here than we need, almost eliminating any need to seek for additional content.


The interactive activity is part of the lesson plan: to teach our AI to recognise garbage from sea animals, and get rid of only those that do not belong anywhere but the recycle bin (or the trash bin, for that matter!).



Generating a lesson plan...


How about we start using some of the things we learned earlier? Of course we can use the outstanding material of the code.org website. But how about we generate our own? Let's go to ChatGPT and create a lesson plan to teach about plastic thrown in the oceans!

“Create a lesson plan to teach primary school students about plastic garbage in the oceans and integrate an AI activity that teaches students how to use machine learning to clean plastic but not sea animals and plants!”



So, now we have our lesson plan in which we can integrate our interactive activity from code.org! How about we also create an introductory presentation on the Great Garbage Patch? We shall use Tome.app: <https://tome.app>

“Create a presentation with 5 slides to introduce primary school students to the problems caused by the Great Pacific Garbage Patch!”



In mere seconds we can have a professionally looking presentation to use as an introduction to our lesson! Of course we can edit it to our heart's content!

Creating written text...

Creating images is a great way to encourage discussion within the classroom - what does the image show? How is it related to our subject? What do you think the animals in the pictures are experiencing/feeling? What would they (the animals) say to us if they could talk - at least in a language we understand? By creating images with Bing, with the proper description, we can use them to encourage students to discuss their relation to our context.

Additionally, we can ask our students to produce written text, to show their ideas of how to clean the oceans.

This text can be produced as part of collaborative work, discussed in the classroom and then used with an image generator to create the picture. We can ask our students to refine their text until the final image is closer to what they want to present. An example could be:

“Tiny robots that roam across the surface of the sea, collecting garbage for recycling.”





Photo stories using images

How about we create a photo story then? We can ask our students to work in groups. Each group can create text with descriptions for a short photo story on how a small (or a team of small) robot(s) can clean the oceans from garbage!



Students can create the text for generating images.

Each small paragraph (we can set the limit) can generate a single image. They can then cut and paste (we can even use glue and markers!) the text and the images, to create a nice poster showing a photo story of how each team can solve the garbage problem!

Of course, the main idea here is to encourage our students to (a) produce written text, (b) suggest ideas on how to clean the oceans, (c) generate images in accordance to their story and (d) present their ideas, as a poster, to the rest of the classroom.

We should always keep in mind that a problem to be solved, we must first identify its source- and the source is people disposing their trash in the oceans, or in places where they make it to the oceans (like rivers). We could extend our lesson(s) to include ways of making sure that the problem stops at its source. It is a good idea to use multiple sources of information so that students can research these sources and suggest solutions.


Creativity to the max!

Whenever the use of digital tools arises, [some] people tend to think that one approach might lead to the total disposal of all the others. Using an image generator can have a positive impact on how students write text descriptions to create their images.. This does not mean that we discourage students from creating their own original artwork!



What we learned so far...

In **chapter 12 "Save the oceans!"**, we used an activity from code.org to learn more about machine learning, and how to use it to solve real problems. A scenario, along with detailed lesson plans and worksheets, has been created online, along with an interactive activity that introduces students to what machine learning is, and how it can be used to clean oceans from plastic pollution. We then proceeded to extend the activity, by using tools we learned in previous chapters of this ebook, to create more content, including a photo story by students.



Further Activities:

Oceans cover 71% of the surface of our planet, the only known celestial body in the galaxy to sustain life (at least for now). Time to learn more about water on our planet!

Activity 1: Ecosystems of the oceans

- Ask students to work in groups and assign one ocean to each group of students.
- Students can use multiple sources of information, including ChatGPT, for their project.
- Give specific instructions to students, to help them focus on their project: what is the name of their ocean? What is the extend of it? What ecosystems does it include? Ask them to present at least one of the ecosystems. What is the main threat to that ocean. What can we do about it? When was it first explored? Think of more questions but also ask students to come up with their own (this can be a part of a whole-class discussion).
- Students can use presentation tools for their project or create a poster.
- Ask them to use image generators to create their images.

Activity 2: Water on other planets

- Ask students to explore the possibility, based on scientific evidence and theories, of water on other planets or moons in our solar system (ie Venus, Mars).
- Students can search for information on the web and/or ChatGPT.
- They can use image generators to create pictures of how water might looked like in the past on Venus or Mars.

