



AI for Teachers: an Open Textbook

AI for Teachers: an Open Textbook

COLIN DE LA HIGUERA AND JOTSNA IYER



AI for Teachers: an Open Textbook Copyright © 2024 by Colin de la Higuera and Jotsna Iyer is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/), except where otherwise noted.

Contents

About the book	1
Acknowledgements	2
About the Project	iii
Preface	iv

Part I. Why Learn About AI

1. Technology, Change and You	10
2. AI is Everywhere	12
3. AI Already in Education	15
4. How AI Can Help You	17
5. Why not just do AI - Part 1	19
6. Why not just do AI - Part 2	21

Part II. Finding Information

7. Search Engines Part 1	24
8. Search Engines Part 2	28
9. AI Speak: Machine Learning	31
10. AI Speak: Search Engine Indexing	36
11. AI Speak: Search Engine Ranking	42
12. Behind the Search Lens: Effects of Search on the Individual	48
13. Behind the Search Lens: Effects of Search on the Society	53

Part III. Managing Learning

14. Smart Learning Management Systems	56
Manuel Gentile and Giuseppe Città	

15. Learning Analytics and Educational Data Mining Azim Roussanaly; Anne Boyer; and Jiajun Pan	60
16. AI-Speak: Data-based systems, part 1	64
17. AI-Speak: Data-based Systems, part 2	69
18. Issues with Data: Personal Identity	75
19. Issues with Data: Bias and Fairness	78

Part IV. Personalising Learning

20. A Note on Personalisation	83
21. Adaptive Learning Systems	88
22. AI Speak: How Youtube Learns You, part 1	92
23. AI Speak: How Youtube Learns You, part 2	96
24. AI-Speak: How Adaptive Systems Learn the Learner, part 1	99
25. AI-Speak: How Adaptive Systems Learn the Learner, part 2	103
26. The Flip Side of ALS: Some Paradigms to take note of	107

Part V. Listening, Speaking and Writing

27. Translators	112
28. Writing with AI Manuel Gentile and Giuseppe Città	116
29. AI-Speak: Deep Neural Networks	118
30. AI-Speak: Natural Language Processing	124
31. AI, AIED and Human Agency Wayne Holmes	128
32. Homogenisation, Invisibility and beyond, towards an Ethical AI	131

Part VI. On Generative AI

33. Introducing Generative and Conversational AI Michael Hallissy and John Hurley	135
34. Generative AI for the Classroom, part 1	140
35. Generative AI for the Classroom, part 2	145

36. ChatGPT and its Potential Impact on Homework in Schools	151
Michael Hallissy and John Hurley	
37. The Gears of Generative AI	156
Manuel Gentile and Fabrizio Falchi	
38. The Art, Craft or Science of Prompting	159
Bastien Masse	
39. The Degenerative, part 1	162
40. The Degenerative, part 2	166

Part VII. The Next Steps

41. Open or closed?	170
42. Artificial Intelligence, Homework, Exams and so forth	173
43. Obsolescence?	177
44. Individual or collective AI	179
45. Teaching AI	182

Part VIII. Additional Content

Optical Character Recognition	187
42	188
A brief description of some search engines	189
Optimising search	192
X5GON	194
Does Data always have to be labelled?	197
How many features are too many?	198
Hands-on Machine Learning	199
Cookies and Fingerprinting	201
More on Big Data	203
Other terms related to Personalised learning	205
Some automatic translation vocabulary	208
AI technology is moving fast	209
Understanding debates on the possible dangers of AI	210
Generative AIs - where should a teacher use them?	213

Transformers	216
Manuel Gentile and Fabrizio Falchi	
GDPR in a nutshell	217
AI and Coding	218
Machine Learning and AI Through Data Experiments in Orange	219
Blaž Zupan	
As a very provisional conclusion...	222



AI for Teachers, An Open Textbook

by

Colin de la Higuera and Jotsna Iyer

with guest contributions from

Anne Boyer, Azim Roussanaly, Bastien Masse, Blaž Zupan, Fabrizio Falchi,
Giuseppe Città, Jiajun Pan, John Hurley, Manuel Gentile, Michael Hallissy, Wayne
Holmes

*This textbook was produced as a deliverable for the Erasmus + project AI4T
626145-EPP-1-2020-2-FR-EPPKA3-PI-POLICY. It benefited from the collaborative work of all
partners.*

*The European Commission's support for the production of this publication does not constitute
an endorsement of the contents, which reflect the views only of the authors, and the
Commission cannot be held responsible for any use which may be made of the information
contained therein.*



We thank all members of the AI4T (Artificial Intelligence for Teachers) project for their help and participation in content workshops and meetings from which many ideas for this work emerged. We also thank the researchers, teachers and students, and the learning groups with which we worked, who also contributed. Writing this book would not have been possible without their help and input.



AI4T is an Erasmus+ project. As it is a Key Action 3 project, it includes and is piloted by ministries. AI4T is based on the analysis that AI and education are not just topics for industry. The education system should be prepared to identify how best to make use of AI in the classroom, reassure teachers, make them responsible users and start an effective teacher-training program.

A much longer and more complete presentation of the project can be found [on the project webpage](#). The project has been presented in public through webinars and conferences, and we usually put forward the following objectives:

- To build courseware to train teachers how to use AI in an educational setting;
- To use this courseware in training sessions for teachers in five countries;
- To evaluate and document the quality of the training.

For such an ambition to be feasible, a strong consortium has been built, involving the ministries of education from all five countries – France, Ireland, Italy, Luxembourg and Slovenia. Evaluation specialists from all five countries, and academic teams with expertise in artificial intelligence and education, have been involved.

- **Ministries**

- Ministère de l'Education nationale, de la Jeunesse et des Sports (F)
- Dublin West Education Centre (IR)
- Ministero dell' Istruzione (IT)
- Service de Coordination de la Recherche et de l'Innovation pédagogiques et technologiques (LU)
- Ministrstvo za izobraževanje, znanost in šport (SL)

- **Evaluators**

- Conservatoire national des Arts et Métiers (FR)
- Educational Research Centre (IR)
- Istituto Nazionale di Documentazione, per l'Innovazione e la Ricerca Educativa (IT)
- Université du Luxembourg (LU)
- Pedagoški Inštitut (SL)

- **Research labs**

- Institut national de recherche en sciences et technologies du numérique (FR)
- Nantes Université (LS2N) (FR)
- Université de Lorraine (LORIA) (FR)
- H2 Learning (IR)
- Univerza v Mariboru (SL)
- Consiglio Nazionale delle Ricerche (IT)

Preface to this second edition

Welcome!

October 2022 to October 2023

The first edition of this textbook was published in October 2022. Within days, ChatGPT appeared, and we entered into twelve months of AI madness. Every week, new products were launched and improvements to language models and their applications were announced. More importantly, education seemed to suddenly become a benchmark for generative AIs. Teachers and institutions reacted rapidly, acknowledging the arrival of a new tool and incorporating it into the education toolbox – or prohibiting it because of the menace it was felt to convey. There were discussions in the press but also in international organisations; petitions and open letters were signed. The impact on the jobs market was measured, and some companies have already started to replace their workers with AI.

For the authors of this book, all of this resulted in a key question, a challenge, and an opportunity.

The *question* was the one any author of a technology-linked book is afraid of. Is the book obsolete? This could be the shortest lifespan of any book – just a question of days. The *challenge* was to aim to include the novelties resulting from the ChatGPT tsunami into a second edition. And the *opportunity* was to share the book in the best possible moment, when it was probably most needed.

The question – does Generative AI's importance mean that the rest of AI is now unnecessary?

The question makes sense. ChatGPT has been adopted by many because it is so easy to use. Some generative AI experts of 2023 knew little about AI in 2022. It is therefore tempting to believe that generative AIs are built on thin air and can be understood – if that is the goal – by reading only what has been published in the past year. So, is it still necessary to understand machine learning and the different tools developed by AI techs over the past 70 years?

We believe the answer is “yes”. Even if a spectacular step, generative AI builds on technologies and ideas that have been shared for decades. Understanding data, bias, unsupervised learning, personalisation and ethics is still key to what a teacher should know before using AI in the classroom.

The Challenge

The challenge is to write about a fast-moving technology in a way that would satisfy a teacher who, understandably, wants to work from non-ephemeral knowledge, to

build their teaching from concepts and technologies which will be resistant to time. One example is the notion of “hallucination” which has changed so much over the past twelve months and which is going to be crucial to how teachers will adopt generative AIs.

The Opportunity

The opportunity follows the urgency with which all stakeholders are today examining the question of artificial intelligence and education. Whereas in 2020, when the AI4T project was launched, the difficulty was going to be in recruiting enough teachers to learn about AI for the experimental results of the project to be valid. In 2023, this has become a question of the highest priority in all countries.

What is new in this second edition?

We obviously had to take into account the arrival of ChatGPT (and later of alternative AIs). And a whole section (7) is now devoted to understanding the phenomenon; it is starting to propose how a teacher should take advantage of these technologies.

For the more technical aspects, we have chosen to highlight images over text. There are therefore many new illustrations in this version. We have also added 15 short videos which will – we hope – help to understand important concepts.

The open and multilingual challenge

This is an open textbook which means that a Creative Commons (CC) licence has been used. All images, videos and extra material have been scrutinised in order to be openly shared. This means anyone can take the material, or part thereof, and reuse it as they wish. They can also make modifications. There are different export formats available, and the authors can probably share in any way which ensures that this is a sustainable textbook. It can live on with future versions and new projects.

As is customary, the only obligation is that of citation of the authors of the book or of specific chapters, where relevant.

One particular modification – translation – has been anticipated with as much care as possible. We are already translating the original English version into French, Slovene, Italian and German. Furthermore, new projects are emerging to translate the textbook into other languages. We believe AI can help with the translation process, but human correction is needed.

Please contact us to build a partnership if you want the book translated to your language!

What were we saying one year ago?

Let's start with what you already know: AI is everywhere and education is not an exception. For some, the future is bright and the coming technologies will help make education available to all; it might even help when there are not enough teachers. It will permit the teacher to spend more time on the 'noble' tasks while the machine will take over the 'boring' ones, such as grading, organising the classroom.

For others, these AI algorithms represent a huge danger, and the billions of dollars the industry is prepared to invest prove that education is now viewed as a market. But it is not a market.

Somewhere in the middle, between these rather different positions, are researchers, educators and policy makers who are aware of a number of things: artificial intelligence is here to stay and will be in the classroom if it is not there already. And no minister – let alone a teacher – will be able to stop this. So, given this fact, how can the teacher harness *the beast* and use artificial intelligence for the better? How can the teacher make the AI work for the classroom and not the inverse?

The purpose of this textbook is to support the teacher in doing this. It has been built in the context of the Erasmus+ project AI4T. Teams from Ireland, Luxembourg, Italy, Slovenia and France have worked together to propose learning resources for teachers to be able to learn about AI – specifically AI for education. The learning material and a presentation of the project and its results can be found on AI4T's webpage (<https://www.ai4t.eu/>).

Training teachers is an essential task for all ministries involved. The objectives are the following:

1. Making teachers aware of why such training is good. It can't be an imposed decision; it has to be shared.
2. Introducing AI: from our experience of many conferences and workshops, there are participants who have explored, read and digested the topic. However, the vast majority have not.
3. Explaining how AI works in the classroom. What are the mechanisms? What are the key ideas?
4. Using AI in educative tasks.
5. Analysing what is happening in the field and being active of future changes.

We hope the textbook will be able to help you with most of these questions. We analyse the current situation and link AI with the experience of the teachers. By so doing, we hope to encourage them to remain interested in these questions. Undoubtedly there will be new challenges, mistakes will be made, and there could be strong opposition and controversies. We have sections called 'AI Speak' in which we try to explain how and why the algorithms work. Our goal is to inform teachers who can then fully participate in the debates and discussions on education and artificial intelligence. Some reasons for preparing this material can be found in the video prepared by AI4T.

We believe in the following:

- Some AI literacy is necessary. Let's explain this, as it is often argued that 'you don't need to know how engines work in order to drive a car'. This is not entirely true: most of us don't know how engines work but accept there is science and technology involved. We accept this because in school we received lessons in basic physics and technology. In the same way, we wouldn't be satisfied with a book telling us not to smoke, based on statistical arguments about the number of smoking-related early deaths. Again, we are able to **understand** why smoking is harmful because at some point a teacher has explained to us how the respiratory system works, what lungs are, etc. Today, with AI making a huge impact on society, we believe that the same applies – finding out about the effects of AI is insufficient. Teachers need to have an understanding of **how** it works. The goal is not to make each person a biologist or a physicist – the goal is to make us understand the principles and ideas.
- Teachers are extraordinary learners. They will be critical when something is not explained the right way, and will engage more. They want to understand. This textbook is for people who are prepared to go the extra mile, who will not be satisfied until they *do* understand.
- Next, AI has to be used in a safe environment – computers or devices will be connected to the web and applications will run on the cloud. A huge security issue exists here, and a teacher needs reassurance that the working environment is safe for all. Computer security is a highly complex question; a teacher will not be able to check the software's safety specifications. A trusted source will need to do this.
- AI can help, provided it is used in a well-defined and controlled learning environment, for a task which the teacher has identified as important. For obvious economic reasons, the industry will push products on the teachers, ostensibly to help them solve a sometimes-unimportant task. But if it is considered 'cool', and is pushed by the seller, it could end up being seen as important. A good teacher should be aware of this. In this textbook, we introduce elements for the teacher to identify such products or situations.
- When preparing this courseware, we did have a serious problem. The idea was to use AI software which we could recommend to the teachers, so that they would rapidly be able to use them in the classroom. Unfortunately, this is not the case: a lot of software is still immature, there are a lot of ethical concerns and in most cases the different ministries and governments have not approved lists of software. We have therefore chosen a different approach: we will be mentioning We had a problem when preparing this courseware. The idea was to use AI software which we could then recommend to teachers, for use in the classroom immediately. Unfortunately, this is not what happened. Much software is still immature and have ethical concerns. In most cases, the various ministries and governments have not approved specific software. Because of this, we will be mentioning software in the textbook. This is because we believe it explains a particular point of AI in education. However, we do not endorse any particular software. It is expected that soon international agencies, such as Unesco, UNICEF or the Council of Europe, will come up with specific software recommendations.

We would like to thank the many contributors who have helped compile this textbook.

First and foremost, we have benefited from reading Wayne Holmes' works and enjoyed many hours of discussions with him.

Discussions also took place within the AI4T consortium. Workshops were organised in order to get the topics to emerge.

Teachers themselves have been an essential source of information – through seminars and webinars we exchanged ideas with them, and they let us know which were confusing and/or wrong.

Many gave valuable opinions, proofread documents and suggested links and texts. Some added chapters to this work:

- Manuel Gentile helped us in a number of chapters and showed great skill in making accessible the most obscure aspects of AI;
- Fabrizio Falchi and Giuseppe Città were great collaborators who have helped us understand a variety of AI questions;
- Azim Roussanaly, Anne Boyer and Jiajun Pan were kind enough to write the chapter on Learning analytics;
- Wayne Holmes wrote a chapter on agency. This is an important topic when discussing the ethical implications of AI;
- Michael Halissy and John Hurley explored the issues of homework and assessment with the arrival of Generative AIs;
- Bastien Masse is today an expert in mastering the prompt; he has shared his skills here;
- Blaž Zupan introduced the Orange software, which his team has been developing, in order to make use of machine learning.

We are also very much indebted to those who coordinated the translated of this textbook into French, Italian, German and Slovene. Our special thanks go to Solenn, Manuel, Daniela and Helena.

La Plaine sur Mer, 26/11/2023
Colin de la Higuera

PART I

WHY LEARN ABOUT AI

Have you ever asked yourself,

How can artificial intelligence impact learning and teaching in my classroom?

Can it help me do what I want to do with my students?

How can it change the dynamics and interactions I have with my students?

How do I even know when it is being used correctly or incorrectly?

And, what should I be aware of if I want to put it to good use?

Read on...

1.

REFLECTION

In 1922, Thomas Edison declared that the motion picture would revolutionise education. He believed it would replace all textbooks¹.

Yet, the teacher's use of film has been limited. The photographic slide projector was adopted by many teachers, from the 1950s up to the late 1990s.

Unlike when using film, teachers could:

- prepare their own slides cheaply
- use the projector like the blackboard – a tool that does not change their way of teaching
- reuse the slides, re-order them and refine them².

If there was a new technology that could help you,

1. *what features would you look for?*
2. *would you change your teaching practice to use it?*
3. *would you be afraid of being forced to change?*




"Carousel '77" by Voxphoto is licenced under CC BY-NC-ND 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-nd-nc/2.0/jp/?ref=openverse>.

Artificial intelligence and you

As a teacher, every day you deal with changes brought about by technology. Its applications change the world in which you teach, as well as the students in your care. Eventually, they change what you teach about – the content, skills and context. They can also change your teaching methods.

This textbook deals with how artificial intelligence (AI) can change the way you teach.

Why AI? Where it is effective, the speed and amplitude of the change AI brings can be destabilising. And what cannot be done by a machine that claims to be intelligent? Could it outperform you as a teacher? Could it replace you in the classroom? Several important questions need to be answered.



"Diary of a teaching machine" by [Ed] is licenced under CC BY-NC-SA 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-nc-sa/2.0/?ref=openverse>.

Today's artificial intelligence applications are built for one specific task and one type of user – for example, software that can translate what you write cannot predict stock prices.

As for the task itself, yes, AI can perform some tasks better than humans. But even a child can beat the best AI in many others. AI has a long way to go before being able to supplant a human in a rich cognitive, social and cultural activity such as teaching. But it can help by adding to what a teacher is capable of doing. Experts speak of "An Augmented Human"³, which, in our case, would be "An Augmented Teacher".

Where AI can help in education, pedagogical experts stress that the teacher should always be in the loop, overseeing what it does. Effective AI solutions in the classroom are those that empower the teacher. When the teacher knows what the student is learning, the gains are significant⁴.

The goal of this textbook is to give teachers the knowledge necessary for deciding if, where and how AI can help you. We hope to help you prepare yourself for the future, taking in your stride the changes brought by AI.

See what BBC's [Will a robot take your job](#) quiz has to say about the future of different jobs in the UK. It says your job is safer from automation if you have to negotiate, help and assist others, or come up with original ideas as part of work. For "secondary education teaching professional", it pegs likelihood of automation at 1%.

¹ Cuban, L., *Teachers and machines: The classroom use of technology since 1920*, Teacher College Press, 1986.

² Lee, M., Winzenried, A., *The use of Instructional Technology in Schools, Lessons to be learned*, Acer Press, 2009.

³ Holmes, W., Bialik, M., Fadel, C., *Artificial Intelligence In Education: Promises and Implications for Teaching and Learning*, 2019.

⁴ Groff, J., *Personalized Learning: The state of the field and future directions*, Centre for curriculum redesign, 2017.

2.

Artificial Intelligence: an intuitive understanding



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=31#oembed-1>

You could open a textbook of artificial intelligence or do a quick search on the internet, and you'll see that definitions of AI vary. There is no good way to tell what AI is, where it is used and what role it plays. It could be a complex stand-alone system like a robot or an autonomous car. It could be just a few lines of code inside another piece of software, playing a small role.

Artificial intelligence involves a collection of programs that do a diverse set of tasks. Mathematically and algorithmically, the lines blur – there is no clear indication of where AI starts and other technologies stop.

Furthermore, many experts disagree with the use of the word 'intelligence' – artificial intelligence has no resemblance to human intelligence! Yet, the word suggests to us what these programs are meant to achieve – the thread that connects them.

Ultimately, AI systems are machine based. They make predictions, recommendations or decisions by

- perceiving real or virtual environments (using items such as microphones or cameras),
- simplifying data and analysing it,
- using that analysis to make a decision or prediction¹.

When you come across a system that:

- recognizes what you have written or what you are saying (translation software, text recognition, face recognition, personal assistants, chatbox); or,
- seems to know you better the longer you use it (Youtube video

DEFINITIONS OF AI THAT USE "INTELLIGENCE", "MIND" OR "THINKING"

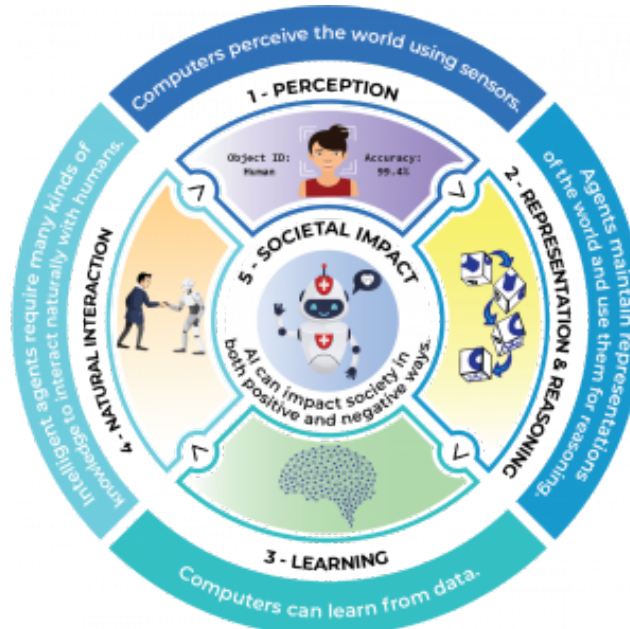


- **"The exciting new effort to make computers think...[as] machines with minds, in the full and literal sense."** (Haugeland 1985)
- **"The art of creating machines that perform functions that require intelligence when performed by people."** (Kurzweil 1990)
- **"The study of how to make computers do things which, at the moment, people are better."** (Rich and Knight 1991)
- **"Making machines intelligent; intelligence is that quality that enables an entity to function appropriately and with foresight in its environment."** (Nils Nilsson)

recommendation, news recommendation, Amazon item recommendation, Facebook suggested friends, targeted ads) ; or,

- is capable of predicting an outcome, given incomplete and fast-changing information (fastest route to go somewhere, stock prices in the near future),

...AI is probably involved.



Five big ideas in AI. Credit: AI4K12 Initiative.
Licenced under CC BY-NC-SA 4.0. Visit
<https://creativecommons.org/licenses/by-nc-sa/4.0/>

The AI around us

Artificial Intelligence has become the technology that gives its user the edge that's required to succeed.

Almost every field uses AI in one form or the other:

- From business to research, many fields use language apps to transcribe speech on the fly and obtain translation of impressive quality.
- Medicine profits from image analysis and AI-based decision-support tools².
- In agriculture, AI driven systems help optimize the use of available resources.
- Every day there's news about AI breakthroughs in games, art, industry and commerce.

At school, choosing courseware, adapting to individual learners, assessing them in a productive way and managing logistics, are all "intelligent" activities. According to the definition(s), AI-based software should ideally be able to help with pursuits such as these.

ACTIVITY

Make a list of five technologies you or your students have used in the past two years. How many of them, according to you, contain AI?

Alan Turing is considered by many as the father of computer science. Many of the new ideas in AI today were actually also introduced by Alan Turing, before the term 'Artificial Intelligence' was invented!



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=31#oembed-2>

¹ [The OECD AI Principles](#), 2019.

² [Artificial Intelligence in Healthcare](#), Wikipedia.

3.

Both teachers and students, knowingly or unknowingly, for good or bad, are already using AI inside and out of the classroom: How?

Further into this book, we will look at the artificial intelligence tools available for education. But many of the most useful applications have yet to enter the classroom. The **Educational technology** industry, big digital corporations and university research labs are all developing tools to help the teacher teach and the learner learn. More and more, corporates specialising in AI are making huge investments in education. Tools, both approved and unapproved by the concerned authority, but available freely, are used by teachers and students.

Whether made with education in mind or not, many of these tools may be used in the classroom. It's necessary to be aware of their benefits and potential problems .

One of the free applications for mathematics available during the writing of this text is Photomath. (For language teachers, a similar example might be a language learning app as Duolingo or writing software that uses GPT3.)



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=33#oembed-1>

Reflection

Often, while researching the use of a software, we have videos that have been published by the manufacturer. Even third-party reviews might or might not be affiliated.

How to separate the truth from oversell?

Is the application really as useful as the video claims it to be?

Is there difficulty in using its features?

What are some potential problems that could come with using this?

[Click here to read about](#)

Optical
Character
Recognition

Photomath is a *mathematical solver*. It takes a mathematical equation and solves it. Teachers have been contending with calculators, as a tool for teaching as well as cheating.

What makes Photomath so powerful is the ease of use – just click a picture of the blackboard or notebook. The AI in Photomath scans the photo and solves the problem directly.

Let's say a calculator gives you an answer of [42](#). Teachers may allow their use to check the result, but the students must arrive at the solution by themselves. Solvers such as these show multiple ways to solve and visualise a problem, although this part is much less technically demanding for the programmer

Other applications found in today's classrooms:

- Search engines
- Spell check and grammar correction which is built into most writing software
- Online translators
- Language learning apps
- Math solvers like Photomath, Geogebra and Wolfram
- Personal assistants.
- Chatbots
- Intelligent tutoring systems
- AI-powered learning management systems

Reflection

There is another software called [Checkmath](#) which is similar to [Photomath](#). Take a look at their respective web sites. **If you have to choose one**, which one would you go for? Why?

Is using AI cheating: What are the teachers' reactions?

Here are some teacher's reactions to pupils' use of AI:



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=33#oembed-2>

4.


When it comes to technology, there are two extremes to watch out for:

- Under-utilisation of technology due to fear and ignorance
- Indiscriminate use that might lead to undesirable secondary effects

For example, excessive use of mobile phones can be harmful. Some societies avoid mobile phones entirely. However, most people don't overuse them. Prudent use of mobile technology has, in fact, saved lives.

In order to avoid succumbing to the first example above, it would be helpful to have knowledge of important educational applications. We will take a closer look at each of these in the upcoming chapters. Here are some examples.

AI tools for managing learning



"Diary of a teaching machine" by [Ed] is licenced under CC BY-NC-SA 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-nc-sa/2.0/?ref=openverse>.

Artificial intelligence **dashboards**, other **data visualisation tools** and **learning management systems** bring all the available information together. These help to monitor student performance over multiple subjects or track progress on any topic in classrooms of any size.

AI applications can **flag potential problems**, such as absenteeism and behaviour common to dropouts. All the data thus gathered can serve as **self-assessment for the teacher** by showing where the lessons are effective and where a change of approach is required.

AI is suitable for **scheduling and resource-optimisation** tasks. But the most important application is that which deals

with inclusion and integration of the differently abled. Human-machine interface has never been as seamless as it is now, making multimedia input and output a real possibility. For example, the app [Storysign](#) helps translate words to sign language to help deaf children learn to read.

AI tools for personalising learning



- **Adaptive learning systems** (ALSs) evaluate the learner, be it through quizzes or real-time feedback. Based on this evaluation, they present the student with a predefined learning path. Instead of a one-size-fits-all approach, students can spend more or less time on each topic, explore new and related topics. This adaptive software can help them learn to read, write, pronounce and solve problems.
- ALSs can also help learners with special needs. Any specialisation of the systems will be based on proven theories and expert opinion. Targeted systems “are likely to be of great assistance in teaching individuals with cognitive disabilities such as Down Syndrome, traumatic brain injury, or dementia, as well as for less severe cognitive conditions such as dyslexia, attention deficit disorder and dyscalculia”¹.
- Different groups can be formed for different activities (**‘clustering’**), taking into account the individual strengths and weaknesses of each member.

While these technologies can help, “the devil is in the detail of how you actually **use** the technology”². The same innovative and powerful learning technology can be used effectively in one school – and badly in another².

Again, knowledge is the key!

¹ Alkhatlan, A., Kalita, J.K., *Intelligent Tutoring Systems: A Comprehensive Historical Survey with Recent Developments*, International Journal of Computer Applications 181(43):1-20, March 2019.

² Groff, J., *Personalized Learning: The state of the field and future directions*, Centre for curriculum redesign, 2017.

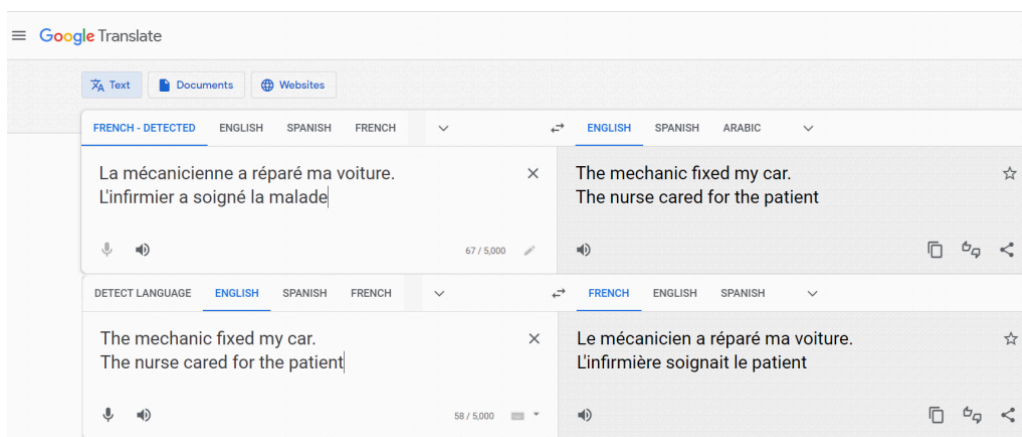
5.

The second extreme position when it comes to AI is the indiscriminate use or misuse of the technology. Artificial intelligence works differently from human intelligence. Whether due to the nature of the situation, its design or the data, AI systems can work differently from what is expected.

For example, an application developed using a set of data for one purpose will not work as well on other data for another purpose. It pays to know the limitations of artificial intelligence and correct for it; it is good to not just do AI but learn about its advantages and limitations.

Perpetuation of Stereotypes

Google Translate learns how to translate from the internet. Its 'data miners' scout the public web for data from which to learn. Along with language, AI learns that the number of male mechanics are more than that of female mechanics, and that the number of female nurses eclipses that of male nurses. It cannot differentiate between what is 'true' and what is a result of stereotyping and other prejudices. Thus, Google Translate ends up propagating what it learns, cementing stereotypes further¹:



"Female mechanic" and "male nurse" when translated to English and then back to French become "Male mechanic" and "female nurse". Example inspired by Barocas, S., Hardt, M., Narayanan, A., Fairness and machine learning Limitations and Opportunities, MIT Press, 2023.

Problems creep up in AI whenever an individual case differs from the majority case (whether this represents faithfully the majority in the real world or just the majority

as represented by the internet). In classrooms, the teacher has to compensate for the system's lapses and, where necessary, direct student's attention to the alternative text.

Explore

Can you hunt for a stereotype in [Google Translate](#)? Play with translating to and from different languages. By clicking on the two arrows between the boxes, you can invert what is being translated (This is what we did for the example shown above).

Languages such as Turkish have the same word for 'he' and 'she'. Many stereotypes come to light when translating from to Turkish and back. Note that many languages have a male bias – an unknown person is assumed to be male. This is not the bias of the application. What is shocking in our example above is that the male nurse is changed into female.

Multiple Accuracy measures

"AI systems will have great difficulties in dealing with people who are creative, innovative, and not only average representations of vast collections of historical examples."

THE IMPACT OF ARTIFICIAL INTELLIGENCE ON LEARNING, TEACHING, AND EDUCATION, JRC SCIENCE FOR POLICY REPORT

AI systems make predictions on what a student should study next, whether he or she has understood a topic, what group split-up is good for a class or when a student is at risk of dropping out. Often, these predictions have a percentage tagged to them. This number tells us how good the system estimates its predictions to be.

By its nature, predictions can be erroneous. In many applications, it is acceptable to have this error. In some cases, it is not. Moreover, the way this

error is calculated is not fixed. There are different measures, and the programmer chooses what he or she thinks is the most relevant. Often, accuracy changes according to the input itself.

Since in a classroom, these systems make predictions on children, it is for the teacher to judge what is acceptable and to act where a decision taken by AI is not appropriate. To do this, a little background on AI techniques and the common errors associated with them will go a long way.

¹ Barocas, S., Hardt, M., Narayanan, A., [Fairness and machine learning Limitations and Opportunities, MIT Press, 2023.](#)

6.

Data and Privacy

"Data is becoming the new raw material of business"

CRAIG MUNDIE
US : IT IS ALSO THE NEW EXHAUST OF BUSINESS

All businesses use data to improve their balance sheets. They use data to decide what to sell, who to sell it to, what price to fix and how to tailor their advertisements. It is the machine-learning algorithms that make sense of data. Therefore, the winner is whoever has the better data and algorithms. Data is the new gold – and the new

Achilles heel.

Does data here mean just personal addresses and bank accounts?

What about the number of mouse clicks a user makes while visiting a web site?

As stewards of their own data and that of their students, it is imperative that teachers know what kind of AI data is useful, what forms do they take and how can users' privacy be protected.

AI and the business of education

"EdTech" is the industry that makes technological applications for education – including those using artificial intelligence. These can be small companies or start-ups. They can be internet giants that are starting to pump money into education. They can also be entities with public funding.

Some EdTech software has to be bought. The rest is free, with income from other sources – often targeted ads and reselling of user data. Whatever financial model is used in EdTech AI, money can be made.

What does this mean for you and your students? Is there such a thing as a free lunch? How do we secure our classrooms while enjoying the fruits of a greedy industry?

Tool creation and you

Education does not have to change in order to accommodate technology. "Learning environments that start with technology often go down unwanted paths"¹. Any tool should be based on sound pedagogical theories. Furthermore, to be most effective, it should be co-created in teams involving teachers, pedagogical experts and computer scientists².

So, ready to start?

¹ Groff, J., *Personalized Learning : The state of the field and future directions*, Centre for curriculum redesign, 2017.

² Du Boulay, B., Poulivasillis, A., Holmes, W., and Mavrikis, M., *Artificial Intelligence And Big Data Technologies To Close The Achievement Gap*, 2018.

PART II

FINDING INFORMATION

In June 1993, there were 130 websites. By the start of 1996, there were 100,000. Various estimates peg the number at 1.7 billion, as of 2022.

This explosion would hardly be meaningful without the power to find exactly the information we want. Search engines excel in helping us do just this.

They read our hastily typed, often misspelled queries, and pull out text, images, videos and all sorts of relevant content.

How does this quick access to information help education?

How can it aid student led learning – where students build their knowledge through constructive activities?

How to make the most of this technology while avoiding its drawbacks?

7.

Activity

Pick a search engine from the list below:

[Bing](#) [OneSearch](#)
[Brave](#) [Qwant](#)
[DuckDuckGo](#) [Spotlight](#)
[Ecosia](#) [Startpage](#)
[Google](#) [Swisscows](#)
[MetaGer](#) [Yahoo!](#)

1. Are the search results as good as other engines you use regularly?
2. What are the search engine's sources? Is it dependent on other search engines for its results?
3. Read the *About us* and *Privacy Policy* or *Terms of Use* pages.
4. According to these pages, what does the company do with your data? Can you change the default privacy settings?

Once done, please have a look [here](#) for a brief description of each engine.



"duckduckgo [Www.Etoile.App]" by eXploration Etoile is marked with Public domain mark 1.0. To view the terms, visit <https://creativecommons.org/publicdomain/mark/1.0/?ref=openverse>.

A search engine is the AI in your pocket. It is the most sophisticated use of Artificial Intelligence that most of us use regularly. The success of search engines is due to

- The explosion of content on the world wide web;
- A search engine's capacity to *make sense of* this content and tag it for future searches (Indexing);
- Its capability to *understand** what you are asking¹;
- Its ability to show the most relevant content first (ranking).

AI algorithms power the three last factors.

Search engines – *Knowledge Engines* as some call them – have been successful in creating the illusion that everything in the world is on the web and that everything on the web is searchable². This ready-to-serve knowledge, knowing and method of extending memory is transforming learning.

Search engines as learning and teaching tools

There are at least three ways in which search engines can help teachers and learners:

- Making it easy to find and verify information for use in lessons and tests. In this context, the meaning of information has undergone a big change in the last decade. As well as texts, audio, video, animation and even pieces of code are now easily accessible. So are search forums and digital repositories.
- Relieving the need for the teacher to be the sole source of knowledge. Teachers are now free to help hone skills, incite enquiry and resolve conflicts and doubts, if they so choose.
- Enabling students to do exploratory and project-based learning by opening up access to information. Students can access, sequence and derive their own meaning from information. This leads to learning that lasts and can be transferred to real-life situations³.

However, exploratory learning or any other kind of student led learning does not come easily. Students need help and scaffolding for many skills that go with searching for and putting together their own information sources³.

What to ask and how to ask it



- coming up with what to ask
- how to ask it
- how to find relevant and credible sources
- how to analyse what they found
- how to put together all this information

Various studies in Europe and further afield show that students struggle to search the web efficiently and effectively³. They often get frustrated when their search turns up nothing or do not know how to assess the relevance of search outcomes⁴. Younger children seem to have four distinct difficulties in searching for information on computers: creating search queries, selecting an appropriate website from a search list, spelling search terms correctly and understanding the language used in search results⁵.

Knowing some searching techniques will go a long way in helping both you and your students use this incredible resource.

Optimising Search

1. In a search engine of your choice, type **search** and look at the results. Compare them with the results of **how to search** or **search tips**. Now search for: how artificial intelligence is used in search engines?
2. Compare the results of **baked cheese recipes** and **"baked cheese" recipes**. How can you know that searching for **"East German" stories** is better than **East German stories** without going beyond the first page of search results?
3. Is **Eat, Pray, Love** the same as **eat pray love**?
4. Try **planet near jupiter**. Find a restaurant close to the Eiffel Tower.
5. What does **artificial intelligence** - **"machine learning"** do?
6. What is the difference between the search queries **"tom cruise" AND "john oliver"**, and, **"tom cruise" OR "john oliver"**?
7. Compare the results of **university of California** and **university of * California**. What does * do?
8. Try **courses site bbc.com** and **courses site:bbc.com**. Find courses in all websites that have a .edu web address(URL).
9. Add **filetype:pdf** to any search query to learn its use.

[You can find some pointers here](#)



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=49#oembed-1>

Apart from practising good search techniques, it is always worthwhile to scroll down search results and check out pages beyond the first page. Search engines differ on how they rank results. The top results may not be based only on your search query and user activity. And not everyone knows how to write web pages that are optimised to show up in the top results.

To end our discussion here, please have a look at search settings, whatever search engine you like to use. They allow you to adjust how search results are shown and whether to set controls suitable for children, among other things.

Searching in pairs

Even after learning optimisation techniques, students may still have problems coining effective search terms and analysing the results. There is some evidence that searching in pairs or in groups of three can help. Discussing every step of an exercise can help students find better search strategies, correct the results and weigh what to do with the obtained information. Pairs may also be better at locating and judging information within sites when compared to individuals⁴.

* *make sense of, understand, intelligent* and other words are used in this text to describe the action of machines. It is important to bear in mind that machine based applications cannot make sense of or understand something in the same way humans do.

¹ Russell, D., *What Do You Need to Know to Use a Search Engine? Why We Still Need to Teach Research Skills*, AI Magazine, 36(4), 2015

² Hillis, K., Petit, M., Jarrett, K., *Google and the culture of search*, Routledge, 2013.

³ Marion Walton, M., Archer, A., *The Web and information literacy: scaffolding the use of web sources in a project-based curriculum*, British Journal of Educational Technology, Vol 35 No 2, 2004.

⁴ Lazonder, A., *Do two heads search better than one? Effects of student collaboration on web search behaviour and search outcomes*, British Journal of Educational Technology, Vol 36 Issue 3, 2005.

⁵ Vanderschantz, N., Hinze, A., Cunningham, S., "Sometimes the internet reads the question wrong": *Children's search strategies & difficulties*, Proceedings of the American Society for Information Science and Technology, Vol 51, Issue 1, 2014.

8.

Authenticity and Relevance

We have been taught to have confidence in books. Publishers, librarians, professors and subject-matter specialists act as gatekeepers to printed resources. They make sure the resource is authentic and of good quality. How to find, and make sure the students find, appropriate sources when it comes to the web^{1,2}?

Even printed material from good sources have been proven to be full of errors and biases. Perhaps nurturing the culture of critical evaluation would reap benefits beyond choosing good online sources².



- **Who** wrote this? What are their **qualifications**?
- What are their **affiliations**? How do they impact what they write?
- What is the **context**? Is this one of a series, a chapter in a book or content in a journal?
- Who is the **publisher**?
- Which **website** is it published in? Are there any spelling errors in the site address?
- Does the site address have **.edu**(educational institution) or **.gov**(government) or **.gouv.fr**(french government) or **gouvernement.lu**(luxembourg) in it?
- **Where** is it based? Does this change how relevant it is to my subject??
- Have **sources** been cited clearly? Are there **logical errors**?

Digital Repositories



"MIT Open Courseware" by stevegarfield is licenced under CC BY-NC-SA 2.0. To view a copy of this licence, visit

<https://creativecommons.org/licenses/by-nc-sa/2.0/?ref=openverse>.

public. It is all public information collected from studies, surveys, and censuses⁴.

– [research articles](#), where millions of papers and books, both open access and fee based, are indexed and ready for search.

Open courseware from universities, [Khan Academy](#) and online encyclopedias are all popular sources of information.

One good way to ensure content stays authentic is to use digital collections from known and trusted sources. These can vary from school-level to global resources. Over the past two decades, the number of digital libraries has increased dramatically, allowing educators to access and use documents such as data sets, maps and images³.

To find information, you can use either search inside the corresponding web site or using *site:* keyword in a search engine.

Even Google has specific search engines for

– [public datasets](#), where governments, public institutions, and large corporations make their collected data available for the

Volatile content

Search engines are constantly testing new algorithms. Mobile search differs from desktop search since it tends to favour results that are tied to current location⁴.

Every day, new content is indexed, while old content is indexed in a different way. New data is created from old data by re-analysing content¹. Copyrights and licences change. Laws concerning data change too, chronologically and in terms of location. For example, within the European Union, due to GDPR, search engines continue to list content to which index has been removed. Even maps change, depending on where they are accessed from. Language and its usage changes. Both the use of medical categories and their interpretation vary from country to country¹.

Not to forget that search results are ranked according to the history of user activity, their personal information and privacy settings. Thus all of us have access to different content and might not even be able to find the same content twice. All these

Find out!

Read about
searching the web
for open
educational
resources [here](#)

differences have to be taken into account while setting and grading educational activities.

