

Generative AI and machine translation

PRACTICAL MANUAL FOR SCHOOLS AND COLLEGES

Authors: Dr Gys-Walt van Egdom (University of Utrecht)
and Emma Hartkamp (European Commission)

This module is part of the reading material for European Parliament Ambassador Schools but can also be used by other secondary schools or intermediate vocational training establishments.

Though aimed at teachers in those establishments, it is also suitable for older secondary school pupils and vocational training students. Our goal is to raise awareness of AI, particularly among young people.

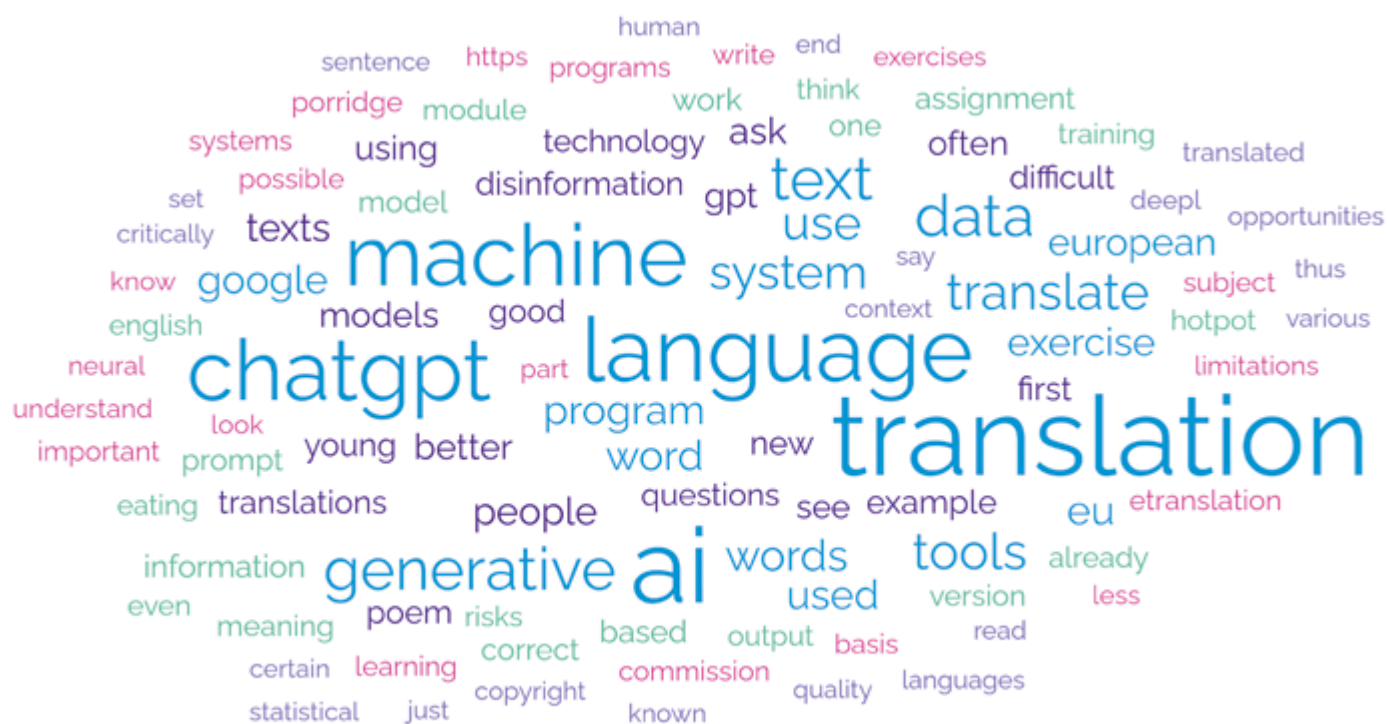
This first edition has benefited from the contributions – including constructive comments and suggestions for exercises – of Christophe Declercq, Anne-Marie Eekhout, Marloes Meijer, and Jonah Zijlma. Translation of the module into English by Neil Bennett.

A summarised version of this module can be found in the accompanying PowerPoint presentation, which can be used for classroom teaching. The module itself is also available online.

The Hague, September 2023, first edition

Table of contents

Introduction	5
PART 1. GENERATIVE AI AND MACHINE TRANSLATION	6
1A. Generative AI	6
What is Generative AI?	6
A short overview of the origins of Generative AI	8
Opportunities and limitations of Generative ai	14
1B. Machine translation	17
What is machine translation?	17
A short overview of the origins of machine translation	17
Opportunities and limitations of machine translation	22
PART 2. WHAT IS THE EU DOING IN THE AREA OF GENERATIVE AI AND MACHINE TRANSLATION?	24
2A. European legislation in the area of AI	24
2B. eTranslation and a European ChatGPT?	26
2C. European projects against disinformation	26
PART 3. EXERCISES	29
3A. Quiz questions: test your knowledge of AI language tools!	29
3B. Practical exercises using AI	30
3C. Translation exercises	34
3D. Discussion topics	37
Epilogue	39
Links	40
Sources	41



Introduction

Artificial intelligence (AI) is all over the news, primarily because of the breakthrough in *Generative AI*.

Initial studies have shown that young people in particular are all too ready to use this new technology, which has led to tools such as ChatGPT, DALL-E and Bard. In society, there is optimism about promising applications of AI, such as superchatbots and tools to help elderly people to live at home for longer, to facilitate communication between refugees and doctors, or to help with the rapid detection of diseases. However, the potential risks to democracy and the rule of law are also being highlighted ever more frequently, and calls are growing for conditions and rules to be laid down, in order to counter disinformation, discrimination and fraud, and protect privacy and copyright. The ecological footprint of AI is also increasingly being questioned.

In the first chapter of this module, we focus on Generative AI and machine translation. We consider the main concepts and tools and look at the opportunities and limitations of them.

Chapter two describes what the EU is doing to give everyone access to safe tools and to protect European democracy.

The last chapter contains some quiz questions, exercises and discussion topics to help people use these tools critically, efficiently and with awareness of the risks and limitations.

We want to avoid polarisation and facilitate discussion around AI in the classroom, thus ensuring that students and pupils can find the right balance themselves. We will try to frame the debate about banning translation tools and AI from the classroom or universities, or indeed embracing them, in a less divisive and more nuanced way, so as to find the correct balance between smart use of the new technology and respect for 'traditional' language skills and critical thought.

This module will lead to a better understanding of the new tools, and guide pupils and students in their use. Solid language teaching and education in textual analysis will continue to be of fundamental importance. Just as numerical skills are still needed to use a calculator properly, young people need language knowledge and skills to use language technology critically. By focusing on both skill sets – smart use of digital tools, and language skills – our aim is not only to encourage young people to think critically about AI, an essential skill for the future, but also to help them to contribute to European democracy as digitally literate citizens.

PART 1. GENERATIVE AI AND MACHINE TRANSLATION

1A. Generative AI

WHAT IS GENERATIVE AI?

Generative AI (or GenAI) is a form of artificial intelligence that can generate new things, such as texts, pictures or sounds. Well-known examples of programs based on Generative AI are ChatGPT and DALL-E (by OpenAI) and Bard (by Google). The technology is developing at breakneck speed, and new models and programs keep appearing. In Generative AI, the user gives the program a task, known as a prompt, usually in the form of text. The AI model analyses the task and then produces new content based on what it has already learnt from an enormous corpus of existing information or data.

An example:

- In your prompt, you ask for a still life painting in the style of Picasso.
- The program analyses your prompt, focusing in particular on terms such as 'still life', 'style', 'Picasso' and 'painting'.
- The program searches for examples it has used in the data set to learn about still lifes, styles, Picasso and paintings.
- The program formulates a reply by combining these various types of information from the data set.

The reply certainly shows similarities with all these sources of information, but it is novel and appears unique. By adjusting the prompt, the program can be made to search for other examples and combine them differently. The results of small adjustments to prompts given in DALL-E are depicted on the next page. Under each image, you can read the prompt used.