Activity 3: Saving the seas

- As a school project, find more about sea pollution, its consequences to sea life -and in extend, human life- and what world governments are doing about it.
- Students can select specific seas (ie Mediterranean) to focus on more specific ecosystems.
- Find which ones (seas) face immediate threat(s).
- Suggest ways to solve the problems by (a) stopping/ changing human activity and (b) cleaning the seas.
- Identify more resources on the web to support your project.

Chapter 13: Machine Learning for Kids

"Children, Computers &
Powerful Ideas!"

(S. Papert)

In this chapter...

In **chapter 13 “Machine Learning for Kids”**, we will use a versatile and simple tool to create and train classes in projects that incorporate images, audio and poses. We will:

- Learn how to create accounts for our school
- Create projects that include images, text or audio
- Create models for use with Scratch 3.0 to create interactive games and activities



It's for kids...

So far we have worked with activities that teach us how to create classes with data and train them to create a model. We then tested it within the environment of Teachable Machine and as a separate weblink. We also worked with an inspiring activity that teaches both what Machine Learning is (and how we train AI) and also how to clean our oceans from plastic waste.

However, all the previous activities were more geared towards going from point A to point B. For example, even with Teachable Machine -though it can be used in much more creative projects and processes- we barely scratched the surface of what we can achieve. And, it does require more skilled teachers and trainers with experience in serious programming to make the most of it.

In this chapter, we will use an amazing tool named Machine Learning for Kids. Through this, we will create projects in which we will add data from our computer or web camera and we will then use them to train new models that can be shared with other people. These models will be




incorporated in a custom version of Scratch 3.0 that allows the fast development of applications and video games - we can even dissect a Pacman video game to see how it was created with the help of AI!

And the most amazing thing about Machine Learning for kids? It's completely free! In the next pages of this chapter, we will learn how to create an account and start producing models for use with Scratch! Enjoy!


Machine Learning for Kids

Machine Learning is not necessarily an incomprehensible subject. In fact, Machine Learning for Kids (<https://machinelearningforkids.co.uk>) is trying to do just that- simplify what it is and how it works!




At first glance, the website is unassuming - a simple introductory web page with 3 simple steps that enable the computer, through Scratch activities, to recognise items.

The first time we visit the site, we have the option of registering for Machine Learning for Kids. However, can just click on the "Try it now" button and start using it!




The website is very simple and offers great tutorials and other resources, such as Worksheets (Menu links).




Book on Machine Learning

The website has links to a complete, outstanding book on Machine Learning, complete with examples and activities. We can access the description of the book from the website menus, and read the full book from the link provided (image below).



This is one of the most comprehensive, easy to understand books on the subject, with examples written specifically for the novice teacher and younger students (primary education). As a matter of fact, we can even examine the entire (392 pages!) book from the links provided!



The book is broken down into chapters that explain the basics of what Machine Learning is, and how we can train our computer to recognise and interact with items. It uses Scratch (with AI extensions) as the basis for most of its examples, clearly explained so that anyone can follow them. This is probably the best place to go, immediately after our (brief!) introduction to Machine Learning.


Printed book

If you are serious about your commitment to Machine Learning as a teacher, you should invest in purchasing the specific book for your personal needs or your school. The sheer number of examples, in addition to the clear and precise instructions give, make this a worthwhile investment!




Sign up...


Even though you can use (almost) all resources available in Machine Learning for Kids without registering, you might want to have the full spectrum of features and capabilities available to you. Not to mention the capability to save your project and create more than one! You can click on the link “Why register” to find the differences between “guest” and registered access (image below).



To create a new account, you click on the “Sign up” button (image above) to start the process. It must be stated that the entire website is open with free access, so no credit cards are involved!



Depending on the type of user, you select the account you prefer. Our suggestion is to create an unmanaged account.



Unmanaged account

Machine Learning for Kids

machinelearningforkids.co.uk/#/signup

Create an unmanaged class account

Username *

alexandroskoffe

We need a unique username to identify you.

15 / 15

Email address *

alexandros.kofferos@gmail.com

Intended use

School

☒ By checking this box, you accept responsibility for any fees or charges incurred on IBM Cloud.