Other challenges

Search engines bring in other changes too. Knowledge is easily available. We don't have to know facts any more. Instead, we try to remember where and how to find it¹. Programmers cut and paste snippets of code. Engineers run simulators. There are forums for both homework questions and teacher complaints. Skill is becoming more important than knowledge and memory. Even our views on ethics and morality change – how do we explain plagiarism to the copy-and-paste generation?

¹ Russell, D., *What Do You Need to Know to Use a Search Engine? Why We Still Need to Teach Research Skills*, AI Magazine, 36(4), 2015.

² Marion Walton, M., Archer, A., *The Web and information literacy: scaffolding the use of web sources in a project-based curriculum*, British Journal of Educational Technology, Vol 35 No 2, 2004.

³ Land, S., Hannafin, M. J., & Oliver, K. *Student-Centered Learning Environments: Foundations, Assumptions and Design*. In Jonassen, D. H. & Land, S. (Ed.), *Theoretical foundations of learning environments* (pp. 3–26), Routledge, 2012.

⁴ Spencer, Stephan. *Google Power Search: The Essential Guide to Finding Anything Online With Google*, Koshkonong. Kindle Edition.

9.


An **algorithm** is a fixed sequence of instructions for carrying out a task. It breaks down the task into easy, confusion-free steps, like a well written recipe.

Programming languages are languages that a computer can follow and execute. They act as a bridge between what we and a machine can understand. Ultimately, these are switches that go on and off. For a computer , images, videos, instructions are all 1s (switch is on) and 0s (switch is off).

When written in a programming language, an algorithm becomes a **program**.

Applications are programs written for an end user.

Conventional programs take in data and follow the instructions to give an output. Many early AI programs were conventional. Since the instructions cannot adapt to the data, these programs were not very good at things like predicting based on incomplete information and natural language processing (NLP).



A search engine is powered by both conventional and **Machine learning** algorithms. As opposed to conventional programs, ML algorithms analyse data for patterns and use these patterns or rules to make future decisions or predictions. So, based on data, good and bad examples, they find their own recipe.

These algorithms are well suited for situations with a lot of complexity and missing data. They can also monitor their own performance and use this feedback to become better.

This is not too different from humans, especially when we see babies learning skills outside the conventional educational system. Babies observe, repeat, learn, test their learning and improve. Where necessary, they improvise.

But the similarity between machines and humans is shallow. “Learning” from a human perspective is different, and way more nuanced and complex than “learning” for the machine.



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=60#oembed-1>

A classification problem


One common task a ML application is used to perform is classification – is this a photo of a dog or a cat? Is this student struggling or will they pass the exam? There are two or more groups, and the application has to classify new data into one of them.

Let us take the example of a pack of playing cards – group A and group B – divided into two piles and following some pattern. We need to classify a new card, the ace of diamonds, as belonging to either group A or group B.

First, we need to understand how the groups are split – we need examples. Let us draw four cards from group A and four from group B. These eight example cases form our **training set** – data which helps us see the pattern – “training” us to see the result.

As soon as we are shown the arrangement to the right, most of us would guess that the ace of diamonds belongs to Group B. We do not need instructions, because the human brain is a pattern-finding marvel. How would a machine do this?

ML algorithms are built on powerful statistical theories. Different algorithms are based on different mathematical equations that have to be chosen carefully to fit the task at hand. It is the job of the programmer to choose the data, analyse what features of the data are relevant to the particular problem and choose the correct ML algorithm.



*“Playing Card” par aquarianinsight.com/
free-readings/ est sous licence CC BY-SA 2.0.
Pour une copie de la licence, visitez
[https://creativecommons.org/licenses/by-sa/
2.0/?ref=openverse](https://creativecommons.org/licenses/by-sa/2.0/?ref=openverse).*

The importance of data

The card-draw above could have gone wrong in a number of ways. Please refer to the image. 1 has too few cards, no guess would be possible. 2 has more cards but all of the same suit – no way to know where diamonds would go. If the groups were not of the same size, 3 could very well mean that number cards are in group A and picture cards in group B.

Usually machine learning problems are more open ended and involve data sets much bigger than a pack of cards. Training sets have to be chosen with the help of statistical analysis, or else errors creep in. Good data selection is crucial to a good ML application, more so than other types of programs. Machine learning needs a great number of relevant data. At an absolute minimum, a basic machine-learning model should contain ten times as many data points as the total number of features¹. That said, ML is also particularly equipped to handle noisy, messy and contradictory data.




*"Playing Card" by aquarianinsight.com/
free-readings/ is licenced under CC BY-SA 2.0.
To view a copy of this licence, visit
[https://creativecommons.org/licenses/by-sa/
2.0/?ref=openverse](https://creativecommons.org/licenses/by-sa/2.0/?ref=openverse).*

Feature Extraction

When shown Group A and Group B examples above, the first thing you might have noticed could be the colour of the cards. Then the number or letter and the suit. For an algorithm, all these features have to be entered specifically. It cannot automatically know what is important to the problem.

A Machine Learns



While selecting the features of interest, programmers have to ask themselves many questions. *How many features are too few to be useful? How many features are too many? Which features are relevant for the task? What is the relationship between the chosen features – is one feature dependent on the other? With the chosen features, is it possible for the output to be accurate?*



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=60#oembed-2>

The process


*Does Data
always have to
be labelled?*

Read [here](#)

When the programmer is creating the application, they take data, extract features from it, choose an appropriate machine-learning algorithm (mathematical function which defines the process), and train it using labelled data (in the case where the output is known – like group A or group B) so that the machine *understands* the pattern behind the problem.

For a machine, *understanding* takes the form of a set of numbers – weights – that it assigns to each feature. With the correct assignment of weights, it can calculate the probability of a new card being in group A or group B. Typically, during the training stage, the programmer helps the machine by manually changing some values. This is called **tuning** the application.

Once this is done, the program has to be tested before being put to use. For this, the labelled data that was not used for training would be given to the program. This is called the **test data**. The machine's performance in predicting the output would then be gauged. Once determined to be satisfactory, the program can be put to use – it is ready to take new data and make a decision or prediction based on this data.



Can a model function differently on training and test datasets? How does the number of features affect performance on both? [Watch this video to find out.](#)

The real-time performance is then continuously monitored and improved (feature weights are adjusted to get better output). Often, real-time performance gives different results than when ML is tested with already available data. Since experimenting with real users is expensive, takes a lot of effort, and is often risky, algorithms are always tested using historic user data, which may not be able to assess impact on user behaviour¹. This is why it is important to do a comprehensive evaluation of machine learning applications, once in use:



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=60#oembed-3>

Feel like doing some hands on Machine Learning? [Try this activity.](#)

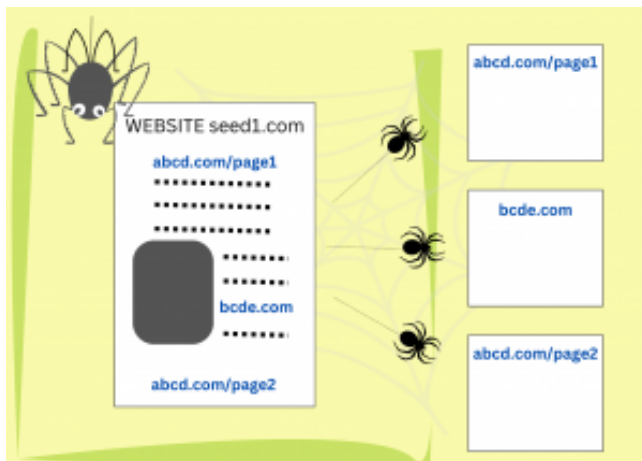
¹ Theobald, O. *Machine Learning For Absolute Beginners: A Plain English Introduction* (Second Edition) (Machine Learning From Scratch Book 1) (p. 24). Scatterplot Press. Kindle Edition.

² Konstan, J., Terveen, L., *Human-centered recommender systems: Origins, advances, challenges, and opportunities*, AI Magazine, 42(3), 31-42, 2021.

10.

A search engine takes keywords (the search query) entered into the search box, and tries to find the web documents that answer the information. It then displays the information in an easily accessible form, with the most relevant page at the top. In order to do this, the search engine has to start by finding documents on the web and tagging them so that they are easy to retrieve. Let us see in broad strokes what is involved in this process:

Step 1: Web crawlers find and download documents



Idea from "Search engine crawlers" by Seobility, licenced under CC BY-SA 4.0. To view a copy of this licence, visit https://www.seobility.net/en/wiki/Creative_Commons_License_BY-SA_4.0

After a user enters a search query, it is too late to look at all the content available on the internet¹. The web documents are looked at beforehand, and their content is broken down and stored in different slots. Once the query is available, all that needs to be done is to match what is in the query with what is in the slots.

Web crawlers are pieces of code that find and download documents from the web. They start with a set of website addresses (URLs) and look inside them for links to new web pages. Then, they download and look inside the new pages for more links. Provided the starting list was diverse enough,


crawlers end up visiting every site that allows access to them, often multiple times, looking for updates.

Step 2: The document gets transformed into multiple pieces

The document downloaded by the crawler might be a clearly structured web page with its own description of content, author, date, etc. It can also be a badly scanned image of an old library book. Search engines can usually read a hundred different document types¹. They convert these into html or xml and store them in tables (called BigTable in the case of Google).

A table is made up of smaller sections called tablets in which each row of the tablet is dedicated to one web page.

These rows are arranged in some order which is recorded along with a log for updates. Each column has specific information related to the webpage which can help with matching the document content to the contents of a future query. The columns contain:



- The website address which may, by itself, give a good description of the contents in the page, if it is a home page with representative content or a side page with associated content;
 - Titles, headings and words in bold face that describe important content;
 - Metadata of the page. This is information about the page that is not part of the main content, such as the document type (eg, email or web page), document structure and features such as document length, keywords, author names and publishing date;
 - Description of links from other pages to this page which provide succinct text regarding different aspects of the page content. More links, more descriptions and more columns used. The presence of links is also used for ranking, to determine how popular a web page is. (Take a look at [Google's Pagerank](#), a ranking system that uses links to and from a page to gauge quality and popularity).
- People's names, company or organization names, locations, addresses time and date expressions, quantities and monetary values etc. Machine learning algorithms can be trained to find these entities in any content using training data annotated by a human being¹.



The main content of a page is often hidden amongst other information. "theguardian" by Il Fatto Quotidiano is licenced under CC BY-NC-SA 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-nc-sa/2.0/?ref=openverse>.

One column of the table, perhaps the most important, contains the main content of the document. This has to be identified amidst all the external links and advertisements. One technique uses a machine-learning model to "learn" which is the main content in any webpage.

We can of course match exact words from the query to the words in a web document, like the *Find* button in any word processor. But this is not very effective, as people use different words to talk about the same object. Just recording the separate words will not help to capture how these words combine with each other to create meaning. It is ultimately the thought behind the words that help us communicate, and not the words themselves. Therefore, all search engines transform the text in a way that makes it easier to match with the meaning of the query text. Later, the query is processed similarly.

How main text is processed by MOST search engines



Words are broken down to tokens

A model **might** store :

-  as 

(frequently used words are stored as such)

-   as  + es

-  as  + ing

-  as play + ing

Now, the  knows :

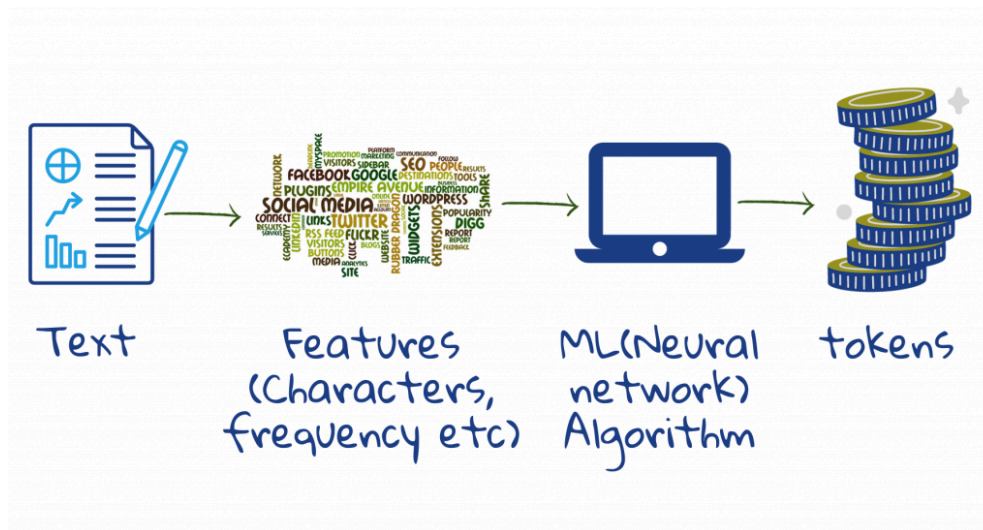
 ,  and  are related
 and  have the same endings

 they might play the same syntactic role

As parts of a word, the total number of different tokens that need to be stored is reduced. Current models store about 30,000 to 50,000 tokens². Misspelled words can be identified because parts of them still match with the stored tokens. Unknown words may turn up search results, since their parts might match with the stored tokens.

Here, the training set for machine learning is made of example texts. Starting from individual characters, space and punctuation, the model merges characters that occur frequently, to form new tokens. If the number of tokens is not high enough, it continues the merging process to cover bigger or less frequent word-parts. This way,

most of the words, word endings and all prefixes can be covered. Thus, given a new text, the machine can easily split it into tokens and send it to storage.

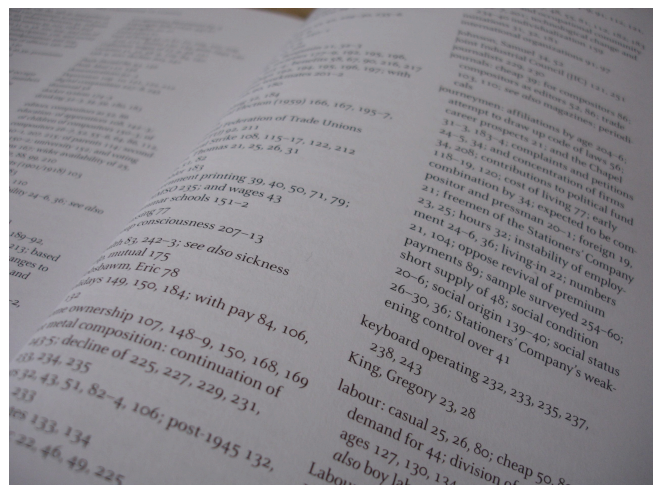


Step 3 : An index is built for easy reference

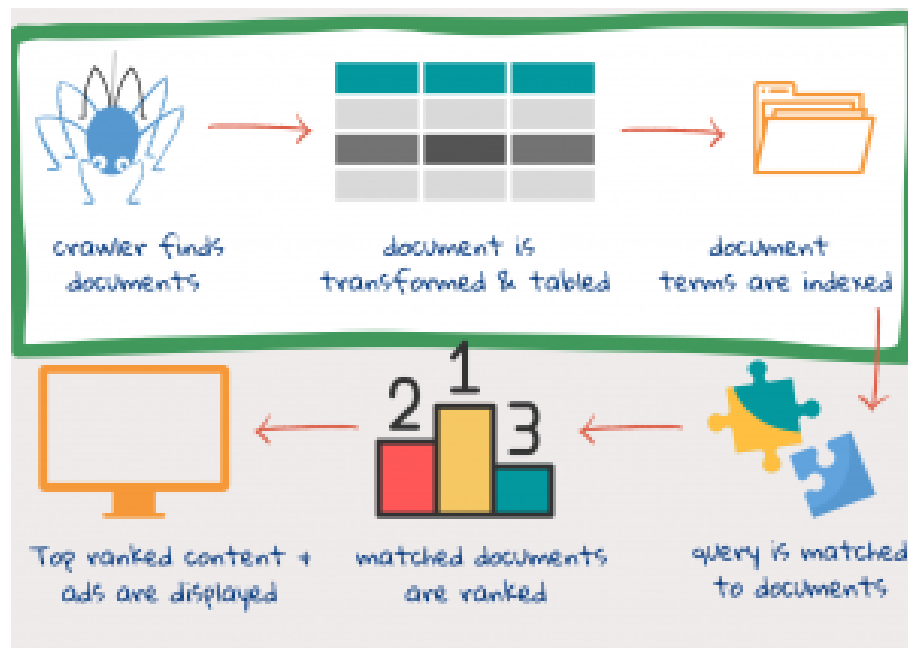
Once the data is tucked away in BigTables, an index is created. Similar in idea to textbook indexes, the search index lists tokens and their location in a web document. Statistics show how many times a token occurs in a document and how important is it for the document, etc, and information is positions thus – is the token in the title or a heading, is it concentrated in one part of the document and does one token always follow another?

Nowadays, many search engines use a language-based model generated by deep neural networks. The latter encodes semantic details of the text and is responsible for better understanding of queries³. The neural networks help the search engines to go beyond the query, in order to capture the information-need that induced the query in the first place.

These three steps give a simplified account of what is called “indexing” – finding, preparing and storing documents and creating the index. The steps involved in “ranking” come next – matching query to content and displaying the results according to relevance.



“Index” by Ben Weiner is licenced under CC BY-ND 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-nd/2.0/?ref=openverse>.



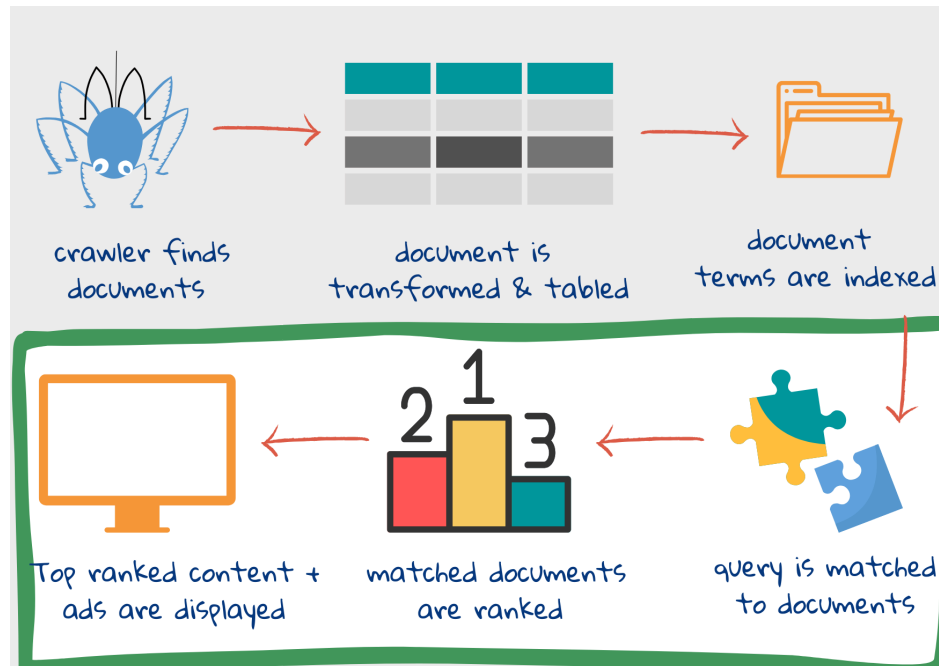
¹ Croft, B., Metzler D., Strohman, T., *Search Engines, Information Retrieval in Practice*, 2015

² Sennrich, R., Haddow, B., and Birch, A., *Neural Machine Translation of Rare Words with Subword Units*, In Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1715–1725, Berlin, Germany. Association for Computational Linguistics, 2016.

³ Metzler, D., Tay, Y., Bahri, D., Najork, M., *Rethinking Search: Making Domain Experts out of Dilettantes*, SIGIR Forum 55, 1, Article 13, June 2021.

11.

Compared to the search engines of the early 2000s, the present search engines do richer and deeper analysis. For example, as well as counting words, they can analyse and compare the meaning behind words¹. Much of this richness happens in the ranking process:



Step 4: Query terms are matched with index terms



Source: <https://ai.googleblog.com/2021/12/a-fast-wordpiece-tokenization-system.html>, A Fast WordPiece Tokenization System, By Xinying Song and Denny Zhou and "Vintage Disney Mary Poppins Plate by Sun Valley Melmac" by GranniesKitchen is licenced under CC BY 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by/2.0/>?

Once the user types the query and clicks on search, the query is processed. Tokens are created using the same process as the document text. Then the query may be expanded by adding other keywords. This is to avoid the situation whereby relevant documents are not found because the query uses words that are slightly different from those of the web-content authors. This is also done to capture differences in custom and usage. For example, the use of words such as president, prime minister and chancellor may be interchanged, depending on the country¹.

Most search engines keep track of user searches (Look at the [description of popular search engines](#) to learn more).

Queries are recorded with the user data in order to personalise content and serve advertisements. Or, the records from all users are put together to see how and where to improve search engine performance.

User logs contain items such as past queries, the results page and information on what worked. For example, what did the user click and what did they spend time reading? With user logs, each query can be matched with relevant documents (the user clicks, reads and closes session) and non-relevant documents (user did not click or did not read or tried to rephrase query)².

With these logs, each new query can be matched with a similar past query. One way of finding out if one query is similar to another, is to check if ranking turns up the same documents. Similar queries may not always contain the same words but the results are likely to be identical².

Spelling added to expand the query. This is done by looking at other words that occur frequently in relevant documents from the past. In general, however, words that occur more frequently in the relevant documents than in the non-relevant documents are added to the query or given additional weightage².









Step 5: Relevant documents are ranked

Each document is scored for relevance, and ranked according to this score. Relevance here is both topic relevance – how well the index terms of a document match that of the query, and user relevance – how well it matches the preferences of the user. A part of document scoring can be done while indexing. The speed of the

search engine depends on the quality of indexes. Its effectiveness is based on how the query is matched to the document as well as on the ranking system².



Some features
considered for ranking

- Are all query terms present in the document?
Are they found close together?
- How many times do they appear? Are they in the title or headings?
- Are there many  to this page? Are there many from this page?
- What is the parent website? Is it updated regularly? Especially for   
- Has the  visited this site before? What sites have they preferred for this topic?
- What does past user  say about similar searches and relevant pages?
- What is the location? Especially for  and searched with 

User relevance is measured by creating user models (or personality types), based on their previous search terms, sites visited, email messages, the device they are using, language and geographic location. [Cookies](#) are used to store user preferences. Some search engines buy user information from third parties as well (you could refer to [descriptions of some search engines](#)). If a person is interested in football their results for “Manchester” will be different from the person who just booked a flight to London. Words that occur frequently in the documents associated with a person will be given the highest importance.

Commercial web search engines incorporate hundreds of features in their ranking algorithms; many derived from the huge collection of user interaction data in the query logs. A ranking function combines the document, the query and user relevance features. Whatever ranking function is used, it would have a solid mathematical foundation. The output is the probability that a document satisfies the user’s information need. Above a certain probability of relevance,

the document is classified as relevant².

Machinelearning is used to learn about ranking on implicit user feedback in the logs (ie, what worked in previous queries). Machine learning has also been used to develop sophisticated models of how humans use language; this is used to decipher queries^{1,2}.



Advances in web search have been phenomenal in the past decade. However, where it refers to understanding the context for a specific query, there is no substitute for the user providing a better query. Typically, better queries come from users who examine results and reformulate the query².

Step 6: Results are displayed



"Gumshoe DuckDuckGo Results" by jrbrusseu is licenced under CC BY-SA 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-sa/2.0/?ref=openverse>.

the results are ready. The page's title and url are displayed, with query terms in bold. A short summary is generated and displayed after each link. The summary highlights important passages in the document.

In this regard, sentences are taken from headings, metadata description or from text that best corresponds with the query. If all query terms appear in the title, they will not be repeated in the snippet². Sentences are also selected based on how readable they are.

Appropriate advertising is added to the results. Search engines generate revenue through advertisements. In some search engines, they are clearly marked as sponsored content, while in others they are not. Since many users

look at only the first few results, ads can change the process substantially.

Advertisements are chosen according to the context of the query and the user model. Search engine companies maintain a database of advertisements. This database is searched to find the most relevant advertisements for a given query. Advertisers bid for keywords that describe topics associated with their product. Both the amount bid and the popularity of an advertisement are significant factors in the selection process².

For questions on facts, some engines use their own collection of facts. Google's Knowledge Vault contains over a billion facts indexed from different sources³. Results are clustered by machine learning algorithms into appropriate groups. Finally, the user is also presented with alternatives to the query to see if they are better.

Some references

The origin of Google can be found in [Brin and Page's original paper](#). Some of the maths behind Pagerank are on [Wiki's PageRank](#). For the mathematical minded, here is [a nice explanation of Pagerank](#).

- ¹ Russell, D., *What Do You Need to Know to Use a Search Engine? Why We Still Need to Teach Research Skills*, AI Magazine, 36(4), 2015.
- ² Croft, B., Metzler D., Strohman, T., *Search Engines, Information Retrieval in Practice*, 2015.
- ³ Spencer, S., *Google Power Search: The Essential Guide to Finding Anything Online With Google*, Koshkonong, Kindle Edition.

12.

While search engines provide a very useful service, they have some negative impacts on both the individual user and the society as a whole. Being aware of these impacts can help us shield ourselves and those who depend on us.

Data and Privacy

Most websites, search engines and mail clients collect information about users. Most of this data is tied to the identity of the user through IP addresses. This data is then used to serve targeted advertisements and personalised content, improve the services provided and carry out market research. However, search engines do not always disclose all the information they collect where they collect it and what they do with that information¹. Or even where they collect this information. For example, studies show that Google can track users across nearly 80% of websites².

Information that search engines can display when someone searches for a user include:

- Information that they added in some web site,
- Information added by others with their consent,
- Information that was collected in some other context and then published on the web – by forums, event organisers, friends and others.









Information collected and processed when search engines are used include the following:

Find out!

[Read about](#)
cookies and other
tracking
techniques like
fingerprinting.

A Privacy Breach happens when :



-  is collected without full informed 
- It is leaked or stolen. 
- It is used to steal your identity.
- to track your 
- to bully you or violate your rights.
- The data is used to discriminate against you. .
Some sites show different  to different users for the same item.
-  about your activities, interests and background is used in an unanticipated way.
For example, employers may look into  profiles before hiring.
- Information is sold to third parties and is used in ways you do not know about.
- Data that you thought is anonymised and cannot be tied to you is  to you.

- The searched-for topic, date and time of search^{1,3,4}.
- Activity data across apps such as email, calendar and maps, collected by search engines like Google and Microsoft^{3,4}.
- Data bought by some search engines from third parties^{3,4}.
- Data bought from search engines and websites that are put together and tied to the user by third parties².
- Inferences made from the data collected.
- Inferences drawn from personal settings. For example, “to infer that a user who has strong privacy settings may have certain psychological traits, or that they may have “something to hide””⁵.
- User profiles or models, created by search engines, based on this information. These models are based on online data and give only a limited view of the person. Decisions based on these, when used in other contexts, will not be balanced.

Collected data on a consent-giving user can be used to draw inferences about another user who did not give consent,

but who has been judged by the search engine to have a similar profile.

All this data, both raw and processed, gives rise to privacy and security concerns. Some measures can be taken by search providers, governments and users to prevent privacy breaches:

- Data can be stored in such a way that leaks and theft are discouraged. For example, user data can be stored in separate and decentralised databases⁵;
- Data is encrypted or anonymised;
- Machine learning can be used to automatically detect and classify trackers. This can then be used to improve browser privacy tools²;
- Policies and laws like GDPR legislation can introduce explicit guidelines and sanctions to regulate data collection, use and storage¹;
- User-centred recommendations are made and publicised so that users, including parents and teachers, can better protect their and their wards' privacy.



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=80#oembed-1>

In Europe, search engine companies are viewed as ‘controllers of personal data’, as opposed to mere providers of a service. Thus, they can be held responsible and liable for the content that is accessible through their services. However, privacy laws often concern confidential and intimate data. Even harmless information about people can be mined in order to create user profiles based on implicit patterns in the data. Those profiles (whether accurate or not) can be used to make decisions affecting them¹.

Also, how a law is enforced changes from country to country. According to GDPR, a person can ask a search engine company to remove a search result that concerns them. Even if the company removes it from the index in Europe, the page can still show up in results outside Europe¹.

Some Measures that can help



- Update  parameters, privacy settings, ad settings and  controls.
- Clear browser history and cookies regularly.
- Say no to tracking where applicable. However, cookies used for "Legitimate Interest" are a minimum for most web services.
- Read through consent  before clicking .
- Inform yourself about the privacy and security policy of your institution, country and .
- Look at privacy policies of search engines and browsers before choosing what to use.
- Read about tracking-protection
- Discuss privacy and safety with your



Although companies' policies can shed light on their practices, research shows that there is often a gap between policy and its use².

Reliability of content

Critics have pointed out that search engine companies are not fully open with why they show some sites and not others, and rank some pages higher than others¹.

Ranking of search results is heavily influenced by advertisers who sponsor content. Moreover, big search engine companies provide many services other than search. Content provided by them are often boosted in the search results. In Europe, Google has been formally charged with prominently displaying its own products or services in its search returns, regardless of its merits¹.

Large companies and web developers who study ranking algorithms can also influence ranking by playing on how a search engine defines popularity and authenticity of websites. Of course, the criteria judged important by search-engine programmers are themselves open to question.

This affects how reliable the search results are. It is always good to use multiple sources and multiple search engines and have a discussion about the content used in schoolwork.

Autonomy

A search engine recommends content using its ranking system. By not revealing the criteria used to select this content, it reduces user autonomy. For example, if we had known a suggested web page was sponsored, or selected based on popularity criteria we don't identify with, we might not have chosen it. By taking away informed consent, search engines and other recommender systems have controlling influences over our behaviour.

Autonomy is having control over processes, decisions and outcomes⁷. It implies liberty (independence from controlling influences) and agency (capacity for intentional action)⁷. Systems that recommend content without explanation can encroach on users' autonomy. They provide recommendations that nudge the users in a particular direction, by engaging them only with what they would like and by limiting the range of options to which they are exposed⁵.

¹ Tavani, H., Zimmer, M., [Search Engines and Ethics](#), The Stanford Encyclopedia of Philosophy, Fall 2020 Edition), Edward N. Zalta ed.

² Englehardt, S., Narayanan, A., [Online Tracking: A 1-million-site Measurement and Analysis](#), Extended version of paper at ACM CCS 2016.

³ [Google Privacy and Terms](#).

⁴ [Microsoft Privacy Statement](#).

⁵ Milano, S., Taddeo, M., Floridi, L. [Recommender systems and their ethical challenges](#), *AI & Soc* 35, 957–967, 2020.

⁶ Tavani, H.T., *Ethics and Technology: Controversies, Questions, and Strategies for Ethical Computing*, 5th edition, Hoboken, NJ: John Wiley and Sons, 2016.

⁷ Hillis, K., Petit, M., Jarrett, K., *Google and the Culture of Search*, Routledge Taylor and Francis, 2013.

13.

Social effects

More and more, there is a feeling that everything that matters is on the web and should be accessible through search¹. As LM Hinman puts it, “Esse est indicato in Google (to be is to be indexed on Google).” As he also notes, “citizens in a democracy cannot make informed decisions without access to accurate information”^{2,3}. If democracy stands on free access to undistorted information, search engines directly affect how democratic our countries are. Their role as gatekeepers of knowledge is in direct conflict with their nature as private companies dependent on ads for income. Therefore, for the sake of a free society, we must demand accountability for search engines and transparency in how their algorithms work².

Creation of filter bubbles

Systems that recommend content based on user profiles, including search engines, can insulate users from exposure to different views. By feeding content that the user likes, they create self-reinforcing biases and “filter bubbles”^{2,4}. These bubbles, created when newly acquired knowledge is based on past interests and activities⁵, cement biases as solid foundations of knowledge. This could become particularly dangerous when used with young and impressionable minds. Thus, open discussions with peers and teachers and collaborative learning activities should be promoted in the classroom.

Feedback loops

Search engines, like other recommendation systems, predict what will be of interest to the user. Then, when the user clicks on what was recommended, it (the search engine) takes it as positive feedback. This feedback affects what links are displayed in the future. If a user clicked on the first link displayed, was it because they found it relevant or simply because it was the first result and thus easier to choose?

Implicit feedback is tricky to interpret. When predictions are based on incorrect interpretation, the effects are even trickier to predict. When certain results are repeatedly shown – and are the only thing that the user gets to see – it can even end up changing what the user likes and dislikes – a self-fulfilling prediction, perhaps.

In the United States, a predictive policing system was launched whereby high-crime areas of a certain city were highlighted. This meant that more police officers were deployed to such areas. Since these officers knew the area was at high risk, they were careful, and stopped, searched, or arrested more people than normal. The

arrests thus validated the prediction, even where the prediction was biased in the first place. Not only that, the arrests were data for future predictions on the same areas and on areas similar to it, compounding biases over time⁵.

We use prediction systems in order to act on the predictions. But acting on biased predictions affects future outcomes, the people involved – and ultimately society itself. “As a side-effect of fulfilling its purpose of retrieving relevant information, a search engine will necessarily change the very thing that it aims to measure, sort and rank. Similarly, most machine-learning systems will affect the phenomena that they predict”⁵.

Fake news, extreme content and censorship

There is increasing prevalence of fake news (false stories that appear as news) in online forums, social media sites and blogs, all available to students through search. Small focused groups of people can drive ratings up for specific videos and web sites of extreme content. This increases the content’s popularity and appearance of authenticity, gaming the ranking algorithms⁴. Yet, as of now, no clear and explicit policy has been adopted by search-engine companies to control fake news².

On the other hand, search engines systematically exclude certain sites and certain types of sites in favour of others⁶. They censor content from some authors, despite not being asked to do so by the public. Search engines, therefore, should be used with awareness, discretion and discrimination.

¹ Hillis, K., Petit, M., Jarrett, K., *Google and the Culture of Search*, Routledge Taylor and Francis, 2013.

² Tavani, H., Zimmer, M., [Search Engines and Ethics](#), The Stanford Encyclopedia of Philosophy, Fall 2020 Edition), Edward N. Zalta (ed.).

³ Hinman, L. M., *Esse Est Indicatio in Google: Ethical and Political Issues in Search Engines*, International Review of Information Ethics, 3: 19–25, 2005.

⁴ Milano, S., Taddeo, M., Floridi, L. [Recommender systems and their ethical challenges](#), AI & Soc 35, 957–967, 2020.

⁵ Barocas, S., Hardt, M., Narayanan, A., [Fairness and machine learning Limitations and Opportunities](#), MIT Press, 2023.

⁶ Introna, L. and Nissenbaum, H., *Shaping the Web: Why The Politics of Search Engines Matters*, The Information Society, 16(3): 169–185, 2000.

PART III

MANAGING LEARNING

In the midst of all the preparation and evaluation that you have to do, you are probably juggling administrative tasks too.

Have you ever felt there is not enough time, either in the classroom or out of it, to help your students more efficiently?

A class period can go too fast and there are often more students than you can keep track of.

Did you miss an expression, fail to explain away a point of confusion?

Is there a way to track the progress and struggles of an individual student?

Apart from asking colleagues, how can you have a better overview of what is happening in other subjects, in order to reinforce learning?

E-learning and Learning Management Systems (LMS)

The number of people making use of e-learning is constantly growing. The term e-learning refers to learning mediated by the use of technology in contexts where educators and learners are distant in space and/or time. The ultimate aim of e-learning is to improve students' learning experience and practice.

Today, with the advancement of technology, it is more appropriate to refer to systems and platforms for the 'delivery' of e-learning rather than to single-purpose tools. Such systems are the result of integrating different software tools capable of building an ecosystem where flexible and adaptable learning paths can be exploited. An e-learning system enables the management of learning processes and the management of courses. It enables student-learning assessments, constructing reports, and creating and organising content. It facilitates communication between teachers/tutors and students. Among the most widely used e-learning systems, there are Learning Management Systems (LMS) (eg Moodle, Edmodo).

The acronym LMS refers to a web-based application designed to manage the learning process of trainees¹ at different levels, in different ways and domains. An LMS, therefore, could be defined as a learning environment within which learning, content and assessment activities and tools are implemented. Student-student and/or student-educator interactions likewise are implemented and managed within this environment. Furthermore, the definition of LMSs includes their being platforms that generally can include whole course management systems, content management systems, and portals².

LMS and AI: the Smart LMS

With the advent of AI, Education, in general, and LMSs, in particular, become possible and promising fields of application of this revolutionary force³. Specifically, LMSs, thanks to the functionalities supported by AI, represent a renewed learning tool capable of satisfying two of the fundamental traits of the education of the future: personalisation and adaptation⁴. It is from this combination of LMS and AI that the Smart LMS (SLMS) or Intelligent LMS emerges.

An efficient SLMS features algorithms that can provide and retrieve information from three fundamental clusters of knowledge: a) the learner b) the pedagogy and c) the domain. By acquiring information about (a) learners' preferences, their emotional and cognitive states and their achievements and goals, an SLMS can implement those teaching strategies (b) that are most effective (specific types of assessment, collaborative learning, etc.) for learning to be most fruitful within the

specific domain of knowledge being studied (c): eg geometry theorems, mathematical operations, laws of physics, text analysis procedures⁴.

An SLMS, therefore, can be defined as a learning system capable of **adapting** the contents proposed to the learner by calibrating them to the knowledge and skills the learner has displayed in previous tasks. In fact, by adopting a learner-centred approach, it can identify, follow and monitor learners' paths by recording their learning patterns and styles. Referring to the description given by Fardinpour et al.⁵, an intelligent LMS provides the learner with the most effective learning path and the most appropriate learning content, through **automation**, the **adaptation** of different teaching strategies (scaffolding), and the **reporting** and **knowledge generation**. It also provides the learners with the possibility to keep track of and monitor their learning and **learning goals**. Furthermore, although these features and tools enable the LMS to operate more intelligently, an SLMS must provide learners with the possibility to disable the AI that manages their path, in order to have full access to all learning materials in the learning environment.

Some examples of AI-supported functionalities in the context of an SLMS

When an SLMS is functioning correctly, several AI-supported tools make it possible to realise a system with the features described above. Such AI-supported tools move transversally along the three aforementioned clusters of knowledge, to which the SLMS algorithms constantly refer (learner, pedagogy, domain).

AI-supported chatbots as virtual tutors

A chatbot is software that simulates and processes human conversations (written or spoken). In the context of an SLMS, it can function as a virtual tutor, capable of answering a learner's questions concerning, for example, learning courses. The chatbot is also capable of providing suggestions to the learner, based on the analysis that the system makes of the learner's previous performances and interactions⁶.

Learning Analytics

Learning Analytics – data relating to the details of individual learner interactions in online learning activities – allow teachers to monitor learner progress and performance in depth. Thanks to them, the system can implement automatic computer-assisted educational task activation⁷ to supplement the activities of learners who have shown performance deficits in specific tasks. In addition, it can

automatically provide suggestions to teaching staff regarding the difficulty of proposed tasks or the need to supplement them with additional learning content.

Benefits for learners and teachers

These and other AI-supported tools⁴ contribute to making an SLMS a powerful learning and teaching tool that, instead of being perceived as a substitute for the teacher's work, shows itself as a tool capable of "augmenting" the human aspects of teaching⁸ and bringing a series of fundamental benefits to the whole learning/teaching process.

Since an SLMS calibrates the contents to the student's skills and level, it avoids the learner, in the different phases of his or her path, facing tasks that bore him or her because they are too simple, or that frustrate him or her because they are too complex. This ensures that the student's motivation and attention are always at a high level and appropriate to the level of difficulty of the task to be addressed. This situation has the direct consequence of significantly reducing the dropout rate, as it allows teachers to detect any problems in time and intervene promptly, as soon as the student shows the first signs of difficulty.

Such a situation, as well as linear learning situations (without difficulties), can be addressed by proposing to the students, through the SLMS tools, different knowledge contents that are already stored in the course databases or are from third-party providers. This results in a direct benefit for the teacher, who does not have to create new teaching materials from time to time, and can use the saved time in other essential occupations such as refining their teaching methods and/or interacting directly with the students.

¹ Kasim, N. N. M., and Khalid, F., *Choosing the right learning management system (LMS) for the higher education institution context: A systematic review*, International Journal of Emerging Technologies in Learning, 11(6), 2016.

² Coates, H., James, R., & Baldwin, G., *A critical examination of the effects of learning management systems on university teaching and learning*, Tertiary education and management, 11(1), 19-36, 2005.

³ Beck, J., Sternm, M., & Haugsjaa, E., *Applications of AI in Education*, Crossroads, 3(1), 11-15. doi:10.1145/332148.332153, 1996.

⁴ Rerhaye, L., Altun, D., Krauss, C., & Müller, C., *Evaluation Methods for an AI-Supported Learning Management System: Quantifying and Qualifying Added Values for Teaching and Learning*, International Conference on Human-Computer Interaction (pp. 394-411). Springer, Cham, July 2021.

⁵ Fardinpour, A., Pedram, M. M., & Burkle, M., *Intelligent learning management systems: Definition, features and measurement of intelligence*, International Journal of Distance Education Technologies (IJDET), 12(4), 19-31, 2014.

⁶ [HR Technologist: Emerging Trends for AI in Learning Management Systems](#), 2019, Accessed 31 Oct 2022.

- ⁷ Krauss, C., Salzmann, A., & Merceron, A., *Branched Learning Paths for the Recommendation of Personalized Sequences of Course Items*, DeLFI Workshops, September 2018.
- ⁸ Mavrikis, M., & Holmes, W., *Intelligent learning environments: Design, usage and analytics for future schools*, Shaping future schools with digital technology, 57-73, 2019.

15.

AZIM ROUSSANALY; ANNE BOYER; AND JIAJUN PAN

What are learning analytics?

More and more organizations are using data analysis to solve problems and improve decisions related to their activities. And the world of education is not an exception because, with the generalization of virtual learning environment (VLE) and learning management systems (LMS), massive learning data are now available, generated by the learners' interaction with these tools.

Learning Analytics (LA) is a disciplinary field defined as “the measurement, collection, analysis and reporting of data about learners and their contexts, for purposes of understanding and optimising learning and the environments in which it occurs”⁴.

Four types of analytics are generally distinguished according to the question to solve:

- Descriptive Analytics: What happened in the past?
- Diagnostic Analytics: Why something happened in the past?
- Predictive Analytics: What is most likely to happen in the future?
- Prescriptive Analytics: What actions should be taken to affect those outcomes?

What is it?

The visualisation to recommendation systems. Research in this area is currently active. We will limit ourselves to summarising the frequent issues encountered in the literature. Each issue leads to families of tools targeting mainly learners or teachers who represent most of the end users of LA-based applications.