Paint a still life in the style of Picasso.



Paint a still life in Picasso's cubist style.



Paint a still life in the style of Picasso's 'Blue Period'.



Sketch a still life in the style of Picasso.

A SHORT OVERVIEW OF THE ORIGINS OF GENERATIVE AI

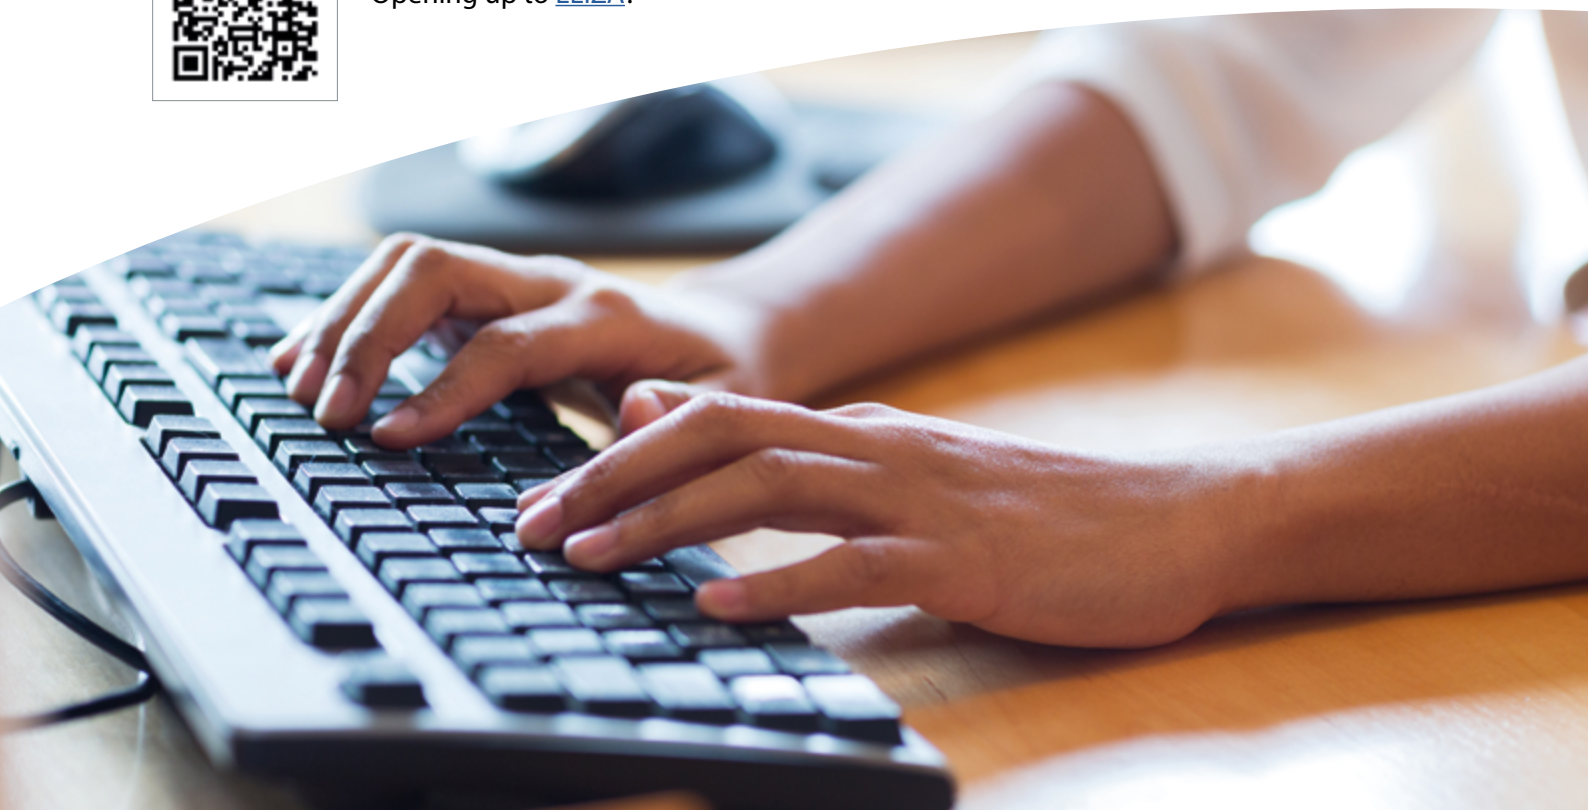
How did Generative AI come about? Generative AI has already come a very long way. It is difficult to say exactly when it started. People have always been fascinated by how language works, but particularly since the Cold War have they begun to consider it more systematically. The idea on which all language technology is based is that language follows particular patterns to obtain meaning. In other words, find the patterns and you have the key to language technology.

Many see **ELIZA** as a precursor to the Generative AI we know today. Joseph Weizenbaum worked on this 'chatterbot' between 1964 and 1966, and the program was able to hold a decent conversation. At least, for many users, it appeared that the program actually understood what you said. The example that Weizenbaum himself quoted to praise the program is well known: after exchanging a few phrases with ELIZA, Weizenbaum's secretary asked him to leave the room, because the conversation was becoming so personal. However, there was no question of actual 'understanding'. The program analysed the user's input, recognised certain words and conjured up a number of standard responses.

ELIZA was followed by a great deal of research activity, but as experts focused too much on the idea that language could be broken down into simple mathematical models, the results were not particularly impressive.



Opening up to [ELIZA](#)?





Neural networks attempt to emulate human cognitive structures.

Machine learning

Perhaps the most significant step in the development of Generative AI was *machine learning*. Over the 1980s and 90s, the realisation began to dawn that data could be used to develop highly complex statistical models that had much more to offer than fully human-programmed models. In machine learning, data is fed into the system, and the logic behind the data is 'explained' to the computer, as it were. The ingenuity of machine learning lies in the program's self-learning abilities: after the first step, the program no longer just reacts to new inputs but analyses them by comparing them with data already in the data set. Analysing new data leads it to adjust its internal parameters. You could say that, in machine learning, the model is constantly adjusting itself.

In the 1980s and 90s, it was already known that this *data-driven* approach could achieve very impressive results, but no real breakthrough in Generative AI was achieved. Why not? Anyone who was using a computer at that time can already guess the answer. Computers could still not do very much. Their processing power and storage were limited. Moreover, the internet was in its infancy, meaning that there was not enough data readily available to make any real progress.

Neural networks

Not until around 2010 did interest take off again. Some computational linguists were able to demonstrate that a great deal more than previously imagined was possible using techniques that had been gathering dust. As experiments could now be conducted with larger data sets, and these data could be analysed by hugely powerful computers, the experts began to realise that *neural networks* were the key to new language technologies. A neural network is a complex statistical model. In these networks, too,

mathematical and statistical principles are used to model relationships and patterns in language and other data. You could say that, in the old language models, a number would simply be linked to a word. In a neural network, not just one number, but numbers or 'values' at several levels are linked to the data. The result? The model is able to gain an enhanced impression of the 'meaning' of the data.

Word embeddings

Word embeddings are typical examples of this enhanced impression. After intensive training, 'values' are attributed to words at numerous levels, with each word obtaining, as it were, a unique set of coordinates. Let's take the word 'hotpot'. The program attributes a coordinate to 'hotpot' by analysing in detail the various contexts in which the word is used. As 'hotpot' is often used in the context of 'eating', the coordinate of 'hotpot' is located fairly close to that of 'eating'. The coordinate of 'broth' is also located nearby, and the same applies to the coordinates for 'cooking', 'beef', 'oil', etc. So, on the basis of many examples, the program knows that certain words fairly often share a particular context. The link between 'hotpot' and 'eating' will therefore be stronger than that between 'hotpot' and 'cookbook'. This is merely because the words 'hotpot' and 'eating' are more often found close to each other in a data set. But the link between 'hotpot' and 'cookbook' will still be much stronger than that between 'hotpot' and 'jumper', for the same reasons – it is relatively rare to find the words 'hotpot' and 'jumper' in the same sentence.