Setting up this tool to allow your students to train machine learning models will require some access codes (called "API keys") from IBM Cloud.

"Do I have to pay for that?"

- No.
- Creating an account on IBM Cloud is free.
- There is a limit to the number of free API keys you can create.
- Watson API keys you can create for free have usage limits. You can choose to pay IBM if you want to increase these limits. It is these charges that are your responsibility.
- An advisory price list is included in the [setup instructions](#) (section 4).

"Will my students need to do this?"

- No.
- This is a one-off setup task that you will need to do as the teacher / group-leader.
- Students in your group will not need to create an account with IBM Cloud.
- Students in your group will not need to visit IBM Cloud.

"How will I know how to do that?"

- Step-by-step instructions with screenshots are available to walk you through this.


"Do all projects require IBM Cloud API keys?"

- No. Machine learning projects to recognize images, numbers or sounds do not require API keys. Projects using pre-trained models do not require API keys.
- Only machine learning projects to recognize text require IBM Cloud API keys.


If you're concerned about any of this, or have any questions, please [get in touch](#).

CREATE CLASS ACCOUNT CANCEL

To create an unmanaged class account, we complete the form above and we click on the "CREATE CLASS ACCOUNT" button on the bottom of the page (image above). A password will be created and displayed on this page, and an email will be sent to the account you typed above. Verify your account from the email link sent to you.



Once you log into your account, you can see the above options. Keep in mind that you can have most of the functionality even without registering for API Keys.




IBM API Keys

API Keys are needed when training learning models for text projects. Otherwise, you can skip this test!


Starting a new project

Projects allow us to “teach” our AI to recognise shapes, sounds or text through feeding data. We do not have to use huge amounts of data, but the more we provide, the better the results (see previous chapter and activity “AI for Oceans” for an example on machine learning).

From the menu “Projects” (image below) we can create a new -you guessed it!- project! Actually, this is the page where all our projects will appear!



First, we click on the button “Add a new project”.




Our project will rely on training our AI to recognise data in one of the forms shown in the image above (text, images, numbers, sounds). Once we choose the type of data, we click on “Create” to start training our AI. Remember: the more data we use, the better our AI will be in recognising our preferred data type.

Proper naming

Giving a proper name to our projects is very important: once we start building -hopefully- a library of projects, generic names such as “MyProjectOne” or “JustAnotherProject” might not do. Try to think of names that describe in a simple sentence what your project’s all about – that way, it will make more sense and help you identify past (or present) projects.




All projects we create, will appear in our “Projects” page. We can click on their rounded rectangle box to edit them, or on the trash can icon to delete them.




Machine Learning for Kids is amazingly simple- once we click on a project (image above), we are presented with 3 different rounded rectangles: Train, Learn & Test, and Make. All three of these are essential in the process of creating and refining our project.


We can revisit any project again and again and improve it, almost indefinitely! In the next pages, we will train our AI to recognise pictures of cats and fish, and also to distinguish between them.




Our first step is to train our computer to recognise images of cats and fish! We start by clicking on the “Train” button (image above).




We will “pour” buckets of data to our model. Let’s click on the “Add new label” and start naming our items.




There are four ways to add images to your "bucket" - the first is to drag and drop an image inside the rounded rectangle we created and named "Fish". We can do that with an image file from our computer, or directly from an image on the web (images below).




Alternatively, we can right click on an image in a web browser, and select "copy image address". Then, we can click on the "www" button in our bucket (image top left) and paste the address (image below).



Keep in mind that only jpg and png files are supported!



We can also draw an image and add it to our bucket!"



Another way to input images is by using your webcam. I have tried it with my cat, he was not very into it, but it is a very convenient way of inserting images of actual persons (more on that later!). By far the easiest way to add images is by drag-and-drop via another web browser window (as shown in the previous page).


Regardless of the method of adding items to your bucket, you need to gather as many as possible to help train your AI to recognise fish from other animals (and later plants).

How many are enough?




Obviously, the more the better, especially when you add items that have great differences between them. However, a sample of at least 10 is a good start. For more information on the sample size, you can check this link:

<http://tinyurl.com/aisamplesize>




So, it's time to add another "bucket". This time, we will add cat photos to our bucket. Once we create a new label (for our item bucket), its name also appears on the header of the page. This is extremely helpful as it shows exactly what we are training our AI to recognise.