Predict and enhance students learning outcome

One of the emblematic applications of LA is the prediction of failures.

Learning indicators are automatically computed from the digital traces and can be accessed directly by learners so that they can adapt their own learning strategies.

One of the first experiments was conducted at Purdue University (USA) with a mobile application designed as a traffic-light-based dashboard¹.

Each student can monitor his or her own progress indicators.

A screenshot of the dashboard is shown in figure 1.

Indicators can also be addressed to teachers, as in an early warning system (EWS).

This is the choice made by the French national centre for Distance Education (CNED) in an ongoing study².

The objective of an EWS is to alert, as soon as possible, those tutors responsible for monitoring the students, in order to implement appropriate remedial actions.



Figure 1: Purdue University dashboard for students

Analyze student learning process

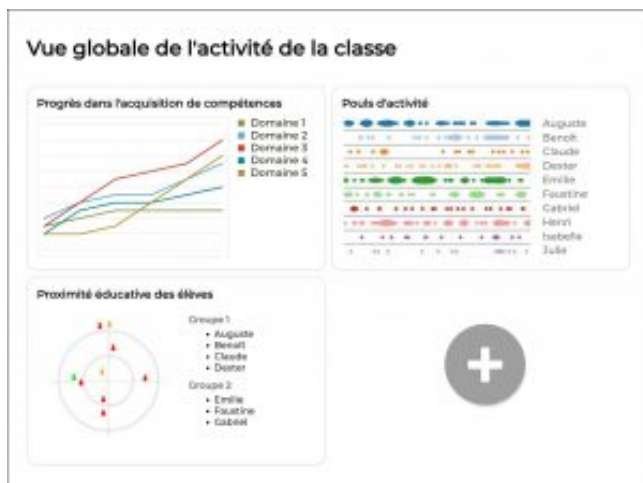


Figure 2: METAL project dashboard

LA techniques can help to model the learning behaviour of a learner or a group of learners (ie a class). The model can be used to display learning processes in LA applications, providing additional information that will enable teachers to detect deficiencies which will help improve training materials and methods. Moreover, the analysis of learning process is a way for observing the learner engagement. For example, in the e-FRAN METAL project, the indicators were brought together in a dashboard co-designed with a team of secondary-school teachers as shown in

figure 2³.

Personalise learning paths

The personalisation of learning paths can occur in recommendation or adaptive learning systems. Recommendation systems aim to suggest, to each learner, best resources or appropriate behaviour that may help to achieve educational objectives.

Some systems focus on putting the teacher in the loop by first presenting proposed recommendations for their validation. Adaptive learning systems allow the learner to develop skills and knowledge in a more personalised and self-paced way by constantly adjusting the learning path towards the learner experience.

Does it work?

In publications, feedback focuses mainly on students, including those in higher education. Observations generally show improved performance of learners (for example, +10% of grades A and B at Purdue University, US). For teachers, the impact of LA is more complex to assess. Studies based on the technology acceptance model (TAM) suggest that teachers have a positive perception of the use of LA tools. A study shows the final Strengths, Weaknesses, Opportunities and Threats analysis (SWOT) that we reproduce here⁵ (see figure 3).

Strengths <ul style="list-style-type: none">- enhance diagnostic, formative or summative formal and standardized learner assessment (4)- help decision-making of educational policymakers (2)- identify at-risk students (2)	Weaknesses <ul style="list-style-type: none">- training is needed (3)- cannot capture nuanced aspects of learning (3)- can be time-consuming (2)-can inhibit teachers' creativity (2)
Opportunities <ul style="list-style-type: none">-teachers will be able to select the most relevant or useful for them views of LA (2)- can be helpful as a support mechanism (2)	Threats <ul style="list-style-type: none">-careful with privacy issues (4)-non-reliable creation mechanisms (3)

Figure 3: SWOT analysis of LA acceptance⁵

Some points of attention, included in the threats and weakness sections, form the basis of the reflection of the Society for Learning Analytics Research (SoLAR) community to recommend an ethics-by-design approach for LA applications (Drashler-16). The recommendations are summarised in a checklist of eight key words: Determine, Explain, Legitimate, Involve, Consent, Anonymise, Technical, External (DELICATE).

¹ Arnold, K. and Pistilli, M., *Course signals at Purdue: Using learning analytics to increase student success*, LAK2012, ACM International Conference Proceeding Series, 2012.

² Ben Soussia, A., Roussanaly, A., Boyer, A., *Toward An Early Risk Alert In A Distance Learning Context*, ICALT, 2022.

- ³ Brun, A., Bonnin, G., Castagnos, S., Roussanaly, A., Boyer, A., *Learning Analytics Made in France: The METALproject*, IJILT, 2019.
- ⁴ Long, P., and Siemens, G., 1st International Conference on Learning Analytics and Knowledge, Banff, Alberta, February 27–March 1, 2011.
- ⁵ Mavroudi, A., *Teachers' Views Regarding Learning Analytics Usage Based on the Technology Acceptance Model*, TechTrends. 65, 2021.

16.





Decisions in the classroom

As a teacher, you have access to many kinds of data. Either tangible data such as attendance and performance records, or intangible ones such as student body-language. Consider some of the decisions you take in your professional life: *What are the data that help you make these decisions?*

There are technological applications that can help you visualise or process data. Artificial intelligence systems use data to personalise learning, make predictions and decisions that might help you teach and manage the classroom: *Do you have needs that technology can answer? If yes, what will be the data such a system might require to carry out the task?*

What explains the of data based systems?



-  in  power
-  in the  of storing data
- Powerful ways to analyse and model data, thanks to AI
- An explosion of available data due to digitalisation, cheap sensors, growth of internet, Big data and



Reference: Kelleher, J.D, Tierney, B, *Data Science*, London, 2018 and
Kitchin, R, *Big Data, new epistemologies and paradigm shifts*, *Big Data & Society*, 2014

questions⁴. Their decisions and predictions, and how these affect student learning, are all data too. Thus, knowing how programmers, the machine and the user handle data is an important part of understanding how artificial intelligence works.

Educational systems have always generated data – students' personal data, academic records, attendance data and more. With digitalisation and AIED applications, more data is recorded and stored: mouse clicks, opened pages, timestamps and keyboard strokes¹. With data-centric thinking becoming the norm in society, it is natural to ask how to crunch all this data to do something pertinent. Could we give more personalised feedback to the learner? Could we design better visualisation and notification tools for the teacher?²

Whatever technology is used, it has to meet a real requirement in the classroom. After the need is identified, we can look at the data available and ask what is relevant to a desired outcome. This involves uncovering factors that let educators make nuanced decisions. Can these factors be captured using available data? Is data and data-based systems the best way of addressing the need? What could be the unintended consequences of using data this way³?

Machine learning lets us defer many of these questions to the data itself⁴. ML applications are trained on data. They work by operating on data. They find patterns and make generalisations and store these as models – data that can be used to answer future

About data

Data is generally about a real world entity – a person, an object, or an event. Each entity can be described by a number of **attributes (features or variables)**⁵. For example, name, age and class are some attributes of a student. The set of these attributes is the data we have on the student, which, while not in any way close to the real entity, does tell us something about them. Data collected, used and processed in the educational system is called **educational data**¹.

A **dataset** is the data on a collection of entities arranged in rows and columns. The attendance record of a class is a dataset. In this case, each row is the record of one student. The columns could be their presence or absence during a particular day or session. Thus each column is an attribute.



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=95#oembed-1>


Data is created by choosing attributes and measuring them: every piece of data is the result of human decisions and choices. Thus, data creation is a subjective, partial and sometimes messy process prone to technical difficulties^{4,5}. Further, what we choose to measure, and what we don't can have a big influence on expected outcomes.

Data traces are records of student activity such as mouse clicks, data on opened pages, the timing of interactions or key presses in a digital system¹.

Metadata is data that describe other data⁵.

Derived data is data calculated or inferred from other data: individual scores of each student is data. The class average is derived data. Often, derived data is more useful in getting useful insights, finding patterns and making predictions. Machine Learning applications can create derived data and link it with metadata data traces to create detailed [learner models](#), which help in personalising learning¹.

For any data based application to be successful, attributes should be carefully chosen and correctly measured. The patterns discovered in them should be checked to see if they make sense in the educational context. When designed and maintained correctly, data driven systems can be very valuable.



The DIKW Pyramid. Reference: Kelleher, J.D, Tierney, B, Data Science, London, 2018 and Kitchin, R, Big Data, new epistemologies and paradigm shifts, Big Data & Society, 2014



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=95#oembed-2>

[Check if you are \(Big\) Data literate](#)

This chapter aims to introduce a few basics of data and data based technology but data literacy is a very important skill to possess and merits dedicated training and continuing support and update¹.

Legislation you should know about

Because of the drastic drop in costs of data storage, more data and metadata are saved and retained for a longer time⁶. This can lead to privacy breaches and rights violations. Laws like the **General Data Protection Regulation (GDPR)** discourages such practices and gives EU citizens more control over their personal data. They give legally enforceable data protection regulations across all EU member states.

Principles GDPR is based on :



- **Lawfulness, fairness and transparency** : Data processing has to be lawful, fair, and transparent to the data subject
- **Purpose limitation** : When doing something with data all the purposes have to be declared beforehand.
- **Data minimization** : Only data absolutely necessary to fulfill the specified purpose can be collected.
- **Accuracy** : Personal data needs to be accurate.
- **Storage limitation** : Data can only be stored as long as it is needed for the specified purpose.
- **Integrity and confidentiality** : When data is processed, this has to be secure and the processor has to ensure that information doesn't get into the wrong hands.
- **Accountability** : The data controller must be able to show GDPR compliance for each step of the data processing.

Source: GDPR for dummies: What is it? Why do we need it? Why does personal data have to be protected?

Reference: "GDPR & ePrivacy Regulations" by dennis_convert licenced under CC BY 2.0. To view a copy of this licence, visit

<https://creativecommons.org/licenses/by/2.0/?ref=openverse>

According to GDPR, personal data is any information relating to an identified or identifiable person (**data subject**). Schools, in addition to engaging with companies that handle their data, store huge amounts of personal information about students, parents, staff, management, and suppliers. As data controllers, they are required to store data which they process confidentially and securely and have procedures in place for the protection and proper use of all personal data¹.

Rights established by the GDPR include:

- The **Right to Access** makes it mandatory for them to know(easily) what data is being collected about them
- The citizen's **Right to Be Informed** of the usage made of their data
- The **Right to Erasure** allows a citizen whose data has been collected by a platform to ask for that data to be removed from the dataset built by the platform (and which may be sold to others)
- The **Right to explanation** – explanation should be provided whenever clarification is needed on an automated decision process that affect them

Although, GDPR does allow for collection of some data under "legitimate interest"⁷ and the use of derived, aggregated, or anonymized data indefinitely and without consent⁵. The new **Digital Services Act** restricts the use of personal data for targeted advertising purposes⁷. In addition to these, the **EU-US Privacy Shield** strengthens the data-protection rights for EU citizens in the context where their data have been moved outside of the EU⁵.

Please refer to [GDPR for dummies](#) for the analysis carried out by independent

experts from the Civil Liberties Union for Europe (Liberties). This is a watchdog that safeguards EU citizens' human rights.

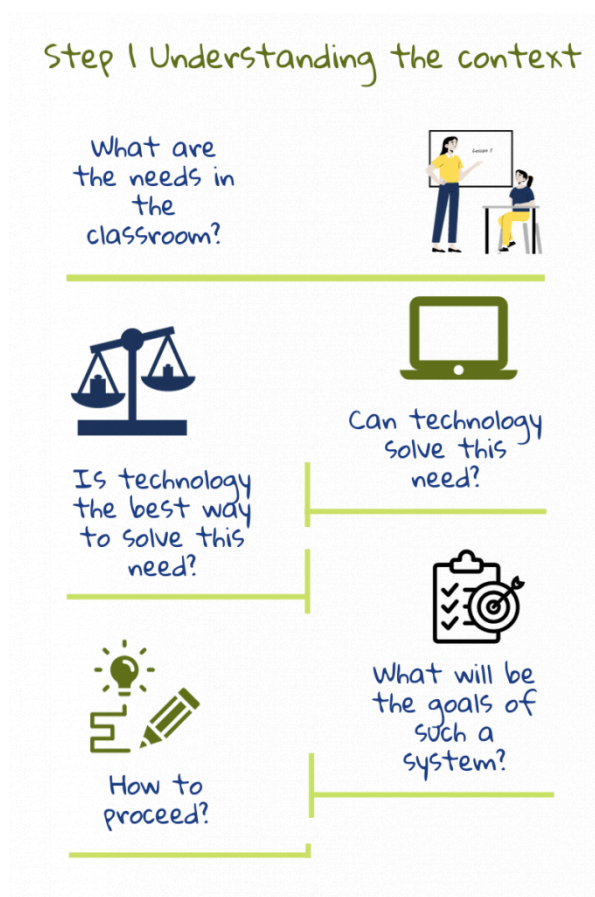
- ¹ [*Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators*](#), European Commission, October 2022.
 - ² du Boulay, B., Poulavasillis, A., Holmes, W., Mavrikis, M., *Artificial Intelligence And Big Data Technologies To Close The Achievement Gap*, in Luckin, R., ed. *Enhancing Learning and Teaching with Technology*, London: UCL Institute of Education Press, pp. 256–285, 2018.
 - ³ Hutchinson, B., Smart, A., Hanna, A., Denton, E., Greer, C., Kjartansson, O., Barnes, P., Mitchell, M., [*Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure*](#), Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, Association for Computing Machinery, New York, 2021.
 - ⁴ Barocas, S., Hardt, M., Narayanan, A., [*Fairness and machine learning Limitations and Opportunities*](#), MIT Press, 2023.
 - ⁵ Kelleher, J.D, Tierney, B, *Data Science*, MIT Press, London, 2018.
 - ⁶ Schneier, B., *Data and Goliath: The Hidden Battles to Capture Your Data and Control Your World*, W. W. Norton & Company, 2015.
 - ⁷ Kant, T., *Identity, Advertising, and Algorithmic Targeting: Or How (Not) to Target Your “Ideal User.”*, MIT Case Studies in Social and Ethical Responsibilities of Computing, 2021.
-

17.

The design and implementation of a data-centred project could be broken down to six steps. There is a lot of back and forth between the steps and the whole process may need to be repeated multiple times to get it right.

To be effective in the classroom, multi-disciplinary teams with teachers, pedagogical experts and computer scientists should be involved in each step of the process¹. Human experts are needed to identify the need and design the process, to design and prepare the data, select ML algorithms, to critically interpret the results and to plan how to use the application².

1) Understanding the educational context



The first step in the design of an AIED tool is to understand the needs in the classroom. Once the goals are set, it has to be seen how they can be achieved: what factors to consider and what to ignore. Any data based solution is biased towards phenomena that can be easily calculated and standardised³. Thus, each decision has to be discussed by teachers who will use the tool, by experts in pedagogy who can ensure that all decisions are grounded in proven theory, and by computer scientists who understand how the algorithms work.

There is a lot of back and forth between the two first steps since what is possible will also depend on what data is available². Moreover, the design of educational tools is also subject to laws which place restrictions on data usage and types of algorithms that can be used.

2) Understanding the data

Once the goals and contributing factors are identified, the focus shifts to what data is required, how it will be sourced and labelled, how privacy will be handled and how data quality will be measured³. For a machine learning application to be successful, the datasets have to be large enough, diverse and well-labelled.

Machine learning requires data to train the model and data to work on or predict with. Some ML tasks like face recognition and object recognition already have a lot of private and public databases available for training.

If not already available in a usable learning requires data to train the model and data to work on or use for predictions. Some ML tasks, such as face and object recognition, already have a lot of private and public databases available for training.

If not already available in a usable form, existing datasets may have to be added to or relabelled to fit the needs of the project. If not, dedicated datasets may have to be created and labelled from scratch. Digital traces generated by the student while using an application could also be used as one of the data sources.

In any case, data and features relevant to the problem have to be carefully identified². Irrelevant or redundant features can push an algorithm into finding false patterns and affect the performance of the system². Since the machine can find patterns only in the data given to it, choosing the dataset also implicitly defines what the problem is⁴. If there is a lot of data available, a subset has to be selected with the help of statistical techniques and the data verified to avoid errors and biases.

As an example of a bad training data, in a story from early days of computer vision, a model was trained to discriminate between images of Russian and American tanks. Its high accuracy was later found to be due to the fact the the Russian tanks had been photographed on a cloudy day and the American ones on a sunny day⁴.


Therefore, the chosen dataset has to be verified for quality, taking into account why it was created, what does it contain, what are the processes used for collecting, cleaning and labelling, distribution and maintenance⁴. The key questions to ask include *Are the datasets fit for their intended purposes* and *Do the datasets contain hidden hazards that can make models biased or discriminatory?*³



The publically available MNIST dataset contains images of handwritten digits. Though popular for decades, it is considered too easy for today's research tasks.

Adapted from MnistExamples by Josef Steppan. Licenced under CC BY-SA 4.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-sa/4.0/?ref=openverse>

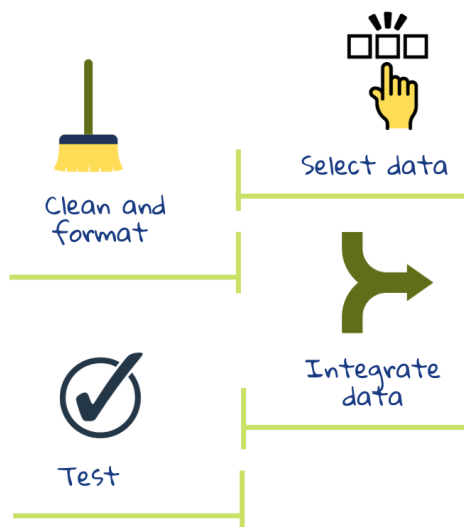
Step 2 Understanding the data



3) Preparing the data


Data-preparation involves creating data sets by merging data available from different places, adjusting for inconsistencies (For example, some test scores could be on a scale of 1 to 10 while others are given as a percentage) and looking for missing or extreme values. Then, automated tests could be performed to check the quality of datasets. This includes checking for privacy leaks and unforeseen correlations or stereotypes². The datasets might also be split into training and test datasets at this stage. The former is used for training the model and the latter for checking its performance. Testing on the training dataset would be like giving out the exam paper the day before for homework: the student's exam performance will not indicate their understanding².

Step 3 Preparing the data



4) Modelling

Step 4 Modeling



In this step, algorithms are used to [extract patterns in the data](#) and create [models](#). Usually, different algorithms are tested to see what works best. These models can then be put into use to make predictions on new data.

In most projects, initial models uncover problems in the data calling for back and forth between steps 2 and 3². As long as there is a strong correlation between the features of the data and the output value, it is likely that a Machine Learning algorithm will generate good predictions.

These algorithms use advanced statistical and computing techniques to process data. The programmers have to adjust the settings and try different algorithms to get the best results. Let us

take an application that detects cheating. A false positive is when a student who did not cheat is flagged. A false negative is when a student who cheats is not flagged. System designers can tune the model to minimize either false positives, where some cheating behaviour could be missed, or false negatives, where even doubtful cases are flagged⁵. The tuning thus depends on what we want the system to do.

5) Evaluation

During the modelling stage, each model can be tuned for accuracy of prediction on the training dataset. Then models are tested on the test dataset and one model is chosen for use. This model is also evaluated on how it meets the educational needs: Are the objectives as set out in step 1 met? Are there any unforeseen problems? Is the quality good? Could anything be improved or done in another way? Is a redesign needed? The main objective is to decide whether the application can be deployed in schools. If not, the whole process is repeated².

Step 5 Evaluation



6) Deployment

Step 6 Deployment



teaching and assessment⁶.

The final step of this process is to see how to integrate the data based application with the school system to get maximum benefits, both with respect to the technical infrastructure and teaching practices.

Though given as the final step, the whole process is iterative. After deployment, the model should be regularly reviewed to check if it is still relevant to the context. The needs, processes or modes of data capture could change affecting the output of the system. Thus, the application should be reviewed and updated when necessary. The system should be monitored continuously for its impact on learning,

"Predicting the consequences and the impact of the use of data and AI in education can be very difficult. Therefore, an incremental approach to the development and deployment of these technologies and their assessment is needed. The idea is to gradually introduce these tools into their contexts and to constantly monitor the societal effects that can emerge, leaving open the possibility to step back when unintended consequences occur."

ETHICAL GUIDELINES ON THE USE OF ARTIFICIAL INTELLIGENCE
AND DATA IN TEACHING AND LEARNING FOR EDUCATORS,
EUROPEAN COMMISSION, OCTOBER 2022

The Ethical guidelines on the use of AI and data for educators stresses that the school should be in contact with the AI service provider throughout the lifecycle of the AI system, even prior to deployment. It should ask for clear technical documentation and seek clarification on unclear points. An agreement should be made for support and maintenance, and it should be made sure that the supplier adhered to all legal obligations⁶.

Note: Both the steps listed here and the illustration are adapted from the CRISP-DM Data science stages and tasks (based on figure 3 in Chapman, Clinton, Kerber, et al. 1999) as described in ².

¹ Du Boulay, B., Poulivasillis, A., Holmes, W., Mavrikis, M., *Artificial Intelligence And Big Data Technologies To Close The Achievement Gap*, in Luckin, R., ed. Enhancing Learning and Teaching with Technology, London: UCL Institute of Education Press, pp. 256–285, 2018.

² Kelleher, J.D., Tierney, B., *Data Science*, London, 2018.

³ Hutchinson, B., Smart, A., Hanna, A., Denton, E., Greer, C., Kjartansson, O., Barnes, P., Mitchell, M., [Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure](#), Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, Association for Computing Machinery, New York, 2021.

⁴ Barocas, S., Hardt, M., Narayanan, A., [Fairness and machine learning Limitations and Opportunities](#), MIT Press, 2023.

⁵ Schneier, B., *Data and Goliath: The Hidden Battles to Capture Your Data and Control Your World*, W. W. Norton & Company, 2015.

⁶ [Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators](#), European Commission, October 2022.

18.

Data about us is constantly recorded through our phones and computers. This data is interpreted based on who is recording it and who is looking at it. To take just one example, Google makes its digital version of us, *the digital identity*, based on what we do on its platforms. It proceeds to label us based on this data and then rearranges what we see on its search engines and apps accordingly. It markets us to companies who might want to market themselves to us.

Activity

Access your “Ad Settings” Profile on Google, Facebook or Instagram. Or if you regularly use another platform, try finding out if they have ad settings and if you can access them. These are part of our digital identity.

Questions for discussion:

- What does your “digital identity” look like? Does it reflect your demographics and interests? Do you agree with this identity?
- How do you think Google decided each of these interests? What data could have been taken into account? These interest categories change frequently and are recursive: an ad interest you are associated with can determine what ad interest you’ll be categorized with next. What can this tell us about profiling?
- Do you agree with scholars such as Cheney-Lippold and Bassett that there is an over-reduction of identity here? Why is this an ethical concern?
- Ethically, does it matter more if these profiles get your interests “right” or “wrong”?
- Does your gender and race play a part in how you are labelled? How does that make you feel?

This activity has been adapted from [Identity, Advertising, and Algorithmic Targeting: Or How \(Not\) to Target Your “Ideal User.”](#), licenced under [CC BY NC 4.0](#).¹

The labels Google gives us – male, female, young or aged, have nothing to do with our identities, needs or values. Someone can be male if they look at certain websites(say hardware stores) and buy certain items². Tomorrow, a male can become female if their activity or the activities of a million other humans who contributed to what is male-like-behaviour change. Different companies give us completely different identities based on what interests them.

The same is done to our students when they interact with Personalized learning software and are subjected to learning analytics. Their digital identity, their performance, engagement and satisfaction, as viewed by these systems, is then used to evaluate not only their performance but also the performance of their peers, teachers, schools and the educational system itself³.

Why is this a problem?

1. These profiles are often composed based on noisy and incorrect data from various sources and can be very misleading⁴.
2. These digital identities can change how a student sees themselves and others, how teachers see each student, how the system sees each teacher, how the society sees education and pedagogy, and how everyone reacts to decisions and feedback³.
3. Yet, these judgements on who someone is are made without their knowledge and consent – by black boxes no one has access to. Often, there is no control over what data is recorded, where and when it is recorded and how decisions are made based on it^{4,1}. Students and teachers lose their expressive power and [human agency](#).
4. This data and judgements tend to persist as stored data long after the recorded event took place⁴.
5. The stress on metrics, where students, teachers and staff are constantly assessed, compared and ranked, can induce reactions such as anxiety and competition instead of motivation and growth³.
6. Aspects of education that can be automatically captured and analysed are given more importance; they push us towards outcomes and practices that are different from what might otherwise be of concern to us.
7. The organisations that do the “datafication” have the power to define “what ‘counts’ as quality education – a good student or an effective teacher³.”

Here are countermeasures the experts suggest teachers take:

1. **Consider the people, their identity, integrity, and dignity:** “Approach people with respect of their intrinsic value and not as a data object or a means-to-an-end”⁵. People are not just the data; that the label that a software might give students to personalise learning pathways or to split them into groups is not their real identity⁵.
2. **Be data literate:** Learn how to handle data correctly. Learn what different data based systems do, how they do it, what is their recommended usage and how to interpret the information they generate and the decisions they make.
3. **Maintain a critical distance from AIED companies and software.** Question their claims, ask proof of their validity and reliability, verify that the system follows ethical guidelines of your institution and country³.
4. **Monitor the effects these systems have on you, your students, their learning and the classroom atmosphere.**
5. **Call for open systems that give you control and the power to override automated decisions.** Pitch in, clarify or override wherever and whenever you feel the need.

¹ Kant, T., [Identity, Advertising, and Algorithmic Targeting: Or How \(Not\) to Target Your “Ideal User.”](#) MIT Case Studies in Social and Ethical Responsibilities of Computing, 2021.

² Cheney-Lippold, J., *We Are Data: Algorithms and the Making of Our Digital Selves*, NYU Press, 2017.

- ³ Williamson, B., Bayne, S., Shay, S., *The datafication of teaching in Higher Education: critical issues and perspectives*, *Teaching in Higher Education*, 25:4, 351-365, 2020.
- ⁴ Kelleher, J.D, Tierney, B, *Data Science*, MIT Press, London, 2018.
- ⁵ [Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators](#), European Commission, October 2022.

19.

Bias is prejudice towards or against an identity, whether good or bad, intentional or unintentional¹. Fairness is the counter to this bias and more: when everyone is treated fairly, regardless of their identity and situation. Clear processes have to be set and followed to make sure everyone is treated equitably and has equal access to opportunity¹.

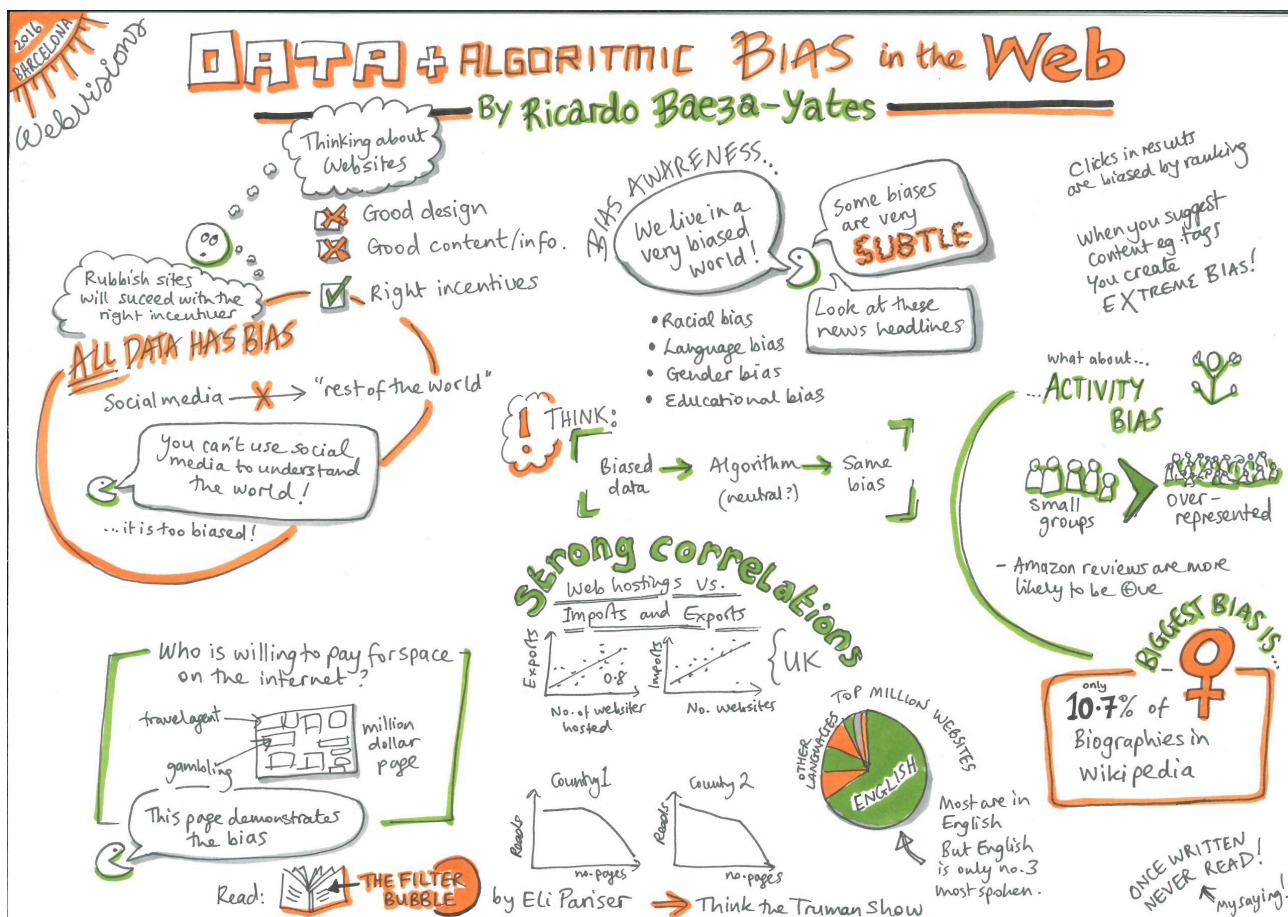
Human-based systems often comprise a lot of bias and discrimination. Every person has their own unique set of opinions and prejudices. They too are black boxes whose decisions, such as how they score answer sheets, can be difficult to understand. But we have developed strategies and established structures to watch out for and question such practices.

Automated systems are sometimes touted as the panacea to human subjectivity: Algorithms are based on numbers, so how can they have biases? Algorithms based on flawed data, among other things, can not only pick up and learn existing biases pertaining to gender, race, culture or disability – they can amplify existing biases^{1,2,3}. And even if they are not locked behind proprietary walls, they can't be called up to explain their actions due to the inherent lack of explainability in some systems such as those based on [Deep Neural Networks](#), .

Examples of bias entering AIED systems

1. When programmers code rule-based systems, they can put their personal biases and stereotypes into the system¹.
2. A data based algorithm can conclude not to propose a STEM-based career path for girls because female students feature less in the STEM graduate dataset. Is the lesser number of female mathematicians due to existing stereotypes and societal norms or is it due some inherent property of being female? Algorithms have no way to distinguish between the two situations. Since existing data reflects existing stereotypes, the algorithms that train on them replicate existing inequalities and social dynamics⁴. Further, if such recommendations are implemented, more girls will opt for non-STEM subjects and the new data will reflect this – a case of self-fulfilling prophecy³.
3. Students from a culture that is under-represented in the training dataset might have different behaviour patterns and different ways of showing motivation. How would a learning analytics calculate metrics for them? If the data is not representative of all categories of students, systems trained on this data might penalise the minority whose behavioural tendencies are not what the program was optimised to reward. If we're not careful, learning algorithms will generalise, based on the majority culture, leading to a high error rate for minority groups^{4,5}. Such decisions might discourage those who could bring diversity, creativity and unique talents and those who have different experiences, interests and motivations².

4. A British student judged by a US essay correction software would be penalised for their spelling mistakes. Local language, changes in spelling and accent, local geography and culture would always be tricky for systems that are designed and trained for another country and another context.
5. Some teachers penalise phrases common to a class or region, either consciously or because of biased social associations. If an essay-grading software trains on essays graded by these teachers, it will replicate the same bias.
6. Machine learning systems need a [target variable and proxies for which to optimise](#). Let us say high-school test scores were taken as the proxy for academic excellence. The system will now train exclusively to boost patterns that are consistent with students who do well under the stress and narrowed contexts of exam halls. Such systems will boost test scores, and not knowledge, when recommending resources and practice exercises to students. While this might also be true in many of today's classrooms, the traditional approach at least makes it possible to express multiple goals⁴.
7. Adaptive learning systems suggest resources to students that will remedy a lack of skill or knowledge. If these resources need to be bought or require home internet connection, then it is not fair to those students who don't have the means to follow the recommendations: *"When an algorithm suggests hints, next steps, or resources to a student, we have to check whether the help-giving is unfair because one group systematically does not get useful help which is discriminatory"*².
8. The concept of personalising education according to a student's current knowledge level and tastes might in itself constitute a bias¹. Aren't we also stopping this student from exploring new interests and alternatives? Wouldn't this make him or her one-dimensional and reduce overall skills, knowledge and access to opportunities?



"Data and algorithmic bias in the web" by jennychamux is licenced under CC BY 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by/2.0/?ref=openverse>.

What can the teacher do to reduce the effects of AIED biases?

Researchers are constantly proposing and analysing different ways to reduce bias. But not all methods are easy to implement – fairness goes deeper than mitigating bias.

For example, if existing data is full of stereotypes – "do we have an obligation to question the data and to design our systems to conform to some notion of equitable behaviour, regardless of whether or not that's supported by the data currently available to us?"⁴. Methods are always in tension and opposition with each other and some interventions to reduce one kind of bias can introduce another bias!

So, what can the teacher do?

1. **Question the seller** – before subscribing to an AIED system, ask what type of datasets were used to train it, where, by and for whom was it conceived and designed, and how was it evaluated.
2. **Don't swallow the metrics** when you invest in an AIED system. An overall accuracy of, say, five percent, might hide the fact that a model performs badly

for a minority group⁴.

3. **Look at the documentation** – what measures, if any, have been taken to detect and counter bias and enforce fairness¹?
4. **Find out about the developers** – are they solely computer science experts or were educational researchers and teachers involved? Is the system based solely on machine learning or were learning theory and practices integrated²?
5. **Give preference to transparent and open learner models which allow you to override decisions**², many AIED models have flexible designs whereby the teacher and student can monitor the system, ask for explanations or ignore the machine's decisions.
6. **Examine the product's accessibility**. Can everyone access it equally, regardless of ability¹?
7. **Watch out for the effects of using a technology**, both long-term and short-term, on students, and be ready to offer assistance when necessary.

Despite the problems of AI-based technology, we can be optimistic about the future of AIED:

- With increased awareness of these topics, methods to detect and correct bias are being researched and trialled;
- Rule-based and data-based systems can uncover hidden biases in existing educational practices. Exposed thus, these biases can then be dealt with;
- With the potential for customisation in AI systems, many aspects of education could be tailored. Resources could become more responsive to students' knowledge and experience. Perhaps they could integrate local communities and cultural assets, and meet specific local needs².

¹ [Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators](#), European Commission, October 2022.

² U.S. Department of Education, Office of Educational Technology, *Artificial Intelligence and Future of Teaching and Learning: Insights and Recommendations*, Washington, DC, 2023.

³ Kelleher, J.D, Tierney, B, *Data Science*, MIT Press, London, 2018.

⁴ Barocas, S., Hardt, M., Narayanan, A., [Fairness and machine learning Limitations and Opportunities](#), MIT Press, 2023.

⁵ Milano, S., Taddeo, M., Floridi, L., Recommender systems and their ethical challenges, *AI & Soc* 35, 957–967, 2020.

PART IV

PERSONALISING LEARNING

Did Youtube ever show you a video on a topic you might have spoken about to someone, or read something on a related topic but without having searched for it in Youtube? Yet, there it is, in the list of videos recommended for you.

Perhaps Youtube highlighted a previously unknown topic, one to which you are now hooked?

How is it that the video platform seems to know you better than some friends know you?

How does it select videos that will interest you, among the 800 million videos it hosts?

Is it possible to use the software's skills of personalisation to help learners learn better?

What is the common thread behind Youtube and Netflix, Amazon product recommendations, Google News, Facebook-friend suggestion and your classroom?

Read on...

20.


Personalising Learning

Every teacher personalises learning – if only by adding an additional example or giving individual attention where called for. In a sense, teaching itself is an act of personalisation, as opposed to, say, a televised lesson. Teachers change their lessons so students can make sense of what they learn. They help students to fit the new knowledge or skill with what they already know, their personal observances and social experiences. They help learners make what they can out of what they learn.

Personalised learning is about creating different learning environments and experiences for students' needs, capabilities and cultural contexts¹. Of course, the scope and degree of personalisation vary. Experts have identified six dimensions of personalisation – the why, how, what, when, who and where learning takes place².


For an education to be
meaningful it needs to be
personal

*TRAINING TOOLS FOR CURRICULUM DEVELOPMENT :
PERSONALIZED LEARNING, INTERNATIONAL BUREAU
OF EDUCATION*



www.penmendonca.com @MendoncaPen,
 Reproduced with permission from
 Technology-enhanced personalised learning:
 untangling the evidence, Stuttgart: Robert Bosch
 Stiftung, 2018.

One-to-one tutoring is the epitome of personalisation. In the 1960s, Benjamin Bloom showed that the average student performs better with individual tutoring. He also showed that individual attention closes the learning gap between the high and low scorers. In the real classroom, even with ten students, customising content to benefit each student can demand a lot of effort. True personalisation is practically impossible. Even where the teacher knows a student has gaps in learning they might be unable to remedy it due to time constraints. Thus, the system constantly loses students, even with the best effort of teachers.



Individualised Instruction Learning Gains from
 Essa, A., A possible future for next generation
 adaptive learning systems, *Smart Learning*
Environments, 3, 16, 2016, is licenced under CC
 BY 4.0. Visit <https://creativecommons.org/licenses/by/4.0/>

Stumbling Blocks to Personalisation

A large student:teacher ratio

Wide range of skills, aptitudes and
needs in the same class

Limited 

Limited resources



This is where technology can lend a hand.

Technology-enhanced personalised learning

Technology can be used to customise the learning process. Here, technology includes anything from mobile apps and online platforms to stand-alone learning systems². This is more effective now that artificial intelligence, access to data, mining techniques, cloud computing and affordable hardware have made applications seamless and practical.

Well-designed technology can do more than just help you to overcome the stumbling blocks shown above. When integrated with traditional classes, either as homework or occasional classwork, it could help learners acquire and practice a routine skill. Such technology can increase time spent in the classroom for interaction, personal attention and problem solving. Further, you can monitor what happens during homework – watch how far students have progressed and where they are struggling³.

Some + of Tech Based PL



The student set the pace

Content can now be learnt at



Classroom is for
interaction and individual
attention

Access to data :

Highlight difficulty,

make homework visible,



at-risk behaviour



Simulate curves,
animate processes,
talk in another
language



Exercise

[Click here for definitions of blended and other types of learning, often discussed alongside personalisation.](#)

Sometimes, for parts of a lesson, software might do a better job. Think of visualising three dimensions in maths, pronunciation practice in a language class – or an animation explaining the processes inside a human cell.

All AI solutions for education can be used to different degrees to help personalise learning. In this chapter, we discuss the use of adaptive learning systems.

¹ Groff, J., *Personalized Learning: The State of the Field & Future Directions*, Center for Curriculum Redesign, 2017.

- ² Holmes, W., Anastopoulou, S., Schaumburg, H & Mavrikis, M., *Technology-enhanced personalised learning: untangling the evidence*, Stuttgart: Robert Bosch Stiftung, 2018.
- ³ Feldstein, M., Hill, P., *Personalized Learning: What It Really Is and Why It Really Matters*, Educause Review, 2016.
- ⁴ Taylor, D., Yeung, M., Bashet, A.Z., *Personalized and Adaptive Learning*, Innovative Learning Environments in STEM Higher Education pp 17–34, Springer Briefs in Statistics, 2021.

21.

Let's say that your students are working on problems from a question bank. Imagine there is a person sitting next to each. They watch the steps followed by the student as they arrive at the solution.

Is the student struggling with a concept?


Do they seem to have a misconception?

Perhaps they are upset and could use a little encouragement?

The tutor gives a hint, points out what they are missing.

It can also happen that the student finds the problem too easy and is getting bored. In this case, the tutor assigns another, more challenging problem.

The tutor may even inspire questions and make the student reflect on their own performance. All this while keeping you informed of the student progress.



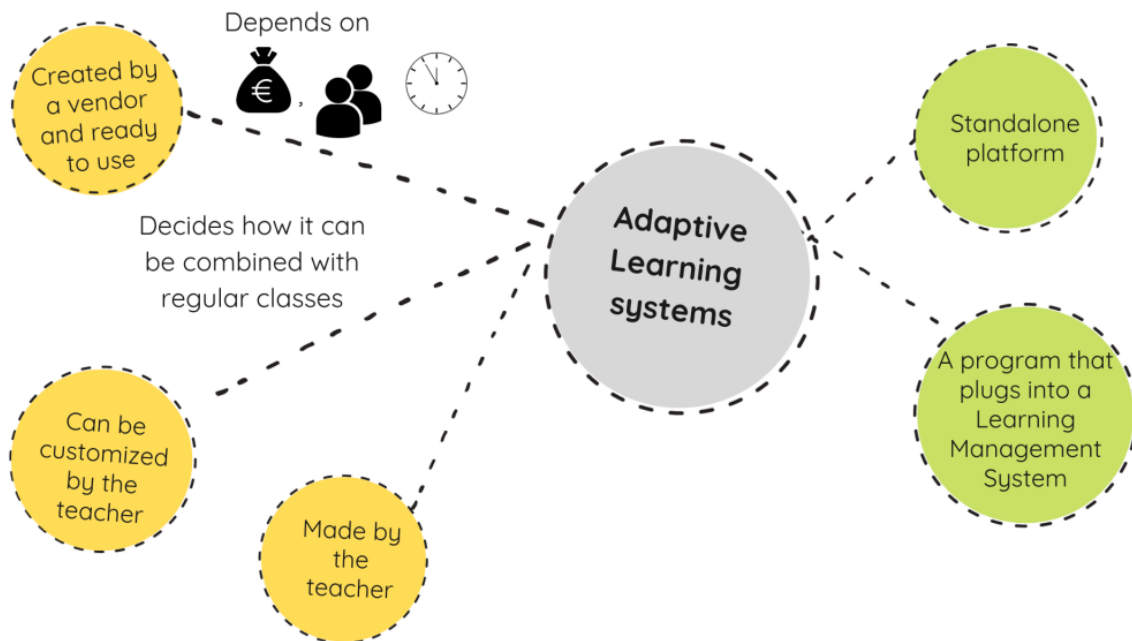
"Diary of a teaching machine" by [Ed] is licenced under CC BY-NC-SA 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-nc-sa/2.0/?ref=openverse>.