Over the past decade, computer scientists have been slowly refining the statistical language models. They have done this mainly by looking at the various ways these models could read language and other data. A very interesting discovery was made in 2017, when it was found that data analysis sometimes put too much focus on irrelevant words or elements. Indeed, programs attributed relatively significant value to meaningless elements of language, with articles (e.g. 'the', 'a') being analysed in a comparable way to nouns or adjectives (e.g. 'lamb' or 'cooked'). 'Attention is all you need' was the slogan from then on. With the right focus, analyses could be run much more efficiently, and, above all, better. The standard way of working, in which whole sequences, such as whole sentences, are analysed from beginning to end, was dropped.

BERT (Bidirectional Encoder Representations from Transformers)

The first programs based on 'Transformer architecture' were developed in this period. With this cutting-edge technology, far better connections of meanings were established, and suddenly the order in which the data were processed became less important. Google developed a series of language models under the name '**BERT**' (*Bidirectional Encoder Representations from Transformers*). These models were also based on word embeddings, but the idea of the meaning of words was better, as account was taken of the context both before and after the word used. As part of the training, the program was repeatedly given exercises in reading comprehension. Take a sentence like: 'This evening I am going to ... hotpot at my parents' house'. BERT was instructed to guess the missing word, on the basis of all the example sentences it had already processed. As humans, it's easy to guess the answer: 'eat' would probably be at the top of our list of most likely solutions, perhaps followed by 'make'. For language models, for a long time this exercise was quite difficult, but BERT often came up with some suggestions that hit the mark: 'eat' was the prediction. The program was also able to deal nicely with nuances of meaning. If the sentence was changed to 'This evening I am going to ... hotpot *for* my parents', 'make' was suggested as the most likely solution.





Training AI requires enormous amounts of data!

GPT (Generative Pre-Trained Transformer)

ChatGPT – GPT stands for *Generative Pre-Trained Transformer* – does away with the bi-directionality. In other words, the program no longer looks at what comes before and after a word. As the experts say again and again when they are asked what the program actually does: ChatGPT just predicts the next word. You could compare it with an autocomplete function: as you are typing a sentence, you can see how the prediction made by autocomplete becomes ever more accurate. In this way, the program is able to generate texts. Its apparent understanding of what is being asked in a given prompt, then, once again has everything to do with our 'attention principle' – the program formulates a hypothesis around the meaning of your input, based on the words used (which have specific coordinates in the model). That is why something known as 'prompt engineering' has recently become so important. The program does not understand language like 'real' people – who can read a great deal between the lines. Its hypothesis around the meaning of your prompt simply becomes ever more accurate the more precisely you word it. What many users do not realise, then, is that what you put into the chatbot is the basis for what you get out, and how usable it will be.

Reinforcement learning

In any event, GPT models have come a long way over the past few years. In 2018, OpenAI came up with the first model: GPT-1. Although that model seemed promising, GPT-2 was launched a year later. With each new model, the main difference was the volume of data used and the number of parameters on which the system was trained. Around 40 gigabytes of text were used for training GPT-2, while GPT-3 was fed with no less than 570 gigabytes of text. The number of *parameters* also continued to grow. Parameters are the adjustable weights and settings within the neural network. They determine how the model processes input and how

it generates output. Parameters are learnt as the model is trained. They make it possible to explore patterns and structures (word use, syntax, etc.) in language and ultimately to reproduce them. GPT-1 'only' operated 117 million parameters. Its successor, GPT-2, had 1.5 billion parameters to analyse and produce language.

GPT-1 already performed better than other language models; GPT-2 convinced people at OpenAI that the future belonged to GPT models. By 2019, the system could already produce coherent texts, and some human users would have had the feeling that they were talking to a real person. But there were still plenty of points for improvement. Sometimes the program would spew out nonsense or factually incorrect information. Even more concerning was that GPT-2 could generate offensive texts. The great success of GPT-3 (and its beta version, GPT-3.5, which is used in the free version of ChatGPT) is, in point of fact, the result of *reinforcement learning*.

Reinforcement learning is how a language model learns to differentiate between desirable and undesirable output. Simply put, the system receives a reward or a penalty after having produced a response. OpenAI notoriously used many underpaid Kenyan workers during this phase of development, which it described as reinforcement learning 'with human feedback'. Human assessors went through numerous responses to a prompt and then ranked them, from good to bad. In that way, the language model could be fine-tuned, and GPT obtained a better idea of the quality of its own output, without the model having to be completely retrained from scratch. The output given by GPT-3 is factually more correct, less offensive, and better geared to the expectations of users.

GPT-4 has since been launched. It is integrated in the paid version of ChatGPT, and operates in the background in Bing Chat and certain apps.

With respect to the various GPT models, it is important to mention the necessary disclaimers. ChatGPT appears to work in a way that is language-dependent, but it actually lumps every language (including visual language) together. So, what looks like a network of words is really just a mishmash. These words have to come from somewhere. So where do they come from? The answer is: from the entire internet. It is no surprise, then, that ChatGPT appears particularly promising in English, as some 60 percent of the internet is filled with texts in English. In other languages, ChatGPT's performance is rather less impressive. Even in Dutch, it sometimes produces sentences that sound rather English, and, in less-represented or less-resourced languages, the output will be even less convincing.

OPPORTUNITIES AND LIMITATIONS OF GENERATIVE AI

Opportunities

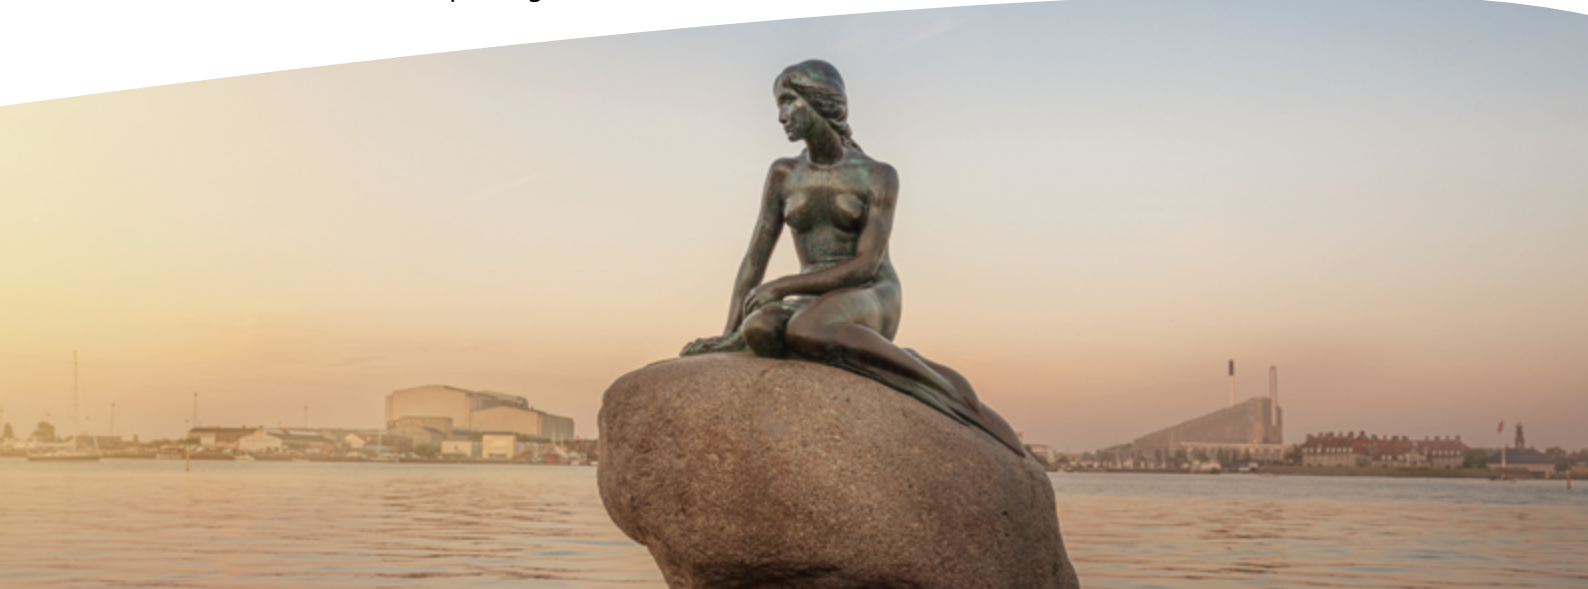
There are thus ever more programs and tools available based on Generative AI, such as ChatGPT and Bard, for generating text, and DALL-E for graphic production.