Learn & Test


We have created the Fish and Cats labels ("buckets") and we have populated each with 20 images. Now it's time to train the computer to recognise images!



We click on the "Learn & Test" button (image above) to start the training process:




If we need to add more images, to improve the recognition of fish and cats, we can go to the previous screen, click on "Train" and drag and drop as many images as we like. Otherwise, we can click on "Train new machine learning model" (image bottom left) to start the procedure.




Once the training is complete, we will be able to test it - and we will do just that in the next page! Alternatively, we can click on the "Delete this model" to -well!- delete it, or train a new machine learning model.


So, let's see if we have adequately trained our model! We will begin with a rough drawing of a fish! Just hold it in front of your camera and click on the "Test with webcam" button. Wow! It worked! (Image below).



Our model is starting to recognise images we use - if we see that there is a strong percentage of error, we can go back and add more images. But, in the above example, we managed to successfully recognise our rough drawing as a fish -which is correct (hah! tell that to my Art Teacher in secondary school who said nobody's gonna recognise my drawings!).




We can also paste the URL of an image (either of a fish or a cat) in the field next to the "Test with www" button (image above).



We can always test our drawings as well!

Let's make a game!

Now that we have trained and tested our model, let's use it to create a simple game using Scratch! Let's go back to our previous screen and click "Make" (image below).




We have 3 great options: too use our learning model with Scratch, Python or App Inventor. For this ebook, we will go with Scratch 3!


Coding in Python



In this ebook we focus mostly on Scratch 3.0. There are many pretrained models to choose from, apart from the ones we are creating. However, if you are a more experienced user or just want to try learning how to type code in Python, you can do so by visiting the amazing tutorials offered through Machine Learning for Kids or other websites. As a starting point for Python, we suggest our previous ebook. You can download it for free from the page <https://mathsis.org>.




By clicking on the Scratch 3 button (image above) we are presented with the blocks that will be added:



We click on "Open in Scratch 3" (image above) to continue.

Scratch-ing code


Scratch 3.0 is a visual programming language and online platform designed to teach coding to beginners, especially children. Developed by the MIT Media Lab, Scratch allows users to create interactive stories, games, and animations by dragging and dropping code blocks rather than typing out text-based code. The platform uses a block-based interface where users can snap together colorful code blocks representing different programming commands.



With Scratch 3.0, users can code by stacking and connecting these blocks, which simplifies the process of


learning programming concepts like loops, conditionals, and event handling. It's a user-friendly tool that promotes creativity and problem-solving skills in a playful environment. Scratch also has an incredibly active online community where users can share their projects, collaborate, and learn from one another.

There are many free resources that help teachers and students learn to use Scratch, including dozens of books. One useful page is <https://scratch.mit.edu/educators>.




Scratch & new model

If you are not familiar with Scratch, we suggest you take a look at the learning resources from the previous page, or some of the many helpful videos online. Once we launch Scratch 3.0 from within our model, we will be greeted with an extra set of tools (image below).



As we can see, there is an extra category of commands called "Cats Fish". They are based on the model we trained earlier! These are automatically generated into Scratch commands that appear in this (slightly) modified albeit full-featured Scratch version. We can use these with some of

the commands already available in Scratch or with some of the extensions provided by Machine Learning for Kids.




From the Scratch environment, we can click on the "Project templates" menu to view all the options available. We can use some of these templates along with our own model or models to extend the capabilities of what Scratch 3.0 can do. And, of course, we get to learn more about creating AI-powered applications and video games.

For a more thorough look on Scratch 3.0, you can visit the official website <https://scratch.mit.edu>.


Pre-trained models

Machine Learning for Kids allows not only the training of our own models but the use of pre-trained ones. We also have the option of sharing our own models, thus helping in the creation of an online community that caters to the needs of almost any educator wishing to use AI.



In the pre-trained page (accessed using the links on the top of the web page) we can find short descriptions of each available model.

Some of these pre-trained models allow the detection of parts of our face ("Face detection") or our pose ("Pose detection"), while others can recognise speech!