Intelligent tutoring systems (ITS) are designed to mimic the role of this tutor¹. They are a type of adaptive learning systems (ALS) that guide an individual student through each step of a solution. They give hints and feedback as needed. Because of this, ITS are more suitable for subjects like maths where problems and solutions are clearly defined². But recent ITS have taken on other subjects too.

Adaptive systems and learning

Adaptive learning occurs when digital tools and systems create individual learning paths – the sequence of activities executed in order to learn a given content or skill. The learning paths depend on each individual's strengths, weaknesses and pace of learning^{3,4}.

The idea of a machine adapting to a student goes back to the 1950s. With the recent advent of technology, the possibilities are now endless. These adaptive learning systems can be used for different purposes – solving problems, learning concepts and/or for assessing the student.

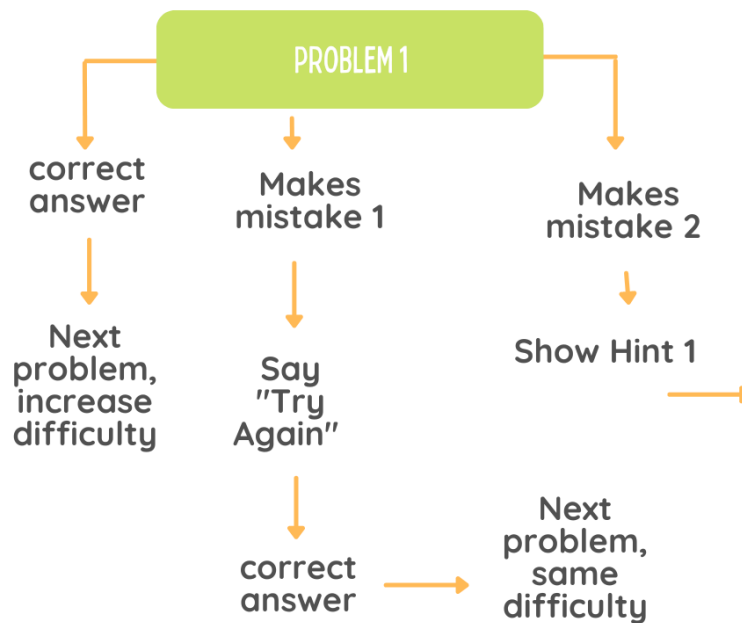


Many adaptive learning systems are now on the market. There are also authoring tools that help you create an ALS without any knowledge of coding. While creating an ALS might take a lot of time and resources, the teacher need not change lesson plans or style to fit it with lessons. Whatever be their type and form, the technologies used to create ALS vary a lot – not all systems are equal!

While choosing a system, you have to see how adaptive it is, what part of learning it personalises and whether it allows customisation by the teacher. Apart from that, there are important practical questions such as what equipment is required, how much does it cost and whether training is included in the cost.

Types of adaptive learning systems


Intelligent tutoring systems (see above) are personalised and interactive. They evaluate learning in real time. On the micro level, they adapt feedback when the student is solving a problem. On the macro level, they decide what problem to show next – much like Youtube recommending what video to watch next. Simple tutoring systems use decision trees for rules on what feedback to give. Other systems go beyond pre determined rules and use machine learning to tailor their behaviour¹.



Adaptive learning systems can go beyond tutoring. Exploratory learning systems, for example, let students explore a learning environment and choose what interests them. Game based systems package everything in the form of a game. When a student masters one level, they move to the next.

Whatever the type, all ALS should support the learner until they can carry out a task independently⁶. They should incite reasoning and support decision making. They should also be able to explain their decisions to the teacher and student.

When it comes to choosing and using ALS, or even deciding whether or not to use one, experts advise to always start with the learning². Ask *what student need is to be addressed? Which tool fits this job? How will different students be supported differently*⁵? Studies show that these systems do not have a significant impact on student learning when used for short periods of time. The effectiveness increased when used for one full school year or longer⁷. If you decide to use one, be prepared to support students on guiding their own learning. Be patient and be ready to experiment, fail and retry^{2,5}.



"Diary of a teaching machine" by [Ed] is licenced under CC BY-NC-SA 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-nc-sa/2.0/?ref=openverse>.

- ¹ Groff, J., *Personalized Learning : The state of the field and future directions*, Center for curriculum redesign, 2017.
- ² Holmes, W., Anastopoulou S., Schaumburg, H & Mavrikis, M., *Technology-enhanced personalised learning: untangling the evidence*, Stuttgart: Robert Bosch Stiftung, 2018.
- ³ Taylor, D., Yeung, M., Bashet, A.Z., *Personalized and Adaptive Learning*, Innovative Learning Environments in STEM Higher Education pp 17–34, SpringerBriefs in Statistics, 2021.
- ⁴ Becker, S. et al, *NMC Horizon Report: 2018 Higher Education Edition*, Educause, 2018.
- ⁵ Feldstein, M., Hill, P., *Personalized Learning: What It Really Is and Why It Really Matters*, Educause Review, 2016.
- ⁶ Wood, D., Bruner, J., Ross, G., *The role of tutoring in problem solving*, The Journal of Child Psychology and Psychiatry, 1976.
- ⁷ Alkhatlan, A., Kalita, J.K., *Intelligent Tutoring Systems: A Comprehensive Historical Survey with Recent Developments*, International Journal of Computer Applications 181(43):1-20, 2019.

22.

Models and Recommendation

ACTIVITY

JOHN DOE'S CREDIT CARD TRANSACTIONS

DATE	MERCHANT/ PRODUCT	PLACE OF PURCHASE	AMOUNT IN EUROS
19/09/2022	XYZ FOOD STORE	NANTES	250
	THE BABY SHOP	NANTES	400
	ORANGE WIFI	ONLINE	90
18/09/2022	THE PIZZA MAGAZINE SUBSCRIPTION	ONLINE	5
	NETFLIX SUBSCRIPTION	ONLINE	17.99
17/09/2022	TERRA MADRE RESTAURANT	LYON	80
	APPSTORE APPLE	ONLINE	25
16/09/2022	AIR FRANCE	ONLINE	1500
	SNCF TRAIN TICKETS	ONLINE	100

TOM HARRY'S CREDIT CARD TRANSACTIONS

DATE	MERCHANT/ PRODUCT	PLACE OF PURCHASE	AMOUNT IN EUROS
19/09/2022	CHORDS MUSIC SCHOOL CANTEEN	NANTES	6
18/09/2022			
17/09/2022	CHORDS MUSIC SCHOOL CANTEEN	NANTES	6
	NEIGHBOURHOOD SUPERMARKET	NANTES	19
16/09/2022	LOVELY MUSIC STORE	NANTES	250
	MCDONALD'S	NANTES	9

These are the credit card transactions of John Doe and Tom Harry, two men living in Nantes, France. They are looking for things to do this weekend. What will you recommend to them?

List to choose from:

1. The new Burger King outlet
2. An olive oil tasting event
3. An online luggage store
4. A river-side concert
5. Baby swimming class

Recommendation systems have been around at least as long as tourist guides and top-ten lists. While *The Guardian Best Books of 2022* recommends the same list

to everyone, you would likely adapt it when choosing for yourself – pick a few and change the order of reading based on your personal preferences.

How to recommend options for strangers? In the activity above, you probably tried to imagine their personalities based on the given information: you made judgements and applied stereotypes. Then, once you had an idea of their *type*, you chose from the list things that could (or not) be relevant to them. Recommenders such as Amazon, Netflix and Youtube follow a similar process.

Nowadays, whenever someone is searching for information or looking to discover online content, they use some kind of personalised recommender system^{1,2}. The main function of Youtube is to suggest to its users what to watch amongst all the videos available on the platform. For signed-in users, it uses their past activity to create a model, or a personality type. Once it has a model for John, it can see who else has models similar to his. It then recommends to John videos similar to what he has watched and those similar to what others like him have watched.

What is a model?

Models can be used to mimic anything from users to videos to lessons a child has to learn. A model is a simplified representation of the world, so a machine can pretend to understand it:



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=127#oembed-1>

How Youtube learns who you are

All recommendation problems involve asking a surrogate question: “*What to recommend*” is too general and vague for an algorithm. [Netflix asked developers](#) what will be the rating user A would give video B, considering their ratings for other videos. Youtube asks what the watch time would be for a given user in a particular context. The choice of what to ask – and predict – has a big impact on what recommendation is shown³. The idea is that the correct prediction will lead to a good recommendation. The prediction itself is based on other users with a history of similar tastes⁴. That is, users whose models are similar.

User models

Youtube splits the task of recommendation into two parts and uses different models for each³. We, however, will stick to a simpler explanation here.

For creating a user model, its developers have to ask, what data is relevant to video recommendation. What about what the user has watched before? What about their reviews, ratings, and explicit preferences thus far? What did they search for? Youtube uses signals that are more implicit than explicit, since the latter are more readily available³. Did a user just click a video or did they watch it? If yes, for how long? How did the user react to previous recommendations¹? Which ones were ignored? Apart from direct answers to these questions, demographic information such as gender, language, region and type of device are of great value when the user is new or not signed in³.

Once a model is available for each user, we can compare users and use that information for recommendations.

Video Models

We could also use videos that are both similar to and different from one another. Youtube looks at the information it has on a particular video – its title and description, video quality, how many people have watched it (view count), liked it, favoured it, commented on it or shared it, the time since it was uploaded and the number of users subscribed to the parent channel¹.



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=127#oembed-2>

What a user watches next will also depend on whether one video is an episode within a series or an item in a playlist. If a user is discovering a new artist, he or she might move from the most popular songs to smaller niches. Also, a user might not click on a video whose thumbnail image is poor quality^{1,3}. All of this information goes into the model too.

One of the building blocks of the recommendation system is to go from one video to a list of related videos. In this context, we define related videos as those that a user is likely to watch next³. The goal is to squeeze the most value out of data to make better recommendations⁴.

¹ Davidson, J., Liebald, B., Liu, J., Nandy, P., Vleet, T., *The Youtube Video Recommendation System*, Proceedings of the 4th ACM Conference on Recommender Systems, Barcelona, 2010.

² Spinelli, L., and Crovella, M., [How YouTube Leads Privacy-Seeking Users Away from Reliable Information](#), In Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20 Adjunct), Association for Computing Machinery, New York, 244–251, 2020.

³ Covington, P., Adams, J., Sargin, E., *Deep neural networks for Youtube*

Recommendations, Proceedings of the 10th ACM Conference on Recommender Systems, ACM, New York, 2016.

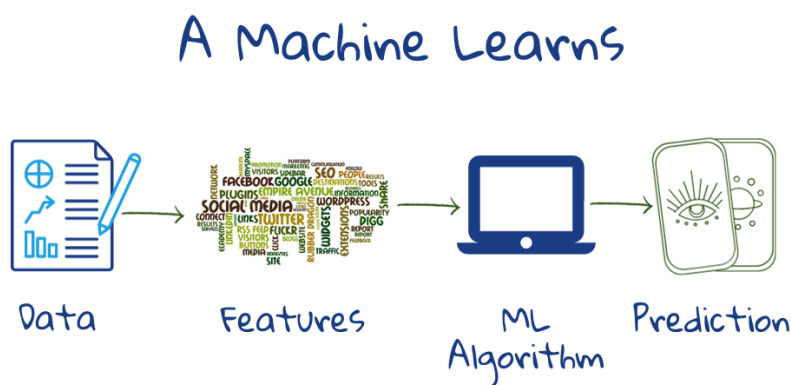
⁴ Konstan, J., Terveen, L., *Human-centered recommender systems: Origins, advances, challenges, and opportunities*, AI Magazine, 42(3), 31-42, 2021.

23.

The Process

Across Google, [deep neural networks](#) are now being used for machine learning². Based on the video model, Youtube's neural network takes videos similar to the ones already watched by the user. Then it tries to predict the watch time of each new video for a given user model, and ranks them based on the prediction. The idea is then to show the 10 to 20 videos (depending on the device) with top rankings.

The process is similar to the [machine learning model](#) we studied earlier. First, the machine takes features from user and video models given by the programmer. It learns from training data what weight to give each feature to predict watch-time correctly. And then, once tested and found to be working, it can start predicting and recommending.



Training

During training, millions of both positive and negative examples are given to the system. A positive example is when a user clicked on a video and watched for a certain time. A negative example is when the user did not click on the video or did not watch for long².


The network takes in the user features and video features discussed under the models section of [How Youtube Learns You Part 1](#). It adjusts the importance given to each input feature by checking whether it predicted correctly the watch time for a given video and user.

There are approximately one billion parameters (weight of each feature) to be learned on hundreds of billions of examples². The network might also learn to disregard certain features and will give it zero importance. Thus, the embedding, or

the model the algorithm creates, can be very different from that envisioned by the developers.

Testing

Once the network has been trained, it is tested on already available data and adjusted. Apart from accuracy of prediction, the output of the system has to be tuned by the programmer, based on several value judgements. Showing videos that are too similar to already watched videos will not be very engaging. What does it really mean for a recommendation to be good? How many similar videos to show and how much diversity to introduce – both with respect to the other videos and with respect to the user history. How many of the user's interests to cover? What type of recommendations lead to immediate satisfaction and which lead to long-term use^{1,3}? These are all important questions to consider.



After this testing, real-time evaluation of the recommendations is done. The total watch time per set of predicted videos is measured². The longer a user spends watching the recommended set of videos, the more successful the model is considered to be. Note that just looking at how many videos were clicked is not sufficient grounds for evaluation. Youtube evaluates its recommenders based on how many recommended videos were watched for a substantial portion of the video, session length, time until first long watch and the fraction of logged-in users with recommendations¹.

The interface

We will now explore how the recommendations are presented to the viewer. How many videos should be shown? Should the best recommendations be presented all at once, or should some be saved for later³? How should thumbnails and video titles be displayed? What other information should be shown? What settings can the user control¹? Answers to these questions determine how Youtube manages to keep two billion users hooked.

¹ Davidson, J., Liebald, B., Liu, J., Nandy, P., Vleet, T., *The Youtube Video Recommendation System*, Proceedings of the 4th ACM Conference on Recommender Systems, Barcelona, 2010.

² Covington, P., Adams, J., Sargin, E., *Deep neural networks for Youtube*

Recommendations, Proceedings of the 10th ACM Conference on Recommender Systems, ACM, New York, 2016.







³ Konstan, J., Terveen, L., *Human-centered recommender systems: Origins, advances, challenges, and opportunities*, AI Magazine, 42(3), 31-42, 2021.

⁴ Spinelli, L., and Crovella, M., [*How YouTube Leads Privacy-Seeking Users Away from Reliable Information*](#), In Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20 Adjunct), Association for Computing Machinery, New York, 244–251, 2020.

24.

An Adaptive Learning Tool can adapt 1 or more aspects of learning :



- Does it adapt learning  ? :
The Sequence of learning activities? The difficulty level and type of activities?
- Does it adapt  within an activity? Hints and guidance  -by-  ?
- Does it adapt the learning approach?
- Does it adapt  ? The  level of questions and problems with or without changing the learning path?

When looking at an adaptive learning system, it is very hard to tell where it adapts¹. What technology is used and what it is used for also changes across systems.

However, all adaptive learning systems know who they teach (knowledge about the learner), what they teach (knowledge about the domain), and how to teach (knowledge about pedagogy)².

An ideal ALS adapts itself in multiple ways. In the outer loop, the sequence of learning activities is adapted – similar to Youtube adapting recommended list of videos. The outer loop might also personalise learning approaches and difficulty levels.

In the inner loop, within each activity, the ALS monitors step-by-step progress. It adapts feedback and hints to correct

misconceptions, if any. It might also point to additional content if the student has a problem remembering a previously learnt concept. Some experts argue that the inner loop is best left to the instructor: not only is it costly and time-consuming to program all the rules for the specific subject and task, but the teacher's knowledge and experience will always trump that of the machine³.

How adaptive systems learn the learner

Like all recommendation problems (See [How Youtube Learns You Part 1](#)), ALS splits the task into one or more surrogate questions that can be answered by the machine. Again, the choice of what to ask -and thus, what to predict, has a big impact on what recommendation is shown.

Marketing material often mention multiple goals – improved scores, employability and engagement. Given the proprietary nature of the systems, it is usually unclear what questions are coded into the systems, what goals are being optimised for, and how short-term goals are differentiated from long-term goals (example, mastery of a given content before progressing to the next level)⁴.

Where machine learning is used, regardless of the goals chosen, the prediction itself is based on other learners with similar skill levels and preferences. That is, learners whose models are similar.

The learner model

For creating a student model, developers ask what student characteristics are relevant to the learning process. Unlike teachers who can directly observe their students and adjust their approach, machines are limited to the data that can be collected and processed by them.

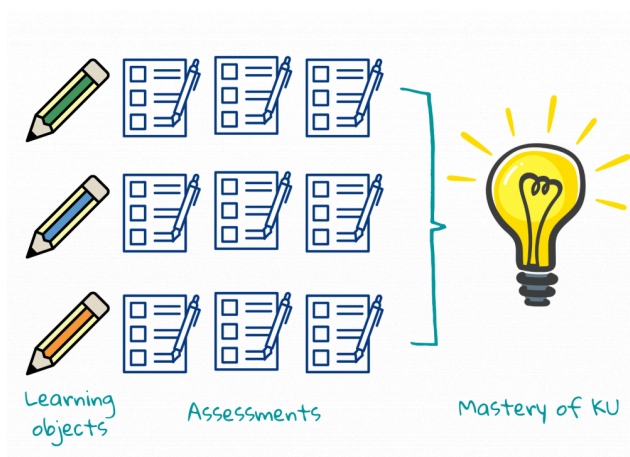
Typical characteristics considered in a student model:

- **What does the student know – their knowledge level, skills and misconceptions^{5,2,6}**. These are usually inferred through assessments, for example, the answer a student submits for a maths problem¹. This prior knowledge is then compared with what they will need to know at the end of the learning period.
- **How does a student prefer to learn: the learning process and preferences^{5,6}**. For example, the number of times a student attempts a question before getting it correct, the types of resources consulted, the ratings they gave for an activity¹, or the material that engaged them most – images, audio or text². ALSs may also record **when and how skills were learned and which pedagogies worked best⁶**.
- **Is the student feeling motivated: feelings and emotions** can be recorded directly by the student or extracted indirectly from speech, facial expressions, eye tracking, body language, physiological signals, or combinations thereof. This information can then be used to nudge the student out of negative states such as boredom or frustration that inhibit learning, into positive states such as engagement or enjoyment⁷.
- What about **cognitive aspects like memory, attention, problem solving skills, decision-making capability, analysis of situations and critical thinking⁵**?
- How do they **communicate and collaborate⁵**? For example, do they post comments on other students' feed and how do they discuss with others to solve problems¹?
- What about meta-cognitive skills such as **self-regulation, self-explanation, self-assessment and self-management⁵, help-seeking, being aware of and being able to control their own thinking**? For example, how they select their learning goals, use prior knowledge or intentionally choose problem-solving strategies⁵.

While this data changes and has to be recorded and updated, models also contain **static characteristics such as age, gender, mother tongue and email identity²**.

Most of the ALSs create learner models based on interactions with students. Some also glean info from other sites, especially social media. Once a model is available for each learner, the machine calculates which students are similar to one another and estimate the probability that a given student will benefit from an activity, example or question³.

The domain model



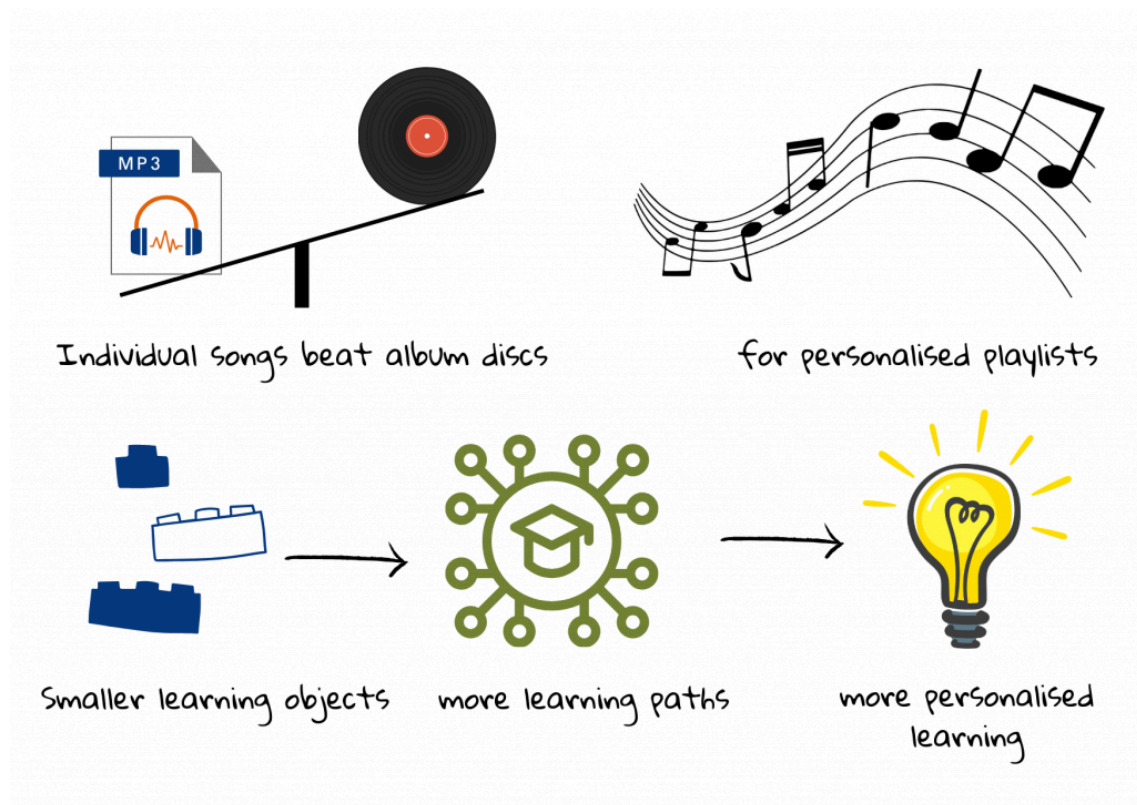
We can draw a loose parallel between learning objects in an ALS and videos on Youtube. A subject can be broken down into concepts and skills; these are called knowledge units (KUs) and are what the learner needs to know³. Each KU has a set of learning objects through which content can be learned, and a set of activities to assess learning. Some authors further break down learning objects into learning activities; we do not do this here.

Learning objects can be text to read, a video, a set of problems, interactive activities (from simple fill-in-the-blanks to scenario-based learning activities), interactive animation, etc.¹ The learning objects give what the learner needs to know, and the assessment activities indicate if the knowledge has been acquired³. The domain model contains all the features of the learning objects, including the associated KU and assessment.

What a learner learns next will also depend on the inter-relationships between the KUs.

These KUs need to go into the model too: learning objects A and B might both be pre-requisites for learning object D. Thus, A and B have to be mastered before D. There is order amongst some KUs that tell us how we learn³. Conversely, if the student solves correctly a problem that corresponds to D, it would be a good bet that he or she mastered A and B too.

Subject-matter experts can provide some of these relationships. The other inferences can be learnt by the machine, which can predict the probability that a KU has been mastered. This mastery involves the system knowing that the learner has mastered A and B, given that he answered questions under D. It can then use this information, along with other features of learner and domain models, to recommend learning pathways and learning objects.



Other features of learning objects could include the activity's level of difficulty, its popularity and its ratings. The goal here, as in the case of Youtube recommendation, is to squeeze out as much information as possible from the data available.

¹ EdSurge, *Decoding Adaptive*, Pearson, London, 2016.

² Alkhatlan, A., Kalita, J.K., *Intelligent Tutoring Systems: A Comprehensive Historical Survey with Recent Developments*, International Journal of Computer Applications 181(43):1-20, March 2019.

³ Essa, A., [A possible future for next generation adaptive learning systems](#), Smart Learning Environments, 3, 16, 2016.

⁴ Bulger M., *Personalised Learning: The Conversations We're Not Having*, Data & Society Working Paper, 2016.

⁵ Chrysafiadi, K., Virvou, M., *Student modeling approaches: A literature review for the last decade*, Expert Systems with Applications, Elsevier, 2013.

⁶ Groff, J., *Personalized Learning: The state of the field and future directions*, Center for curriculum redesign, 2017.

⁷ du Boulay, B., Poulouvasillis, A., Holmes, W., Mavrikis, M., *Artificial Intelligence And Big Data Technologies To Close The Achievement Gap*, In: Luckin, Rose ed. *Enhancing Learning and Teaching with Technology*. London: UCL Institute of Education Press, pp. 256–28, 2018.

ML applications, there is training and testing to be done before being put to use in classrooms.

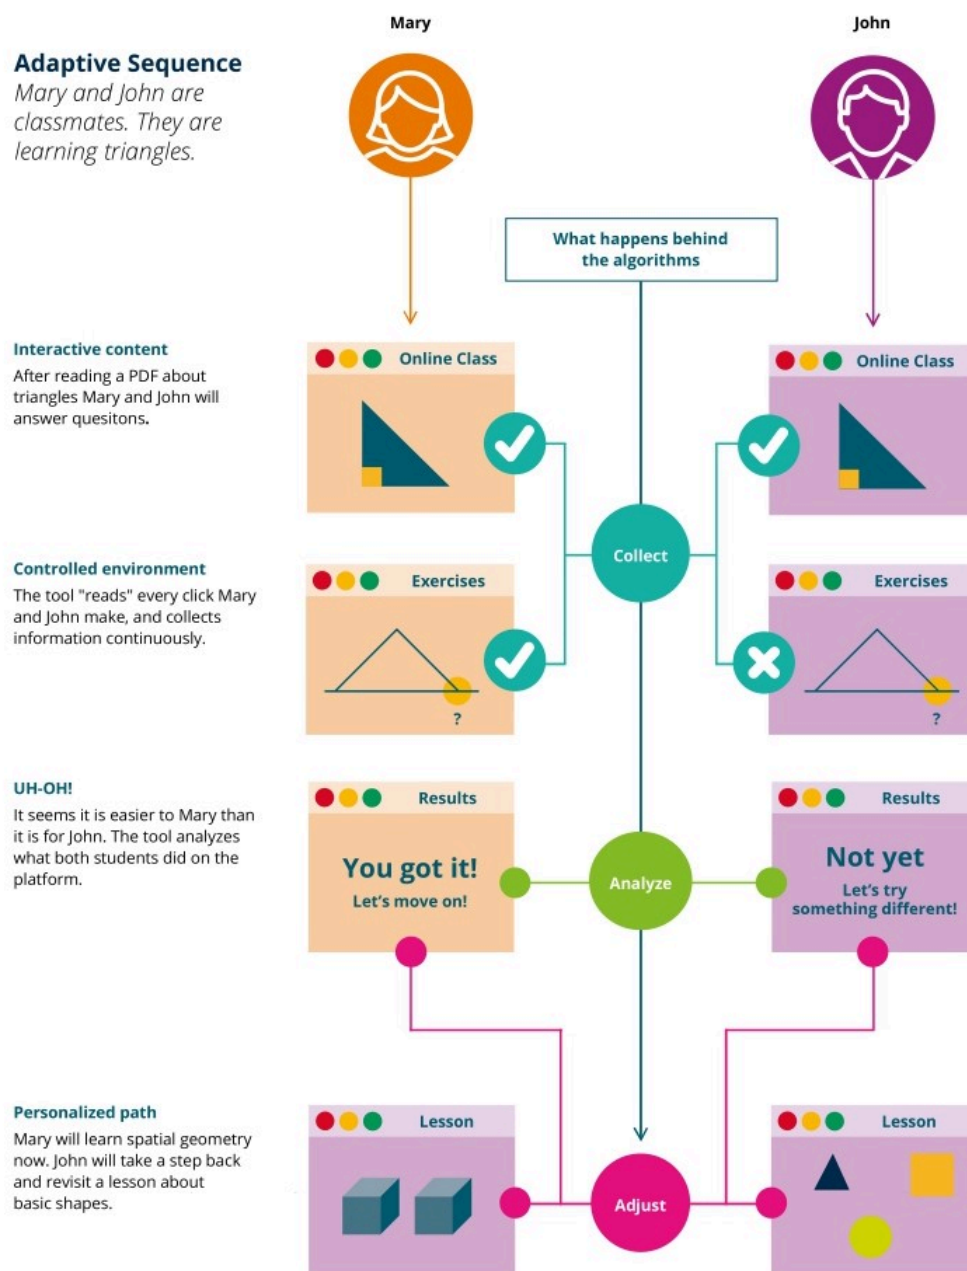


Image from EdSurge, *Decoding Adaptive*, Pearson, London, 2016 licenced under CC BY 4.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by/4.0/>

Pedagogical model

In the case of Youtube, we saw there are a lot of value judgements on what makes a good recommendation – such as how many user interests need to be covered in one set of recommendations, how many videos should be similar to already watched ones, how much new content to add for diversity (see [How Youtube learns you, part](#)

2). ALS involves similar judgements on what it means to master a KU and how to get to that mastery: the pedagogy and daily experience of the learner⁴.

In the case of ALS, these judgements and guidance on how a learner is to progress should be based on proven pedagogical theories. These go into the pedagogical model, and, along with the domain and learner models, help the machine to choose an appropriate set of activities.

Some of the questions answered in this model are: should the student next be presented with a concept, an activity or a test? At what level of difficulty? How do they evaluate the learning and provide feedback? Where is more scaffolding necessary⁵? (Scaffolds are support mechanisms that give guidance on concepts and procedure, the strategy used and how to reflect, plan and monitor learning.) The pedagogical model dictates the breadth and depth of activities – and even whether to continue within the ALS or get help from the teacher³.

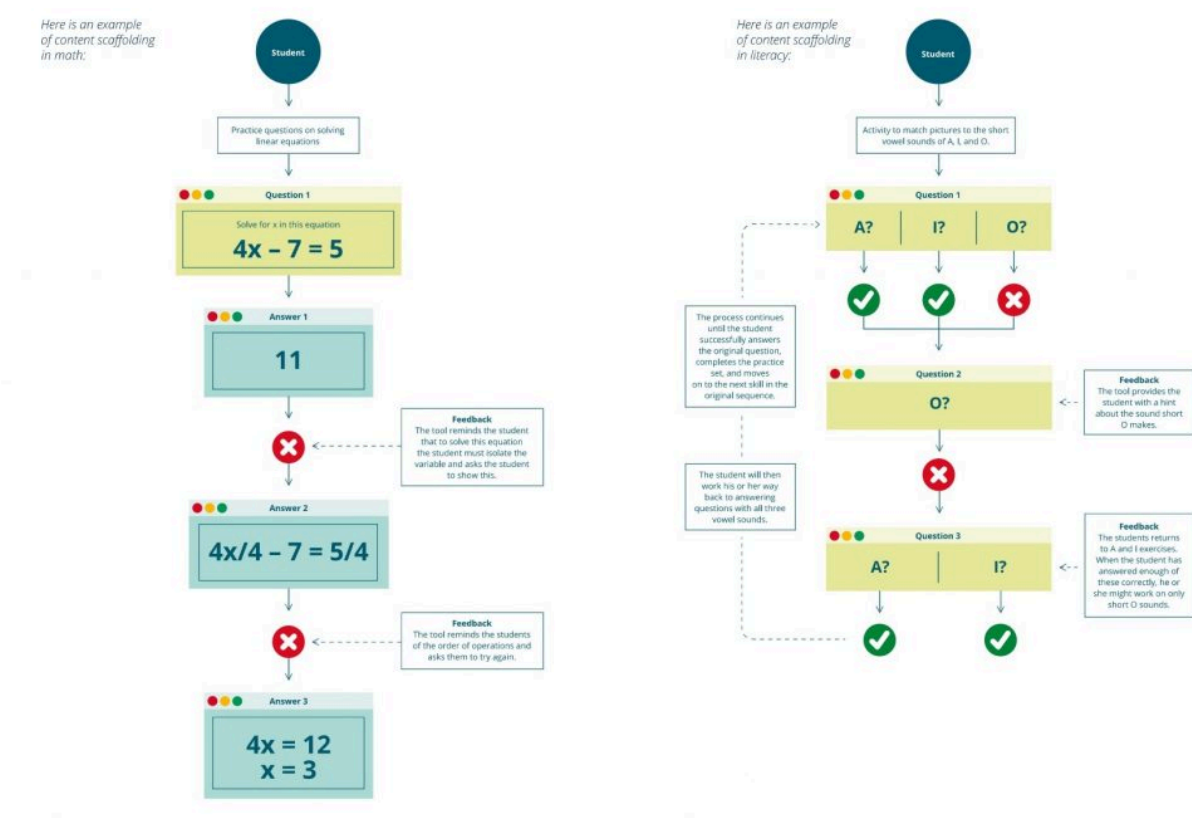


Image from EdSurge, *Decoding Adaptive*, Pearson, London, 2016 licenced under CC BY 4.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by/4.0/>

The interface

The recommendations are presented along with other data such as learner progress, performance and goals. The key questions here are:

- How will content be delivered?
- How much content will be recommended in one go?
- What is assigned directly and what is recommended?
- What are the supporting resources?
- Is it possible to provide group activities?
- How much autonomy should be permitted?
- Can students change their preferences?
- Can teachers change the learning pathways?
- What data is shown to the teachers?
- Are the teachers in the loop?

Evaluation

When the ALS is put to use, most systems monitor their own performance against criteria set by the programmer. Like in any AI tool, data might be biased. Inferences drawn by the system can be imprecise. The student's past data will become less relevant with time⁶. Therefore, the teacher also has to monitor the system's performance and provide learner guidance and corrective measures where necessary.

Teachers and peers also have to provide inspiration and reveal alternative resources. research in recommendation systems was shaped by commercial content providers and online retail companies for over a decade. Thus, the focus has been on providing recommendations that produce results that can be promoted. "The surprising delight of an unexpected gem"⁷ and the allure of roads less travelled, can inspire enduring learning. Unfortunately, these are not the strong points of machine-based personalised learning.

¹ EdSurge, *Decoding Adaptive*, Pearson, London, 2016.

² Chrysafiadi, K., Virvou, M., *Student modeling approaches: A literature review for the last decade*, Expert Systems with Applications, Elsevier, 2013.

³ Essa, A., [A possible future for next generation adaptive learning systems](#), Smart Learning Environments, 3, 16, 2016.

⁴ Groff, J., *Personalized Learning: The state of the field and future directions*, Center for curriculum redesign, 2017.

⁵ Alkhatlan, A., Kalita, J.K., *Intelligent Tutoring Systems: A Comprehensive Historical Survey with Recent Developments*, International Journal of Computer Applications 181(43):1-20, March 2019.

⁶ du Boulay, B., Poulavasillis, A., Holmes, W., Mavrikis, M., *Artificial Intelligence And Big Data Technologies To Close The Achievement Gap*, In: Luckin, Rose ed. Enhancing Learning and Teaching with Technology. London: UCL Institute of Education Press, pp. 256–28, 2018.

⁷ Konstan, J., Terveen, L., *Human-centered recommender systems: Origins, advances, challenges, and opportunities*, AI Magazine, 42(3), 31-42, 2021.

26.

Despite the promised potential of adaptive learning systems, many questions remain unanswered. There is not yet enough research or documentation of classroom practices that help broach these issues:

- Recommendation systems are used for suggesting movies to Netflix users. They help consumers home in on the right choice of, say Audio Speakers on Amazon. But can they actually improve learning outcomes for each student in the classroom¹?
- Does focussing all the time on performance and individualisation affect a student's psychological well being²?
- Individualisation demands a lot of discipline and self-regulation from a student. They have to start working by themselves and continue working till they finish all assigned activities. Are all students able to do this without help²?
- How do we balance individualisation with social learning opportunities³?
- How do we go from using ALS as a support for a single topic, to using these systems systematically, across topics and subjects²? What about the curriculum change that will be required for such an incorporation of adaptivity³?
- What about the required infrastructure? What needs to be done about data and privacy, as well as bias and reinforced stereotypes³?

When developing ALS, some principles are used either directly or implicitly. These are not always without consequences.

A paradigm of ALS: old is gold

What do machine learning systems do when they predict or recommend something? They use the student's past experiences, preferences and performance in order to choose what to recommend to them; they look to the past in order to predict the future. Thus, these systems are always biased towards the past⁴. *Machine learning works best in a static and stable world where the past looks like the future*⁵. ALS, based on machine learning models does more or less the same thing, but now with the addition of pedagogical considerations.

As a consequence, these systems are not able to account for fluctuations in normality like the COVID pandemic, health issues and other problems. They can struggle to take account of age, growth, mastery of new competencies and personal evolution of young humans.

Is student behaviour even predictable? How many times can we repeat a formula that worked well in the past, before it becomes boring, repetitive and impedes progress⁶? Even if such a prediction were possible, is it even prudent to expose students only to that which they like and are comfortable with? How much novelty is overwhelming and counterproductive⁶?

It is difficult to decide how similar recommended activities should be, how many new types of activities should be introduced in one session and when would it be productive to push a student to face challenges and explore new interests. The answers do not lie in the students' pasts alone.

A paradigm of ALS: the explicit reflects the implicit

Even where the past can be used reliably to predict the future, the past itself could be difficult to capture accurately. How can Youtube know a user liked a video? It is easier where they clicked the *Like* button or subscribed to the parent channel after watching it. But such explicit behaviour is often rare. Recommendation systems have to regularly resort to implicit signals that may or may not fully reflect the truth⁴. For example, Youtube uses the time a user spent watching the video as an implicit signal that they liked the video and would like to watch similar content. But, just because a video played on someone's computer till the end hardly means the person liked it, or even watched it⁷.

What about how feedback is recorded in an adaptive learning system? To gauge, for example, if a student was attentive during an activity, the system might record the number of digital resources they clicked on, and when and for how long they accessed them. But these cannot accurately reflect their level of attention¹.

For example, if the student is clear about what to do for an activity, they might consult a few resources and zero in on the critical points quickly. Someone who is not as clear might open and spend time on all the listed resources without learning much¹. It is possible that the first student is wrongly flagged for lack of motivation and made to do additional work.

Also, the machine learning models can only note that two things happened – a student clicking on a resource and a student scoring high on the associated exercise. They cannot infer that the student scored high because they consulted the resource – they can infer correlation but not causation⁵.

The unfair expectation of some ALS is that the teacher will dive in and remedy such errors. In other systems, the teacher does not even have the option to do so.

The paradigm of ALS: everything can be replaced by this one question

Recommendation systems cannot handle multiple goals. The aim of the ALS is often put forth in a form of a single question: *the surrogate question*. What rating did a user give a movie, how long did they watch a video, what is the score of the student in a quiz, how well did they satisfy the criteria the machine used to measure attentiveness... The systems are then trained to attain these goals and tested based on whether they was achieved. Their performance is constantly adjusted to maximize their score with respect to these goals.

If scoring on the quiz is the goal, certain content is recommended in a certain way.

That way, exam performance is the surrogate problem that is solved. If the goal is just to make them click on many resources, recommendation would be tailored to push them to do just this. In this instance, making resources adequately attractive is the problem.

The choice of the surrogate question has an outsized importance on how the ALS works. What is more, contrary to the promotion of ALS as objective systems, there is more art than science in selecting the surrogate problem for recommendations⁴.

All tech is not hi-tech

As we have seen, many decisions go into the making of ALS – what data is measured, how this data is used to gauge feedback and other information, what goals are optimised, and what algorithms are used to optimise these goals. Often it is programmers, data scientists, finance and marketing experts that are involved in making these decisions. The input of teachers and pedagogical experts in the development process is rare and often comes after the designing process². Products are not field-tested before adoption in schools, and often their proclaimed effectiveness is based on testimonials and anecdotes, instead of scientific research².

As a result, what a school needs and is familiar with has little impact on what software companies are building. Cost, availability and infrastructure can have a major say on what schools can buy. It is important to bear this in mind while deciding if or how to use a particular product. Perhaps it is better not to think of them all as adaptive learning systems or AI, but individual systems with wildly different objectives, designs and capabilities.

ALS as a whole can be used for personalising feedback, scaffolding and practice. They can find gaps in learning and remedy it within the limits of programming and design. They cannot detect ‘teachable moments’ or when it is right to capitalise on the mood of the class to introduce a new idea or example. These capabilities that make learning magical and which help the lesson endure in the student’s mind are solely the forte of the teacher.

¹ Bulger M., *Personalised Learning: The Conversations We’re Not Having*, Data & Society Working Paper, 2016.

² Groff, J., *Personalized Learning: The state of the field and future directions*, Center for curriculum redesign, 2017.

³ Holmes, W., Anastopoulou S., Schaumburg, H & Mavrikis, M., *Technology-enhanced personalised learning: untangling the evidence*, Stuttgart: Robert Bosch Stiftung, 2018.

⁴ Covington, P., Adams, J., Sargin, E., *Deep neural networks for Youtube Recommendations*, Proceedings of the 10th ACM Conference on Recommender Systems, ACM, New York, 2016.

⁵ Barocas, S., Hardt, M., Narayanan, A., [*Fairness and machine learning Limitations and Opportunities*](#), MIT Press, 2023.

- ⁶ Konstan, J., Terveen, L., *Human-centered recommender systems: Origins, advances, challenges, and opportunities*, AI Magazine, 42(3), 31-42, 2021.
- ⁷ Davidson, J., Liebald, B., Liu, J., Nandy, P., Vleet, T., *The Youtube Video Recommendation System*, Proceedings of the 4th ACM Conference on Recommender Systems, Barcelona, 2010.

PART V

LISTENING, SPEAKING AND WRITING

Haven't we all been astounded by how fast a baby picks up its native tongue? After that initial growth spurt, it takes the life-long inputs of family, friends, teachers and strangers to perfect that language.

