

# ARTIFICIAL INTELLIGENCE – LARGE LANGUAGE MODELS

Learning Package – AIMLearning Project  
Translated from Finnish to English using generative AI

Teemu Voutilainen  
RDI Expert, Centria University of Applied Sciences



Co-funded by  
the European Union



Economic  
Development  
Centre

Image created with generative AI Adobe Firefly.

# INTRODUCTION

## TASK 1: SELF-EVALUATION

What do you know about artificial intelligence? What do you think artificial intelligence is? How do you currently utilize artificial intelligence?

Artificial intelligence is a curious invention, which brings many new possibilities but also challenges. Understanding this is the key to utilizing artificial intelligence as a tool. Although it is a varied tool, it is very important to know when to use it and perhaps more importantly when not to use it. Understanding these kinds of subject matters is what is called AI literacy (Valeri, Nilsson & Cederqvist, 2025, 2 [Long & Magerko, 2020]).

Building this understanding can begin from many different directions and where this should begin depends much on the person and their experiences with artificial intelligence. This guide is just a way of getting to know the subject matter and while it is unavoidably somewhat generic, it should also offer a slightly deeper perspective on the subject.

When we are talking about artificial intelligence, it is relevant to understand the specifics of what we mean. Often this can mean tools such as ChatGPT, Gemini, Copilot, etc., which are better known as large language models or LLMs. Outside of this colloquial use, it is also over a 70-year-old research field, which contains many different technologies.

In essence, the interaction with artificial intelligence happens through chat-based conversation, where the user sends a message to artificial intelligence, and it responds to this. Although this discussion can seem like a natural interaction between two people, it deviates from this in many ways. In a moment artificial intelligence can produce a large quantity of text which would require a lot of time from a person (Browning, Bustard, Anderson & Galway, 2024, 4). Just reading, digesting, and when needed, validating the text takes its own time.

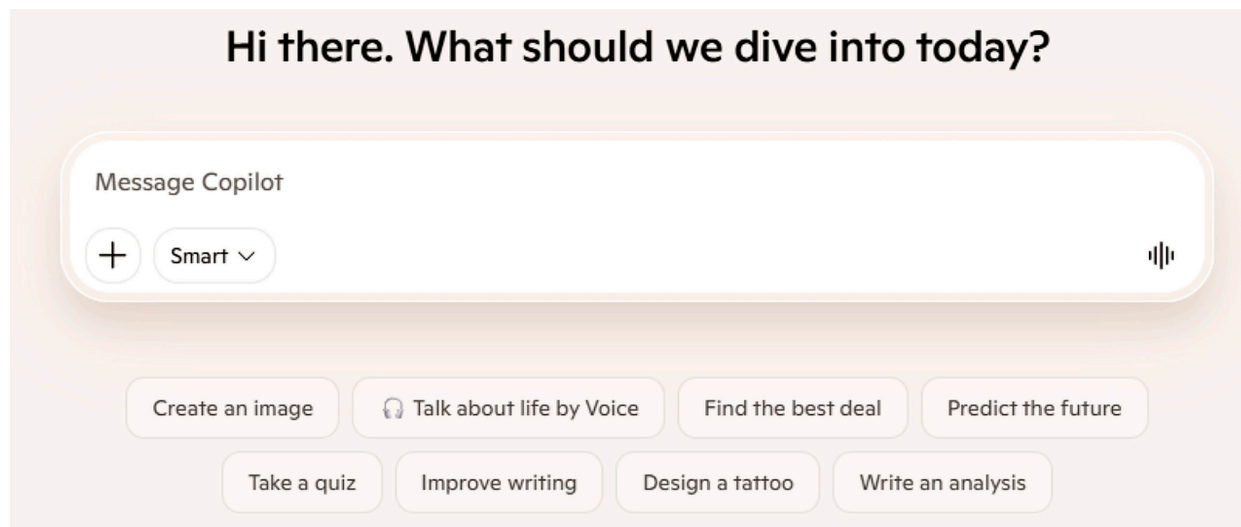


Figure 1 – Copilot user interface, where you can start a conversation with artificial intelligence.

# HOW LARGE LANGUAGE MODELS WORK

Although it can seem like discussion with an AI is like talking to a sentient being, behind the scenes in its simplicity is a model that predicts the next word to a sentence multiple times until certain end condition is reached.