This version of Scratch 3.0 includes all the standard extensions, plus all the extra ones from Machine Learning for Kids (labeled in the image above as ML for Kids). We can select one of the pre-trained models to test it with some sample code.


Sharing is caring

Most of the functionality found in pre-trained models used to be the work of dedicated software and hardware that cost hundreds if not thousands of euros just a few years ago. Today, everyone with some basic understanding of Scratch and a webcam can easily create games that recognise gestures! This is the true power of sharing with others!




Machine Learning Projects

To help us get started, we can view some of the projects already available on the web site, under the “Worksheets” link (top of the page).



Under every screenshot of the project, we can find information on the programming environment (Scratch, Python, App Inventor), the difficulty level and the type of information that is recognised (text, numbers, images, sounds). By clicking on the “Download” button under each one, we can view teacher guides and project worksheets.


These are highly detailed documents that contain not only a complete lesson plan, but also precise information for the students, including necessary screenshots and explanations.



By studying the projects available, we can improve our own understanding on how to use AI models more effectively in designing our own games and applications. Some of the available examples include games like Pac-Man or Rock Paper Scissors (or you can extend this game by adding **Lizard** and **Spock** - for fans of “Bing Bang Theory”)!

What we learned so far...

In **chapter 13 “Machine Learning for Kids”**, we created an account on the Machine Learning for Kids web tool. With Machine Learning for Kids, we can create and train learning models consisting of many classes that can then be used to create AI enabled projects in Scratch, Python and App Inventor. The models we create allow us to recognise images, audio and text. We also used the templates and examples that are included with the Scratch 3.0 variant created for Machine Learning for Kids.



Further Activities:

Machine Learning for kids is the perfect place to start training models and using them with your own Scratch-based creations!

Prerequisites: for the activities below, we assume that students already have basic knowledge and experience with Scratch 3.0, they know how to create and use variables, conditions and loops to create simple games.

Activity 1: Fruit bats!

- Fruit bats are a very peculiar kind of bat: they are smaller than their other cousins, and they mostly eat rotten fruit. They are also almost extinct in most of the world!
- Train your model to recognise between rotten and fresh fruit. You can focus on just one kind like apples, for example.
- Create a background where you have fruit randomly appearing on the ground and on the branches of a tree.
- Click only on the fruit that your fruit bat can eat!

- Make your game more educational: provide information on the fruit bat as you play along! You can even include photos of actual fruit bats in the game.

Activity 2: The Recycle game

- Following on the path of the previous chapter, let's re-create part of the interactive activity we used for cleaning the oceans.
- Create a set of classes that differentiate between trash and sea life. Make sure you include both fish, mammals and other sea life, like invertebrates, all in their respective class.
- Include classes for different types of trash we can (unfortunately) find in the seas: plastic bottles, shoes, clothes, tyres.
- Train your model and use it with the Scratch 3.0 variant of Machine Learning for Kids.
- Assign variables to a simple game: -1 when you click on a sea animal, +1 when you click (to remove) trash.

Chapter 14: RAISE Playground


"The Playground is where
imaginations run wild
and dreams come true"

(T. W. Thomas)

In this chapter...

In **chapter 14 “RAISE Playground”**, we use an MIT-developed variant of Scratch and its extensions that enable us to develop AI games and applications. We will:

- Study the AI extensions of a variant of Scratch 3.0
- Create a simple game that uses the webcam to control a sprite
- Use facial expressions to control the movement of an object on screen



Let's raise playgrounds!

In the previous chapter, we focused on Machine Learning to create models that we then used in a special version of Scratch 3.0. The focus of Chapter 12 was exactly on creating the models and then using Scratch.

In this chapter, the final chapter of our ebook, we focus on the RAISE Playground, a new platform that supports AI and robotics and its aimed for students and people who are newcomers in the world of programming.

For this chapter, as in the previous, we assume that the reader has at least some basic knowledge of programming in Scratch. Therefore, we will proceed almost immediately to create our first simple games using AI extensions. We will also do a small comparison between this version of Scratch and the one used by Machine Learning for Kids. Which one you prefer to use is up to you - which one satisfies your needs the most or is more approachable to your level of understanding?