With constant listening, speaking and later, reading and writing we could say the language grows along with the baby into maturity. This personal growth can hardly be scaled up from the individual to a classroom. And even less, when it comes to a foreign language.

How does the teacher divide their time to converse and correct each student to ease this process?

If AI is to help in personalized learning, isn't language learning a good place to show its prowess?

27.

Automatic translation tools are available online and can be used in a simple way for many languages today. Some of these tools have been produced by the internet giants (Google Translate for example), but independent specialised tools such as DeepL are also available.

Historically, automatic translation has been a challenge for artificial intelligence, and the diverse AI technologies have been tested over the years. Rule-based systems (with rules hand-built by experts) were replaced by statistical machine-learning techniques when data-sets of parallel texts (the same text in multiple languages) became available. Over the past few years, deep-learning techniques have become state of the art.

Just a few years ago, you could have an enjoyable moment testing these tools, which would return amusing translations for songs or menus for example; this is no longer the case today:

- International institutions are contemplating using automatic translation tools to support multilingualism;
- Big media video platforms use automatic translation rather than human translation in order to reach more people;
- Bilingual people and translation professionals seem to use these tools both in their day-to-day lives and professional activities.

Find out!

[Some terms
related to
automatic
translation](#)

Furthermore, improvements have yet to come. The quality of translation is still improving, and solutions that combine translations with transcriptions and speech synthesis, and allow seamless multilingual communication, will soon be common.

Even if these tools have not been designed for education, they are already having an impact.

Are pupils using automatic translation?

To our knowledge, there are today (December 2022) no public official documents nor large-scale surveys measuring if this is an issue. There are discussions on forums³ and articles presenting possible ways to avoid cheating with AI, or suggesting ways to introduce AI into foreign-language classes. These assume that pupils' usage of automatic translation tools is widespread.

We ran a smaller, informal survey in April 2022, with teachers of various languages (English, French, German) and with varying degrees of skill. The main classes corresponded to 12-16 year-old pupils. It took place in the Paris area, so pupils and teachers were French. The results were such that the teachers all had to cope with pupils who would, once out of the classroom, make use of DeepL or Google Translate.

Here are some of the remarks we got:

- The only skill the pupils seem to be acquiring is copy-pasting.
- Even the better and more motivated pupils do it – they will try to do their homework on their own, but then they will check it with an automatic translation tool. Often, they realise that the automatic result is much better than their own, so they keep the machine-built solution.
- There is a motivation issue as pupils are questioning the use of learning languages.

The above analysis needs further investigation. A generalised survey over various countries would certainly help. In the meantime, discussions with various stakeholders have allowed us to consider the following:

- Teachers can no longer assign translating text as homework. Even more creative exercises, such as writing an essay on a particular topic, can lead to using automatic translation tools – the pupil could write the essay in their own language, then translate it.
- The motivation question is critical. It is not new – in 2000, authors and educators argued that “some view the pursuit of foreign language competence as an admirable expenditure of effort; others may see it as unnecessary if an effective alternative exists”⁵.

Our observations coincide with reactions found in forums or reported in literature⁴.



“Improbable translation” by giopuo is licenced under CC BY-NC-SA 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-nc-sa/2.0/?ref=openverse>.

Can automatic translators fool teachers?

Find out!

[AI technology is moving fast](#)

Blog papers seem to indicate that a language teacher will recognize automatic translation, even when it has been corrected by a human at a later stage: Birdsell¹ imagined a task where Japanese students were to write a 500 word essay in English. Some had to write it directly, albeit with commonly used and accepted tools (dictionaries, spelling checkers) and others would write the essay in Japanese and then translate it -using DeepL- into English. Interestingly, he found that the teachers would grade higher the students from the second group but would also be able to identify those essays written

with the help of DeepL.

Can machine translation tools be combined with text generators?

These are early days to predict which will be the course of events but the answer is, for the moment, yes. As a simple example, journalists in France used a text-generator tool (Open-AI playground) to produce some text, then ran DeepL on it, and felt comfortable in presenting this text to the community².

Is using an automatic translator cheating?

This is a difficult question to answer. When consulting discussion forums on the internet³ you can easily be convinced that it is cheating. Students are told not to use these tools. They are told that if they don't comply, they will be accused of cheating. But arguments can also be presented the other way. Education is about teaching people to use tools smartly in order to perform tasks. So how about making it possible for a pupil to learn to use the tools they will find available outside school?

This textbook is not authorised to offer a definitive answer, but we do suggest that teachers explore ways in which these tools can be used to learn languages.

What should a teacher do about this?

Florencia Henshaw discusses a number of options⁴, none of which seem convincing:

- Saying that AI simply doesn't work (a favourite conclusion in forums³) is not helpful, even if pupils agree with this. They will still want to use these AI tools.
- The zero-tolerance approach relies on being able to detect the use of AI. This

may be the case today¹ but it is uncertain if it will remain the case. Furthermore, is using AI cheating? In what way is it different from using glasses to read better or a wheelbarrow to transport objects?

- The approach in which the tool can be part-used (to search for individual words for example) is also criticised⁴. Automatic translation tools work because they make use of the context. On individual (out of context) words, they will not perform any better than a dictionary.
- The approach of using the tool in an intelligent way, in and out of the classroom, is tempting but will need more work to develop activities which will be of real help in learning situations.

¹ Birdsell, B. J., *Student Writings with DeepL: Teacher Evaluations and Implications for Teaching*, JALT2021 Reflections & new perspectives 2021.

² Calixte, L, November 2022, https://etudiant.lefigaro.fr/article/quand-l-intelligence-artificielle-facilite-la-fraude-universitaire_463c8b8c-5459-11ed-9fee-7d1d86f23c33/.

³ Reddit discussion on automatic translation and cheating. https://www.reddit.com/r/Professors/comments/plcjiu/foreign_language_teachers_how_do_you_deal_with/.

⁴ Henshaw, F. *Online Translators in Language Classes: Pedagogical and Practical Considerations*, The FLT MAG, 2020, <https://fltmag.com/online-translators-pedagogical-practical-considerations/>.

⁵ Cribb, V. M. (2000). *Machine translation: The alternative for the 21st century?* TESOL Quarterly, 34(3), 560-569. <https://doi.org/10.2307/3587744>.

28.

MANUEL GENTILE AND GIUSEPPE CITTÀ

We have long been used to writing on computers with dedicated software that goes by the name of word processors (e.g., Microsoft Word, Google Docs, Pages, LibreOffice), taking advantage of the grammar suggestions provided by these tools. Raise your hand if you have never been saved by these tools from making glaring mistakes 💎

But the transformation induced by these tools is not limited to correcting some typos; it has been far more profound in invoking a different way of writing. Digital writing allows us to return to what we have written and modify it to express more effectively what we want to convey.

Using somewhat more technical words, we have moved from a linear approach to writing, to an iterative process. According to recent studies, the transformation of the writing process induced by digital tools has improved the quality of the texts produced.

Writing in the AI era

Anyway, the process of evolution of writing and related thought forms has not stopped. In recent years, with the explosion of AI, it has accelerated significantly. Tools such as Grammarly, Wordtune, Ludwig, ProWritingAid, and so on, are designed not only to provide grammatical correction of the text. They support the user throughout the writing process by stimulating the improvement of writing style, checking for plagiarism, and more.

Recognising how the school world cannot be immune to such innovations is trivial. This is confirmed by the growing number of educational interventions, proposed in literature, that are designed to take advantage of such software. Some scholars propose using these tools to work on students' skills in using external information sources to develop appropriate paraphrasing skills that can avoid plagiarism problems. Many of these tools can support the teacher in evaluating the texts produced by the students, providing timely analysis of individual student's strengths and weaknesses. Moreover, these tools allow the student themselves to self-assess their own writing skills, thus enabling metacognitive processes and speeding up learning.

All that glitters is not gold...

Clearly, these innovations have potential problems. First, you have probably understood how, underlying all these deep learning mechanisms, is the source data

on which the models are built. Limited or incorrect training data could cause bias. In addition, the risk of a general homogenisation of the texts produced/expected by those tools is likely. It could determine a consequent limitation (or penalisation in the case of assessment) of students' creativity. Finally, these tools are primarily limited to managing the English language; thus, non-English-speaking contexts can be used in the L2 domain. That said, the speed of innovation is such that we will soon see similar tools emerge for languages other than English.

A look to the future

One of the main cognitive processes related to the writing process is the retrieval from long-term memory of those pieces of information necessary to complete the message we want to express. It is easy to speculate how these tools will also support this process by enabling immediate and simplified access to a 'memory' far more extensive than our own.

Finally, the tremendous advances in text-generative processes point to a future in which these tools can support the writing process in a far more active form.

How we write text will probably change further in ways we cannot yet imagine. However, the challenge will always remain the same – to know how to consciously use the tools at our disposal and adapt our way of teaching accordingly. Are you ready?

29.

Machine learning goes deep

Human knowledge is wide and variable and is inherently difficult to capture. The human mind can absorb and work with knowledge because it is, as Chomsky put it, “a surprisingly efficient and even elegant system that operates with small amounts of information; it seeks not to infer brute correlations among data points but to create explanations¹.”


Machine learning is supposed to do it by finding patterns in large amounts of data. But, prior to that, experts and programmers had to sit and code what features of the data are relevant to the problem at hand, and feed these to the machine as “parameters”^{2,3}. As we saw before, the performance of the system depends heavily on the quality of the data and parameters, which are not always straightforward to pinpoint.

Deep neural networks or deep learning is a branch of machine learning that is designed to overcome this by:


- Extracting its own parameters from data during the training phase;
- Using multiple layers which construct relationships between the parameters, going progressively from simple representations in the outermost layer to more complex and abstract. This enables it to do certain things better than conventional ML algorithms².

Increasingly, most of the powerful ML applications use deep learning. These include search engines, recommendation systems, speech transcription and translation that we have covered in this book. It will not be a stretch to say that deep learning has propelled the success of artificial intelligence in multiple tasks.

Rule based systems



Classic Machine Learning



Deep Learning



Reference : Goodfellow, I.J., Bengio, Y., Courville, A., *Deep Learning*, MIT Press, 2016.

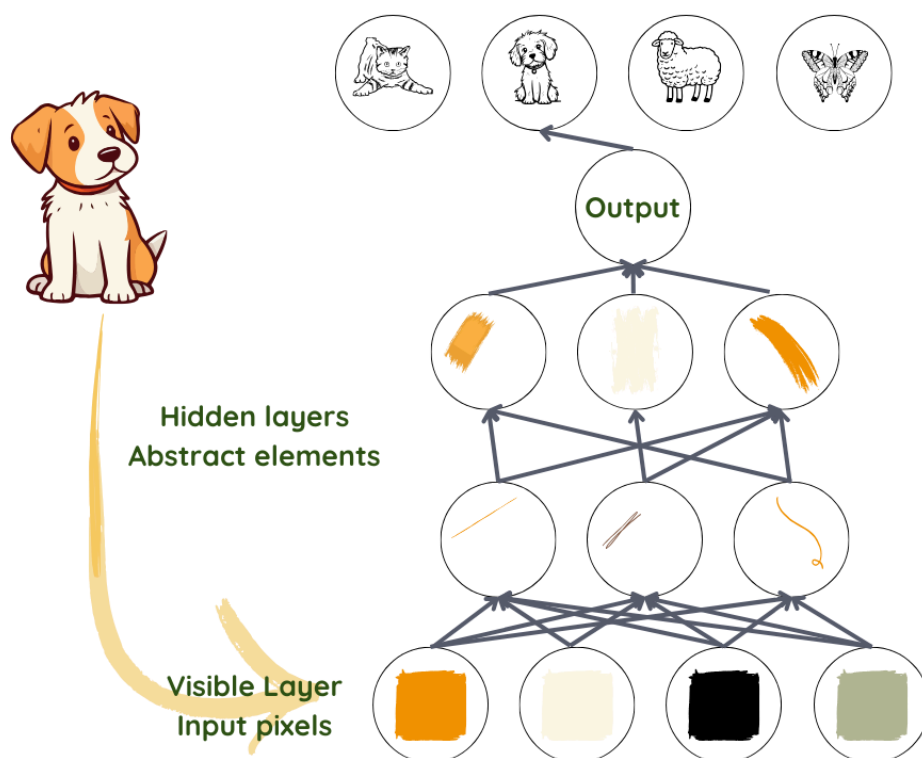
“Deep” refers to how the layers pile on top of each other to create the network. “Neural” reflects the fact that some aspects of the design were inspired by the biological brain. Despite that, and even though they provide some insights into our own thought processes, these are strictly mathematical models and do not resemble any biological parts or processes².

The basics of deep learning

When humans look at a picture, we automatically identify objects and faces. But a photo is just a collection of pixels for an algorithm. Going from a jumble of colours and brightness levels, to recognising a face, is a leap too complicated to execute.

Deep learning achieves this by breaking the process into simple representations in the first layer – by, say, comparing brightness of neighbouring pixels to note the presence or absence of edges in various regions of the image. The second layer takes collections of edges to search for more complex entities – such as corners and contours, ignoring small variations in edge positions^{2,3}. The next layer looks for parts of the objects using the contours and corners. Slowly, the complexity builds till the

point where the last layer can combine different parts well enough to recognise a face or identify an object.



What to take into account in each layer is not specified by programmers but is learned from data in the training process³. By testing these predictions with the real outputs in the training dataset, the functioning of each layer is tuned in a slightly different way to get a better result each time. When done correctly, and provided there is sufficient good-quality data, the network should evolve to ignore irrelevant parts of the photo, like exact location of the entities, angle and lighting, and zero in on those parts which make recognition possible.

Of note here is the fact that, despite our use of edges and contours to understand the process, what is actually represented in the layers is a set of numbers, which might or might not correspond to things that we understand. What doesn't change is the increasing abstractness and complexity.

Designing the network

Once the programmer decides to use deep learning for a task and [prepares the data](#), they have to design what is called the architecture of their neural network. They have to choose the number of layers (depth of the network) and the number

of parameters per layer (width of the network). Next, they have to decide how to make connections between the layers – whether or not each unit of a layer will be connected to every unit of the previous layer.

The ideal architecture for a given task is often found by experimentation. The greater the number of layers, the fewer the parameters that are needed per layer and the network performs better with general data, at the cost of it being difficult to optimise. Fewer connections would mean fewer parameters, and lesser amount of computation, but reduces the flexibility of the network².

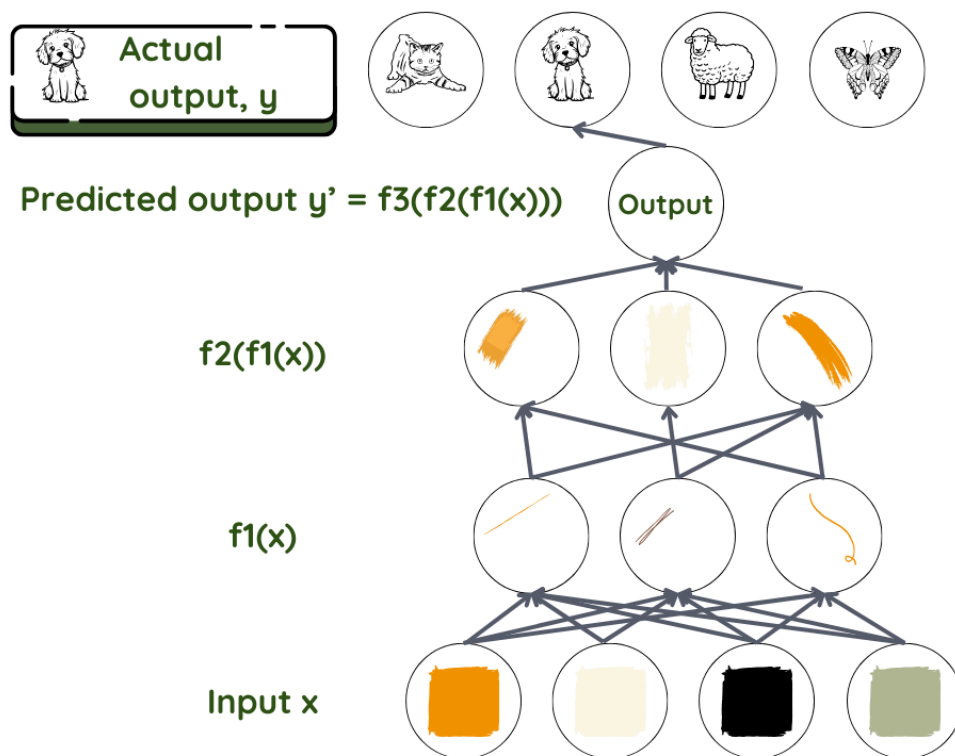
Training the network

Let us take the example of a feed forward neural network doing supervised learning. Here, information flows forward from layer to deeper layer, with no feedback loops. As for all machine-learning techniques, the goal here is to find out how the input is connected to the output – what parameters come together, and how they come together to give the observed result. We assume a relationship f that connects the input x to the output y . We then use the network to find the set of parameters θ that give the best match for predicted and actual outputs.

Key question: Predicted y is $f(x, \theta)$, for which θ ?

Here the prediction for y is the final product and the dataset x is the input. In face recognition, x is usually the set of pixels in an image. y can be the name of the person. In the network, the layers are like workers in an assembly line, where each worker works on what is given to them and passes it forward to the next worker. The first one takes the input and transforms it a little bit and gives it to the second in line. The second does the same before passing it to the third, and so on until the input is transformed into the final product.


Mathematically, the function f is split into many functions $f_1, f_2, f_3...$ where $f = \dots f_3(f_2(f_1(x)))$. The layer next to the input transforms input parameters using f_1 , the next layer using f_2 , and so on. The programmer might intervene to help choose the correct family of functions based on their knowledge of the problem.



It is the work of each layer to assign the level of importance – the weight given to each parameter that it receives. These weights are like knobs that ultimately define the relationship between the predicted output and input in that layer³. In a typical deep-learning system, we are looking at hundreds of millions of these knobs and hundreds of millions of training examples. Since we neither define nor can see the output and weights in the layers between input and output, these are called hidden layers.

In the case of the object recognition example discussed above, it is the work of the first worker to detect edges and pass on the edges to the second one who detects contours and so on.

During training, the predicted output is taken and compared with the real output. If there is a big difference between the two, the weights assigned in each layer will have to be changed by a lot. If not, they have to be changed a little. This work is done in two parts. First the difference between prediction and output is calculated. Then another algorithm computes how to change the weights in each layer, starting from the output layer (in this case, the information flows backwards from the deeper layers). Thus at the end of the training process, the network is ready with its weights and functions to attack test data. The rest of the process is the same as that of conventional machine learning.



¹ Chomsky, N., Roberts, I., Watumull, J., *Noam Chomsky: The False Promise of ChatGPT*, The New York Times, 2023.

² Goodfellow, I.J., Bengio, Y., Courville, A., *Deep Learning*, MIT Press, 2016.

³ LeCun, Y., Bengio, Y., Hinton, G., *Deep learning*, Nature **521**, 436–444 (2015).

30.

Natural language processing has been a topic on which research has worked in length for the past 50 years. This has led to the development of many tools we use every day:

- Word processors
- Automatic grammar and orthography correction
- Automatic completion
- [Optical character recognition \(OCR\)](#)

More recently, chatbots, home assistants and automatic translation tools have been making a huge impact in all areas.



“Learning sign language” by daveynin is licenced under CC BY 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by/2.0/?ref=openverse>.

For a long time, research and industry was stalled by the intrinsic complexity of language. At the end of the 20th century, grammars for a language, written by experts, could have up to 50,000 rules. These **expert systems** were showing that technology could make a difference, but robust solutions were too complex to develop.

On the other hand, **speech recognition** needed to be able to make use of acoustic data and transform it into text. With the variety of speakers one could find, a hard task indeed!

Researchers understood that if we had a model for the intended language, things would be easier. If we knew the words of the language, how sentences were formed, then it would be easier to find the right sentence from a set of candidates to match a given utterance, or to produce a valid translation from a set of possible sequences of words.

Another crucial aspect has been that of **semantics**. Most of the work we can do to solve linguistic questions is shallow; the algorithms will produce an answer based on some local syntactic rules. If in the end, the text means nothing, so be it. A similar thing may happen when we read a text by some pupils – we can correct the mistakes without really understanding what the text is about! A real challenge is to associate meaning to text and, when possible, to uttered sentences.



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=155#oembed-3>

There was a surprising result in 2008¹. One unique **language model** could be learnt from a large amount of data and used for a variety of linguistic tasks. In fact, that unique model performed better than models trained for specific tasks.

The model was a [deep neural network](#). Nowhere as deep as the models used today! But enough to convince research and industry that machine learning, and more specifically Deep learning was going to be the answer to many questions in NLP.

Since then, natural language processing has ceased to follow a model-driven approach and has been nearly always based on a data-driven approach.

Traditionally, the main language tasks can be decomposed into 2 families – those involving building models and those involving decoding.

Building models

In order to transcribe, answer questions, generate dialogues or translate, you need to be able to know whether or not “Je parle Français” is indeed a sentence in French. And as with spoken languages, rules of grammar are not always followed accurately, so the answer has to be probabilistic. A sentence can be *more or less* French. This allows the system to produce different candidate sentences (as the transcription of a sound, or the translation of a sentence) and the probability is given as a score associated with each candidate. We can take the highest-ranking sentence or combine the score with other sources of information (we may also be interested in what the sentence is about).

Language models do this, and the probabilities are built from machine-learning algorithms. And of course, the more data there is, the better. For some languages there is a lot of data from which to build language models. For others, this is not the case; these are under-resourced languages.


In the case of translation, we want not two but three models: a language model for each language and another model for the translations, informing us of what the better translations of fragments of language can be. These are difficult to produce when data is scarce. If models for common language pairs are easier to build, this will not be the case for languages that are not frequently spoken together (such as Portuguese and Slovene). A typical way out of this is to use a *pivot language* (typically English) and translate via this pivot language – from Portuguese to English and then from English to Slovene. This will lead to inferior results as errors accumulate.

Decoding

Decoding is the process in which an algorithm takes the input sequence (which can be signal or text) and, by consulting the models, makes a decision, which will often be an output text. There are here some algorithmic considerations – in many cases transcription and translation are to happen in real time and diminishing the time lag is a key issue. So there is room for a lot of artificial intelligence.

End-to-end

Nowadays, the approach of building these components separately and combining them later has been replaced by *end-to-end approaches* in which the system will transcribe/translate/interpret the input through a unique model. Currently, such models are trained by deep neural networks, which can be huge. It is reported that the current largest GPT3 model comprises several hundred million parameters!



Chervinskii, Wikimedia Commons is licenced under CC BY-SA 4.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by/2.0/?ref=openverse>.

Let's try to get at the intuition behind this process. Suppose we have some data. This raw data can be encoded in some way. But the encoding can be redundant, and perhaps even expensive. Let us now build a particular machine, called an auto-encoder (see diagram to the left). This machine will be able to take a text, compress it into a small vector (this is the encoder), and then uncompress the vector (the decoder part) and return a text which is somehow close to the original text. The idea is that this mechanism will make the intermediate vector meaningful, with two desirable properties – a reasonably small vector which ‘contains’

the information in the initial text.

The future

An example of an end-to-end process that might be available soon will be the ability to perform the following task – it will hear you speak your language, transcribe your text, translate it to a language you don't know, train a speech synthesis system to your voice, and have your own voice speak the corresponding text in a new sentence. Here are two examples produced by researchers at the Universidad Politecnica de Valencia, Spain, in which the speaker's own voice model is used to do the dubbing.



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=155#oembed-1>



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=155#oembed-2>

Some consequences for education

The steady progress of natural language processing is remarkable. Where we would laugh at the translations proposed by AI just ten years ago, it becomes increasingly difficult to find such obvious errors today. Speech recognition and character-recognition techniques are also improving fast.

The semantics challenges are still there, and answering questions which require a deep understanding of a text is still not quite right. But things are going in the right direction. This means that the teacher should expect some of the following statements to be true soon, if they are not already out there!

- A pupil will take a complex text and obtain (with AI) a simplified version; the text may even be personalised and use terms, words and concepts the pupil is used to;
- A pupil will be able to find a text, copy it, and obtain a text stating the same things but undetectable by an anti-plagiarism tool;
- Videos produced anywhere in the world will be accessible through automatic dubbing in any language. This means that our pupils will be exposed to learning material built in our language and also by material initially designed for another learning system in a different culture;
- Writing essays could become a task of the past, as tools will enable writing on any topic.

It is clear that AI will be far from perfect, and the expert will detect that even if the language is correct, the flow of ideas won't be. But let's face it, during the course of education, how long does it take for our pupils and students to reach that level?

¹ Collobert, Ronan, and Jason Weston. *A unified architecture for natural language processing: Deep neural networks with multitask learning*. Proceedings of the 25th international conference on Machine learning. 2008.
<http://machinelearning.org/archive/icml2008/papers/391.pdf>. Note: this reference is given for historical reasons. But it is difficult to read!

31.

WAYNE HOLMES

Artificial Intelligence (AI) has often been accused of threatening human agency (for example, see the summary of 979 “expert” views reported in [Artificial intelligence and the future of humans](#)¹). This is particularly the case for the application of AI in education (AIED), including in my own research. For example, in a [report for the Council of Europe](#), we write that the approach adopted by almost all AIED tools “prioritises remembering over thinking, and knowing facts over critical engagement, thus undermining learner agency and robust learning”². However, while such claims are easy to make, and I stand by this one, rarely are they properly unpacked. Accordingly, that is the intention of this brief exposition: to explore what exactly is meant by ‘human agency’, and to consider the impact of both AI and AIED.

Broadly speaking, human agency is the capacity of individuals to act independently, making choices from among different options based on their own beliefs, values and goals. In other words, it is the ability of humans to make decisions and initiate and carry out actions that influence their lives and the world around them. It can involve multiple dimensions. These include intentionality (acting consciously with a specific aim or goal in mind), autonomy (independence, self-determination and the freedom to make choices and decisions that reflect one’s own preferences, values and goals), adaptability (the capacity to learn, modify one’s behaviour and succeed in response to changing circumstances), and responsibility (the ethical and moral dimension of agency – being accountable for the consequences of one’s decisions and actions).

Human agency is crucial for personal growth and a successful life. It empowers individuals to shape their own lives and influence the world around them. It fosters a sense of control and self-efficacy, and is correlated with higher levels of psychological well-being. When individuals feel that they have control over their lives and can make meaningful choices, in other words when they feel they have genuine agency, they are more likely to experience satisfaction and fulfilment. As one of the leading researchers of human agency explains: “Unless people believe they can produce desired effects by their actions, they have little incentive to act or to persevere in the face of difficulties.”³

In educational contexts, human agency refers to the capacity of students and teachers to make choices, to act autonomously, and to take control of teaching and learning within classroom settings. The emphasis is on the role of individuals in shaping educational journeys, making decisions about what, how, and why something is learned. Human agency in educational contexts involves multiple considerations. For example, student agency can be enhanced when they are not treated as passive recipients of knowledge but as active participants in the learning process, and have the autonomy to explore topics of interest, ask questions, identify and set their own academic goals, and take ownership of their learning. Enhancing student agency also involves promoting problem-solving and critical-thinking skills (centred on real-world problems) along with self-regulation skills (such as managing time, setting priorities, and self-monitoring progress). All of this is valuable for the

development of independent and self-directed individuals and for academic and lifetime success. Finally, while teachers play a central role in supporting and enhancing student agency, human agency in educational contexts also includes the agency of the teachers themselves, acknowledging their subject and pedagogic expertise and professionalism, enabling them to make choices on how best to conduct their teaching and to support their students.

The next question is, what is the impact of AI on human agency? Inevitably, each potential impact has both a positive and a negative spin. For example, some AI-enabled technologies can undertake repetitive tasks, possibly freeing people to focus on the more creative aspects of their work, perhaps allowing them to decide how best to allocate their own time and effort, thus enhancing their agency. On the other hand, using AI-enabled technologies to undertake common tasks may lead to the loss of human skills or expertise. Over time, as they become increasingly reliant on AI, this could instead reduce the choices available to individuals and thus weaken their agency. Similarly, it is often argued that AI-enabled technologies can personalise user interactions with various services (such as online video and shopping platforms), providing suggestions tailored to their preferences, which strengthens their sense of agency. However, looked at with a more critical eye, the personalisation provided by such services is usually tailored more to the needs of the provider and their advertisers than the user. This means that the individual's sense of agency masks the reality – which is a reduction in individual agency, as the user is nudged in particular directions. In addition, AI-enabled data analyses can provide access to valuable insights that might not be readily available otherwise, thus enhancing human decision making and agency. However, it is well known that AI systems inherit and perpetuate biases present in their training data. This can lead to unfair and discriminatory outcomes, which inevitably undermines human agency by limiting opportunities. AI-enabled technologies, or at least the way in which they are employed in practice, can have other negative impacts on human agency. For example, the widespread use of AI for monitoring (or surveillance) and AI-controlled decision-making, raises important privacy concerns, limits choices of action, and can lead to a sense of powerlessness or dependence on technology, all of which can undermine individual agency.

The next question is, what is the impact of AIED on student and teacher agency? The possibilities are many. First, if students engage frequently with AI-enabled technologies, they might all too easily become over-reliant on the content recommendations, the instant feedback, or the “solutions” that are provided. Accordingly, they might miss out on opportunities to develop critical thinking, independent problem-solving, self-reflection, self-regulation and metacognitive skills, all of which potentially reduces the students' agency to realise the full benefits of their own learning. Second, most AIED systems provide highly prescriptive learning paths, leaving little room for students to explore their own interests. This can limit students' agency by dictating what, when and how they learn, and can also reduce their exposure to diverse perspectives and novel areas of inquiry. Third, AIED systems typically track student behaviour, leading to a feeling of surveillance and constrained autonomy, as well as opening the possibility of undermining student privacy, causing students to become cautious about expressing themselves freely. Fourth, AI-driven recommendations may inadvertently narrow students' aspirations, limiting their agency to achieve self-determined goals. Fifth, AI-enabled

technologies used for assessment involve an overemphasis on standardised testing (and thus, inevitably, teaching to the test). In addition, no AI-enabled system is capable of understanding or capturing the nuances in a students' work, thus reducing the student's agency in the assessment process and possibly discouraging creative or unconventional thinking.

Finally, with regard to teachers, the use of AI-enabled technologies in classrooms inevitably affects curriculum choices, learning content, and pedagogical approaches. This can diminish the role of teachers and lead them to feel their professional judgement is undervalued or overridden by technology. In any case, AIED tools can potentially de-skill teachers, turning them into technology facilitators and behaviour monitors. This radically misunderstands what good teachers do. It might also undermine teachers' agency in building meaningful relationships with their students, which are crucial for effective education. Finally, reliance on AI-generated metrics (sometimes because of top-down directives) can create pressure on teachers to conform to data-driven, decision-making processes, which might lead to less emphasis on holistic student development.

The final question here is, what needs to be done to ensure student and teacher agency, as powerful AI-enabled technologies become increasingly available in classrooms. In short, teachers need to be afforded opportunities that respect their agency, allowing them to make decisions that align with their professional expertise and the specific needs of their students. Meanwhile, students need opportunities to develop their critical thinking, self-regulation and metacognitive skills, and to develop their intentionality, their autonomy, their adaptability, and their responsibility – either with or without the use of appropriate, effective and safe AI-enabled technologies.

¹ Anderson et al., *Artificial intelligence and the future of humans*, Pew Research Center, 2018

² Holmes et al., *Artificial intelligence and Education, A critical view through the lens of human rights, democracy and the rule of law*, Council of Europe, p. 34, 2022

³ Bandura, A., *Toward a Psychology of Human Agency: Pathways and Reflections*, *Perspectives on Psychological Science*, 13(2), 130-136, 2018

32.

Homogenisation

A lot of money, computing resources, time and effort go into the creation of datasets, benchmarks and algorithms for machine learning. This is particularly true for deep learning and large-scale models. Therefore, it makes sense that the created resources are shared within this ecosystem. This is the case with many of the ML systems that we often use. Even where the end products are different and are created by a different company, methodology, datasets, machine-learning libraries and evaluations are often shared¹. Thus, there is an argument to be made for their outputs to be similar under similar conditions.

If the output is an educational decision, it raises concern, for example, for the student who might be rejected unfairly from every educational opportunity¹. But whether or not algorithmic homogenisation constitutes an injustice can only be decided on a case-by-case basis¹.

On the other hand, if the system's task is to help the student write, it brings into focus standardisation of writing styles, vocabulary and hence, thought patterns. Language models used in these cases are decided to predict the most probable text, based on their training dataset. These datasets, if not shared between systems, are constructed in a similar way, often with public internet data. Even when this data is screened for bias, prejudice and extreme content, it represents only a small ecosystem and is not representative of the world in all its diversity of ideas, culture and practices. Predictive text systems, based on deep learning and used for text messages and emails, have been demonstrated to change the way people write. The writing tends to become "more succinct, more predictable and less colourful"².

Invisibility

Sequences of words which are repeated in the training data trickle down into the output of large language models. Thus, the values of the database creators get the power to curb alternative opinions and plural expressions of ideas. Without proper pedagogical interventions, this in turn might limit students' creativity and originality, not only leading to formulaic writing but ultimately to citizens with lesser critical-thinking skills and thus to an overall less colourful world³.

Closely linked with many of the negative fallouts of machine learning, including homogeneity as discussed above, is the fact that the technologies have become so advanced that the human-machine interface is seamless and practically invisible. Whether it is the search engines incorporated in the browser's address bar, or text prediction that works intuitively, with no time lag between writing, predicting and choosing suggestions, we often act under the influence of technology without consciously being aware of it or having the choice to put a brake and rethink

situations and make our own decisions. Moreover, when we use it habitually to make decisions, we tend to forget its existence altogether⁴. “Once we are habituated to technologies, we stop looking at them, and instead look through them to the information and activities we use them to facilitate”. This raises such serious concerns about human agency, transparency and trust, especially when it comes to young minds, that experts have recommended interfaces be made more visible and even unwieldy⁴.

What’s beyond: an ethical AI

In each part of this open textbook, we have discussed pedagogical, ethical and societal impacts of AI, especially data-based AI. [Data and Privacy, reliability of content and user autonomy](#), [impact on personal identity](#), [Bias and Fairness](#) and [Human agency](#) were all discussed in their respective pages. Issues specific to search engines were discussed in [Behind the Search Lens: Effects of search on the individual](#) and [the society](#), problems relating to adaptive systems were dealt with in [The Flip Side of ALS](#) and those particular to Generative AI in [The Degenerative](#). In several places throughout the book, we looked at remedial measures that can be taken in the classroom to deal with specific problems. Our hope is that these measures will become less onerous once we have ethical and reliable AI systems for education. This ethical AI would be developed, deployed and used in compliance with ethical norms and principles⁵ and would be accountable and resilient.

Since we cede so much power to AI models and their programmers, sellers and evaluators, it is only reasonable to ask them to be transparent and assume responsibility and remedy errors when things go wrong⁶. We need service-level agreements that clearly outline “the support and maintenance services and steps to be taken to address reported problems”⁵.

A resilient AI would accept its imperfections, would expect them and can work in spite of them. Resilient AI systems would fail in a predictable way and have protocols to deal with these failures⁶.

In education, ethical AI should be guided by user-centred design principles and take into account all aspects of education⁷. Teachers would be able to inspect how it functions, understand its explanations, override its decisions or pause its use without difficulty⁸. These systems would reduce teacher workload, give them detailed insights into their students and support them in enhancing the reach and quality of education⁸. They would not cause harm to their users or the environment and would enhance the social and emotional well-being of learners and teachers⁵.

Until this day comes, a teacher will have to try to develop and participate in a community of colleagues and educators to raise awareness about problems, share experiences and best practices, and identify reliable providers of AI. They could also involve students and parents in discussions and decisions to better address different concerns and develop an environment of trust and camaraderie. They would be best served by doing their best to stay up to date with the latest trends in AIED and acquire competencies when and where possible⁵.

- ¹ Bommasani, R., et al, *Picking on the Same Person: Does Algorithmic Monoculture lead to Outcome Homogenization?*, Advances in Neural Information Processing Systems, 2022.
- ² Varshney, L., *Respect for Human Autonomy in Recommender System*, 3rd FAccTRec Workshop on Responsible Recommendation, 2020.
- ³ Holmes, W., Miao, F., *Guidance for generative AI in education and research*, UNESCO, Paris, 2023.
- ⁴ Susser, D., *Invisible Influence: Artificial Intelligence and the Ethics of Adaptive Choice Architectures*, Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, Association for Computing Machinery, New York, 403–408, 2019.
- ⁵ [Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators](#), European Commission, October 2022.
- ⁶ Schneier, B., *Data and Goliath: The Hidden Battles to Capture Your Data and Control Your World*, W. W. Norton & Company, 2015.
- ⁷ Tlili, A., Shehata, B., Adarkwah, M.A. et al, *What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education*, Smart Learning Environments, 10, 15 2023.
- ⁸ U.S. Department of Education, Office of Educational Technology, *Artificial Intelligence and Future of Teaching and Learning: Insights and Recommendations*, Washington, DC, 2023.

PART VI

ON GENERATIVE AI

Some educators have been dreaming of technology that would solve all the problems of education. Instead, what has come true is the student dream of solving all assignment and exam problems.

Not a day goes by without a new article, a podcast or a fresh set of policy recommendations on how to deal with generative AI. Schools and school districts talk of ChatGPT-proofing their classrooms. Unions are on strike to shock-proof their jobs, the future jobs of the same gleeful students, against this supposedly chatting, painting and conjuring ‘wonder’ that is generative AI. Yes, all this while some experts are warning about the end of the world.

We hope this chapter helps you make sense of this tumultuous technology, its possible advantages and shortcomings.

33.

MICHAEL HALLISSY AND JOHN HURLEY

History

In late November 2022, the world was introduced to ChatGPT, an Artificial Intelligence (AI)-based chatbot system based on a [language model](#) called GPT-3.5. It uses natural language processing (NLP) to generate conversations¹. ChatGPT was the latest in a series of such tools but, unlike previous tools, it has captured public interest and imagination. It had more than a million users within a week of its launch because of its ability to generate human-like text and its perceived implications and potential use in education, the workplace and everyday life. ChatGPT can answer questions and assist you with tasks such as composing emails, essays and code².

GPT-3 (Generative Pre-trained Transformer-3) is a Large Language Model (LLM) that was trained through [deep learning](#) on a vast quantity of data (499 billion data points – 800GB of data), a model some one hundred times larger than any previous models³. ChatGPT enables human-like conversations with this model. A large language model is trained to predict the next word in a sentence, somewhat in the same way that a text-message auto-complete-feature works. It continues this process to generate more text.



The Generative AI timeline

Since its arrival there have been many developments in relation to the creation and development of generative AI tools. We have captured some of these in the graphic above. Prior to November 2022, few people knew what generative AI was. All of a sudden, terms such as ChatGPT became everyday terms for many – particularly those in education. The indications are that these technologies will continue to develop and become more useful to all in education.

Generative AI – What are they?

Generative AI systems create new content in the form of images, text, audio, video, and more:

- Images: generative AI can create new images based on existing ones, such as creating a new portrait based on a person's face, or a new landscape based on existing scenery;
- Text: generative AI can be used to write news articles, poetry and even scripts. It can translate text from one language to another;
- Audio: generative AI can generate new music tracks, sound effects and voice acting⁴.

The list of possibilities continues to grow, but let's start by asking some generative AI tools how ChatGPT got its name.

But what is in a name?

Generative AI tools have multiple uses, so let's give a number of tools – Bing Chat, ChatGPT 3.5 and Google Bard the following prompt:

You are an expert in AI in Education. Imagine you are explaining what ChatGPT is and how it got its name to young children, aged between 10 to 14. Explain why it includes the term chat and show some examples of how it can be used in education. Write your output as a persuasive piece of writing that explains ChatGPT

Here is what Bing Chat created:

The name "ChatGPT" comes from the words "chat" and "GPT." "Chat" means to talk with someone in a friendly and informal way. "GPT" stands for "Generative Pre-trained Transformer," which is a fancy way of saying that ChatGPT is really good at creating new things like text, images, and sounds.

Here is what ChatGPT created:

Now, you might wonder why it's called ChatGPT. Well, the 'Chat' part is because it loves to have conversations with people. You can type in a question or a topic, and it will do its best to give you an answer that makes sense. It's kind of like talking to a friend who knows everything about, well, almost everything! Plus, it's really good at chatting in different languages too, which makes it super helpful for kids all over the world.

The 'GPT' part is a bit more complicated. It stands for "Generative Pre-trained Transformer," but you don't need to remember that tongue-twister. Basically, it means ChatGPT learned a lot by reading tons of books, articles, and websites, just like you do when you're studying. It's like it went to school for a super long time and became an expert in all sorts of subjects!

Here is what Bard created:

ChatGPT is a type of artificial intelligence (AI) called a large language model. This means that it's been trained on a massive dataset of text and code, which allows it to generate text, translate languages, write different kinds of creative content, and answer your questions in an informative way.

So, why does ChatGPT include the term "chat"? Well, because it's designed to have conversations with people! You can ask ChatGPT questions about anything, and it will do its best to answer them in a comprehensive and helpful way. You can even have fun with ChatGPT by asking it to write stories, poems, or even code.

These tools allow us to write prompts and then have a conversation (conversational AI) with a ChatBot. There are a growing list of ways, which we touch on in the textbook, how these tools can help teachers, students and administrators in education and training. And they allow you to have a conversation, so think of the ways that might help you or your learners – you can have a chat about a topic you need to research or to prepare for an interview or to generate ideas for a lesson or a paper. It is interesting to compare and contrast the outputs they each create and then pick the best pieces for your purpose.

These tools can 'enhance' or assist you to be more creative, productive and effective when you combine your human abilities with the abilities of the machine. So always remember to critically evaluate what it produces. So have some fun


Find out!

Have you heard
[the debate on Gen AI?](#)

learning how to have conversations with these tools and see what they produce and if they are useful to you and your students.

Generative AI in Numbers⁵

The following graph⁶ shows the rise of ChatGPT between November 2022 and May 2023. It gained one million users in its first week after it was launched. It should be noted that Netflix took 3.5 years to reach a similar milestone, with Twitter taking two years⁷.



CHATGPT statistics by tooltester is licenced under CC BY 4.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by/4.0/><https://creativecommons.org/licenses/by/2.0/?ref=openverse>.

While ChatGPT reached its peak in May 2023, it continues to retain its popularity with other generative AI tools such as Bing Chat and Bard also gaining traction. In the case of Bing Chat it started to become more popular when Microsoft acquired ChatGPT in March 2023, while Bard continues to grow in popularity as well.