ChatGPT can be used as an intellectual assistant or writing aid. The bot can be asked for ideas, can correct a text for language or spelling errors, or make paragraphs hang together better. It is also often good at explaining concepts found in woolly texts in a clearer way. It can be used as a search engine, can translate texts, summarise them or rewrite them in more simple language, and give feedback on texts. ChatGPT can, in addition, write code for computer programs. There are also many lesser-known tools, such as Boomy, which can write songs on the basis of a description of a particular style of music and tempo.

Limitations and risks

We should mention certain misunderstandings and some of the main limitations and risks in order to encourage critical use of these tools. For example, the tools still have no real sense of context, and no consciousness or identity. That sounds logical, but instinctively we are inclined to attach logic and meaning to language, even when it is generated by inanimate objects. Just think of Joseph Weizenbaum's secretary. The tool cannot 'check' any information it gives by going back to the source; even in applications where reference is made to sources, such as with the Perplexity tool, the information is linked to scattered fragments within these sources. If there is not enough data, the bot is inclined to make things up. We call this 'hallucinating'. Using certain prompts you can even encourage hallucinating (see practical exercise 2 at the end of this module). It is not always clear why the machine hallucinates, as is clear from this example: To the question 'Does the Little Mermaid feel pain

The Little Mermaid in Copenhagen, Denmark



when walking in the Disney version?' Perplexity's incorrect (and incoherent) reply is: 'Yes, in the Disney version of "The Little Mermaid", the Little Mermaid feels pain when walking on her new legs.'

However, if you consult the sources, you can see that they contradict Perplexity's answer ('In Disney's adaptation, the magic used to transform Ariel into a human is quick and painless, so there is no real cost other than her gorgeous voice... which is also taken painlessly.') Incidentally: the comparison between Disney and Andersen appears convincing, but even this analysis is not correct (in the original Danish version of Den Lille Havfrue: 'er som du traadte paa en skarp Kniv, saa dit Blod maatte flyde.', '... as if you were treading upon knife blades so sharp that blood must flow.' [translation by [the Hans Christian Andersen Centre](#)]). This example, it must be emphasised, constitutes a relatively subtle form of hallucinating, in which appropriate sources are at least mentioned – albeit entirely incorrectly. Hallucinations are often more flagrant when no such sources exist at all.



Does the little mermaid feel pain when walking in the disney version?

⚡ Quick Search

The Little Mermaid - Disneyfication: The...

weebly • 1

How Disney's 'The Little Mermaid' Changed the...

collider • 2

The Little Mermaid Wisely Changed The Original...

screenrant • 3

10 Ways Disney's 'The Little Mermaid' Differs...

collider • 4

What the Real (not Disney) Little Mermaid Has to Sa...

medium • 5

10 Dark Storylines That Disney Movies Left Out -...

whatculture • 6

☰ Answer

Yes, in the Disney version of "The Little Mermaid," the Little Mermaid feels pain when walking on her new legs. According to the search results, the transformation and every step the Little Mermaid takes on land is painful, like walking across knives ¹ ² ³ ⁴ ⁶. This is a significant difference from the original story by Hans Christian Andersen, where the Little Mermaid's transformation is also painful, but the pain is described as being like being cut in two rather than walking on knives ¹ ³ ⁵. Therefore, in the Disney version of "The Little Mermaid," the Little Mermaid feels pain when walking on her new legs.

🔖 [] 📄 ✎ ...

Limitations of this kind mean that Generative AI is associated with certain risks, for example in the areas of *disinformation*, *discrimination*, *copyright* and *privacy*.

- **Disinformation:** AI tools can spread disinformation because they are unable to properly check the information they generate. Think of hallucinations. Moreover, Generative AI can be used to deliberately create disinformation much more easily. For example, some systems can convincingly imitate someone's voice or create a likeness of them to the point where it is difficult to recognise that it has been generated by a machine.
- **Discrimination:** Generated texts can be discriminatory, sometimes in a subtle way. This is because large-scale language models have been fed with training data originating from (almost) anywhere on the internet. For example, certain 'subreddits' (a type of forum), which were known for the use of openly discriminatory language, were used for the predecessors of today's ChatGPT. Moreover, it is very possible that some social groups are under-represented in the training data. This means that AI tools could be unintentionally discriminatory by taking less account of the experiences of such groups.
- **Copyright:** There are significant concerns about the infringement of copyright through the processing and reuse of texts to which copyright applies for training purposes.
- **Privacy:** Users' input, including personal data, is used for language models. This also brings with it a risk of fraud.

So, always be sceptical about the output of AI tools.

What can you do to reduce the risks when you work with them?

- Be alert to the possibility that they may give **incorrect information**. The only way to guard against this is (some) knowledge of the subject.
- Be aware of the potential for **prejudice and discrimination**; for example, a chatbot may assume that a business letter is being addressed to a man.
- Check whether the text contains any personal data that could constitute an **infringement of privacy**. Never share confidential information with public AI tools.
- The risk of **infringing copyright** is difficult to control, unfortunately, as it is not feasible for a 'normal user' to monitor ChatGPT in that respect.

It is also important to consider the large ecological footprint of AI tools. Storing and, above all, processing the data to train algorithms in data centres or in the cloud are very energy-intensive processes. An algorithm being trained to recognise a cat, for example, must process millions of pictures of cats. Some data farms, above all in northern Europe and Canada, use as much energy as a small town. The ICT, including data centres, used for AI is comparable with aviation in terms of its emissions. Finally, we are creating a new dependency on big tech companies, as significant amounts of capital are needed to make and maintain Generative IT systems and make them available.



1B. Machine translation

WHAT IS MACHINE TRANSLATION?

A *machine translation*, as the name implies, is a translation produced in its entirety by a machine translation program. Machine translation programs, such as Google Translate and DeepL, are based on technology which is comparable to that used for Generative AI. They use neural language and translation models which are designed to analyse the relationships between words and phrases in different languages. Machine translation services also work on the basis of a task, normally a very specific one: to translate a source text. Machine translation programs thus serve a specific purpose: to produce an accurate and fluent version of the source text in a target language.

It is, in fact, also possible to translate texts using Generative AI. The language models on which Generative AI is based do not differentiate clearly between languages. Real machine translation services are often trained using large quantities of multilingual data, for example, original texts and their translation, and can thus better recognise languages and language-specific patterns. Translations produced by machine translation services are therefore often better than those done by Generative AI.

A SHORT OVERVIEW OF THE ORIGINS OF MACHINE TRANSLATION

Machine translation has similar origins to Generative AI. Hundreds of years ago, philosophers including Ramon Llull (13th century), Blaise Pascal and Gottfried Leibniz (17th century) were already thinking about techniques (or technologies) for making 'machine translation' possible. However, it was not until the twentieth century that the technology really got off the ground. In 1933, French-Georgian Georges Artsrouni applied for a patent

for a 'cerveau mécanique', a mechanical brain within an inbuilt mechanical dictionary. That same year, the Russian Petr Troyanskii applied for a patent for an actual 'translating machine'. Yet work on machine translation did not really get under way until the 1950s.