Remember, when we are creative, we choose our own tools and get on with our work - not the other way around! So, move on and get creative with RAISE Playground, the latest gift from MIT!


Game time!

Scratch has enabled students and teachers to create simple games in almost no time since its very first versions. With RAISE Playground, using AI to create interactive games is incredibly easy. By far, this is the easiest to use tool that exists when it comes (for newcomers) to create AI content!




Raise Playground


Since this is an MIT educational website, it follows the same pattern as other similar creations: simple, single-page layout, with the bare minimum set of information. In a nutshell: this is a very sophisticated yet uncomplicated online platform - <https://playground.raise.mit.edu!>



By scrolling the page, we find the link to launch the Playground version of Scratch 3.0, as well as links for examples created with it, and curricula that correspond to the capabilities provided by the platform (image top right). It must be stated that all the content of this page, including the examples and the tool itself, are completely free!




The Curricula is a very important part of the platform, since it provides examples and web pages with detailed material.




Getting started...

To access RAISE Playground directly, you can type <https://playground.raise.mit.edu/main/> in your web browser. It will lead you directly to the modified version of Scratch (image below).



On screen, we see the word “RAISE” dominating the upper left part of our window, while gone is the familiar friendly cat and we can see a small robot (BBC Micro:bit based) in its place. Don’t be alarmed, our cat is gone (for now!) But we can always bring him back using our sprite library!



In the next pages, we will use AI extensions to create a simple game with our little robot. By clicking on it (image top) we can see that it has two costumes: one with the (on-board) LEDs set to off, and another with the LED lights creating a red shape. This is a very good example of showing students how to program the array of LEDs on the surface of the actual but also the virtual BBC Micro:bit.


Virtual robots

AI can be used to control robots but also to enable them, via their sensors and programming, to take decisions on the spot. The use of a virtual robot was an important design decision by the creators of the platform, since it shows a clear direction towards teaching robots.




RAISE PLAYGROUND

There are subtle, yet significant differences between the two variants of Scratch 3.0 used in RAISE Playground and Machine Learning for Kids. Apart from the standard set of commands, they add more extensions and templates.




RAISE Playground even supports importing trained models from Teachable Machine.




MACHINE LEARNING FOR KIDS

This variant of Scratch 3.0 offers a plethora of Project templates, right from the menu items. These allow the easy and fast creation of games and applications that use pretrained models.




There are less extensions than RAISE Playground, but they are still incredibly versatile and enable creativity.




Choosing extensions

By clicking on the Scratch Extensions (see previous page), we can access all the available ones. Here, we can see even extensions from Teachable Machine, and we can even use the models we created with that tool.




For our example we will use the “Hand Sensing” extension. We just click on the extension to enable it. It requires access to our webcam to work, which we must give in order to work with the specific provided commands. Next, we will use our hands to control the position of the virtual robot on the screen!



As we can see in the image above, the available commands of the extension “Hand Sensing” have appeared and we will use them accordingly!


In the next page, we will create a simple program to send the virtual robot anywhere our finger points!

This is a simple yet intuitive way to use our finger as a video game controller!




Moving around...

By clicking on the Scratch Extensions (see previous page), we can view a list of the available ones - the standard extensions found in every Scratch 3.0 version, as well as some new extensions for this specific variant. Some of the extensions allow interactivity between this Scratch version and other AI tools such as Teachable Machine, allowing us to use models we created earlier.



For our example we will use the "Hand Sensing" extension. We just click on the extension to enable it. It requires access to our webcam to work, which we must give in order to work with the specific provided commands. Next,

we will use our hands to control the position of the virtual robot on the screen!




The pretrained model for this extension works remarkably well - not only it recognises specific fingers, but it also recognises the part of the finger. This allows great flexibility in creating games where we control our sprite with gestures and hand movements.

Video transparency

There are cases where you want to be a part of the actual game, and other cases where you might want to be as subtle as possible. Changing the video transparency is a very flexible way to do just that. Make sure you adjust it accordingly to the requirements or even needs of your game!






By changing the video transparency to 100, we completely “hide” our video feed. This is necessary if we want to have an unobtrusive view of the screen area and also use a background.