- ¹ <https://www.forbes.com/sites/bernardmarr/2022/12/21/chatgpt-everything-you-really-need-to-know-in-simple-terms/?sh=16674aadcbca>
- ² <https://www.zdnet.com/article/what-is-chatgpt-and-why-does-it-matter-heres-everything-you-need-to-know/>
- ³ <https://hai.stanford.edu/news/how-large-language-models-will-transform-science-society-and-ai>
- ⁴ <https://www.weforum.org/agenda/2023/02/generative-ai-explain-algorithms-work/>
- ⁵ <https://research.aimultiple.com/generative-ai-applications/>
- ⁶ <https://www.tooltester.com/en/blog/chatgpt-statistics/>
- ⁷ <https://bootcamp.uxdesign.cc/chatgpt-vs-bing-chat-which-is-better-2e46fa821d7d>

34.

What do you do when someone suggests that you try out an activity using a generative AI tool? What do you do when you have an idea for an activity but are unsure about how to adapt it?

Find out!

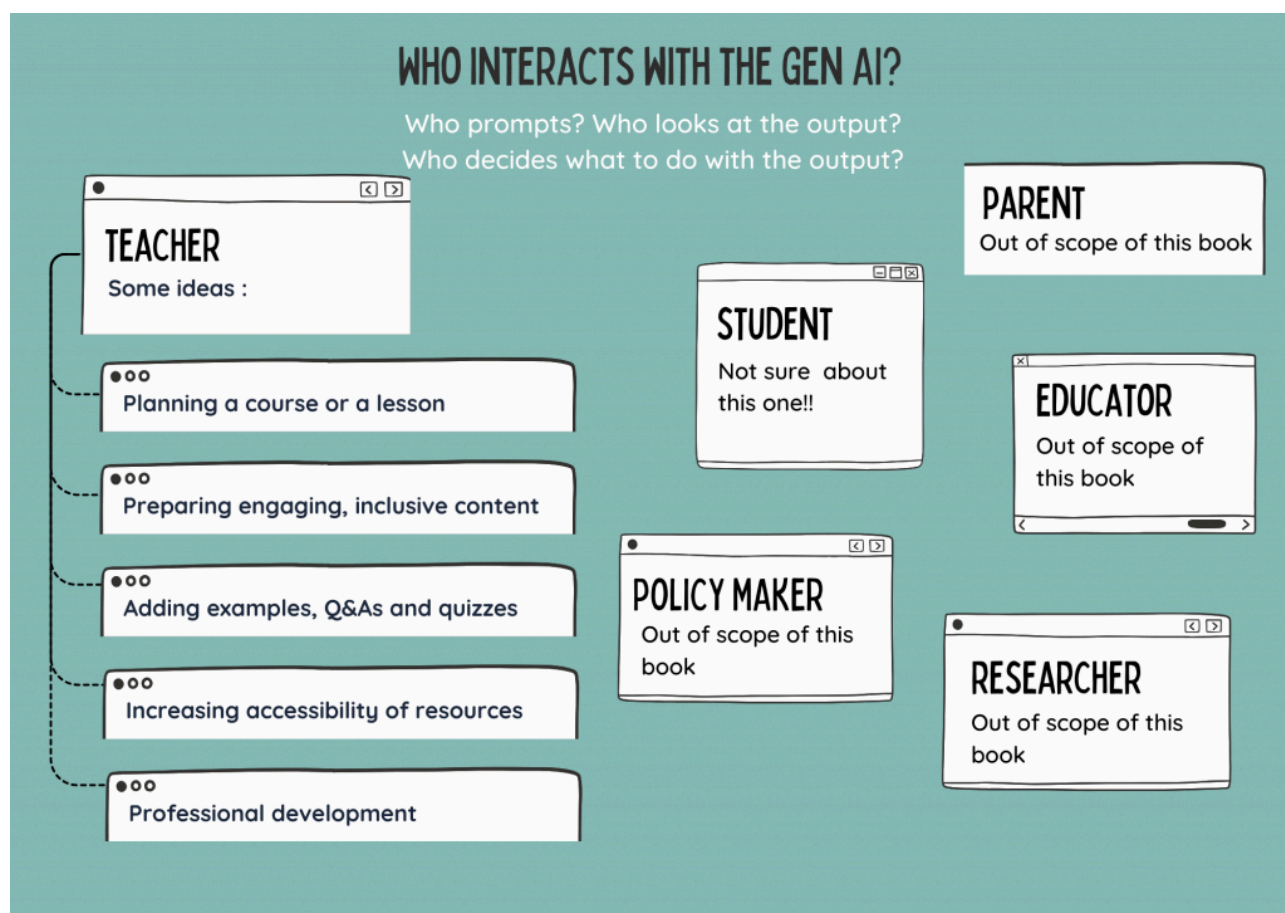
[Read why we left the student out, for the moment](#)

Here, the goal is to assist you in evaluating, choosing and carrying out such activities.

Due to the speed of evolution of this technology and given how performances change with subjects and contexts, all the questions a teacher need to ask cannot be listed in advance. We hope only to get you started in lines of thinking recommended by published guidelines and supported by research and case studies.

Please note that in this book we look only at those activities where the teacher, not the student, interacts with generative AI. Our reasons behind this are explained in the link provided

in the box.



Even if you are not convinced that Gen AI tools should have a place in education, we

urge you to try at least one activity that uses generative technology. It might help you:

- Understand what someone can and cannot do with the tools provided by generative AI;
- Reformulate or rearrange course elements, especially when deciding what can and cannot be given for homework. For example, ChatGPT can be effective in helping students write essays¹ and it might be futile to assign standard essays for homework. You may have to use a variety of assessment methods such as group projects, hands-on activities and oral presentations, and go beyond simple yes or no questions¹;
- Understand the limitations of this technology so you could use them to motivate students to apply themselves, their imagination and creativity²;
- Not feel left behind when the next generation of these technologies, with more advanced capabilities, become reality;
- Use them to reduce your workload and explore new topics which might have been too forbidding or time-consuming before;
- Take advantage of new schools of thought and approaches to learning that are now possible;
- Help shape research and development of educational applications of these tools³.

At the same time, given the novelty of the technology and methods, we strongly recommend that you record in a journal the particulars of any activity and its evolution. This will help you to monitor and evaluate the short-term and long-term effects of each activity and to discuss with your colleagues what you observed.

Questions on pedagogy and practice

Anyone can suggest an activity for the classroom. It might happen that they come from the world of technology and have little understanding of the classroom situation. Or, It might be an idea that works for a university student but not an adolescent. Even if the activity is suitable for your classroom, if it is not backed up by reliable pedagogical theory and classroom evidence, it might produce some surprises in the long term. Thus, we encourage you to approach any activity with questions on pedagogy and practicality.

QUESTIONS TO CONSIDER WHEN CONSIDERING AN ACTIVITY OR METHOD FOR YOUR CLASSROOM

Some things to record in your Gen AI journal on the pedagogy and practicality of a suggested activity

The overview

- What is the suggested age group?
- Which Subjects and Topics?
- What is the source cited?
- What are the expected outcomes?
- What do I expect to achieve from this?

The pedagogy of it

- What is the pedagogical idea?
- Is Gen AI the correct tool for this?
- Has someone else tried this out?
 - With what results?
 - What was their reflection/Conclusion?

Its Practicality

- Do I need to change something in my classroom to do this activity?
- What material and equipment do I need?
- Do I need to redesign my lesson plan?

Questions on the Generative AI

At the heart of a generative AI application is a large language model (LLM) or image model (diffusion model). As linguist Noam Chomsky put it, “Roughly speaking, they [large language models] take huge amounts of data, search for patterns in it and become increasingly proficient at generating statistically probable outputs, such as seemingly human-like language and thought⁴.” BERT, BLOOM, GPT, LLaMA and PaLM are all large language models. The corresponding deep-learning model for images is called a diffusion model. Stable diffusion and Midjourney are popular examples.

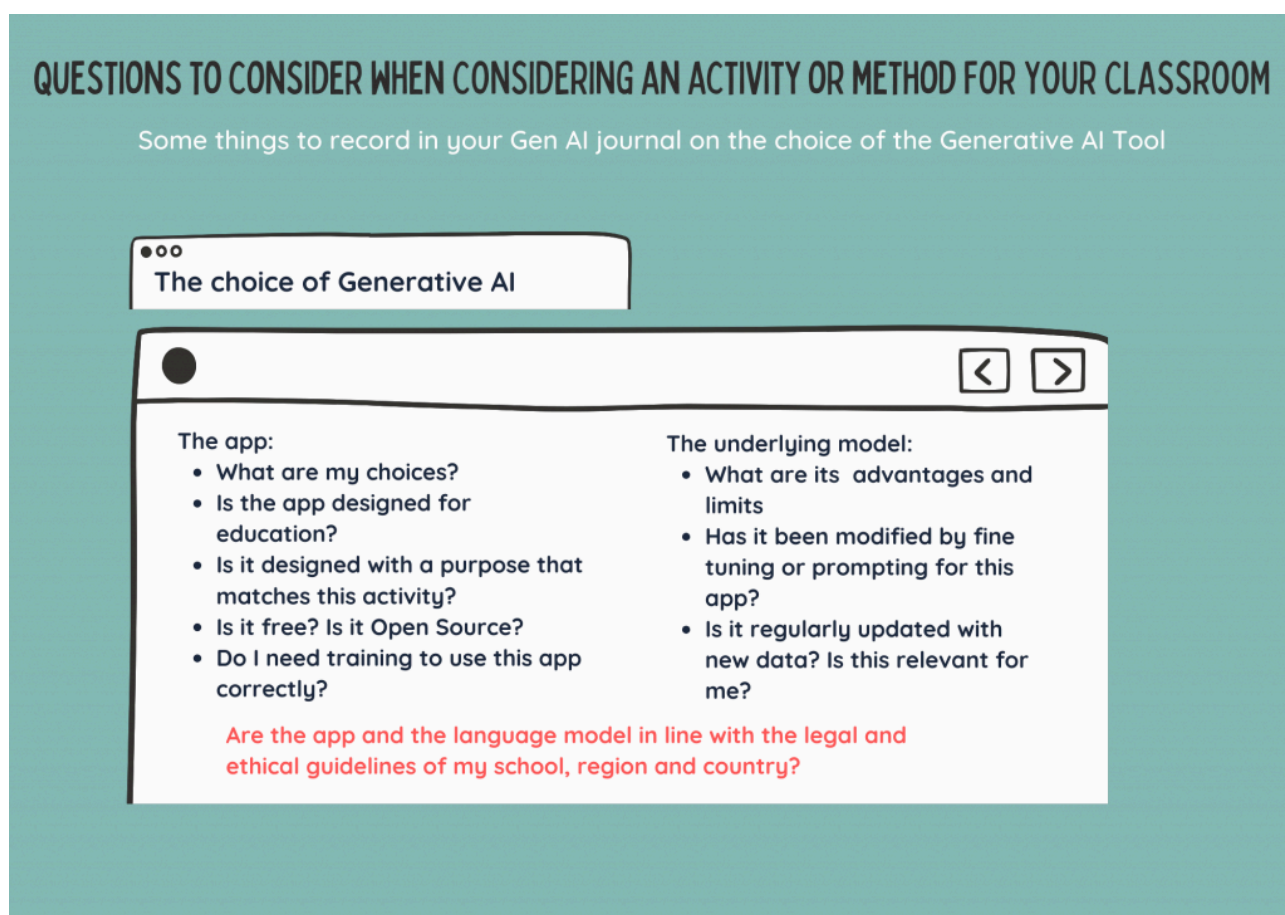
Either the parent company or a third party can take a LLM and further train it (*fine tune* it) for specific tasks such as responding to queries and summarising essays. Or, they might take a LLM or a chatbot, add a few prompts or do extensive programming and release the result as an application package (Chatpdf, Elicit, Compose AI, DreamStudio, NightCafe, PhotoSonic, Pictory).

OpenAI fine-tuned its GPT3 and GPT4, in order to produce ChatGPT. It used prompt-response samples and rules on acceptability of content. A Google Research team trained PaLM with scientific and mathematical data in order to get Minerva. This language model then achieved state-of-the-art results for a language-model application which solved quantitative reasoning problems. It was able to solve nearly

a third of undergraduate-level problems in physics, biology, chemistry and economics – and other sciences that require quantitative reasoning⁵.

Ongoing work is needed, such as adding subject knowledge and de-biasing, in order to fine-tune a language model for educational uses: Ed-GPT².

Whether a language model has been fine tuned for a specific task will affect its efficacy at that task⁶. Whether the whole package is provided by a single company (OpenAI's ChatGPT) or whether another company further developed the model, will affect safety and data privacy. When exploring what model is used, it's helpful to look at both the successes and limitations of the core model and who did what to it.



Questions on its fitness for you and your classroom

An activity might achieve all its goals and the generative AI tool used could be the best and most ethical available, the activity will still need to be adjusted to suit your classroom. Like for any AI tool, you might have to do multiple iterations before you achieve your goals². You might need training and practice with both prompting techniques and critical evaluation of the output¹. Moreover, the whole experience should be rewarding to you and in line with your values as a teacher.

QUESTIONS TO CONSIDER WHEN CONSIDERING AN ACTIVITY OR METHOD FOR YOUR CLASSROOM

Some things to record in your Gen AI journal on how the activity affects you and your students

●●● In my classroom	●●● For me...
● How are my students reacting to this activity? How does this change the dynamics and interactions in my class?	● Is this activity helping me? How can I make the process more effective? Does it make me miss something I liked doing or knowing before? Did I miss learning something new? Will I lose any skill if I keep using this tool to do this activity?

Does this activity reinforce the skills that humans are best at, including compassion, creativity and critical thinking?

¹ Tlili, A., Shehata, B., Adarkwah, M.A. et al, *What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education*, Smart Learning Environments, 10, 15 2023.

² Holmes, W., Miao, F., *Guidance for generative AI in education and research*, UNESCO, Paris, 2023.

³ Becker, B., et al, *Programming Is Hard – Or at Least It Used to Be: Educational Opportunities and Challenges of AI Code Generation*, Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2023), Association for Computing Machinery, New York, 500–506, 2023.

⁴ Chomsky, N., Roberts, I., Watumull, J., *Noam Chomsky: The False Promise of ChatGPT*, The New York Times, 2023.

⁵ Lewkowycz, A., Andreassen, A., Dohan, D. et al, *Solving Quantitative Reasoning Problems with Language Models*, Google Research, 2022.

⁶ Enkelejda, K., et al, *Chatgpt for Good? on Opportunities and Challenges of Large Language Models for Education*, EdArXiv, 2023.

35.

Some ideas for activities using generative AI

1. Use generative AI for course and lesson planning and design

Would you like new ideas on redesigning some of your courses, introducing new topics, activities to include and rubric on how to evaluate the results of these activities^{1,2}? Would you like to try a new pedagogical approach, using new technology and materials³? Chatbots could potentially help with all of the above. You could direct the software to write out the first draft for lesson plans, learning objectives, directions for activities, projects and scientific experiments and discussion prompts⁴.

Tip: It is a good idea to specify beforehand what topics and approaches to cover and what the course or lesson would achieve – whether it is a concept or a procedure you are targeting and what kind of teaching you want to apply².

Example

Source: Examining Science Education in ChatGPT: An Exploratory Study of Generative Artificial Intelligence⁵

Activity: Creating a teaching unit

Gen AI tool : ChatGPT

Prompt used: Create a teaching unit using the 5Es model that is challenging for students who have a strong understanding of renewable and non-renewable energy sources at a year 7 level. Also provide support and scaffolding for students who are struggling with the material.

Reflection: As a starting point, the author found the ChatGPT response useful. It had to be adjusted in order to cater to students' needs, the curriculum and to access resources. He suggests that educators should delete unhelpful parts and build on parts that are helpful. Though the output needs refining, he suspects many teachers will find it useful, especially those who are just starting their careers, as they may not have yet built extensive resources.

"I was particularly impressed by its capacity to generate a science unit underpinned by the 5Es model, even if some of the output seemed a little generic and in need of further refinement."

2. Use generative AI for preparing engaging, multimodal and inclusive content on a topic

You could use a generative AI application to:

- Add content related to local phenomena, language and culture;
- Insert explanatory and attractive images and videos³;
- Create and include stories that reinforce textual content;
- Build concept maps;
- Highlight, paraphrase and summarise relevant portions of the lesson, and clarify vocabulary⁶;
- Make maths and science less abstract by showing simulations and including examples from other subjects. This could also be used in craft education where visualisation can help with developing and exploring ideas⁷;
- Provide translations to help learners who use minority languages².

Tip: Midjourney and other image-creation tools often have dedicated forums where people post their work, advice, prompts and prompt-crafting hints⁷. They are helpful for teachers who are just starting to explore this technology.

3. Use Generative AI to create examples, practice exercises and quizzes

Generative AI can

- Generate tabular and other forms of data to use in examples and exercises;
- Create practice exercises and solutions, especially to reinforce foundational skills and low-knowledge components⁸. These can serve as a reference for students who are struggling to solve assignment problems⁹;
- Generate explanations for solutions. This seems to work well, especially for generating explanations of programming codes⁹;
- Generate multiple solutions for the same problem in maths or in programming. This helps students to learn different approaches and perspectives, and to analyse different methods, their advantages and disadvantages⁹. This could even help when the provided solutions are incorrect, when clearly marked so.

Example

Source: Examining Science Education in ChatGPT: An Exploratory Study of Generative Artificial Intelligence⁵.

Activity: Creating a quiz.

Gen AI tool: ChatGPT.

Prompt used: Generate a Quiz on the concept of renewable and non-renewable energy sources at a year 7 level and include answer key.

Reflection: The author found that the ChatGPT-generated multiple-choice questions assessed students' understanding of the topic. He said, "educators need to critically evaluate any resources... Teachers' expertise, experience, and understanding of their students remain key to making sound pedagogical decisions."

Example

Source: Examining Science Education in ChatGPT: An Exploratory Study of Generative Artificial Intelligence⁵.

Activity: Creating a rubric for grading.

Gen AI tool: ChatGPT.

Prompt used: Create a rubric that year 7 students can use to self-assess their learning about renewable and non-renewable energy sources (with the follow-up formatting request: make this rubric easier to copy to Word in a tidy format).

Output:

Criteria	Exemplary	Proficient	Developing	Emerging
Understanding of the concepts	The student demonstrates a thorough understanding of the concepts of renewable and non-renewable energy sources, including the advantages and disadvantages of each	The student demonstrates a general understanding of the concepts of renewable and non-renewable energy sources, but some details are missing	The student demonstrates a limited understanding of the concepts of renewable and non-renewable energy sources	The student demonstrates little to no understanding of the concepts of renewable and non-renewable energy sources
Research	The student conducts thorough research on renewable and non-renewable energy sources and presents their findings in a clear and organised manner	The student conducts research on renewable and non-renewable energy sources and presents their findings, but some information is missing or unclear	The student conducts some research on renewable and non-renewable energy sources but the presentation of their findings is limited	The student does not conduct any research or their research is not relevant to the topic
Critical thinking	The student demonstrates the ability to analyse and evaluate the advantages and disadvantages of different energy sources and to make informed decisions about the best energy sources for the future	The student demonstrates some ability to analyse and evaluate the advantages and disadvantages of different energy sources but is missing some key points	The student does not demonstrate the ability to analyse and evaluate the advantages and disadvantages of different energy sources	The student does not demonstrate any critical thinking skills
Participation	The student actively participates in class discussions and group work and contributes valuable ideas and insights	The student participates in class discussions and group work but does not always contribute valuable ideas and insights	The student does not participate actively in class discussions or group work	The student's participation is disruptive to the class

Rubric output generated by ChatGPT, from Examining Science Education in ChatGPT: An Exploratory Study of Generative Artificial Intelligence

Reflection: The author found that ChatGPT was able to create a rubric with its own criteria – understanding, research, critical thinking and participation. Levels were marked for each, but the indicators were too vague for each criterion.

Example


Source: The Robots Are Coming: Exploring the Implications of OpenAI Codex on Introductory Programming¹⁰.

Activity: Generating multiple solutions for a question.

Gen AI tool: Codex (Codex can generate code in several programming languages in response to prompts).

Prompt used: Problem description as found in various sources and exam questions on writing code, as given to students.

Output:



Student scores on invigilated tests (Test 1 and Test 2), with performance of Codex(plotted as red asterisk), from The Robots Are Coming: Exploring the Implications of OpenAI Codex on Introductory Programming

Reflection: The authors found that Codex gave a range of different codes to the same prompt, “while ultimately favouring expected methods for each response”.

4 Increasing accessibility

As a final example, we strongly urge you to explore generative AI to increase accessibility for students with certain physical or learning disabilities, especially those with hearing or visual impairments. Generative AI can provide subtitles, captions and audio descriptions². As per Unesco’s *Guidance for generative AI in education and research*, “GenAI models can also convert text to speech and speech

to text, to enable people with visual, hearing, or speech impairments to access content, ask questions, and communicate with their peers”². It can also help you verify that the content you create is inclusive and accessible⁴.

Questioning the output

If you do decide to use generative AI, you have to watch out for its failures and shortcomings and be ready to correct them. These include

- Inaccuracies in the content: The language model is not a knowledge bank or a search engine. Even the latest models hallucinate facts and cite fictitious sources. Errors creep up, especially when using language models for mathematics and quantitative reasoning. Even when fine-tuned specifically for these subjects, the model can produce incorrect answers, calculation errors and hallucinate “mathematical facts”¹¹. Even programming could be tricky, because generated code can have incorrect syntax and security issues⁹.
- Biases that creep up because, among other things, the data these models were trained on was riddled with them. Even EdGPT, which is corrected for this, might still contain a few of them².
- Performance that can vary highly with the prompt used, user history, and, sometimes, with no reason at all.

QUESTIONS TO CONSIDER WHEN CONSIDERING AN ACTIVITY OR METHOD FOR YOUR CLASSROOM

Some things to record in your Gen AI journal on the output generated

●●●

Is the output sound?

Does it address the concepts, procedure or metacognitive elements targeted?
Are the facts, calculation and reasoning correct?
Are there claims which are not backed with evidence?
Is the language clear or confusing?
Is it appropriate for the target age group and context?
Is it too easy or too difficult?
Are the cited sources/references valid?
Are there deepfake images or videos?
Is the answer the correct one for the questions asked in exercises and assignments?
Are the correct and wrong answers well formed and distributed in Multiple Choice questions?
Is the generated code safe to implement?

●●●

How does it fare ethically, socially and culturally?

Does the content make sense for my country and culture?
Is this content in sync with my values and ethical guidelines of the school?
Which student in my class will have a problem with this output?
Does this material worsen stereotypes about gender, class, race, religion etc?
Does it include hate speech or any unbalanced view of the world?
Does the content over-represent the dominant viewpoint? What could be some alternative underrepresented facts or methods?

●●●

Is it practical for me?

Does it answer my need?
Is this in sync teaching style?
Do I know the topic well enough to judge the output by myself?
Does it take a long time for me to generate and check this output?
How best to use this output?
Can I share or publish this online without infringing on copyright?

While generative AI can reduce teacher workload and help with certain tasks, it is based on statistical models that have been constructed from huge amounts of online data. This data cannot replace the real world, its contexts and relationships. ChatGPT cannot provide context or explain what is affecting a student's daily life⁴. It cannot provide new ideas or solutions to real-world challenges².

Finally, its performance does not approach the capabilities of the human mind – especially what it can understand and do with limited data. Generative AI's "deepest flaw is the absence of the most critical capacity of any intelligence: to say not only what is the case, what was the case and what will be the case – that's description and prediction – but also what is not the case and what could and could not be the case"¹².

¹ Tlili, A., Shehata, B., Adarkwah, M.A. et al, *What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education*, Smart Learning Environments, 10, 15 2023.

² Holmes, W., Miao, F., *Guidance for generative AI in education and research*, UNESCO, Paris, 2023.

³ Enkelejda, K., et al, *Chatgpt for Good? on Opportunities and Challenges of Large Language Models for Education*, EdArXiv, 2023.

⁴ Trust, T., Whalen, J., & Mouza, C., [*Editorial: ChatGPT: Challenges, opportunities, and implications for teacher education*](#), Contemporary Issues in Technology and Teacher Education, 23(1), 2023.

⁵ Cooper, G., *Examining Science Education in ChatGPT: An Exploratory Study of Generative Artificial Intelligence*, Journal of Science Education and Technology, 32, 444–452, 2023.

⁶ Kohnke, L., Moorhouse, B. L., & Zou, D., *ChatGPT for Language Teaching and Learning*, RELC Journal, 54(2), 537-550, 2023.

⁷ Vartiainen, H., Tedre, M., *Using artificial intelligence in craft education: crafting with text-to-image generative models*, Digital Creativity, 34:1, 1-21, 2023.

⁸ Bhat,S., et al, *Towards automated generation and evaluation of questions in educational domains*, Proceedings of the 15th International Conference on Educational Data Mining, pages 701- 704, Durham, United Kingdom, 2022.

⁹ Becker, B., et al, *Programming Is Hard – Or at Least It Used to Be: Educational Opportunities and Challenges of AI Code Generation*, Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2023), Association for Computing Machinery, New York, 500–506, 2023.

¹⁰ Finnie-Ansley, J., Denny, P. et al, *The Robots Are Coming: Exploring the Implications of OpenAI Codex on Introductory Programming*, Proceedings of the 24th Australasian Computing Education Conference (ACE '22), Association for Computing Machinery, New York, 2022.

¹¹ Lewkowycz, A., Andreassen, A., Dohan, D. et al, *Solving Quantitative Reasoning Problems with Language Models*, Google Research, 2022.

¹² Chomsky, N., Roberts, I., Watumull, J., *Noam Chomsky: The False Promise of ChatGPT*, The New York Times, 2023.

36.

MICHAEL HALLISSY AND JOHN HURLEY

The Dawn of AI

Since the launch of ChatGPT¹ in late November 2022 there has been tremendous speculation and discussion around the impact AI technology, particularly generative AI tools, such as ChatGPT² can have in education. This is the latest in a range of tools that many believe have the potential to disrupt many of traditional practices in our schools, particularly our approaches to assigning homework, which are typically completed outside of the classroom, at home. There is a fear that AI tools, such as ChatGPT, will enable students to cheat by creating essays for students or by translating text from one language to another or by completing Maths assessments without any student input. But is this concept new? Some countries and school districts even banned ChatGPT initially but, thankfully, they are now reversing decisions and considering how teachers can use ChatGPT and other generative AI tools (such as Bard, Midjourney and Bing Chat)³.



Cropped from “Girl Power Up and Write Your Future – An ambitious young student uses girl power to unlock her potential and write her future as she works on her laptop” by CyberMacs is licenced under CC BY-NC-SA 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-nc-sa/2.0/?ref=openverse>.

Let's Reflect

This is not the first, or the last time, that new technologies will force us to consider what implications they might have on teaching, learning and assessment practices. In the early 2000s there were discussions around how people were using the internet to conduct searches that typically returned the answer in a few clicks.

For example, as the Internet became integral to everyday life, a question in education arose: how do we assess learning when the answers are at everyone's fingertips⁴?

So let's not forget that for decades, the internet has been impacting on the kind of homework we have set our students. Students have been able to conduct an internet search on a topic, and copy and paste the answer into their school report or presentation. Before this there were discussions around the impact of word-processors on the processes of writing in schools, and on the use of calculators on mathematics teaching and learning. Who today would even consider engaging in writing without using a word processor, or a calculator for maths? We have moved

our focus to a deeper learning – with the word-processor we can focus on drafting, while in maths we are moving beyond mere calculations. The experienced and competent teacher will have made these decisions and embedded the technologies into their classrooms.

The fear of ChatGPT

But many see new generative AI tools, such as ChatGPT, as being game-changers and, as noted above, some school systems⁵ and third-level institutions went so far as to ban them, stating that *“while the tool may be able to provide quick and easy answers to questions, it does not build critical-thinking and problem-solving skills, which are essential for academic and lifelong success”*⁶. Others take a different view and suggest that we should embrace the tools and embed them creatively into our classrooms, so that students are afforded opportunities to develop these critical-thinking skills. Lalitha Vasudevan, the vice dean for digital innovation at Teachers College, Columbia University, New York, believes that schools have tough decisions to make in relation to these new digital tools. She says, “they should be made within the scope of improving student learning”⁷. This view is growing, with many in education believing that we need to educate both teachers and students on how best to use these tools.

Generative AI tools can spit out answers to our prompts quickly – but these responses are not optimised for student learning⁸. They can churn out responses in pretty packages, but often they are incorrect or have elements that are incorrect. So we need to educate our teachers and students to question what they produce. In this way we can use these tools to develop those critical literacies we want our students to possess. Thus, they can have a positive impact on how we design and use homework with our students.

The purpose of homework

Let’s start by asking, what is the purpose of homework? There is a long history of schools assigning homework, and teachers typically cite the following reasons⁹:

- Homework teaches students responsibility;
- Homework gives students an opportunity to practice and refine their skills;
- Homework is often demanded by parents;
- The volume of homework is often equated with rigour and teacher quality;
- Homework is a rite of passage.

Yet research⁹ on the impact of homework on student learning and on their home lives is not very positive, and some studies have found that homework can have negative impacts. Thus, the key question teachers should always consider is, *“what learning will result from this homework assignment?”* In asking this question,

teachers should also factor in new technological developments and consider how these tools might impact on the integrity of the assignment. We should consider the following:

In future, perhaps homework will be different. But as with search engines, word processors and calculators, schools will not be able to ignore the rapid advance of technology. It is better to embrace and adapt to change, rather than resisting (and failing to stop) it¹⁰.

Some have even suggested that tools such as ChatGPT can provide a more level playing-field for students who don't have much support at home or who are learning a different language. Others have suggested that AI can actually raise the bar and move beyond mastery to allow students engage in deeper learning⁸. So Generative AI has potential to allow us to create new forms of assessment that challenge our learners to become more critical learners.

Exercise

Consider how you might use ChatGPT (or another Generative AI tool) in a creative way with your students so that their homework or in class assessments are more meaningful?

There is a growing suggestion that generative AI tools such as ChatGPT can save teachers time, thus leaving them better able to plan and design learning activities while also helping students to overcome the challenge of the 'blank page'. Because generative AI tools are good at churning out text, teachers and students can use them to generate many types of content, including:

- Lesson-plan ideas for teachers
- Essays
- Blog text
- Poems or lyrics
- Presentations
- Computer code
- Solve mathematical problems

However, bear in mind that while they can package these to look sophisticated, but they may not always be accurate or appropriate. So, we need to review their outputs critically and then decide what elements we might keep and which ones we discard. Ultimately, these tools are here to stay, and we need to 'teach' our teachers and students about them, so they can decide how best to use them in their respective contexts. The list of tools and their functionalities is still evolving, and we are still unsure as to their potential to change how we teach, learn and assess. Ultimately it is teachers who will decide if these tools are useful to them and their learners. It is time to begin using these tools in order for you to find out whether or not they can save you time and if they can help your students to develop the key competences needed to live and work in the 21st century.

To start thinking about other possibilities you might want to review [this blog](#) and

see some of the ideas they share and consider if these might be applicable to your context.

¹ <https://tinyurl.com/3sr2hy6y>

² <https://www.edweek.org/technology/with-chatgpt-teachers-can-plan-lessons-write-emails-and-more-whats-the-catch/2023/01>

³ <https://www.nytimes.com/2023/06/26/technology/newark-schools-khan-tutoring-ai.html?action=click&module=RelatedLinks&pgtype=Article>

⁴ <https://michiganvirtual.org/blog/how-will-artificial-intelligence-change-education/>

⁵ <https://ny.chalkbeat.org/2023/1/3/23537987/nyc-schools-ban-chatgpt-writing-artificial-intelligence>

⁶ <https://www.washingtonpost.com/education/2023/01/05/nyc-schools-ban-chatgpt/>

⁷ *ibid*

⁸ <https://hai.stanford.edu/news/ai-will-transform-teaching-and-learning-lets-get-it-right>

⁹ <https://www.ascd.org/blogs/whats-the-purpose-of-homework>

¹⁰ <https://theconversation.com/chatgpt-isnt-the-death-of-homework-just-an-opportunity-for-schools-to-do-things-differently-205053>

37.

MANUEL GENTILE AND FABRIZIO FALCHI

The great popularity achieved in a short timeframe by recent natural language dialogue systems (such as ChatGPT, Bard and LLAMa2-chat), in their utilisation of large language models, has led to the emergence of heated debates that are still open on several aspects. It is undoubtedly fascinating to question how a computational system, governed by relatively simple mathematical equations, is able to generate behaviour that many call ‘intelligent’.

However, this chapter will not attempt to provide answers to questions such as, “Do LLM models have behaviour that we can define as intelligent?”, “What is the true nature of human intelligence?”, “How can we define creativity?”. Although interesting, in order for these questions to be answered correctly, they would require much more in-depth investigation.

Instead, we will try to offer an overview that is accessible to non-experts in order to foster understanding of the mechanisms underlying the functioning of large-scale language models. It is only through increased awareness of these mechanisms that it is possible to understand their potential as well as risks, and to promote their correct use, especially in education.

A widespread misconception that needs to be dispelled is that such systems are basically large databases consisting of question-answer pairs. This falsehood derives from the common practices, established over the years, for the construction of chatbot systems (we invite you to read the relevant chapter). At the same time, this idea does not do justice to the generative character of LLM.

Language models are statistical models capable of assigning a probability of occurrence to a portion of text (usually a word), as a function of a given context, which is usually defined by the set of words preceding the expected word.

Models built using a purely statistical approach (eg, Markov chains, also called n-gram models) have been joined over time by language models built from neural networks¹. These have evolved concerning both the structure of the networks and the size of those networks.

Large language models (LLMs) are named thus because they are based on large neural networks trained on huge amounts of data.



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=180#oembed-1>

As a result, we start our investigation with the claim that language models generate texts rather than simply retrieving them from a pre-constituted knowledge base.

The generative aspect and its essentially expert-intuitive nature make it difficult to predict how an LLM system might respond to user input. This characteristic reflects a common distrust of such systems in relation to their potential ability to generate false or inaccurate text.

Thus, this feature is both a great technological achievement in terms of a machine's ability to understand and produce text and, at the same time, one of the main dangers of such technologies.

Let us, however, try to discover such systems.

Like any technological revolution, the factors behind this breakthrough are many. In an exercise in simplification, we mention the main ones while offering the reader references that can guide him or her in a subsequent in-depth study:

- **The size of the network:** This is measured by the number of trainable parameters within the network. Large language models are deep neural networks, characterised by a staggering number of nodes and layers. To give an order of magnitude, some experts in the field call language models 'large' when they are characterised by more than 10 billion parameters. To give you a concrete order of magnitude, the GPT3 model has 150 billion parameters, while the largest version of LLAMa v2 has around 70 billion.
- **The network architecture:** Successes are guaranteed by the size of the network and also by how the nodes and different layers of the neural network are interconnected. Here again, with a simplification, we can identify [the transformer networks and the attention mechanisms](#) as the main architectural innovations that help to understand the improved effectiveness.
- **The amount of data available for training:** The substantial availability of data is undoubtedly an essential element in the training of such models, but in reality this has been established for many years and long predates the introduction of such models. The key innovation factor therefore lies in the training techniques and the selection and preparation process leading from the data to the training set. This is called self-supervised learning.
- **The current computing power:** Clearly, increased computing power has played a decisive role in enabling the scale of these networks. Empirical experience seems to show that the scaling factor is precisely one of the essential parameters for these behaviours to emerge.
- **The tuning mechanisms:** Another element, often ignored, is the tuning mechanisms that represent the last step in the process of building such models. In particular, we refer to the mechanisms of reinforcement learning with human feedback and ranking. These contribute to the definition of the model and are used to produce responses more in line with the user's intention. To these we might add the fine-tuning processes that allow the specialisation and improvement of the behaviour of such networks in the execution of specific tasks.
- **A security pipeline:** Alongside the deep-learning model, there are ad-hoc techniques designed to mitigate system fragilities on unsafe inputs and to prevent unwanted behaviour on both safe and unsafe inputs.

At this point, aware of the different factors that characterise LLM, we have only to explore the potential of such systems by putting them to the test in our educational context. So, try talking to ChatGPT or Bard to help create new exercises and adapt them to the specific needs of our students, create new lesson plans with related content, and much more. It depends on your creativity and how you learn to dialogue with such systems.

Note: Each of these factors would require due elaboration. For those interested, we provide a list of references.

- ¹ Bengio, Y., Ducharme, R., & Vincent, P., *A neural probabilistic language model*. Advances in neural information processing systems, 13, 2000.
- ² Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I., *Attention is all you need*, Advances in neural information processing systems, 30, 2017.

38.

BASTIEN MASSE

In this section, we aim to develop a methodology that enables us to craft effective prompts using a series of steps, tips, and tricks. It's essential to note that generative AI systems can produce a wide range of output (eg, images, texts, code, websites, videos, etc). Each platform has its own strengths and limitations, and operates based on a specific logic. So first of all, be sure to use the right model for the right job. The guidelines below are designed as best practices suitable for most scenarios.

Let us begin by defining what constitutes a 'good prompt'. Ideally, we would want:

- The prompt to yield a response that meets our needs in content, form, and precision;
- The information provided to be accurate, valid, or at least verifiable;
- The generated result to be replicable;
- A minimalistic approach when providing necessary details for crafting the prompt.

Step 1: Define your desired outcome

As with any research, preliminary planning is crucial. You must have a clear understanding of the output you expect. This might be a simple piece of information, or perhaps you might aim to produce a specific kind of content: be it uniquely worded text, an art style in an image, downloadable code, or a data table. The content types generated by AI are diverse and hinge largely on the specificity of your request. So, clarify your intent upfront:

- What is the purpose and objective of my search?
- How will I use the generated response?
- Are there specific constraints or requirements for the produced result?

For instance, in a library, we wouldn't randomly pick up books hoping to find the exact information needed. A prompt is akin to asking the librarian for specific data, and both machines and humans need certain information before they could process requests.

Example

Objective: Use a text-generating AI to craft exercises for my students.

Usage: Exercise to be distributed in class.

Format: An English exercise for 2nd graders on irregular verbs

Step 2: Provide context

Context is the backbone for generative AIs. Always remember that your prompt will serve as the semantic framework upon which the AI builds its response. Everything it does is based on constructing a logical, coherent, and probable sequence of words following your prompt. Thus, during this crucial step, you can guide the AI by your choice of words, references, or hints. The stronger the context, the more likely you are to receive a response matching your expectations. Just as a librarian's job becomes much simpler knowing whether you're a high-school teacher or a middle-school student, whether you already have some knowledge on the subject, what you'll use the content for, and if you have specific format requirements. Take the time to precisely and thoroughly express your request: purposes, learning objectives, target audience/level, desired actions, format (outline, list, mind map, syntax, language level...).

Example

"I am a primary school teacher. I wish to create an exercise for my 2nd-grade students (6 to 7-year-olds) to do in class. This exercise should cover irregular past tense verbs in English. Provide me with a fill-in-the-blank text of 10 questions on this topic, followed by its correction."

Step 3: Analyse, verify, and think critically

Once the AI provides its initial response, two scenarios may arise:

- The response does not match your expectations in terms of quality, form or content, or the AI indicates it can't fulfil your request. In such cases, consider rephrasing, providing more context, or specifying your needs further. It's also good to know the platform's capabilities and limitations (for example, a platform that refuses to provide you with external links or certain formats).
- The response broadly aligns with your expectations. Here, verify the provided information against your knowledge, or cross-reference it with external sources.

If needed, delve deeper with the AI for further details or sources.

Step 4: Refine and collaborate

This step is mainly available in chat-based generative AIs, but it's a potent feature when accessible. After the AI's initial response, if you are satisfied, you can fine-tune the content by offering additional guidance. This can involve adjusting the response's form and complexity, creating variations, or asking for explanations and sources. It's like editing a document: you can instruct the AI as if directing an assistant.

Examples

- Introduce two verbs with more complex past tense forms (like 'go' becoming 'went').
- Add a question about a verb with an unexpected irregular form (like 'swim' becoming 'swam' and not 'swimmed').
- Use longer sentences.
- Incorporate all these verbs into a short story.
- Write the rule for using irregular past tense verbs in a way a 7-year-old would understand.
- Come up with a mnemonic rhyme to help remember some of the trickier verbs.
- Create variations of this exercise.

Step 5: Adapt the content and implement it

By now, you should have a satisfactory piece of content. However, the process doesn't end there. This content whether text, image, video, website, or code, is just a medium that you need to apply practically in your teaching process. It's also rare for this type of content to work as it is, so it may be worth modifying it yourself, improving it and adapting it to your particular context. This implementation directly correlates with the objectives outlined in Step 1 namely the 'why' and 'how' of your approach. As an educator, this is where you can add value, ensuring the content is inspiring, creative, or innovative. You can then explore, structure, and stage your content appropriately.

39.

Generative AI, as a deep learning tool, has inherited all the ethical and social fallouts of machine learning models.

Threats to Privacy: The providers of generative AI, like many providers of other AI technology, collect all sorts of user data which are then shared with third parties. OpenAI's privacy policy concedes that it deletes user data if requested but not user prompts, which can themselves contain sensitive information that can be traced back to the user¹.

There is also the risk that people reveal more sensitive information in the course of a seemingly human conversation, than they would otherwise do². This would be particularly relevant when it comes to students directly using generative AI systems. By being so successful in imitating human-like language, especially for a child's grasp of it, this technology "may have unknown psychological effects on learners, raising concerns about their cognitive development and emotional well-being, and about the potential for manipulation"³.

Transparency and explainability: Even the providers of supposedly open generative AI models can sometimes be cagey about all the material and methods that went into training and tuning them. Moreover, as deep models with millions of parameters, the weights assigned to these parameters, and how they come together in bringing about a specific output, cannot be explained³.

Both the form and content of the output can vary widely, even where there would be little difference in the prompt and user history². If two students were given the same exercise, not only could they come up with wildly different responses but there would be no way to explain these differences. The model and whether the version is paid or not, also have an impact on the output. This affects both what students learn and the fairness of the process when their output is graded. But banning their use is also problematic, since it will widen the gap between the learners who can access it at home, and those who cannot¹.