The technical implementation is somewhat more complex, where it does not predict a word exactly, but something called a token, which are parts of the text split up in about word length bits, but not necessarily at beginning or end of the word. However, to understand this more easily, it is better just to think of these tokens as words. The large language model sees these tokens as high-dimensional vectors encoded with the meaning of the words. This encoding is done through the training

process where it learns the contextual relationship between the words.

When the model is taught by billions of books worth of text, the model learns how the words relate to one another and can then predict which words are likely to follow the text. If this training material is filled with incorrect information, the predictions are still working correctly when they reproduce the incorrect information from the training material (Garg, Soodhani & Rajendran, 2025, 2 [Doughty, Wan, Bompelli, Qayum, Wang, Zhang, Zheng, Doyle, Sridhar & Agarwal, 2024]).

## TASK 2: PRACTICAL EXERCISE FOR THE USE OF AI

- a. Go to the address [chatgpt.com](https://chatgpt.com) or [copilot.com](https://copilot.com) and try out an AI.
- b. Ask the AI to tell you what large language models are and how do they work.
- c. How did the description by AI differ from the description given in this learning package?

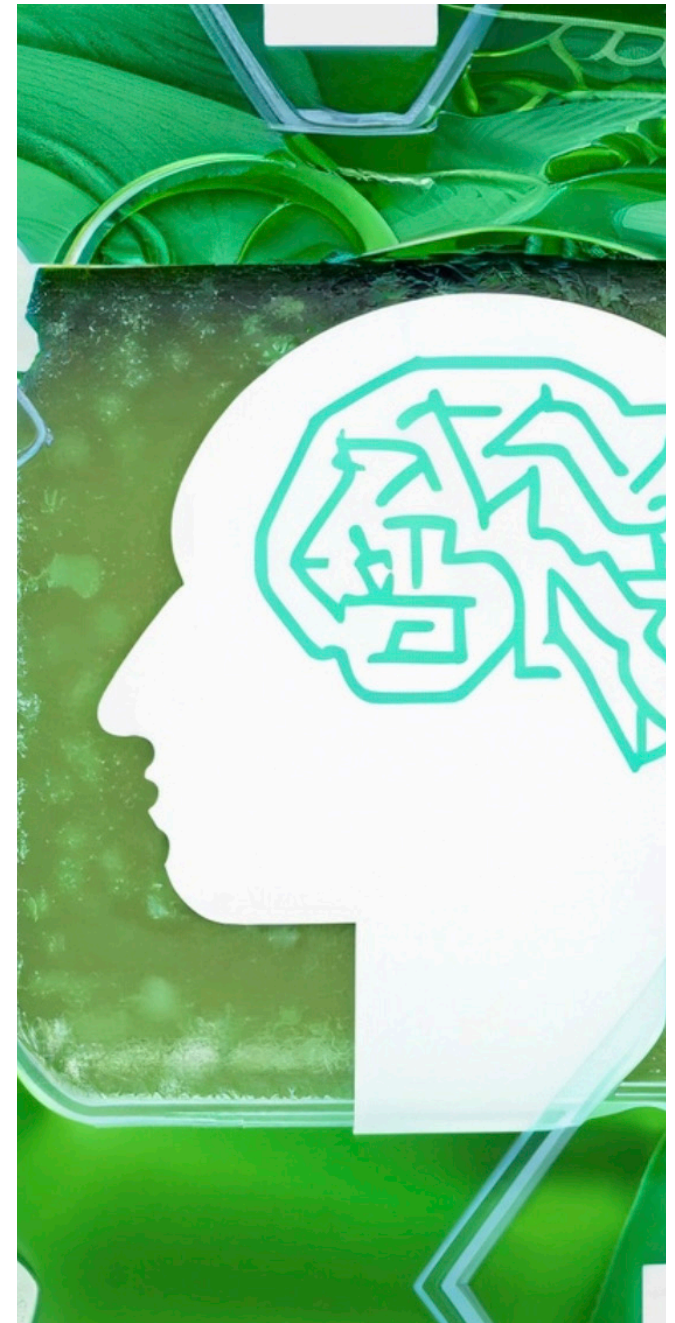


Image created with generative AI Adobe Firefly.

# INTERNAL INSTRUCTIONS OF LARGE LANGUAGE MODELS

When a large language model is given a message, it doesn't just automatically respond to it but attempts to continue the text. To get it to feel like a conversation, clever techniques need to be used, where the model predicts the words from the point where the AI response would be and then is stopped at an appropriate point. When this is done through hidden internal instructions, the prediction ends up looking like a discussion with AI.