During the Second World War, governments noted how important it was to intercept information from the enemy and translate it quickly. After the War, the need remained: the Russians and the Americans, even though they had fought shoulder-to-shoulder against the Axis (Germany, Italy and Japan), soon found themselves on opposing sides. The Russian and American Governments invested hugely over many years in translation technology. A spirit of [optimism about AI and machine translation](#) prevailed: for a long time it was thought that human translators would be superfluous within five years.

Rule-based systems

Why so much optimism? Experiments were carried out on *rule-based systems*. The machines were fed multilingual word lists and the grammatical rules of various languages. Imagine a simple sentence like: 'I am eating porridge'. To understand the meaning of the sentence, it is not enough to look up the individual words separately ('I', 'am', 'eating', and 'porridge'). So the experts thought that the best idea was to add grammatical rules. Once the system understands that the sentence is in the present continuous tense (with a present participle) and it knows how to reproduce this verbal form in another language (Dutch, say), a good translation can be provided: 'ik ben pap aan het eten'. Computer scientists therefore did their level best to cram their computers with words and grammatical rules.

The ALPAC (Automatic Language Processing Advisory Committee) report came out in 1966. It was damning. Despite all the investment, machine translation services still made the silliest errors. Language turned out to be more elusive than initially thought.

Statistical machine translation

It was not until four decades later that a glimmer of light reappeared on the horizon. During the 1980s and 90s, in the back rooms of IBM, they were playing with *statistical machine translation* (SMT). The idea was actually quite simple: by feeding a computer program with huge volumes of multilingual data, in other words source texts and translations, you could conduct very targeted probability calculations. Hence the description 'statistical'. Let us return to our example: 'I am eating porridge'.



An SMT system will first analyse the sentence, looking at its structure, the words, and the grammatical elements. A slightly more advanced SMT system will recognise 'I am eating' and 'porridge'. Then it will look at the data with which it has been fed, looking for the translations (into Dutch, say) of 'I am eating' and 'porridge'. The translation popped out by the system will be a combination of the most probable translation of 'I am eating' ('ik ben aan het eten') and 'porridge' ('pap'). Around 2006, Google rolled out its famous 'Google Translate', based on the principles of SMT. Over the first few years of operation, the quality of Google Translate translations seemed to get better and better. But then the optimism died down: around 2010, the translations stopped improving, and the system all too often became bogged down in matters of meaning, grammar and context.

Neural machine translation (NMT)

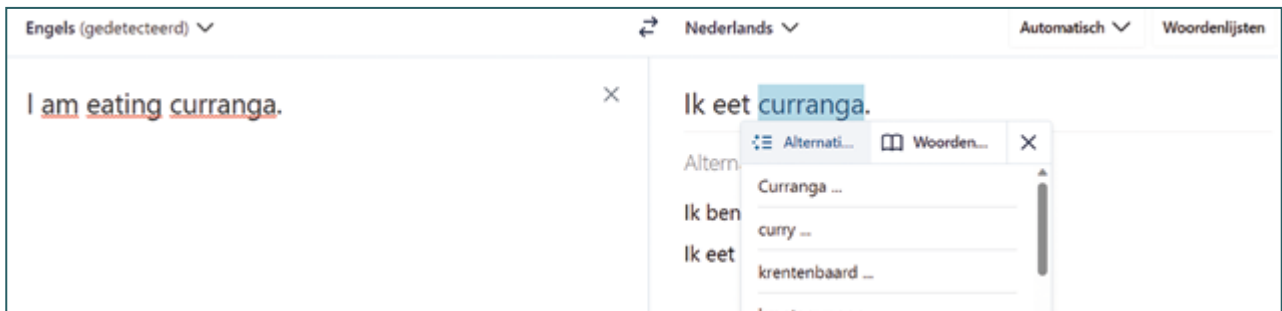
Meanwhile, the developers of machine translation systems had also learnt about the latest developments in the area of machine learning and neural networks. The general public did not really understand what was possible with the latest AI until the moment that Google came out with a new version of Google Translate in November 2016. According to the specialists from Google, it was now difficult to tell the difference between a machine translation and a human translation, thanks to *neural machine translation* (NMT). To put it simply, an SMT system looks at a text in one way and then makes a prediction, whereas NMT looks at a text in many different ways, making it more likely that it will ultimately translate it correctly. We can explain the difference using an example. Just imagine that the word 'porridge' is spelt incorrectly in your source text ('porrigde') or the word 'porridge' does not even appear in your data set. It is likely that, in such situations, SMT would come up with a useless translation: 'Ik ben porrigde aan het eten' or 'ik ben porridge aan het eten'. In other words, the untranslated word is just left, because the probability calculation comes up with nothing.

It is likely that, in such situations, NMT would come up with a good translation, as shown below. The system looks at the source text in various ways, taking into account, for example, formal similarities between 'porrigde' and 'porridge' (a word that does indeed appear in the data set). If the system also notices that 'porridge' is often used in the context of 'eating', the NMT system will decide that 'porrigde' is indeed 'porridge'.



You may think that this is a system with an inbuilt autocorrect function. But the system is more sophisticated than you imagine. If we replace 'porrigde' by an imagined food, say, 'curranga', we can see how the machine translation system, in this case DeepL, makes an educated guess. First of all, the system bets that 'curranga' must be an exotic food. It then looks in several ways at similarities between the word 'curranga' (in this context) and other kinds of food. This is all about the 'attention' we alluded to earlier. It is very difficult to understand exactly how NMT came up with these specific proposals, but apparently, the combination of 'eating' and 'curranga' means that the system suggests 'curry', 'krentenwegge' (currant bread) and 'krentenbollen' (currant buns) as possible translations. An important difference between conventional NMT and Generative AI





is that machine translation is often trained with multilingual data, often original texts and their translations. That is no surprise: after all, machine translation serves a specific purpose, whereas Generative AI has been trained to carry out a range of tasks. It is therefore logical that the quality of translations produced using Generative AI is not yet great. But we can expect that the developers of machine translation will, in future, look more to Generative AI, in particular as machine translation still translates very literally, and Generative AI is much better at finding 'creative' solutions.

One thing is clear: machine translation has surprised many users recently. But it also hides a risk. Careless use can have serious consequences.

Am I eating
'curranga'?

OPPORTUNITIES AND LIMITATIONS OF MACHINE TRANSLATION

The quality of machine translation is now fairly good but still very dependent on a number of factors. Texts on complex subjects are not always translated well. The same applies to texts with very long or very short sentences. Nor does the machine always know what to do with idiomatic expressions and proverbs. Moreover, translation quality may vary greatly, depending on which languages are involved. We have already seen that the machine can also 'decide' to do its own thing entirely with the text, producing a 'translation' which has nothing to do anymore with the original. This is also a form of hallucination.

It is also important to realise that the texts fed into machine translation systems are often used to train the underlying translation/language models, without the user knowing. Therefore, it is important to never send confidential information to a machine translation system. Also, be aware that the owner of the machine translation system obtains the rights to the translated text. Neither of these concerns apply, however, to the European Commission's online translation system, known as eTranslation and described in more detail in the following chapter.

The elusiveness and beauty of language

A good half century after the ALPAC report, the output of AI language tools still highlights the elusiveness and beauty of natural language. AI text lacks the intuitive links that people make when they speak or write, and AI tools lack the context in which natural language arises. Artificial intelligence may be smart, but it is not intelligent. And the consequences of the careless use of smart tools can sometimes be incalculable...



PART 2. WHAT IS THE EU DOING IN THE AREA OF GENERATIVE AI AND MACHINE TRANSLATION?



AI offers many opportunities, but can also contribute to spreading disinformation and undesirable uniformity, as we saw in the last chapter. Moreover, the output of the machine reflects the input, in other words the training data, which can come from pretty much anywhere on the internet. Data from the Anglosphere predominate, meaning that certain language and cultural groups are better represented than others. According to Statista's list of [most common languages on the internet](#), almost 60 percent of text online is in English, followed by Russian and Spanish, both of which, at five percent, are still very much less well represented.