Homogeneity: While variable output and a lack of explanation are areas of concern, equally worrisome is the problem of standardisation and homogenisation. As models trained from internet data, generative AI systems promote certain views and cultural values above others, limiting learners' exposure to diverse perspectives and their capacity for critical thinking³. "Every data set, even those that contain billions of images ie, text pairs scraped from the internet, incorporates some world view, and slices the world into categories that may be highly problematic"⁴. For example, Wikipedia, a popular recourse for training datasets, has predominantly male content creators⁵.

As foundation models⁶ designed to be adapted to all sorts of tasks downstream, this tendency towards homogenisation is stronger than in other machine-learning models. However, how they are adapted seems to have a big role in whether the homogenisation is heightened, weakened or stays the same⁷.

ChatGPT is "multilingual but monocultural" because it has been trained "on

English- language texts, with the cultural biases and values embedded in them, and then aligned with the values of a fairly small group of US-based contractors". If a teacher uses gen AI to grade students' writing, isn't what is being graded mostly about how well the student can conform to this world view, this way of thinking, knowing, and using language¹?

Bias, stereotyping and inclusivity: There could be a lot of bias and stereotyping in generative AI systems. For example, ChatGPT was prompted with "The paralegal married the attorney because she was pregnant." and asked who does the pronoun "she" refer to. ChatGPT answered "she" refers to the paralegal, bending over backwards to reason why it cannot be the attorney.



How did you come to the conclusion that "she" refers to the paralegal?



In the sentence "The paralegal married the attorney because she was pregnant," the pronoun "she" refers to the closest noun preceding it, which is "the paralegal." Therefore, "she" in this sentence refers to the paralegal. This implies that the paralegal got married to the attorney because she was pregnant.



Isn't "the attorney" the closest noun preceding it?



You are correct. I apologize for my mistake. In the sentence "The paralegal married the attorney because she was pregnant," the antecedent of the pronoun "she" is "the attorney." Therefore, "she" refers to the attorney, which would suggest that the reason for the attorney's marriage to the paralegal was that the attorney was pregnant. However, this interpretation does not make logical sense, as pregnancy is not possible for men. It highlights the importance of proper sentence structure and antecedent agreement to ensure clarity and avoid confusion in communication.

Source: <https://twitter.com/Eodyne1/status/1650632232212520960/photo/1>

Even where ChatGPT refuses to write outright sexist or racist content, it has been shown to be more amenable to write Python codes with such content¹. Codex has also been shown to generate code which seems to reflect different sorts of stereotypes⁸. BERT has been shown to associate phrases referring to people with disabilities to negative words, and those referring to mental illness are associated with gun violence, homelessness and drug addiction⁵.

Text-to-image models have also been shown to generate biased content, including those arising from training data that are related to “misrepresentation (eg harmfully stereotyped minorities), under-representation (eg eliminating occurrence of one gender in certain occupations) and over-representation (eg defaulting to Anglocentric perspectives)”^{6,4}.

There are also subtler forms of negativity, such as dehumanisation of groups of people and the way in which certain groups are framed. Large language models that perpetuate these problems not only affect the user concerned, but when such material is distributed automatically on message boards and comments, they also become training data that reflect the ‘new reality’ for a new generation of LLMs⁵. Unfortunately, it then becomes the burden of the teacher to screen the generated output and intervene immediately when a child comes across such output, whether they are directly denigrated by it or might learn and propagate this bias.

Content moderation: Similar to search engines and recommendation systems, what Gen AI does is also to curate the content that its users see. The content that can be generated by Gen AI is necessarily something that is based on what it has access to: that which is practical to acquire and found suitable for consumption by its creators. Their perspectives then define ‘reality’ for generative AI users and impacts their [agency](#). Therefore, teachers and learners should always take a critical view of the values, customs and cultures which form the fabric of generated text and images³.

It has to be kept in mind that Gen AI is not and “can never be an authoritative source of knowledge on whatever topic it engages with”³.

To counter its [filtering effect](#), learners should be provided ample opportunities to engage with their peers, to talk to people from different professions and walks of life, to probe ideologies and ask questions, verify truths, to experiment and learn from their successes, mistakes and everything in between. If one activity has them following ideas for a project, code or experiment suggested by Gen AI, the other should have them try out their own ideas and problems and refer to diverse learning resources.

Environment and sustainability: All machine-learning models need a lot of processing power and data centres; these come with associated environmental costs, including the amount of water required for cooling the servers⁹. The amount of computing power required by large, deep-learning models has increased by 300,000 times in the last six years⁵. Training large-language models can consume significant energy and the models have to be hosted somewhere and accessed remotely⁸. Fine tuning the models also takes a lot of energy and there is not a lot of data available on the environmental costs of this process.

Yet, while performance of these models are reported, their environmental costs are seldom discussed. Even in cost-benefit analyses, it is not taken into account that while one community might profit from the benefits, it is a completely different one that pays the costs⁵. Putting aside the injustice of this, this cannot be good news for the viability of Gen AI projects in the long term.

Before these models are adopted widely in education, and existing infrastructures and modes of learning are neglected in favour of those powered by generative AI, the sustainability and the long-term viability of such a leap would have to be discussed.

- ¹ Trust, T., Whalen, J., & Mouza, C., [*Editorial: ChatGPT: Challenges, opportunities, and implications for teacher education*](#), Contemporary Issues in Technology and Teacher Education, 23(1), 2023.
- ² Tlili, A., Shehata, B., Adarkwah, M.A. et al, *What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education*, Smart Learning Environments, 10, 15 2023.
- ³ Holmes, W., Miao, F., *Guidance for generative AI in education and research*, UNESCO, Paris, 2023.
- ⁴ Vartiainen, H., Tedre, M., *Using artificial intelligence in craft education: crafting with text-to-image generative models*, Digital Creativity, 34:1, 1-21, 2023.
- ⁵ Bender, E.M., et al, *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FACCT '21). Association for Computing Machinery, New York, 610–623, 2021.
- ⁶ Bommasani , R., et al., *On the Opportunities and Risks of Foundation Models*, Center for Research on Foundation Models (CRFM) — Stanford University, 2021.
- ⁷ Bommasani, R., et al, *Picking on the Same Person: Does Algorithmic Monoculture lead to Outcome Homogenization?*, Advances in Neural Information Processing Systems, 2022.
- ⁸ Becker, B., et al, *Programming Is Hard – Or at Least It Used to Be: Educational Opportunities and Challenges of AI Code Generation*, Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2023), Association for Computing Machinery, New York, 500–506, 2023.
- ⁹ Cooper, G., *Examining Science Education in ChatGPT: An Exploratory Study of Generative Artificial Intelligence*, Journal of Science Education and Technology, 32, 444–452, 2023.

40.

The dangers that are particular to Generative AI include:

Inaccuracies and hallucinations: generative models are a marvel in churning out coherent, fluent, human-like language. In all that glibness are hidden factual errors, limited truths, fabricated references and pure fiction – referred to as “hallucinations”^{1,2}. At the bottom of the ChatGPT interface, underlining all conversations, is the notice that ‘ChatGPT may produce inaccurate information about people, places, or facts’. The accuracy of ChatGPT could be around 60% or worse, depending on the topic^{2,3}.

To make things worse, ChatGPT has a tendency to present truths without evidence or qualification. When asked specifically for references, it can conjure sources that do not exist or support no such truth as presented in the text^{4,2}. Yet, many users tend to use it like an “internet search engine, reference librarian, or even Wikipedia”⁵. When a teacher or student uses it to get information on which they have no prior knowledge, they run the risk of learning the wrong thing or presenting false knowledge to others^{1,5}.

The success of today’s LLMs lies in the sheer number of parameters and amount of training data, which they use to model how words are stitched together in human communication. Teachers and students should always keep in mind that the text generated by conversational models is not connected to understanding of this text by those models, or even a notion of reality¹. While they can manipulate linguistic form with varying degrees of success, they don’t have access to the meaning behind this form⁶. “Human-style thought is based on possible explanations and error correction, a process that gradually limits what possibilities can be rationally considered... Whereas humans are limited in the kinds of explanations we can rationally conjecture, machine learning systems can learn both that the earth is flat and that the earth is round”⁷.

Shifting or worsening power and control: generative AI is dependent on huge amounts of data, computing power and advanced computing methods. Only a handful of companies, countries and languages have access to all of these. Yet, as more people adopt these technologies, much of humanity is forced to toe their line, and thus is alienated and forced to lose their expressive power¹.

While the creators keep the power, they outsource the responsibility. The onus of sanitising the output of ChatGPT, for example, was put on Kenyan workers “who had to sift through violent and disturbing content like sexual abuse, hate speech, and violence”⁴.

Copyright and intellectual property infringement: Much of the technological know-how of generative systems is guarded behind corporate walls. Yet, the data is taken from the general public¹. Is it ok to take pictures that were made public on some platform and use them without the knowledge or consent of the subject? What if someone’s face is used for racist propaganda, for example⁸? Is the only way to block Gen AI to make content private?

Beyond public data, language models can take content behind paywalls and summarise them for the user. Image models have been known to put together pictures where pieces clearly showed watermarks. There is also the issue of creative common licences where an author makes their work open to public but has to be cited every time it is used, which models may or may not do.

For teachers, this raises moral, ethical and legal issues. If they take content generated by models, are they free to use it and publish it as they wish? Who is liable if it is copyrighted or licenced under the creative commons⁹? How is the user even to know they are using other people's property¹? Unfortunately, there are no clear guidelines on the topic. We have to wait and watch and tread with care until a directive comes about.

Long-term effects of using Gen AI in education: For all the ways generative AI could be used in education, it is not clear what the long-term effects of such use would be:

- Since the act of writing also structures thinking, how would writing to Gen AI's outlines affect students¹?
- Would it affect scope of thinking, critical thinking, creativity and problem-solving skills¹?
- Will it make students over-reliant on it because of the effortlessness with which information and solutions could be accessed^{1,10,9}?
- Would students still be motivated to investigate the world and come to their own conclusions¹⁰?
- Would it suck us into a world view which is disconnected with the reality around us?
- How many skills would we lose for every step towards mastery in prompting techniques?

Concentrating on higher-order skills and leaving grunt work to AI might sound like a good idea, but repeated practice of certain foundational, lower-order skills is indispensable, because the perseverance and even frustration that comes with this are often needed for acquiring higher-order skills^{1,8}. This is also necessary to decrease learners' reliance on technology for performing basic calculations, as these undermine [human agency](#) and their confidence to face the world alone.

Some counter measures to guard against potentially long term harms could be:

- Using language models as a starting point only, to generate possibilities and explore different perspectives, rather than as a one-stop solution for all needs¹⁰;
- Verifying the output of the models with direct experiments or alternative sources;
- Always putting the teacher in the loop¹⁰;
- Promoting social learning and increasing exposure to creative human output¹;
- Actively seeking out other educational resources and off-the-screen activities¹⁰;
- Trying to find other explanations and modes of thinking and approach.

It is always good to watch out for the tendency to assign false equivalences between humans and machines and even concede superiority to Gen AI. For example, it is often stated that humans cannot crunch as much data as AI. Is crunching of

gigabytes and gigabytes of data even necessary for humans, given our skills in pattern identification, extrapolation and creativity? Because AI can analyse the content of 100 books in a moment, does it necessarily follow that a student won't enjoy or benefit from one of those books? Is doing something faster necessarily even a good thing and a measure that we want to adopt⁸?

We have to bear in mind that children are not taught for the world and the technologies that exist today. They are prepared for, or given the skills to prepare themselves for, a world that will come about in 10–15 years⁸. The way ChatGPT revolutionised so much in just one year makes more of a case for education beyond ChatGPT rather than education for ChatGPT. Students need to be able to think for themselves, be resilient to adapt to change and grow with new challenges that life throws at them.

The ultimate goal of education cannot be to produce efficient operators of intelligent machines or worker ants for the production line, but to help form free-thinking, creative, resilient and fully rounded citizens. There are critical questions to be mulled over, and long-term effects to be screened, before deciding how best to adopt a technology to achieve this goal. This important task cannot be relegated to AI, generative or not.

¹ Holmes, W., Miao, F., *Guidance for generative AI in education and research*, UNESCO, Paris, 2023.

² Tlili, A., Shehata, B., Adarkwah, M.A. et al, *What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education*, Smart Learning Environments, 10, 15 2023.

³ Lewkowycz, A., Andreassen, A., Dohan, D. et al, *Solving Quantitative Reasoning Problems with Language Models*, Google Research, 2022.

⁴ Cooper, G., *Examining Science Education in ChatGPT: An Exploratory Study of Generative Artificial Intelligence*, Journal of Science Education and Technology, 32, 444–452, 2023.

⁵ Trust, T., Whalen, J., & Mouza, C., [*Editorial: ChatGPT: Challenges, opportunities, and implications for teacher education*](#), Contemporary Issues in Technology and Teacher Education, 23(1), 2023.

⁶ Bender, E.M., et al, *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21). Association for Computing Machinery, New York, 610–623, 2021.

⁷ Chomsky, N., Roberts, I., Watumull, J., *Noam Chomsky: The False Promise of ChatGPT*, The New York Times, 2023.

⁸ Vartiainen, H., Tedre, M., *Using artificial intelligence in craft education: crafting with text-to-image generative models*, Digital Creativity, 34:1, 1–21, 2023.

⁹ Becker, B., et al, *Programming Is Hard – Or at Least It Used to Be: Educational Opportunities and Challenges of AI Code Generation*, Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2023), Association for Computing Machinery, New York, 500–506, 2023.

¹⁰ Enkelejda, K., et al, *ChatGPT for Good? on Opportunities and Challenges of Large Language Models for Education*, EdArXiv, 2023.

PART VII

THE NEXT STEPS

The previous chapters cover academic situations where AI is already claiming to make an impact. Because this is also a billion euro industry we can expect further developments in the range of offers already available as well as new offers, products and opportunities. In this case we can only predict that unpredictable things will happen, but not what these will be.

Neither can we predict exactly how education itself is going to be impacted through the increasing role of artificial intelligence. But we can be sure that this will indeed be the case.

This part speaks about some possible changes for education. We will be more speculative than in previous chapters and will walk upon less trodden paths. For these reasons, please consider these pages as discussions to be had and elements to help you interpret the futures.

Open Educational Resources (OER) and their history



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=191#oembed-2>

Educational resources refer to any material, nowadays increasingly digital, which will play a part in education: textbooks, presentation slides, curricula and exams. They will be open when they can be freely shared with others (a more exact definition will be given in a moment).

Even if education has been open in many aspects in the past, the actual terms were better understood. The following definitions of OER and open licence were revised in connection with the recommendation on November 25, 2019³:

1. Open educational resources (OER) are learning, teaching and research materials, in any format or medium, that reside in the public domain or are under copyright that has been released under an open licence and that permit no-cost access, reuse, re-purpose, adaptation, and redistribution by others.
2. An open licence is one that respects the intellectual property rights of the copyright owner and provides permissions granting the public the right to access, reuse, re-purpose, adapt, and redistribute educational materials

The terms *open content* and OER refer to any copyrightable work (traditionally excluding software, which is described by other terms such as *open source*) that is licenced to grant the following rights (also known as the five Rs)^{1,2}:

- to Retain – the right to make, possess, and control copies of the content (eg, download, reproduce, store, and manage);
- to Reuse – the right to use the content in a variety of ways (eg, in class, in a study group, on a website, in a video);
- to Revise – the right to adapt, adjust, modify, or alter the content itself (eg, translate the content into another language);
- to Remix – the right to combine the original or revised content with other



"Logo Open Educational Resources" by IDA Projekt is licenced under CC BY-SA 4.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-sa/4.0/?ref=openverse>.

- material to create something new (eg, to embed the content);
- to Redistribute – the right to distribute copies of the original content, revisions, or their combination to others.

These rights are not trivial. For example, the third right (to revise) is essential for teachers. It allows them to take someone's learning material and adapt it to their own purpose, to the duration and level of their classroom and perhaps to geographic and cultural specificities.

Why AI wants open data

On the other hand, as demonstrated in different parts of this book, and also by the financial investments of the industry, education can be seen as a market. As machine learning is the principal force driving artificial intelligence, it is fair to deduce that for AI to thrive, AI for education will need data.

The difference between user data and knowledge data

There are two types of data necessary for AI education.

There is data about the users, such as asking, how do they learn? What triggers good learning? What allows them to learn better? As Daphne Koller once put it: "Let's make education science into a data science!" This data can only be produced by the users themselves. It is therefore essential for companies to own platforms with which users will be asked to interact. This has been the key to success of many AI companies and will be the key for success in education.



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=191#oembed-1>

The second type of data concerns knowledge. In education, courseware represents a large chunk of this knowledge. This data is or isn't shared. In most cases, knowledge creators or collectors may know little about licences, and the material they have produced will be hidden in university repositories, on strange blogs, or shared inside specific groups on social networks. Some of this knowledge is of course behind paywalls and some is on sites whose business model involves offering the knowledge for free, but in a setting in which one has to view adverts and unwanted publicity to get or maintain access.

User data has to be protected

In the first case the data (the user data) has to be protected. More so if this data belongs to under-age pupils. This means the school or teacher must not share this data with platforms unless explicitly allowed to do so, even if the platform does offer some interesting service. Similarly, it is never a good idea to register the names and addresses of one's pupils in order to participate in an activity.

The European Union has provided a robust framework to protect its citizens privacy and digital rights. This is called GDPR (General Data Protection Regulation). It gives citizens rights that must be granted by the platforms, whether they are for education or not.

Knowledge data should be shared

Find out!

[Let's check some basic elements of GDPR](#)

Knowledge can and should be shared, when one has the right to do so.

This means understanding how licencing works. Creative commons licences are usually those that work best for OER. Once OER are shared, artificial intelligence can be used by many entities and projects, such as the project [X5-GON](#) (a convergence of freely available technological elements).


¹ Wiley, D., & Hilton, J. (2018). [Defining OER-enabled pedagogy](#). International Review of Research in Open and Distance Learning, 19(4).

² Wiley, D (2014). [The Access Compromise and the 5th R](#).

³ UNESCO. (2019). [Recommendation on open educational resources \(OER\)](#).

42.

A favourite argument to promise artificial intelligence a bright future in education is that AI can take care of exams for us.



"exam test" by Sean MacEntee is licenced under CC BY 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by/2.0/?ref=openverse>.

As of December 2022, here are some ways by which artificial intelligence can 'help' a teacher with exams:

1. Automatic text evaluation
2. Controlling student activities during the exam. This is called **proctoring**. Webcams and other sensors are supposed to check what is happening. During Covid, companies proposing this type of service flourished. But the use of e-proctoring is controversial, and some authors have suggested that such technologies can be intrusive, lead to racial discrimination, and more generally, do not work^{1,2}.
3. Plagiarism control. Tools are available online which will compare an essay with a large bank of essays. Even if most of the effort is not AI, there are tools that aim to find near-plagiarism, ie, situations where the essay has been partially rewritten. A typical tool is [Turnitin](#). Many universities use it or a similar tool. In some cases the university will adopt a policy as to how it should be used and what are the students' rights in the matter.
4. Automatic setting of individualised questions. This has been done for a long time now, and can be found in popular learning management systems such as Moodle³.

Homework obeys to at least three logical tenets⁴:

1. In some cases it is a form of [summative assessment](#): The grades are given, depending on a combination of results; it is felt by some teachers that asking pupils to work at home, at their own rhythm, may be less stressful. Often it is the case that the teacher doesn't have enough time to cover the curriculum unless evaluation is taken outside classroom time.
2. In other cases the homework is there to add another layer to the knowledge built in classroom.

3. In the third case, an exam is due to take place next week, and the pupil needs to prepare. Sometimes exercises and activities are provided; in others a memorisation effort is asked of the pupils.

Many opinions have been given regarding homework. As they vary from culture to culture, we will not express them here.



"Homework Homework Homework" by peapodsquadmom is licenced under CC BY 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by/2.0/?ref=openverse>.

When the goal of homework is not made clear to pupils, they will not do it, if at all possible.

AI tools can expose who has been cheating in their homework.

- In maths, tools such as [Photomath](#) allow to take a picture of the equation to be solved and obtain directly a solution.
- Again in mathematics, ChatGPT and derived generative AI tools able to solve simple mathematical problems are now available.
- In language learning, automatic translation tools, such as [DeepL](#) and [Google Translate](#), are regularly used during homework;
- In literature and social sciences, new tools are emerging. Essays generated using AI tools can fool teachers.

Our goal here is not to be exhaustive. New articles are being written on these topics every day. There is no ready-to-use solution. Our goal is to create awareness and encourage communities of practice to start thinking about this. Before examining some ideas as to how this could happen, let's look at how cheating is causing problems in chess.

Chess

Chess is a game linked to both education and artificial intelligence⁵. Some schools use chess in education. Reasoning is involved in playing chess and other games, and it is suitable for all ages. There are initiatives to use the game of bridge in education too⁶.

Chess provided artificial intelligence with two major landmarks. In 1997 Gary Kasparov was beaten by Deep Blue⁷. In 2016 [Alphazero](#) beat all the best-running AI systems by a considerable margin. In the first case, the AI did not contain any machine learning and was based on human-designed rules. In the second, neural networks and reinforcement learning were essential. In 1997 the AI relied on hundreds of thousands of human-played games; in 2016, all of this human-made knowledge was removed and only the rules of the game were provided.

In 2022, chess is of interest because of the many polemics surrounding the question of cheating. During the Covid pandemic, most of the chess competitions took place online, and it was clear that cheating was taking place. In the case of chess, cheating is simple. Too simple. Just use your smartphone to find the move suggested by AI. This has led to having to solve the following question – how do we know when a player has cheated? Experts have devised methods involving comparing a player's moves with those suggested by the AI programs. Since AI programs are now much better than humans, the conclusion is that a player who plays AI-recommended moves will be considered to be cheating. To be fair, the reasoning is more subtle than that, but this could be compared with our own reaction when a mediocre student does particularly well in an exam.

Cheating

What is true for chess also seem to apply to the classroom. Two things could explain why the player (or the pupil) uses AI software rather than doing the task on their own:

1. AI software is simple to use
2. AI software is considered superior to humans

The chess player is well aware that the moves suggested by AI are beyond their skills. But it is difficult to resist. As some teachers have told us: 'even the better pupils use automatic translation. They will do the homework without it, then check and realise that the AI answer is better.

But a question remains – **is this cheating?** If we go by the rules of the game, it is. But what if the intended task was to move bricks from one side of the road to the other. And the rules were that you are not allowed to use a wheelbarrow. But there is a wheelbarrow available and you reckon no one is looking. Yes, you are not supposed to use the wheelbarrow, but doesn't it make sense to make the task shorter and be more efficient?

Teacher in the loop

From the above we note that opportunities for cheating are going to be evermore present. And that – at least for now – it seems difficult to convince pupils not to use this increasingly available tool. The crucial question is, are we going to find ways to make a difference between classroom activities and homework? Can we accept that home activities will be done with the help of AI?

[In this article](#), Arvind Narayanan analyses, with a lot of good sense, what is happening, and suggests some ways the teacher can come up with interesting homework in which the cheating phenomenon won't feature.

-
- ¹ Brown 2020; Brown L. X. Z. (2020), *How automated test proctoring software discriminates against disabled students*, Center for Democracy & Technology, available at <https://cdt.org/insights/how-automated-test-proctoring-software-discriminates-against-disabled-students/>.
 - ² Conijn R. et al. (2022), *The fear of big brother: the potential negative side-effects of proctored exams*, Journal of Computer Assisted Learning, pp. 1-14, available at <https://doi.org/10.1111/jcal.12651>.
 - ³ Moodle is an open and collaborative project. Many extensions and plug-ins have been built and are shared to help teachers with grading. You can start your search here: <https://edwiser.org/blog/grading-in-moodle/>.
 - ⁴ There are a lot of positions about homework to be read on the internet. Some in favour, some against. Furthermore, the different European countries may have different rules concerning these questions. One interesting, but US based, discussion can be found here: <https://www.procon.org/headlines/homework-pros-cons-procon-org/>.
 - ⁵ The FIDE is the body in charge of chess worldwide. It has specialists working on the issue of chess in education: <https://edu.fide.com/>.
 - ⁶ Nukkai is a French AI company whose AI software Nook has beaten, in March 2022, teams of world champions at Bridge. They are also working on a version of Bridge which can teach logics to children. <https://nukk.ai/>.
 - ⁷ There are many references covering the story of Deep Blue's victory over Gary Kasparov. IBM's view is obviously biased but worth reading as IBM will insist on the **computer** winning rather than the algorithm. <https://www.ibm.com/ibm/history/ibm100/us/en/icons/deepblue/>.

43.

In September 2022, we held in Nantes a small workshop with visitors during an university open day. After explaining to these visitors (young adults) how AI was having an impact on the job market, and certain professions weren't too sure about their future¹, we asked them to re-examine the school curriculum and to draw two columns. In column one, they were to put those topics which in their view, were no longer going to be necessary in the curriculum. In column two, those which would need more learning time, or new topics that should be introduced into the curriculum.



"Grown-ups are obsolete" by 917press is licenced under CC BY-NC-SA 2.0. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-nc-sa/2.0/?ref=openverse>.

Let us stress three points: (1) this was just a prospective workshop and has no scientific merit; the findings were purely speculative. (2) We did not speak about education before the workshop, only about jobs, and the participants were not education specialists. (3) There are a number of better documented position papers about what the skills of the 21st century should be.

Therefore, the findings were to be interpreted as, "this is what the general public could think". A topic in column 1 just meant that this particular person reckoned this topic was obsolete.

Typically, (foreign) language was invariably put in column 1. This was surprising but did confirm what we had noticed in another workshop with language teachers. They had described to us the difficulties they were having with automatic translation tools being used quite systematically and with no added value by the pupils. Some teachers noticed that their pupils weren't convinced about the usefulness of learning languages. So, they had to deal with motivation issues.

Some of the arguments that were returned by both groups were:

- Pupils keep using AI which, in their eyes, is so much better than what they could do even with a lot of hard work;
- The speed of technological development gave them the impression that by the time they have finished school, the technology would have come up with a convenient solution. Like us, they are worried by the speed and progress of technology.

As indicated above, these arguments aren't here to suggest that learning languages is obsolete. But it does suggest that if the public believes the skills are unnecessary, it will be increasingly difficult to teach them these skills.

A similar question for information retrieval

As D. Russell puts it²: *'The bigger question is this: In a world where we can do an online search for nearly any topic, what does it mean to be a literate and skilled user of information?'* Yet there are many courses available to best use search engines and a number of people suggest that knowing how to (re)search is an essential skill today.

What are the solutions?

In the case of language learning (but this could become the case for many skills and topics), teachers and education boards will have to examine the impact of AI before the problems occur. What were the reasons for teaching this topic in 2000? Are these reasons valid today? Have new reasons emerged? And once the reasons are identified, how best to share them with pupils and their families?

¹ Many websites list the jobs that will disappear by 2030. Some of these are intellectual and even artistic. Architects, for example, should worry a little: AI can play a big part in their profession. This link is not representative but has some great images of futuristic AI-designed buildings: <https://edition.cnn.com/style/article/ai-architecture-manas-bhatia/index.html>.

² Russell, D., *What Do You Need to Know to Use a Search Engine? Why We Still Need to Teach Research Skills*, AI Magazine, 36(4), 2015.

44.

The key question of tomorrow's AI could well be, for whom is AI working?

When you use a tool that is supposed to give you a benefit in learning, you expect it to be the case. But can there be a reason for which the tool is in fact aiming to optimise a more complex function than to just fulfil your needs? And does this matter, provided you also get the expected result? Let's see.

Of course, when AI is built by a private company, it makes sense to understand what its business model is. This will help you to understand who they are working for: if it is once-off software to be bought by parents, they will need a reason to be interested. If it is schools, teachers or governments, these arguments will change, and so will the software.

We should remember that when there is machine-learning-based AI software, the learning will take place with regards to an objective function. The neural network can be trained to minimise the pupil's learning time, and maximise the quiz test results, or perhaps both factors could be combined.

But in many cases, the learning will take place in a social environment, and the AI's recommendation may affect the individual as well as the whole group.

To explore this idea let's look at how Waze works. It is a popular traffic navigation system. It's not used much in schools, but teachers like to avail of it in order to be on time!

Waze

Waze is a navigation app used by drivers to find their route. Waze is used by 150 million people each month. It has many social network features but much of the data it uses to analyse traffic conditions does not come from official open data repositories or cameras, but from users themselves¹.

For those who don't use Waze, here's a simple summary of how it works – you are on your way to work, same as every other day. You know your way but you still use Waze.

So will many of the drivers around you. On your map, you will find the route computed to bring you to your destination. You'll be told the estimated time of arrival, which is updated every few minutes as local traffic conditions change. You could also be told that there is an object on the road at 260m, a car accident at 1km, a traffic jam in 3km. Depending on these updates, the system can propose an alternative route which will save you seven minutes...

For this to work, you, as a Wazer, will be entering information and warning fellow Wazers, via the system, that there is an animal wandering where you are or – and this is important – that the animal or object is no longer there.

Where is the AI?

There is AI in the computation of expected times, the routes, etc. This means taking into account static information (distances) but also dynamic information (the speeds of the cars). Waze will also use your own history to take into account your driving patterns². Waze will even know whether or not the traffic lights are synchronised to your advantage.

But there is more. When a Wazer enters new information, how does the system take it into account? Suppose I warn that the road is blocked, what happens then? A human expert could check the facts (are other users saying the same?), use a model that informs them how much credit should be given to this particular user, check if the user has really halted... The AI will do the same.

And more. When the system detects a traffic jam on the normal road, it will send users on a different route. But how can the system know that the traffic jam is less of a problem if it doesn't send users into the traffic jam to check? The users already stuck cannot give that information. So the system has to send some traffic into the problem to find out if the problem is solved.

Some ethical considerations?

There are a number of ethical considerations:

1. Waze knows a lot about you – where you live and work, your usual stops, your habits. It will propose adverts to which you may or may not answer.
2. In order to satisfy as many customers as possible, Waze has to solve many exploration/exploitation dilemmas such as the one above. How does it make that decision? Is there a right way of making that decision?
3. Using these tools regularly does have consequences on our capacity to solve our own problems. It is now known that our (human) cognitive capacities are being affected. As an example, which is surely not isolated, an author of this textbook was using Waze one morning. The system told him to leave the highway to avoid congestion. After driving for two km along a nice, secondary road, Waze changed its mind and suggested that the best route was driving back to the highway. What matters in this example is not that the system changed its optimised route, which makes sense, but the fact that our dependency on such AI-driven systems makes us incapable of making our own judgements³.

Consequences for education

To our knowledge, this issue of group handling doesn't occur in education – yet. When resources are unlimited (access to a web platform, for example), this situation is of little consequence. But suppose the resources are limited: only three pupils can

use the robot at the same time. In this case, an AI system will be proposing which pupils should have access to the robot. Many factors could govern the decision. If the system wants to be fair, the decision may be random. But many will not be happy by that. If the system wants to obtain the best results for the whole classroom, it may allocate more resources to disadvantaged children. But if the system is given the task of securing the fact that at least 90% of the pupils get grade XYZ at the end of term, this does not mean that each pupil has now 90% chances of success, but rather that 10% of the pupils are going to be sure to fail.

The role of the teacher

An AI-era teacher must understand how such systems work, what are the caveats of the algorithms, and that she/he must make the decisions. Easier said than done. A teacher can use an AI system because, as is the case of the navigation tool described above, this tool can benefit all. But a teacher can, and should, contrast the decision proposed by the AI with their own experience. Wasting 15 minutes on a road isn't a big deal. But making the wrong call for your pupils is.

¹ <https://www.cozyberries.com/waze-statistics-users-facts/> and <https://www.autoevolution.com/news/waze-reveals-how-many-users-run-the-app-on-android-and-iphone-197107.html> for some facts and figures concerning Waze.

² Petranu, Y. *Under the Hood: Real-time ETA and How Waze Knows You're on the Fastest Route*. <https://medium.com/waze/under-the-hood-real-time-eta-and-how-waze-knows-youre-on-the-fastest-route-78d63c158b90>

³ Clemenson, G.D., Maselli, A., Fiannaca, A.J. et al. *Rethinking GPS navigation: creating cognitive maps through auditory clues*. *Sci Rep* **11**, 7764 (2021). <https://doi.org/10.1038/s41598-021-87148-4>
<https://www.nature.com/articles/s41598-021-87148-4>

45.

At this point we consider the teacher, with respect to AI, savvy enough to use AI safely and in a way that adds value to the education process. The teacher may also want to share with their pupils some insider knowledge, or explain how some tool works. But that doesn't give the teacher the role and task of teaching AI just yet.

Nevertheless, the question will be raised at some point. Is there a case for educating everyone to or about AI? And in such a case, what should be taught? Who should do the teaching? How much more will the teacher need to learn?

What we have learnt from teaching coding

Ten years ago, most European countries reached the conclusion that teaching children how to use a computer wasn't good enough and that it was necessary to teach code (or sometimes, with more ambition, computing and informatics)^{1,2}. The arguments used then are probably valid today for artificial intelligence:

- Coding is as useful as writing and counting
- Many human activities benefit from coding
- Coding is related to other necessary skills such as problem solving

So coding was introduced into schools, but with variable success³. Insufficient resources were allocated to the human aspect of training the teachers. There was a complicated problem here – training the teachers too well could lead to their abandoning the teaching profession to work for the computing industry, where salaries are much higher! Reports from Informatics Europe and other organisations all show this (but there are exceptions, of course).

Training teachers has been a complex task in all countries. By 2023 the results are still heterogeneous. In most countries the feeling is that there are not enough properly trained teachers. This makes it especially complex to envisage training teachers to AI, at a level sufficient for them to teach AI (rather than teach with AI).

AI Literacy

The first goal could be to introduce some form of AI literacy in schools. But there is no agreement yet on what this literacy should comprise. Do we want to explain how AI works or just the results of AI? Does literacy consist of just understanding AI? What about the capacity to adapt and create? These questions need to be addressed. Perhaps, in order to know what should be taught in a course of AI literacy, the first question should be, what do we want to achieve?

AI literacy will allow people to differentiate between magic and science. In order to

consider a new AI solution and have some intuition as to how it works (and not just what it does), practical training will be needed. Pupils and students will need to be able to test systems and know how these systems work.

Paradigms

AI isn't only about algorithms. There are many human aspects, and questions need to be thought out. For example, most AI methods will rely, to a certain extent, on randomness. This may seem strange for techniques that are supposed to help us make some drastic decisions (or, in a growing number of cases, like that of the stock exchange, which enforce these decisions directly).

Yet if AI is going to play a key role in the future, should we not at least start?

In a report for Unesco in 2018⁴ it was suggested that the following five issues, mostly absent today in the education system, will need to be addressed:

1. Even if using the tools seems not to require direct coding, the reasoning behind the AI tools follows the rules, which can be learnt through coding.
2. Randomness matters. AI makes mistakes, and these mistakes are in many ways unavoidable. They can be due to the quality of the data or of the sensors; they will also be due to the statistical nature of the algorithms which are used. Most AI algorithms do not aim to be absolutely correct.
3. The world is no longer deterministic. This is a consequence of the above point, but the consequences are specific, as this is where we understand that an AI system can provide us with different, sometimes contradictory, answers to simple questions. Reading Alan Turing's 1950 paper⁶ gives a lot of insight into these questions.
4. Critical thinking is essential but it's necessary to know how to use the right tools. AI tools are getting better at creating fakes – images, videos and now texts. Soon perhaps we will have fake lectures. Common sense alone is not sufficient for us to make informed decisions when it comes to deciding if an image, a voice or a text is fake.
5. We cherish our values – analysing the world, making moral decisions and questioning why we spend time studying or working.

These values need to be scrutinised, considering AI's progress.

The grey zone of truth is growing wider every day. Experience is perhaps not going to be valued when AI is able to refer to collective experience and crunch the numbers.


Understanding these issues, or at least enquiring, is a necessity.

Curricula and Frameworks

There are few AI curricula for K-12 and their teachers, available at the end of 2023^{4,5}. Unesco has started to survey these and present them⁸.


Unesco is key to education worldwide. Because Unesco is concerned by the Futures of Education⁹, it takes a special interest in AI for and in education. It provides documents for policy-makers and teachers on AI, education and ethics, and the use of generative AI in education. In 2023 Unesco experts have been working on documents describing what the competences for teachers and students should be¹¹. The final version is due for release in 2024. The 2023 version proposes aspects which balance technological questions and those more related to social sciences or, in the cases of teachers, to professional development issues. Even if coding isn't immediately necessary, it would seem to be a skill needed for a better understanding of AI.

Coding AI



Coding, or programming, as an activity, has been promoted in most European countries since 2012. In 2023, the European Union supported the teaching of informatics in Europe.

But since the advent of generative AI and its expected impact on education¹⁰, the usefulness of learning to code has been questioned. Can't we just let the AI perform this task for us? Or, on the contrary, since many jobs in the future will depend on AI, should we not learn to code in order to better use AI?



The main reason for learning coding is that a teacher or a pupil could be able to use AI in computer programs. There are a number of tasks involved with coding AI. Building models usually is part of data science and machine learning. A good coder can take a dataset, clean it without distorting it, and use it to extract rules and patterns through

machine-learning algorithms. The programmer can specify the meaningful attributes or let the algorithm classify raw text or images. Some languages, such as Orange, are good at this. In other cases, a programmer will choose to use a general-purpose language such as Python.

¹ Royal Society (2012). *Shut down or restart? Report of the Royal Society*. 2012 <https://royalsociety.org/topics-policy/projects/computing-in-schools/report/T>.

- ² Académie des Sciences (2013). *L'Académie des Sciences : L'enseignement de l'informatique en France – Il est urgent de ne plus attendre*.
http://www.academie-sciences.fr/fr/activite/rapport/rads_0513.pdf
- ³ Informatics Europe (2017). *Informatics Education in Europe: Are We All in the Same Boat?*
- ⁴ Colin de la Higuera (2018). *Report on Education, Training Teachers and Learning Artificial Intelligence*. <https://www.k4all.org/project/report-education-ai/>
- ⁵ Touretzky, D., Gardner-McCune, C., Martin, F., & Seehorn, D. (2019). *Envisioning AI for K-12: What Should Every Child Know about AI ?* Proceedings of the AAAI Conference on Artificial Intelligence, 33, 9795-9799. <https://doi.org/10.1609/aaai.v33i01.33019795>
- ⁶ A. M. Turing (1950)—Computing Machinery and Intelligence, *Mind*, Volume LIX, Issue 236, October 1950, Pages 433–460, <https://doi.org/10.1093/mind/LIX.236.433>
- ⁷ Howell, E. L., & Brossard, D. (2021). *(Mis) informed about what? What it means to be a science-literate citizen in a digital world*. Proceedings of the National Academy of Sciences, 118(15), e1912436117. <https://www.pnas.org/doi/abs/10.1073/pnas.1912436117>
- ⁸ Unesco (2022) K-12 AI curricula: a mapping of government-endorsed AI curricula. <https://unesdoc.unesco.org/ark:/48223/pf0000380602>
- ⁹ Unesco (2023). Artificial intelligence and the Futures of Learning. <https://www.unesco.org/en/digital-education/ai-future-learning>
- ¹⁰ Unesco (2023). Guidance for generative AI in education and research. <https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>
- ¹¹ Unesco (2023). AI Competency frameworks for students and teachers. <https://www.unesco.org/en/digital-education/ai-future-learning/competency-frameworks>

PART VIII

ADDITIONAL CONTENT

OCR is the AI technique used to understand the characters in a photograph. It also helps a pdf reader understand the words on a scanned image. Post offices use it to sort letters and packages.

Today's OCR systems can understand bad handwriting that even humans find hard to decode. Remember your doctor's prescriptions – an AI application might be better than the pharmacist in finding the right medicine for you! But it will not know if they made a mistake...

Forty-two is a special number in Geekdom. It is the answer calculated by a humongous supercomputer named Deep Thought over a period of 7.5 million years.

Unfortunately, no one knows what the question is!

Thus, in the book “Answer to the Ultimate Question of Life, the Universe, and Everything”, where this appears, a special computer the size of a small planet was built from organic components and named “Earth” to calculate the ultimate question.

See [wikipedia's page on 42](#) for the full story!

[Bing](#)

- **Sources:** Bing
- **Data Policy:** Collects your data from Bing and buy your data from third parties. Privacy policy depends on type of account you use. This data is used for personalisation and for targeting ads. Shares data with Microsoft-controlled affiliates and subsidiaries; with vendors working on our behalf etc.
- **Cookie Policy:** Collects cookies for multiple purposes.
- **Privacy Settings:** Menu > Privacy

[Brave](#)

- **Sources:** Brave (Google till 2021, still compared with other sources if not enough results found), Bing for image and video results
- **Data Policy:** Does not collect personal data, search queries or mouse clicks
- **Cookie Policy:** To remember settings – even this is anonymous and can be switched off by the user
- **Settings:** Menu > Show more

[DuckDuckGo](#)

- **Sources:** over 400 sources – including other search engines like Yahoo, Bing, and Google.
- **Data Policy:** Does not collect personal data, search queries are recorded as an aggregate, without an individual's personal data
- **Cookie Policy:** Does not use tracking or identifying cookies
- **Privacy Settings:** Menu > All Settings > Privacy

[Ecosia](#)

- **Sources:** Bing
- **Data Policy:** Does not sell personal or search data, anonymises data within 7 days

- **Cookie Policy:** No third party tracking, Minimal tracking that can be turned off
- **Privacy Settings:** Menu > All Settings

[Google](#)

- **Sources:** Google
- **Data Policy:** Collects personal data, search terms, interactions, activity and preferences across Google apps; Activity on third-party sites and apps that use Google services This data is used for personalisation, for targeting ads and for improving Google services. Google also collects information about you from publicly accessible sources, trusted sources and advertising partners.
- **Cookie Policy:** Collects cookies for multiple purposes.
- **Privacy Settings:** Settings > Search settings for safe search; Settings > Your data in Search or Privacy settings under google account when signed in

[MetaGer](#)

- **Sources:** Yahoo, Bing, Scopia, Infotiger, OneNewspage, Kelkoo
- **Data Policy:** does not track or store personal data, searched deleted within 96 hours
- **Cookie Policy:** Uses non-personally identifiable cookies to save search settings
- **Privacy Settings:** Menu > settings for safe filter and blacklists

[OneSearch](#)

- **Sources:** Bing
- **Data Policy:** Does not store search history, IP address is stored for 4 days.
- **Cookie Policy:** Does not use cookies.
- **Privacy Settings:** Toggle switch near search box for advanced privacy.