Adding either one of the included backgrounds or creating your own, does not cause any problems -as it shouldn't- with the camera. As soon as we are comfortable with the use of your finger for controlling the sprite (the robot or any other sprite), we can include some more commands, variables and other sprites to make a (more) proper game.

As an example, let's populate the screen with fruit. For each fruit we touch, we receive points.



The above is an overly simple example: we have a score variable, and an item (an Apple). We move our finger (and our robot) to “eat” the Apple and increase its score. Of course we can populate the screen with more fruit appearing in random parts, or even moving around. To spice the game a bit, we can even have “rotten” apples that deduce points if eaten. The choices are infinite!

Video transparency VS turn video off

In the first case (transparency) we just change the way the video appears (or seems to disappear) on the screen. Even with full transparency (100%) the webcam is still on and it detects our movements. However, when we set the video off, it stops completely - and all motion recognition with it. Therefore it is necessary to turn it on and off based on the requirements of a game.



What we learned so far...

In **chapter 14 “RAISE Playground”**, we worked with a special variant of Scratch 3.0 that was developed for AI programming. We compared the RAISE Playground version of Scratch 3.0 with the Machine Learning for Kids variant.

We then used one of the many AI extensions of Playground to create a simple game in which we control a sprite using finger movements.



Further Activities:

RAISE Playground offers a powerful Scratch 3.0 variant, with extensions that allow the fast and easy development of games. For the following proposed activities, we assume that students already have basic knowledge of programming using Scratch, including the use of loops, variables and conditions.

Don't forget: In Playground you can use models trained in Teachable Machine or rely on the pre-trained models.

Activity 1: Create a finger-controlled game

- Create one example for your students - a game in which they must (for some reason) break green balloons but not black ones (they might explode -just think of an excuse!). Show how to play this game using a camera and your fingers.
- Ask your students to come up with their own description of a similar game and allow them time to create it. Allow their peers to test each others game.
- Ask students to suggest improvements/changes to the games their peers created!

Activity 2: Escape the labyrinth!

- Use the extension for controlling your robot on screen with the tip of your index finger.
- Create a labyrinth as a background (maybe add "hidden" treasure for added complexity and excitement!).
- Constraint your robot to horizontal and vertical movement, to move perpendicular to the labyrinth walls.
- Suggestion: maybe you recreate the game using sprites that resemble a mouse and cheese.

Activity 3: Play headball!

- Use the extension that tracks body movement (in this case, the head)
- Create a game where you have to "shoot" the ball in the right basket. For example, red balls might go to the red basket, blue balls to the blue basket.
- Place baskets left and right on the screen.
- Set the balls to fall randomly from the sky, and try to hit them with your head and send them to the right basket!

Short bio...

I was born in Nicosia, Cyprus, on Christmas Eve of 1974. My parents bought my first computer, an Atari 520STFM, in early 1988. I started programming the same year, using the Metacomco BASIC that was included with my machine, and started developing video games a year later with STOS BASIC. I began my studies at the Higher Technical Institute (Electrical Engineering) but moved after one year to the University of Cyprus for a B.A. in Education. During my studies, I bought my first Macintosh computer and with savings from my part time jobs, and a Bank Loan, I purchased Hypercard 2.2 and later Macromedia Studio 3. I continued my studies with an M.A. on Curriculum and Instruction and a PhD on Information and Communication Technologies. I have been working as a teacher since 1999, with a brief hiatus from the classroom between 2009-2011 to work at the Pedagogical Institute, developing educational content for supporting ICT in Education. My personal and research interests include Online Learning Management Systems, video games in education and computer history. With my childhood friend, Nicholas Ktenas, we founded (2014) the first Cyprus Computer History Museum.

Alexandros Kofteros, PhD

alexandros@mathisis.org





An introductory guide for using AI in education

ARTIFICIAL INTELLIGENCE: TEACHING & LEARNING WITH AI

Alexandros Kofteros, PhD

This guide is aimed mainly at teachers wishing to learn more about AI and implement related activities with their students. The (student) target group is upper primary to secondary education.

With this guide:

- We take a glimpse at a short history of AI
- We engage in activities involving Machine Learning
- We generate text, images, charts and presentations using AI
- We develop simple games using pre-trained models
- ...and we get to appreciate robots a bit more...



ISBN 978-9925-8055-1-8