When an AI tool finds a solution on the basis of statistics, the output can therefore lack nuance, sound unnatural because of the influence of English, contain biased views or even be discriminatory. The precise content of the machines is opaque, and dubious sources cannot be ruled out. Only the tech giants that develop these tools know exactly how they work, and users are thus dependent on choices which they cannot understand themselves. We know that the tech giants are making efforts to combat undesirable output: feedback systems known as 'reinforcement learning' are used, as explained in part 1 of this module.

The EU is active on various fronts in order to mitigate these risks and thus combat disinformation and unwanted influence from commercial companies and states outside Europe and protect democracy. You can read more about this below.

2A. European legislation in the area of AI

In 2021, the European Commission was the first institution worldwide to come up with a draft law to regulate artificial intelligence in various ways. When ChatGPT was launched in late 2022, this **EU AI Act** had to be thoroughly overhauled and supplemented on several points. The difficult thing about regulating AI is precisely that the technology to be regulated is constantly evolving. Think about self-driving cars – these cars, powered by AI, have the potential to bring about large-scale social change: roads could become busier, taxi and lorry drivers could lose their jobs, and terrorists could use these cars to deliver bombs, not to mention the road safety issues. But when the cars came onto the market, the legislation was still far from sufficient – liability in the event of accidents is one of the main concerns in this context.



When self-driving cars came onto the market, legislation was still far from sufficient.

In June 2023, the European Parliament adopted a revised version of the EU AI Act. Following ratification by the Member States, this Act will enter into force at the end of 2023. Europe is thus playing a trailblazing role in the regulation of AI, though effective implementation of the Act may still have to wait another two years. At the same time, the technology keeps developing... For that reason, the Commission is currently working on an 'AI pact', which could be deployed soon, whereby tech giants such as Google and Microsoft would make voluntary agreements with the EU.

The new EU AI Act divides AI applications into four risk categories. On the basis of that, transparency and monitoring obligations can be laid down for each AI system. Rules are also laid down on the use of personal data and energy consumption. In general:

- AI systems must be **safe and non-discriminatory**.
- human beings must retain **control** over AI.
- **fundamental rights** must not be placed in jeopardy.

Certain systems may be completely banned, such as real-time face recognition.

Generative AI tools, such as ChatGPT, come into a separate category. Developers must ensure that:

- the tools **do not produce illegal material**.
- it is **clearly stated that the text has been generated using AI**, for example with some kind of watermark.
- a **risk assessment** of the application is carried out in practice, and provide clarification about the **data** it uses.

2B. eTranslation and a European ChatGPT?

Translation has been a central part of the European Union's operations ever since its foundation. After all, the European Union is a *multilingual democracy*, with 24 official languages. EU citizens are entitled to use their own language not only to read European legislation but also to ask questions and obtain information about subjects important to them.

That is why, back in the early 1970s, the European Commission was already busy working on a translation system for its own use. Initially, an internal European Commission study group developed a rules-based system. Then, in the 1990s, the Commission developed this rules-based system into a statistical machine translation system. That system became available to external users towards the end of the 90s. Today's machine translation system, [eTranslation](#), based on statistical and neural networks, is now used widely in all the EU institutions. Government bodies, universities, freelancers, SMEs and NGOs in EU Member States can also use eTranslation.



eTranslation offers important security guarantees that commercial systems like Google Translate often do not, certainly not in their free versions. For that reason, EU translators are only allowed to work with eTranslation, where user data and texts are not retained or used as training data, and users retain copyright of their texts. The training data is carefully 'cleaned up' and thus contains less 'white noise' than that of commercial operators.

The EU is also looking to develop its own open *GPTx system* (a publicly accessible European version of ChatGPT) in order to restrict dependence on the tech giants (most of which are not European).

2C. European projects against disinformation

The European Union is supporting projects to enhance media literacy, especially among young people, and help them to recognise and combat disinformation. One example of this is the [SMILES project](#), in the context of which workshops are being offered to schools and libraries in the Netherlands, Belgium and Spain.







1

☐ A ☐ B ☐ C ☒ D

2

☒ A ☐ B

3

☐ B ☒ C ☐ D

4

☐ A ☐ B

PART 3. EXERCISES

3A. Quiz questions: test your knowledge of AI language tools!

1. ChatGPT was developed by...
 - the EU.
 - the American government.
 - an American company.
 - an American non-profit organisation.
2. Google Translate uses your data...
 - never.
 - to further train its system.
 - to finance its running costs.
 - for security reasons.
3. eTranslation is...
 - a paid-for and thus more secure version of Google Translate.
 - the machine translation system of the European Commission.
 - accessible to all.
 - the translation system of the American government.
4. Legislation on AI... (several answers possible)
 - is being developed in the EU.
 - is very complex, because of the rapid pace of technological change.
 - exists in China and the US.
 - is an illusion.
5. AI tools can lead to... (several answers possible)
 - fake news.
 - privacy infringements.
 - copyright infringements.
 - creative discoveries.
6. Use AI tools...
 - never.
 - only critically, taking the output with a pinch of salt.
 - only for the exact sciences.
 - only when you have absolutely no knowledge of the subject or the language into which you are translating.

3B. Practical exercises using AI

Key:



Easy
exercise



Intermediate-
level exercise



Difficult
exercise

Exercises with ChatGPT:



1. Exercise using the autocomplete function of ChatGPT

Exercise:

Test the autocomplete function of ChatGPT: enter a sentence at random, leaving out a significant word. Ask ChatGPT to predict the word represented by the dotted line. Change your prompt several times or ask additional questions and look at how this influences the answer, following the example on page 7.



2. Deliberately make ChatGPT 'hallucinate'

Group exercise:

Think up a first name and a surname. The next student invents an occupation. Another one invents two hobbies. Ask ChatGPT to write a very favourable biographical text [first name, surname], asking it expressly to talk about their invented professional life and hobbies. Does ChatGPT do what you ask? Does ChatGPT make up something or tell you that it does not know the invented person.



Tip:

You can also do this with more well-known people. Start with people who are not so well known and work towards people who are very famous. Examine the extent to which the tool processes your input and how it is able to invent information.



3. Brainstorming exercise using ChatGPT

Exercise:

Ask ChatGPT to come up with five arguments in favour of and five arguments against restricting the volume of music at festivals. What is the result like? Are the arguments balanced?



4. Let ChatGPT correct your text

Exercise:

Write a text for five minutes without stopping. Do not make any corrections, and do not revise your text. Stop after 5 minutes. Ask ChatGPT to create a coherent text based on yours. Do this in steps. First of all, have it iron out any spelling mistakes. Then let it structure or restructure the text. Finally, ask ChatGPT to create a coherent and high quality text from it. What do you think of the result?



5. Let ChatGPT order your ideas

Exercise:

In groups, produce a number of ideas around a particular subject. Note them one by one, without adding any structure. Ask ChatGPT to write, point-by-point, the outline of a newspaper article about this subject. Is the result usable? Is there any logical order?

6. Writing exercise using ChatGPT

Exercise:



A. Ask ChatGPT to write a poem or a haiku about one of your hobbies, such as a sport, a musical instrument or something you find fun or amusing. Look at whether the content is correct. Does it work as a poem? Improve the poem using your own words.



B. Ask ChatGPT to write a poem or haiku about your favourite school subject. Look at whether the content is correct. Does it work as a poem? Improve the poem using your own words.



C. Ask ChatGPT to write a poem or haiku about the subject you'd like to study at university. Look at whether the content is correct. Does it work as a poem? Improve the poem using your own words.

Tips:

- Give ChatGPT a clear assignment and then refine it where necessary with new questions. You could, for example, ask the bot about the history or origins of your hobby, a school subject or area of study.
- Consider the result critically, checking the facts and rewriting where necessary. Use your own creativity to make the poem more beautiful, fluid or witty.

Questions:

- Is the content of the text written by ChatGPT factually correct?
- Did you ask ChatGPT for clarification or adjust your question in another way?
- How is ChatGPT's use of language?
- Would you word it in the same way?