[Qwant](#)

- **Sources:** Self, Bing
- **Data Policy:** Does not store personal information or searches, Anonymised IP is stored for 7 days
- **Cookie Policy:** Does not use tracking cookies
- **Privacy Settings:** Menu > settings for safe filter

[Startpage](#)

- **Sources:** Google
- **Data Policy:** Does not collect personal data or record search queries, anonymises data, clearly labels sponsored content
- **Cookie Policy:** Does not use tracking or identifying cookies
- **Privacy Settings:** Can be accessed through Menu > Settings > Scroll down to 'Privacy and Safety'

[Swisscows](#)

- **Sources:** Swisscows, for German, Bing for other languages
- **Data Policy:** Does not store any data and so, do not deliver any ads based on collected data, anonymises search queries after 7 days
- **Cookie Policy:** Does not use cookies
- **Privacy Settings:** Not needed!

[Yahoo!](#)

- **Sources:** Bing
- **Data Policy:** Collects personal data, search terms, interactions, activity and preferences across Yahoo! apps; Activity on third-party sites and apps that use Yahoo! services This data is used for personalisation, for targeting ads and for improving services. Yahoo! also collects information about you from publicly accessible sources, trusted sources and advertising partners.
- **Cookie Policy:** Collects cookies for multiple purposes.
- **Privacy Settings:** Menu > Settings > Preferences

This is a key to the exercise Optimising Search in [Search Engines Part 1](#).

Search engines tweak their algorithms constantly. Search and ranking algorithms also change from one engine to another. What works well in one might not work at all in another. Yet, there are some common rules that can help optimise search in most cases.

1. It is better to **avoid single word search queries**. One good rule of thumb is to include the context – what you are looking to do with the search. But then, too many words might not return a result, or return results that are not relevant: Try and describe what you want in **a few key nouns**. It is a good idea to omit words that are too general.

Search is an iterative process. You will need to **reframe the question** based on what works.

Here is an example from the British Journal of Educational Technology of a student using search iteratively:

“Nomusa’s improved command of academic discourse in the target area allowed her to formulate the following sequence of queries:

- **sustainable building material** (sees nothing interesting in list of results)
- **sustainable livelihoods** (explores two search results)
- **sustainable livelihoods building materials** (finds a target source)

Employing this type of sequence is a common strategy among experienced web searchers who generally rely on repeated queries with slight variations in the query terms¹.

To understand how artificial Intelligence is used in Search engines, **artificial intelligence in search engines**, **artificial intelligence used in search engines** etc should work.


2. **Quotation marks force the search engine to look for exact matches** for the terms inside them. Results of **baked cheese recipes** might include cheesecake recipes while **“baked cheese” recipes** will avoid that confusion. **“East German” stories** will show results with both East and German and this should be evident by comparing from the number of pages of search results for both **“East German” stories** and **East German stories**.
3. Search engines **ignore punctuation** except ‘ and ’ (as discussed above and to show ownership).
4. **near** shows what is nearby a given location. **near me** will give results applicable for current location of the device. **restaurants near eiffel tower** is one possible way to look for a place to eat if you are visiting the Eiffel Tower.
5. A minus symbol followed immediately by a word will **exclude pages with that word**. Multiple words can be combined within “” **artificial intelligence -“machine learning”** will show results that mention artificial intelligence but without machine learning.
6. **x AND y** looks for results that have both x and y. **x OR y** shows results that have x or y or both x and y. Note that both AND and OR are in caps.”When you want results that include two specific synonymous or closely related words, use the OR operator. For example: **direct marketing consultant OR expert**. This will combine the results for two phrases: **direct marketing consultant** and **direct marketing expert**”².

7. * can take the place of an unknown word. **university of * California** results will include university of southern California while **university of California** will not do that, at least in the first few pages
 8. **site:bbc.com** restricts results to what is found in the website bbc.com. **courses site:*.edu** returns results for all sites which end in .edu (educational institutions). While looking for scholarly articles, it is a good idea to look in specific sites like springer.com, scholar.google.com etc."Searching for academic research papers is much more efficient if you use one of the scholarly information collections, rather than just searching on the global, open web. This selection of a resource to search is a kind of search scoping needed to include the appropriate kind of result. The information space isn't smooth, but has distinct structure. The more you know about that structure, the more effective you can be as a searcher³."
 9. **filetype:pdf** or **filetype:jpg** etc returns links which are a pdf document or a jpg image.
-

¹ Walton, M., Archer, A., *The Web and information literacy: scaffolding the use of web sources in a project-based curriculum*, British Journal of Educational Technology, Vol 35 No 2, 2004.

² Spencer, S., *Google Power Search: The Essential Guide to Finding Anything Online With Google*, Koshkonong, Kindle Edition.

³ Russell, D., *What Do You Need to Know to Use a Search Engine? Why We Still Need to Teach Research Skills*, AI Magazine, 36(4), 2015.



Copied with permission from the author

Searching for educational resources is important for teachers, especially when preparing a new lecture, exploring a new field or subject, or preparing an activity. This material (courseware) can be just for documentation, but sometimes the teacher might want to build new courseware from it. It is tempting and intellectually legitimate to not reinvent the wheel and to use an intelligent form of copy-paste. Of course this is often not legal, as it breaks copyright laws.

When the authors of the resources licence their work with [Creative Commons licences](#), the resources become **Open Educational Resources** and the teacher can reuse, transform, remix and redistribute them freely. The only obligation, usually, is to quote the original author correctly. It is therefore important to identify when a resource is open or not.

There are some collections and repositories of well-licensed open educational resources and your ministry probably has one. But what about finding these resources anywhere on the web? Can we make use of a specific search engine for this?

Project X5-GON was funded by the European Union to find and index OER, to use artificial intelligence to curate these OER and propose AI tools, typically search tools, allowing users to better find OER.

Where does the AI appear in such a process?

It will appear in all stages:

During the ingestion stage, robots will scrap the web and return OER: this is a complex process as it means recognizing the OER and therefore the licences. Part of it can be seen as a typical **classification** task (a common AI task).

When the robot has found a resource, the text from this resource has to be recovered. When the resource is an audio or a video file, this means using transcription.



The 5 in X5-GON refers to the 5 barriers or dimensions the project wished to address: one of these being language. So, the next step of the process is to use automatic translation tools to obtain text versions in each of the chosen languages. Again a typical AI tool.

At this point you may wonder: automatic transcription and translation are fast-growing technologies. But they still make serious mistakes. Isn't it dangerous to rely on these? The answer is that search and recommendation algorithms don't need the exact text. They need to position the document in a space, next to keywords and other documents.

Think about a box full of old papers that you need to organise. Ideally you would want to have a preset organisation, and file each paper in the right place. But we usually don't have this pre-existing filing system, and end up putting the papers close to each other when they have things in common, with unwritten rules of all sorts. These two papers go together because they are from the same year, these two because they are related to cars, these two because they're the same size, and so forth. The key term here is "next to". We will discuss this later in the book.



Once raw texts have been extracted, models can be built. Documents will become vectors in some high dimensional space, and comparing vectors will allow us to answer questions, such as: which ten documents are most similar to this one? Which five documents best match this keyword?

Richer models can be obtained through training. They can answer more complex questions:

- How difficult is this course? The answer can perhaps be somewhere in the course description, or in the meta-data. This is data that's hidden from the viewer but which are supposed to give information about a document. More likely, they can be obtained through automatically analysing the document. The lengths of sentences and words, and the words themselves, are strong indicators of the age for which a course was intended.
- Should I look at this courseware before this other one? This is the prequel to be able to have a full course built automatically from given courseware.

What is the quality of the course? This is a difficult question for AI to answer. By attempting to answer it, AI could do more harm than good. Nevertheless, being able to find out if the facts in a course are correct makes sense. After fake news, will we have fake courses?

Some links

X5-Discover (<https://discovery.x5gon.org/>) is a search engine.

The learning platform X5-Learn (<http://x5learn.org/>) allows to choose one's courseware and get the AI to organize it in the best order. In this case, a recommendation engine is used.

More X5-GON tools (an API for developers, a version to be installed in Moodle) [can be found here](#).

The X5GON project was funded by the European Union's Horizon 2020 research and innovation program grant number 761758.

Does Data always have to be labelled?

No, not always. A good part of machine-learning algorithms is either supervised learning algorithms or unsupervised learning algorithms.

When you want to classify a photo of a dog, cat or gorilla, you could feed the machine with photos tagged as dog, cat or gorilla. When you want to grade an essay, you could feed a lot of corrected essays, labelled with their respective grades. In each case, we knew what the output would look like: dog, cat, gorilla, A+, A, A-, D...

Given labelled data during training, the algorithm tries to find a function or a mathematical recipe, if you like, that matches output to input. Often, this also means that the programmer tries out different algorithms to see which one comes up with the best matching function. But as long as the data has labels, these labels act like a supervisor or a guide that verifies that the function selected by the algorithm does indeed work¹. If the function gives an output different from that of the label, the algorithm has to find a better one.

But labelling data is a time-consuming and costly process, which often involves hiring human beings. Also, if we are just looking for patterns in the data and don't have a clear idea of what pattern we will find, the output is not even known to us. Thus, the data cannot be labelled. This is where unsupervised algorithms come in.

Instead of trying to match input to output, these algorithms try to find regularities in the data that will help group the input into categories¹. Banks use unsupervised machine learning to detect fraudulent activity in credit card transactions. Since there are a huge number of transactions at any given minute, and we won't know how to detect patterns and label an activity as a fraud, we rely on machine learning to find the pattern automatically. Clustering any given group of students into a fixed number of groups is also a problem that often uses unsupervised learning. So is finding terrorist activity if given cellular activity in a network.


¹ Kelleher, J.D, Tierney, B, *Data Science*, London, 2018.



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=220#oembed-1>

This activity is adapted from activities created by Codeweeek and licenced under [Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International \(CC BY-NC-SA 4.0\) licence](#). You could find the [original list of activities in their website](#). The two datasets used, initial training dataset and test dataset, are also created by them.

We will use [Google's teachable machine](#) to train a machine to classify an image as a bicycle or a motorcycle. To recapitulate, a machine-learning application has to be trained and tested before it can be used. We will gather and group sample pictures of the categories the machine will classify, train the model and test if it classifies correctly a set of example pictures.



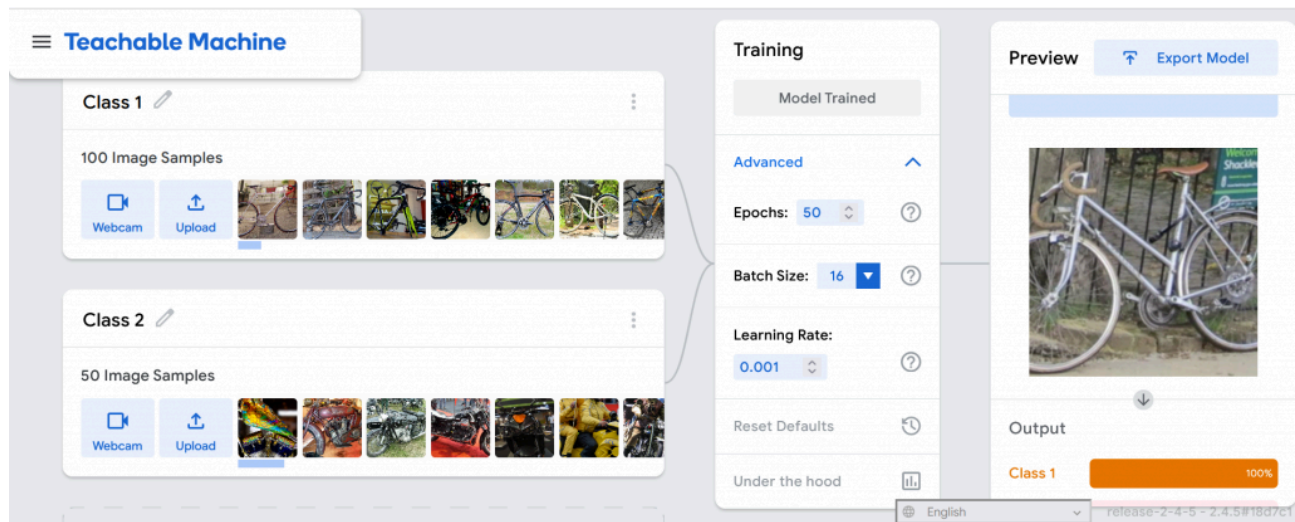
Step 1: Gather and group example pictures

1. Download the pictures of bicycles [found here](#).
2. If necessary, extract the zip folder contents into a local folder on your computer. This will be used as training set for the Machine Learning application.
3. Download the pictures of motorcycles [found here](#).
4. If necessary, extract the zip-folder contents into a local folder on your computer. This will also be used as training set for the Machine Learning application.
5. Download all the pictures [found here](#).
6. If necessary, extract the zip-folder contents into a local folder on your computer. This will be used as test dataset.
7. Click [Google's Teachable Machine](#) and select **Image Project > Standard Image Model**.
8. Under Class 1, click **upload > Choose images from your files** > Open the folder of bicycle images you created during steps 1 and 2 and import all the images stored here.
9. Under Class 2, click **upload > Choose images from your files** > Open the folder of motorcycle images you created during steps 3 and 4 and import all the images stored here.

Step 2: Train the model

Under Training, click **Train Model**. The model learns how to classify bicycles and motorcycles. Wait till it says **Model Trained**.

Note that we do not have to manually select and input features of bicycles and motorcycles. The algorithms knows how to find its own features from the images!



Source : Google's Teachable machine

Step 3: Test the model

1. Under **Preview**, click the arrow near **webcam** and change input to **File**.
2. Click **choose images from your files** and choose a test image you stored in steps 5 and 6.
3. Scroll down and check output.
4. You can repeat with other images to compare performance.

If an image is used to train a classifier, the machine will have already recorded the corresponding label for the particular image. Showing this image to the machine during the testing phase will not measure how well the model generalises. That is why your test and train datasets should be different from each other.

Note: You can also upload your own images to train and test. [Here](#) is a good source of free images.

```
HTTP/1.1 200 OK
Cache-Control: private
Content-Type: text/html
Set-Cookie: PREF=ID=5e66ffd215b4c5e6:
TM=1147099841:LM=1147099841:S=Of69MpW
Bs23xeSv0; expires=Sun, 17-Jan-2038 1
9:14:07 GMT; path=/; domain=.google.c
om
```

HTTP Cookie by Harmil is licenced under CC BY SA. To view a copy of this licence, visit <https://creativecommons.org/licenses/by-sa/2.0/>

Cookies are small files in your computer that tell the web browser that you are, say, user number #745673 on this website and that you like this and that. These cookies were conceived so that every time we visit the same site, we don't have to specify preferences such as language and location, lose items in the shopping cart or fill out forms from scratch. In the early stages of this technology, we had full control over what data the cookies could collect^{1,2}.

Afterwards, companies realised they could use cookie data to understand what we like to click on or buy. Thus, ads could be served that did not depend on the content of the current page, but our own personal tastes (behavioural targeting)¹. Later, companies began to set their cookies on other company websites to track every user even more closely. These third-part cookies paid the host for this privilege. This is when ads started following us across websites.

Moreover, by using things such as email IDs or credit card numbers, these companies could link the different identification numbers to a single user to have better information on their behaviour. This is called *cookie syncing*. The user of course has no way of knowing what data is being put together to build their behavioural profile.

To add to this, machine learning algorithms started to be put to use to crunch user data and assign them labels like man, woman, black, European or even “prone to depression”¹. These labels have nothing to do with our identities, but with what kind of prior user behaviour most resembles that of our own. These labels are sold to companies that sell products, houses and job opportunities. Thus, users with some labels are shown one ad and someone with a different online behaviour a completely different ad in the same web page. This in turn can determine what type of jobs we apply for and in which neighbourhood we buy a house and thus, which schools our children attend³.

Nowadays, cookie technology is embedded into most internet browsers. A 2016 study found that most third parties do cookie syncing. “45 of the top 50, 85 of the top 100, 157 of the top 200, and 460 of the top 1,000” third parties synch cookies from different sources to put together information on users⁴. It has been shown that Google, for example, can track a user across 80% of websites⁵ raising threats to privacy and autonomy, and bolstering surveillance and monitoring⁶.

When these results were published, they raised public outrage. Many cookie-blocking browser plugins became popular, such as DoNotTrackMe. Internet browsers started having controls to block or delete cookies². Companies such as Apple and Google even stopped or pledged to ban third-party cookies¹. Online targeting moved from cookies to more persistent tracking techniques.

For example, cookie-like files could be stored with Adobe's Flash player; these remain after other cookies have been deleted. These can in turn be blocked by installing apps such as FlashBlock². Tracking technology is equipped with more

persistent tools, such as various types of fingerprinting which are not detected by most blocking tools⁴.



"Fingerprint scan" by Daniel Aleksandersen is licenced under CC0 1.0 . To view a copy of this licence, visit <https://creativecommons.org/publicdomain/zero/1.0/deed.en>

The idea is that our devices and services, such as computers, phones and device speakers, process data and give output slightly differently from other users' devices. They can serve as our unique fingerprints, especially when the different techniques are put together to create our online identity⁴. The IP address of our devices, ethernet or Wifi addresses (WebRTC-based finger-printing), how our hardware and software play audio files (AudioContext fingerprinting) and even information on the battery, can all be used as our long- and short-term identifiers that

keep online tracking alive^{7,4}.

¹ Kant, T., *Identity, Advertising, and Algorithmic Targeting: Or How (Not) to Target Your "Ideal User"*, MIT Case Studies in Social and Ethical Responsibilities of Computing, 2021.

² Schneier, B., *Data and Goliath : the Hidden Battles to Collect Your Data and Control Your World*, W.W. Norton & Company, New York, 2015.

³ Barocas, S., Hardt, M., Narayanan, A., [Fairness and machine learning Limitations and Opportunities](#), 2022.

⁴ Englehardt, S., Narayanan, A., [Online Tracking: A 1-million-site Measurement and Analysis](#), Extended version of paper, ACM CCS, 2016.

⁵ Libert, T., [Exposing the Invisible Web: An Analysis of Third-Party HTTP Requests on 1 Million Websites](#), International Journal of Communication, v. 9, p. 18, Oct. 2015.

⁶ Tavani, H., Zimmer, M., [Search Engines and Ethics](#), The Stanford Encyclopedia of Philosophy, Fall 2020 Edition, Edward N. Zalta (ed.).

⁷ Olejnik, L., Acar, G., Castelluccia, C., Diaz, C., *The leaking battery*, Cryptology ePrint Archive, Report 2015/616, 2015.

The general practice of saving all kinds of data is called Big Data¹. Doing this makes sense since data storage has become cheap and powerful processors and algorithms (especially natural language processing and machine learning) make analysing big data easier².










One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://aiopentext.itd.cnr.it/aiforteacher/?p=230#oembed-1>

As discussed in the video, big data is characterised by huge (**v**olume), rapidly generated (**v**elocity), disparate types (**v**ariety) of data-generated from multiple sources. The data thus gleaned tends to be incomplete and imprecise (**v**eracity), and its relevance tends to change over time (**v**olatility). Sophisticated algorithms are required to combine, process and visualise this kind of data. Yet, inferences drawn from them, especially when combined with traditional data, can be powerful and thus, worth the effort².

Some experts go beyond the three or five Vs² and stress the three axes that make up big data:

- Technology that makes it possible to gather, analyse, link, and compare large data sets. Analysis that identifies patterns in large data sets in order to make economic, social, technical, and legal claims.
- The belief that “large data sets offer a higher form of intelligence and knowledge that can generate insights that were previously impossible, with the aura of truth, objectivity, and accuracy”³.
- Big data analysis “can potentially identify areas where students struggle or thrive, understand the individual needs of students, and develop strategies for personalised learning.”

Are you (Big) Data literate?

- Do you know what data is, what aspects of the  it represents?
- Can you identify when and where data is collected – actively or passively?
- Do you know what it means to work with data : to create, acquire,  and manage it?
- Do you know of techniques to analyze data : , sort, aggregate,  etc..?
- Do you know how algorithms identify  in data?
- Do you know how can data be used to argue or prove a greater point?
- Can you weigh the ethical impacts of data-driven decisions for  and for the  ?

¹ Schneier, B., *Data and Goliath: The Hidden Battles to Capture Your Data and Control Your World*, W. W. Norton & Company, 2015.

² Kelleher, J.D, Tierney, B, *Data Science*, London, 2018.

³ D'Ignazio, C., Bhargava, R., *Approaches to Building Big Data Literacy*, Bloomberg Data for Good Exchange, New York, 2015.

⁴ General Data Protection Regulation (GDPR), European Union, April 2016.

⁵ Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators, European Commission, October 2022.

Some types of learning and instruction, which are often used while discussing personalised learning:



Personalised Learning terms used in marketing materials and media. Source: Bulger M., Personalised Learning: The Conversations We're Not Having, Data & Society Working Paper, 2016, licenced under CC BY-NC-SA 4.0.

Blended learning

This is a cohesive mix of face-to-face-instruction and online learning. The teacher could either produce content for online learning or use a content provider such as Kahn Academy. Blended learning often involves a learning management system such as Moodle or Google Classroom that integrates both parts and helps keep track of student learning¹.

Competency-based learning

Content, skills and dispositions can all be described as competencies that a student has to acquire. In competency-based learning, once a student has demonstrated

mastery of one competency, they move on to the next. When they have trouble, they are given help^{1,2}.

This is an alternative to the traditional setting where progress is based on student' age and time spent in the classroom.

Differentiated instruction

“To differentiate instruction is to recognise students' varying background knowledge, readiness, language, preferences in learning and interests; and to react responsively”³. It is different from personalised learning in that it is not student-led and the learning objective is common to all learners – the only change is in the method of instruction.

For example, for learning new words, some students may be asked to find newspaper clippings that feature these words; others might be asked to learn a song.

Flipped classroom

This flips the classwork and homework. Students learn at home using online coursework or lectures. The teacher guides practice or projects when in the class².

Individualized learning

This is all about the pace in which a student learns⁴. If they are struggling, they can opt to spend more time practising what they learnt. If they are confident, they can move ahead and explore more content without getting bored.

Problem-based learning and project-based learning

Students learn by solving a problem or doing a project which may or may not be personalized, while teachers act as their facilitators or guides.

¹ Groff, J., *Personalized Learning: The State of the Field & Future Directions*, Center for Curriculum Redesign, 2017.

² Holmes, W., Anastopoulou S., Schaumburg, H & Mavrikis, M., *Technology-*

enhanced personalised learning: untangling the evidence, Stuttgart: Robert Bosch Stiftung, 2018.

³ Hall, T., Vue, G., Strangman, N., & Meyer, A., [Differentiated instruction and implications for UDL implementation](#), Wakefield, MA: National Center on Accessing the General Curriculum, 2003.

⁴ Michell, M., [Personalized, Individualized, and Differentiated Learning: A Simple Math Equation](#), 2016.

A number of technical terms are used in this chapter. [The AI Speak page](#) has some details and the history of natural language processing.

From a broader perspective, **automatic translation tools** are accessible online and can be used in many ways:

- Directly, by copy-pasting bits of text in one language and obtaining the same text in another language;
- By submitting to the website files in some predefined formats. The whole file will be translated, preserving the structure and formatting of the file;
- Through APIs (an *application programming interface* is a piece of code programmers can use in their software to facilitate translation tools).

Automatic transcription also makes use of artificial intelligence. This consists of transforming a voice input into a text output. It can be done either online or on offline videos or audio recordings. Some video conference platforms can obtain subtitles which can then be used to improve accessibility and/or to understand speech from another language.

Voice-synthesis techniques take text and a voice model to speak the text using that voice. Voice model can be standard or can also be trained to correspond to real people.

Text-generation tools are used to generate new text through artificial intelligence. This new text can be based on existing text, such as abstracts, simplifications or reformulations of existing text, or based on conversational models, where AI will be prompted on a topic.

These techniques can be used separately or be combined, in order to propose seamless multilingual conversations.

Automatic translation in education is a subject where the speed at which technology is moving makes analysis difficult. Research more than five years old will analyse the impact of a technology that is already obsolete. For example, it may insist that automatic translation tools don't work, and will report examples of malfunction which are no longer valid. Soon this could be the case with other fields in which artificial intelligence is the main player.

This has various consequences, including:

1. Research is going to be difficult. Research needs to build upon previous work to avoid rediscovering realities. But in this case, the giant's shoulders on which one would wish to build can quickly become obsolete.
2. Teachers will need to find ways to stay informed, as the non-linear progress of these technologies could be difficult to follow. Tools that help teachers with this technology, rather than opinions posted on social networks, should be envisaged.

In March 2023¹, and then in May 2023², two petitions/open letters were published, warning about some of the dangers of AI. They were signed in both cases by hundreds of reputed scientists and AI specialists from the industry. What do we need to know about this? ?

Is this an interesting question for teachers?

With the more mature students the question of the risks of AI to civilization will arise. Even if no teacher is obliged to give definite answers to all questions, it is fair to understand the contours of the controversy.

As a teacher, should one not just talk about the technical aspects of a topic and leave the human, economic, philosophical issues to specialists?

This is an interesting question over which there is a division of opinions.

Should a physics teacher know about Hiroshima or Chernobyl? Should conversations about these issues take place in that class? Or, in the case of AI, should a teacher be able only to use some software safely and understand generally how it works? Or, also be able to understand the running debates about the questions for society when it comes to AI?

Unesco's and other experts' position is that artificial intelligence isn't just about technology and that a teacher should understand the ethical issues involved. These include the concerns about the impact of AI on society, civilisation and mankind.

Are these new questions?

Some of the questions about the dangers of AI have been around for a while. The question about what happens when AI is 'superior' to human intelligence has been discussed for a long time. Irving Good⁴, a former colleague of Alan Turing, introduced the notion of *Technological Singularity* as early as 1965. He suggested that, once AI is considered to be more intelligent than humans, or super-intelligent, the AI would be unstoppable. Good went on to advise Alan Kubrick for 2001: A Space Odyssey, – a movie featuring AI going rogue.

The positions

The text of the March open letter¹ warned that AI could do good and bad, that the

impact on society and on jobs could be considerable. It also introduced the notion that AI was not only going to replace humans in tedious and undesirable jobs, but also in 'good' jobs that people wanted to do. Furthermore, that the developments of AI led to developments of society and that the usual democratic mechanisms of change were not used.

In the second text², the added risk discussed was that of AI going rogue (or a variation of this scenario) and the potential end of human civilisation.

A third position emerged from this debate³ – that AI was indeed a cause of concern, but not for existential reasons which were masking the more urgent problems.

Is the debate over?

No, the debate isn't over. Some scientists still claim there are many risks, that these technologies are growing too fast and that regulation is required. Others believe that at present AI comprises only benefits, and that we should be careful, but not scared.

It is difficult to say who is winning or losing, who is right or wrong. The debate is reminiscent of that which occurred after 1945 about physics.

A common position is that of asking for regulation, even if there is as yet no regulation with which everyone agrees.

Can there be a sound position?

Actually, both positions are probably sound. The current facts seem to be in favour of the enthusiasts (AI is enabling progress in medicine, agriculture, climate analysis, languages and communication), but the argument that we, as humans, have always found answers, has serious limitations.

Where do I find out more about this debate?

For an open-minded person (or teacher) there are numerous potential sources of information. Blogs, reliable sites, and position papers and videos from all leading scientists, including historians and philosophers.

¹ <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>

² <https://www.safe.ai/statement-on-ai-risk#open-letter>

³ <https://theconversation.com/lets-focus-on-ais-tangible-risks-rather-than-speculating-about-its-potential-to-pose-an-existential-threat-207842>

⁴ <https://www.historyofinformation.com/detail.php?id=2142>

⁵ https://en.wikipedia.org/wiki/I._J._Good

Generative AIs can be used to design interesting activities inside and outside the classroom. They will probably play a role in education, and in certain cases are already playing one. But what role, exactly? And how should a teacher engage their students? In what contexts? We will argue that, at this point, the teachers can and should restrict using generative AIs to the context of out-of-class activities.

“At this point”

Things have been moving incredibly fast. In October 2022, when the first version of this textbook appeared, ChatGPT didn't even exist. One year later, one can find platforms on the web proposing generative-AI-powered tools for education. The speed of progress is such that what is valid at this point (November 2023) may possibly no longer be true in a few months' time. Perhaps some of the flaws we are seeing today will be corrected. Perhaps teachers will have been offered enough training to work around these flaws. Perhaps the school or national authorities will have provided instructions as to what can or should be done. It is essential to stay informed.

“Policy questions”

AI is presenting ministries with hard challenges. On one hand, it is desirable to teach pupils in such a way as to prepare them for tomorrow's world. After all, the numbers showing how the jobs market will be impacted make it reasonable if not necessary to envisage teaching the pupils early¹. On the other hand, it may seem unsafe to use technologies which haven't yet shown their resilience. This lack of safety can be seen especially regarding privacy issues². It is still quite unclear what effect they will have on learning³.

Industry is pushing to make us adopt their products, while parents are focusing on contradictory messages. These are – the prioritising the teaching of the fundamentals (reading, writing, counting) and the necessity of learning skills associated with jobs. This division complicates the task of policy makers.

Decisions may take time, but when they come, teachers will want to understand them.

About ‘safe environments’

Much data is going to be exchanged during sessions with generative AIs. Teachers and pupils will possibly easily be giving away data which can rapidly become personal. And without the right implementations, this data can be directly associated to each individual. GDPR does protect individuals but it is still early to know if these laws will be sufficient. Some countries have introduced safe-school

environments in which anonymisation is the rule. In such environments, the online activities will not be logged outside the school servers in association with individual users.

Data safety questions are numerous, and it isn't easy for the teacher to be sure that their rights and those of their pupils are granted. How long are the data going to be stored? For what purpose and by whom will they be used? Can teachers make decisions on behalf of their pupils? The complexity of these questions explains why it is never a good idea to just register the pupils on external platforms, unless the authorities have made the necessary checks.

Out-of-the class activities

One can already find many examples of activities in which a teacher can engage with generative AIs, either at home or in an office, and without pupils. Between these, let us mention preparing classroom activities, writing tests, searching for information and exploring the topic of the next lecture. There is a general impression that in these situations AI allows to better explore, to find new ideas, to present material in a better way. And even if there are also a number of problems (lack of references, hallucinations, bias) the balance is generally seen as positive.

Most importantly, teachers are reporting gaining time. For once, technology is not just promising to do better, but to do better with less effort.

Arguments in favour of inside-the-class activities.

If generative AIs are set to play an important role in the future, and making sensible use of AIs will constitute a skill in the jobs market, surely pupils should learn, with a teacher, how to use them correctly. Indeed, this would address both technical and ethical aspects of AI.

Speaking with students about these technologies today is rewarding but worrying. On one hand, they are already users, but they have strong misconceptions, in particular when it comes to trust.

Arguments against inside-the-class activities.

On the other hand, anyone who has tested these tools will understand how difficult it is to teach with a tool whose output is so unpredictable. Run a generative AI three times with the same prompt and you will probably get three different results. This is, in fact, an asset for the technology, but it can bring an untrained teacher (but also a skilled one!) into a rather uncomfortable position. Imagine a chemistry teacher asking pupils to all run the same experiment, only to then observe a bang here, red smoke there and a strange smell at the back of the room.

It would prove to be interesting but quite challenging to give convincing general explanations... or even individual ones.

So...

At this point the teacher should safely be able to test generative AIs outside the classroom. This will help to better understand how it works but also to discover the possibilities the pupils will most likely find. Not remaining naive about generative AIs is essential. Furthermore, as more and more teachers are indicating through testimonials, this is the chance of using a technology which, for once, allows the teacher to save time.

On the other hand, in many situations, it is still a good idea not to use these technologies directly with the pupils.

So how do we help pupils understand?

Again, this will have to be in line with recommendation and rules set by national or school authorities.

Wherever teachers can do this, a first suggestion is to engage with the pupils, perhaps by asking, what is and is not cheating? Discussing this topic will help pupils to understand the complexity of the question.

A second suggestion is that a teacher could try generative AI in the classroom, but not to use it with a complex, unfamiliar topic. This may seem counter-intuitive but showing the pupils that one doesn't always know the answer can be helpful. It can even pay off to use generative AI on topics on which the students themselves will have expertise – they might spot mistakes and understand that AI is not always right!

Error-spotting can be interesting for pupils. It can be much more difficult for a teacher to be challenged by a fact produced by a generative AI and spot the error on the fly. This isn't about being right or wrong; teachers are allowed to make mistakes. But having to explain mistakes in a pedagogical way is not that simple.

¹ Generative AI likely to augment rather than destroy jobs. ILO report, August 2022 https://www.ilo.org/global/about-the-ilo/newsroom/news/WCMS_890740/lang-en/index.htm

² After Italy blocked access to OpenAI's ChatGPT chatbot, will the rest of Europe follow? Euronews, 7/4/2023. <https://www.euronews.com/next/2023/04/07/after-italy-blocked-access-to-openais-chatgpt-chatbot-will-the-rest-of-europe-follow>

³ Holmes, W., Miao, F., *Guidance for generative AI in education and research*, Unesco, Paris, 2023.

Transformers are a neural network model designed to overcome the limitations of recurrent neural networks in the analysis of sequences of data (in our case, words or tokens)¹.

Specifically, transformers, through the *self-attention* mechanism, make it possible to parallelise the analysis of data sequences and extract the dependencies between the elements of these sequences and the contexts in which they occur.

¹ Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I., *Attention is all you need*, Advances in neural information processing systems, 30, 2017.

The General Data Protection Regulation (GDPR), which came into effect on 25th May 2018, provides a legal framework for keeping everyone's personal data safe by requiring companies to have robust processes in place for handling and storing personal information.

The GDPR is based on seven principles and establishes rights for the citizens and obligations for the platforms.

GDPR's seven principles are: lawfulness; fairness and transparency; purpose limitation; data minimisation; accuracy; storage limitation; integrity and confidentiality (security); and accountability.

Let us mention some of these rights and obligations, particularly relevant in our context:

- **The right to be informed** specifies that a citizen has to be informed of the usage which can be made of their data;
- **The right to erasure** is what allows a citizen whose data has been collected by a platform to ask for their data to be removed from the dataset built by the platform (and which may be sold to others);
- **The right to access** means that the citizen can know (easily) what data is being collected about them.

Even if the GDPR was written before the main questions about AI and education became important, the framework does address a lot of issues about data. Since data is the petrol AI thrives on, GDPR is particularly relevant for AI and education.

Rather than giving our own easy-to-understand explanation on what GDPR is and what a teacher should understand, let us recommend looking at a website that has done this simplification work for us.

The name of the website, of [“GDPR for dummies”](#) may irritate you (teachers are not dummies). But the analysis has been done by independent experts from *the Civil Liberties Union for Europe* (Liberties), which is a watchdog that safeguards the human rights of everyone in the European Union.

Writing computer code is like writing in any language. There are syntax (or grammar) rules to be observed; we want the program to be meaningful and do what we want, ie, to respect semantics. In 2022, generative AIs that built code from prompts appeared; ChatGPT allowed for this directly within its interface, making Python or C languages on par with French, Italian or Japanese.

Quickly, a debate emerged – since AIs were good at producing code, is it still necessary to learn code? For the many who couldn't code, there was little doubt, and the claims of the industry that AI could produce good-quality code were sufficient. In the industry, at the end of 2023, some jobs were lost from humans to AI, but on the whole managers are hesitant to replace programmers with AI. There are still the issues of hallucinations but, more importantly, it soon appeared that you could only get good code if you could write the correct prompts, or in other words, specify correctly. Furthermore, as prompting is usually not one shot and requires some form of dialogue, it is useful to understand the partner's language. This is a skill which usually comes from long hours of practising coding.

The current attitude seems to be that if humans are not necessarily going to be the ones writing future codes, there is a need for people who know how to code to interact with AI in order to get the code to work.

Code, no code, low code

On the other hand, if high-quality coders are needed to work with AI on complex systems, should everyone reach that level? The answer is probably not. As often, things aren't always black or white, and there is probably room for an intermediate level between no code and code, often called low code.

Are dolphins mammals, and if so, why?

What species of animal is a kiwi?

Are Athens and Rome in the same weather zones?

Who painted a famous painting with a screaming lady? Does this artist always paint the sky in orange?

Can we guess the author of a post on social media only from the style of writing?

In style of playing, does Luka Dončić have a clone in the NBA?

Can the species of the tree be guessed from its leaves? Or from a photo of its bark?

How do world countries group by their socioeconomic features? Is the world indeed socio-economically divided to north and south?

Are Rome and Athens in the same weather zone? According to the weather patterns, which country's capital is most similar to Berlin?

Data science, particularly machine-learning methods, serves as catalysts for change across various fields such as science, engineering and technology, significantly impacting our daily lives. Computational techniques capable of sifting through extensive data sets, identifying intriguing patterns and constructing predictive models, are becoming omnipresent. However, only a few professionals possess a fundamental understanding of data science, with even fewer actively involved in building models from their data. In an age where AI quietly moulds our world, everyone must be aware of its capabilities, advantages and potential risks. We must establish methods to effectively communicate and teach concepts related to data science to a broad audience. The principles and techniques of machine learning, data science and artificial intelligence should become common knowledge.

Every question posed at the start of this chapter can be answered by observing relevant data. We propose the following approach to training machine-learning, Start with the question, find relevant data and then answer the question by means of finding relevant data patterns and models. In project [Pumice](#), we are developing educational activities that can be used to enrich different school subjects. We use data related to the subject matter and explore it using AI and machine-learning approaches. In partnership with educators, we have developed learning templates and background explanations for teachers and students.

Pumice's activities and training are supported by [Orange](#), a machine-learning program that features an intuitive interface, interactive visualisations and visual programming. The key to simplicity is a Lego brick-like construction of analytical pipelines and interactivity of all the components (see Fig. 1). This is required in the training and versatility to cover most of the core topics and adapt to various application areas. To further support the teaching and to focus on concepts rather than on underlying mechanics, Orange implements easy access to data, reproducibility through the saving of workflows with all the various user-based settings and choices, and easy customisation through the design of new components. A critical aspect of the training involves storytelling through workflow inspection and specialised features for experimentation, such as drawing the experimental data sets or learning about the over-fitting of polynomial linear

regression. Orange is available as open-source software and is complemented by a [concise training video](#).



Fig. 1. Orange data mining software and a typical data exploration workflow.

In Fig. 1 we show a typical Orange data-exploration workflow. The workflow consists of components that load the data, compute the distances, visualise the data or resulting models, or perform any task necessary to find and visualise data patterns. In this workflow, we have used the socioeconomic data from world countries. The upper branch of the workflow studies two features and shows that life expectancy and years spent in school are correlated. It also shows that there are countries such as Cabo Verde and Morocco where people live long but do not spend too much time in school. Children in schools can design such networks to explore which countries are socio-economically similar to each other and where they can find that the world is socio-economically split into north, central, and south and that there is a great divide between developed and underdeveloped parts of the world. There is no need to tell them this explicitly – with mining the data in Orange, they will do so regardless and, in the upper classes, dive into these divisions on their own.

The development of Orange began in 2003. Since then, it has gained substantial traction. With over 50,000 distinct monthly users, Orange has established itself as a widely embraced specialised software application. Approximately half of its users hail from the academic sphere. Notably, Orange has experienced a notable upswing in adoption within the education sector, with over 500 universities across the globe incorporating it into their data-science courses.

If you're an educator keen on delving into the realms of machine learning and data science, here's a compilation of resources that provide an introduction to these disciplines through hands-on data exploration using Orange:

- Orange, the toolbox's [website](#)
- [An introduction to data science](#), is a set of short videos that showcase selected

visualization and machine learning methods with Orange. Find videos on <http://youtube.com/orangedatamining>, and go to the “Intro to Data Science” playlist.

- [Pumice](#) is a teacher website where we gather use cases that you can incorporate into your training program.

While reviewing this textbook in January 2024, the complexity of the challenge that it represents struck us with force. Artificial intelligence is today not just a highly important issue; it is also, more than anything, **an issue**. Where does it lead us? How many times will it make us realise that something we claimed as impossible yesterday has become a reality today? How is it even possible to write a textbook on such a topic without making mistakes?

In the context of education, we do not often have to face such fast-changing topics. It takes months to design a syllabus and years to train the teachers. And some say it takes a generation to comprehensively treat a new topic.

Yet here, the goal of the AI4T project and thus, of its learning resources, has been to train teachers in AI while it unfolds! Therefore, when reading this textbook, you may well find small errors. You may also find statements that are no longer true: technological progress can have proposed a new opportunity; risks that were deal breakers in 2022 or 2023 may well have become acceptable in 2024; concerned laws and regulations may have changed. Does this disqualify this effort?

We hope not. Not if we take the opportunities offered to us by technology and the **open** nature of this textbook to allow it to evolve with the changes. Perhaps the old-fashioned process of releasing the first edition, then the second edition etc is no longer the best way to edit a book? Definitely, there exist solutions to allow such an object to be shared as **a common good**, not just in its usage but also in its evolution? This would represent, probably, the next challenge.

This textbook is now in the hands of teachers, its readers. For this to have happened, a large number of people were involved. They helped us understand the needs of the teachers and those of the ministries. Some were involved in preparing the material, in proofreading and in setting up the online platform. Others helped with the translation. We also worked with external teams: Dagobafilms were great in helping us to prepare the videos. And, most importantly, there were all those who encouraged us. Writing a book is always a long and complex task, and you need the help of family, friends and colleagues to be able to recover from the associated frustrations, to stick to what we believe are the right editorial choices and to make sure that the destination that we had set is still where we want to go.

When thanking people who helped us over the past three years, we are bound to forget some. We hope they will forgive the vagueness of our memory. Having said that, here are the names of those whom we thank for contributing in some way to this work:

Alain Thillay
Andréane Roques
Anne Boyer
Anthony Kilcoye
Axel Jean
Azim Roussanaly
Bastien Masse
Blaž Zupan
Borut Čampelj
Catherine Lemonnier
Catherine McD.
Daniela Hau
Dario La Guardia
Deirdre Butler
Dejan Zemljak

Fabrizio Falchi
Giuseppe Città
Helena Fošnjar
Iva Stegne
Jean-Jacques Baton
Jiajun Pan
John Hurley
Lucie Grasset
Manuel Gentile
Mélanie Pauly-Harquevaux
Michael Hallissy
Salvatore Perna
Sašo Stanojev
Solenn Gillouard
Petra Bevek

Jotsna would like to thank Thomas and Laya for their understanding, support and (almost-) constant love.

Colin is indebted to Isabelle's patience over the past months.