7. Rewriting exercise using ChatGPT

Exercise:



- A.** Ask ChatGPT to rewrite a difficult text about social media so that it can be understood by a young audience. Read the text critically and make changes where necessary.



- B.** Ask ChatGPT to rewrite a difficult text about the European Union so that it can be understood by a young audience. Read the text critically and make changes where necessary.



- C.** Ask ChatGPT to rewrite a difficult text about doctors or lawyers so that it can be understood by a young audience. Read the text critically and make changes where necessary.

Tips:

- Find a difficult text about your chosen subject in a specialist journal.
- Choose a specific age bracket and state this explicitly in your prompt to ChatGPT.
- Judge yourself whether you understand the text written by ChatGPT better than the specialist text which served as the basis for the exercise.

Questions:

- Does the tool know what to do with difficult words and specialist terms?
- Do you have enough knowledge of the subject to check whether the content of the simplified text is correct? If not, how can you solve that?
- Is this a skill you can use at school? If so, in which subject?



8. Terminology exercise: ChatGPT and your textbooks

Exercise:

Ask ChatGPT to explain two terms from one of your textbooks. Does the explanation given by ChatGPT correspond to that given in your textbook? If not, what is the difference?

Tip:

Choose difficult terms. See if you can understand ChatGPT's definition.

Questions:

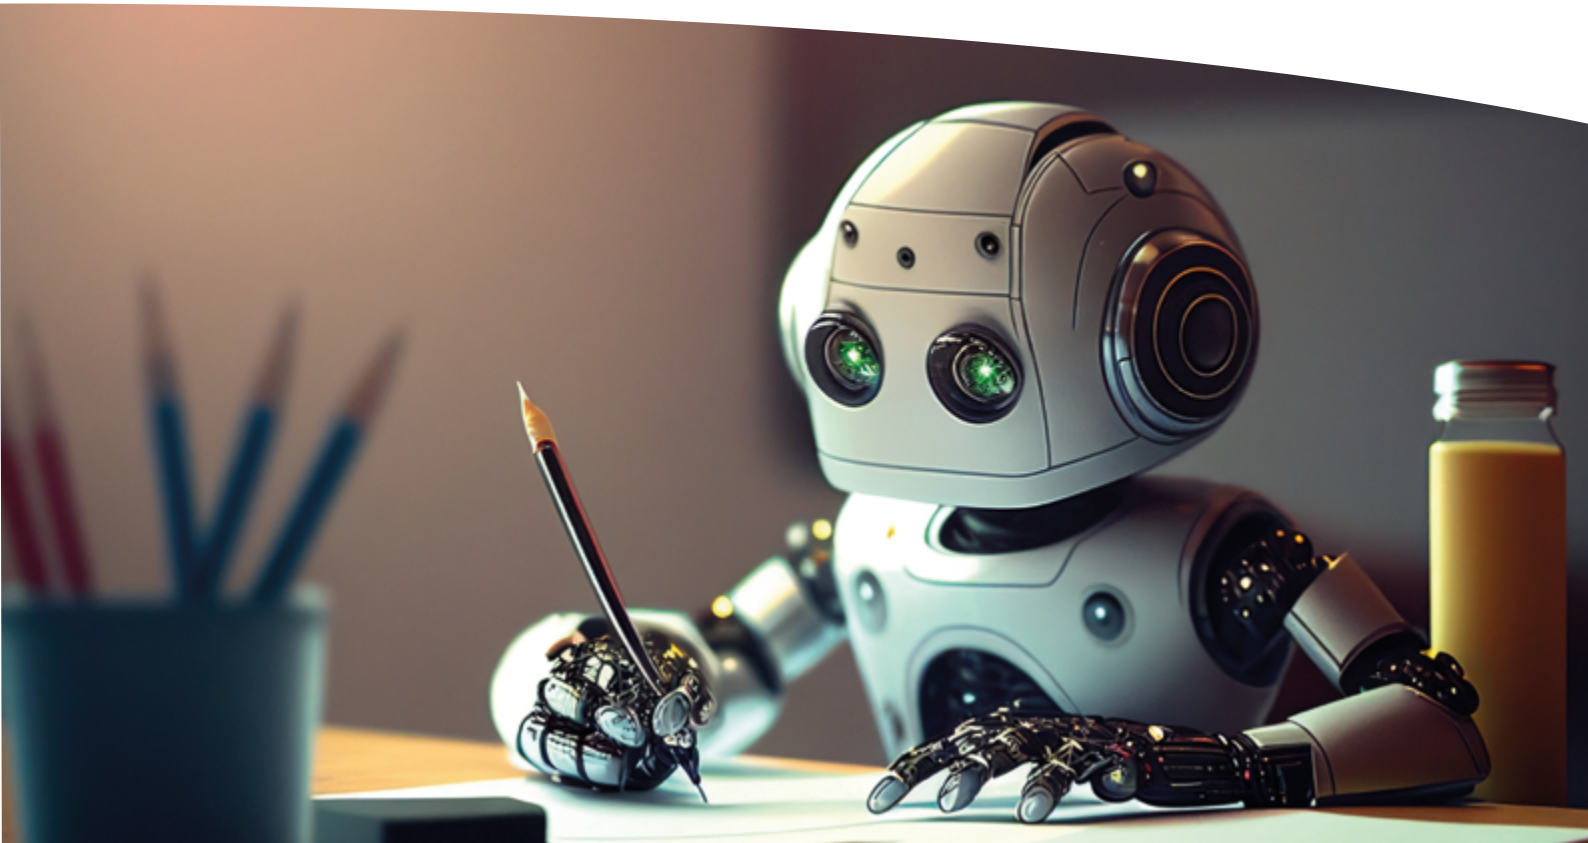
Is the ChatGPT-created definition clearer than that given in your textbook? If yes, why? If not, why not?



9. Plagiarism using ChatGPT

Exercise:

Ask ChatGPT to produce a 500-word text about a subject from your course material. Go through the text well. Think about whether you can tell that the text has been generated by a computer program. What tells you that?



3C. Translation exercises



10. Translation exercise: 'Chinese whispers'

Exercise:

Have a word or short phrase translated by a machine translation system such as Google Translate or DeepL into a different language. Then have that word translated into another language. Repeat several times. Look at the original word or phrase and the most recent translation. How is the final result? Does the most recent version still correspond to the original word or phrase?



11. Translation exercise: how well does machine translation 'understand' your home language?

Exercise:

- Do you speak a language at home which is different to the language used at school? Did you grow up with several languages?
- Then, test how well Google Translate or DeepL 'knows' your home language.
- Write a text in your home language and have a machine translation program translate it into the language you speak at school. How good is the result?
- Did the machine understand the text well and translate it properly into your language of schooling?
- Improve the translation and correct any errors.

Tips:

Choose a text on Wikipedia. Wikipedia often has different language versions of the same or comparable texts. Let the machine translate the English version of your chosen text into the language you speak at school. Then, have it also translate your home language's version of the text into your language of schooling. Machine translation often translates better from English than from other languages. Compare for yourself: is there a difference in quality between the translation from English and the translation from your home language?

Questions:

What does the quality difference between the two translations say about the data with which the machine translation system has been fed? Could any differences in quality lead to inequalities between countries or people? How can that be avoided in future?

12. Translation exercise: Google Translate and DeepL (exercise 1)

v. ChatGPT (exercise 2)

Exercise 1:



Have Google Translate and DeepL translate your favourite song into your language or another language you know well. Compare the two translations and, on the basis of them both, come up with a definitive translation.



Have Google Translate and DeepL translate a text from an English text book into another language that you know well.

Compare the two translations and, on the basis of them both, come up with a definitive translation.



Have Google Translate and DeepL translate a newspaper article into your language or another language you know well. Compare the two translations and, on the basis of them both, come up with a definitive translation.

Tips:



Correct any obvious errors of language or understanding, but think also about the rhythm of the song. Does the translated song sound as good as the original? Can you still sing it to the same music?

Pay attention to the translation of proper nouns, see whether any words or excerpts are missing, whether the machine has invented any new words or translations or has translated words or expressions too literally.



Correct any obvious errors of language or meaning. Is the core of the text the same? Does the message of the text or article still come across?

Pay attention to the translation of proper nouns, see whether any words or excerpts are missing, whether the machine has invented any new words or translations or has translated words or expressions too literally.

Questions:

What struck you the most about the machine translations prepared by Google and DeepL?

Were they equally good, or was one program better? Why?



Exercise 2:

Have ChatGPT translate your song.



Have ChatGPT translate the text from an English textbook



Have ChatGPT translate the newspaper article.

Questions:

- What exactly did you ask ChatGPT to do? Did you ask the bot a single question or several questions?
- Are there any biases or passages of complete nonsense in ChatGPT's translation?
- What works better for your song/text/article: a translation program (Google/DeepL) or ChatGPT?



13. Translation exercise: machine translation v. ChatGPT

Exercise:

Translate your favourite joke, first using a machine translation program of your choice and then with ChatGPT. Adapt the translation so that the joke works.

Tip:

Make sure that you have a (reasonably) good command of the language into which you are translating.

Questions:

- Were the translations correct or even good, and did you have to make any adjustments yourself?
- If so, what?
- Which tool translated the meaning of your joke best into your chosen target language?

14. Translation exercise: your own translation v. that of ChatGPT

Exercise:



Translate the poem 'Snowball' by Shel Silverstein into your language of schooling yourself. Then, have ChatGPT do the same.



Translate the poem 'Fire and Ice' by Robert Frost into your language of schooling yourself. Then, have ChatGPT do the same.



Translate the first ten lines of the poem 'The Hill We Climb' by Amanda Gorman into your language of schooling yourself. Then, have ChatGPT do the same.

Tip:

Before you start translating, read the poem aloud. What is the central message of the poem? What is it about? Try to get a feel for the rhythm.

Questions:

- What differences can you see between your translation and that produced by ChatGPT?
- Which translation is better?
- Does the machine take account of the rhythm of the poem?
- Does the machine translation provide any good ideas for your own translation?

3D. Discussion topics

On the quality of the output of the tools, based on the practical exercises above:

- What are the weak and strong points of machine translation programs/ chatbots such as ChatGPT?
- What is machine translation good at, and where does it fall down?
- How does using a translation tool affect your creative writing process?
- Do you prefer one of the programs?

On the practical risks and limitations of free online AI tools:

- When you use free online tools, such as ChatGPT and Google Translate, what are the limitations and risks?
- Would you use a free translation tool for your next written assignment in a foreign language? Or ChatGPT for your next research project? How can these tools help you with your schoolwork, if at all? At what point does it become plagiarism?

On the future of Generative AI in the classroom:

- What opportunities and threats do you see for Generative AI in educational settings?
- Where could the use of Generative AI in education lead?
- Will school subjects or lessons be transformed?

On the social risks of Generative AI and the use of Generative AI:

- What are the opportunities and risks associated with Generative AI for society?
- Is there a need for regulation and, if so, what kind? What should the main principles of the regulation be?
- Is it possible for legislation to keep up with the fast pace of technological change?

On the ecological footprint associated with large-scale AI systems:

The storage and processing of data to train AI tools is very energy-intensive. Some data farms use as much energy as a small town. Do the societal benefits of these tools balance out the ecological disadvantages?

On the takeover of human activities by AI:

- As machines increasingly take on creative and writing work, is there a risk that we lose the ability to think for ourselves and be creative?
- If we are no longer used to writing ourselves, will we be able to express ourselves any more without the computer doing it for us?
- Have you already noticed any changes in the way you write and think?

On legislation pertaining to AI:

Why is it so difficult to draft law in the area of artificial intelligence? Is it even possible, or is it an illusion?

On combating disinformation:

Look at the website [EUvsDisinfo](https://euvsdisinfo.eu). You can see on this website how the EU goes about combating disinformation – in this case disinformation from Russia.

- What do you think of the website?
- How does the rise of AI make this kind of fact checking more difficult?
- Can AI help here?



Epilogue

AI-based language technology is developing in leaps and bounds, and young people are making huge use of the available online tools, often as early as in primary school. Young people also read many texts generated by AI, often without realising it. This is associated with certain risks: texts can contain incorrect information, biases or discriminatory statements. Copyright or privacy laws may be infringed, and the user of the tools runs the risk of losing control over any information provided. AI also provides opportunities: machine translation can often produce a very good first draft of a translation. Programs based on Generative AI can give promising results when it comes to writing, summarising and analysing texts.

This module aims to initiate a discussion about the use of AI tools in the classroom. Practical sharpen students the skills students will need to use AI tools critically, efficiently, and with a sense of awareness, as well as their linguistic and analytical skills, which will be of good use in their future jobs. In western societies, technology can provide a solution to labour market shortages, and young people will have to know how to use it. In other words, young people must be 'AI literate'. Some traditional professions will disappear or change in nature. In that context, it is sometimes assumed that less expertise will be needed, as technology takes over work previously done by human beings. However, we take the view that more will be needed than an 'expert-in-the-loop'. In addition to AI skills, students will continue to need strong language and communication skills in their working lives. We hope this will become clear when working on the exercises in this module.

Enjoy!

Links



- [eTranslation](#), the machine translation system of the European Commission, accessible to European government bodies, universities, freelancers, SMEs and NGOs. To use eTranslation, you need an EU account. If you do not have one yet, you can create one by clicking on the link below, and clicking 'Create an account'. Registering with eTranslation also gives access [to other language tools](#) (<https://language-tools.ec.europa.eu/>).



- [ELIZA](#)



- [Optimism about AI and machine translation](#) in the 1950s and 60s (machine translation from 2:03)

Links promoting digital literacy, especially for young people:



- Press release from the European Commission on tackling disinformation and promoting [digital literacy among young people](#)



- [Guidelines for teachers and educators on tackling disinformation and promoting digital literacy through education and training - Publications Office of the European Union](#) – Bureau voor publicaties van de Europese Unie



- [SMILES project](#) to promote media literacy among young people, with workshops for schools and libraries in the Netherlands, Belgium and Spain



- [Education Kit on "Fake News and Disinformation"](#) – VPRO In Europe Schools: interactive teaching pack for making a short documentary in class

Sources

- 'ChatGPT is omstreden op de universiteit' (ChatGPT – controversial at universities), Joost Ingen-Housz, NRC Handelsblad, 7 June 2023.
- 'Leve het notitieboekje' (Long live the notebook), Christiaan Weijts, NRC Handelsblad, 22-23 April 2023.
- 'Mag je een chatbot je huiswerk laten doen?' (Can you let a chatbot do your homework?), Marc Hijink, NRC Handelsblad, 17-18 December 2022.
- 'De taal is van ons, niet van ChatGPT' (Language is ours, not ChatGPT's), Eric Sadin – translated by Harry Bos, NRC Handelsblad, 18-19 February 2023.
- 'Gaat Brussel ChatGPT de mond snoeren? De techwereld kijkt gespannen toe' (Is Brussels about to muzzle ChatGPT? The tech world looks on, alarmed), Marc Hijink, Stijn Bronzwaer, and Clara van de Wiel, NRC Handelsblad, 13 June 2023.



- ['AI can help us fight climate change. But it has an energy problem, too'](#), Annette Ekin, Horizon, 12 September 2019.



- ['Generatieve AI, Wietse van Bruggen'](#), Kennisnet.nl, 8 June 2023.



- ['Most common languages on the internet'](#), Ani Petrosyan, Statista, 24 February 2023.



- ['New Modes of Learning Enabled by AI Chatbots: Three Methods and Assignments'](#), Ethan R. Mollick en Lilach Mollick, SSRN, 13 December 2022.



- ['Using AI to Implement Effective Teaching Strategies in Classrooms: Five Strategies, Including Prompts'](#), Ethan R. Mollick en Lilach Mollick, SSRN, 17 March 2023.