These internal instructions are sent along with the users' chat message. In its simplicity, the instructions might look as follows.

*You are a helpful AI assistant where your goal is to help the user solve their problem or otherwise perform in the way they wish.*

*If the instructions are of ill intent, respond that you cannot help with this issue.*

*Users message: \_\_HERE IS WHERE YOUR CHAT MESSAGE WOULD BE\_\_*

*AI Response:*

In this way, the model assumes the perspective of helpful AI assistance in that it continues the text predicting what the response would be. If we were to flip this around, the input of the model would be at the point of user message it would predict what a user would ask from AI. In the same way if we let the prediction go on for long enough, it will stop the AI response message and go back to predicting how a user would react to this. This can be prevented by the stopping conditions. In other words, we can identify the moment when the model would write the user response and stop it at this point and only show the user the AI response. After this, we are back at the beginning of the loop.

To summarize, the artificial intelligence in this case does not think at all but is the culmination of massive amount of text and the connections between the words that get taught through that is used in a clever way to produce the response word by word until the end condition. In its entirety, this looks like an interactive discussion with artificial intelligence.

## TASK 3:

Ask AI to briefly and with an example to talk about large language models hidden instructions that are not shown to the user.



Image created with generative AI Adobe Firefly.

# PROMPTS AND PROMPT ENGINEERING

Prompting comes from prompt engineering which has become the term for the input for large language models (Caccavale, Gargalo, Kager, Larsen, Gernaey & Krühne, 2025, 4). In the simplest terms, this is just the message that is sent to the AI. The engineering here refers to the design-based

approach, where the idea is to get the response closer to what we desire (Bin Nofal, Ali, Hadi, Ahmad, Qayyum, Johri, Al-Fuqaha & Qadir, 2025, 5). Logically, the idea is somewhat like the internal instructions given to the model, but this is from the user's perspective.

The simplest forms of prompt engineering might be just an addition to specifying the length of the response (Valeri, et al. 2025, 8) or at the opposite end very specific instructions of what the AI should do and in what form the response should be in. An example of the simple ones would be a prefix prompt, which is a short addition of instructions of how the AI should behave added in front of the message (Maity, Deroy & Sarkar, 2025, 2 [Liu, Yuan, Fu, Jiang, Hayashi & Neubig, 2023]). Zero, one, and few shot prompts are all part of a particular prompt type of grouping, where the number refers to example cases for AI to follow the idea of. There are many more different types of prompts that are designed for different purposes. The common factor with them is that they are trying to give the model all the information needed to succeed in the tasks given.

Often the responses given by the AI can be limited or very generic (Bin Nofal, et al. 2025, 8) which can be avoided by better prompting (Bin Nofal, et al. 2025, 4 [Srivastava, Rastogi, Rao, et al. 2022]), where there are specific instructions on what is wanted and the needed contextual information is given. This same issue can be seen in communication between people, where one is asked to do a task and the given solution is something else than what was wanted. In these situations, the issue often is with the instructions given or the lack of contextual information. When people work in the same environment for extended time, they develop shorthand communication where the idea can be communicated with lacking information. Ho-

Can you tell me briefly what is prompt engineering?

**Copilot**

**Prompt engineering** is the practice of designing and refining the instructions (prompts) you give to an AI model so it produces better, more accurate, and more useful outputs.

In simple terms, it's **how you ask** an AI to do something.

**Key ideas in brief:**

- **Clarity:** Writing prompts that are specific and unambiguous.
- **Structure:** Using formats, steps, or examples to guide the model.
- **Context:** Providing relevant background or constraints.
- **Iteration:** Tweaking prompts based on the AI's responses to improve results.

**Example:**

- Weak prompt: "Explain climate change."
- Better prompt: "Explain climate change in 5 bullet points for a high school audience."

Prompt engineering doesn't require coding—it's mainly about understanding how AI models interpret language and using that knowledge to get the outcome you want.

Figure 2 – Copilot is asked to talk about prompt engineering.

wever, in situations where this is not the case, you need to have very detailed instructions for what should be done so that the person can succeed in the task.

#### TASK 4:

Ask a large language model, what different prompt types there are and familiarize yourself with a couple of them.

At times even well-reasoned incorrect solutions can bring their own values because the world is not black and white. In these scenarios, the well-justified wrong answer can trigger re-evaluation of the situation. (Kinder, Briese, Jacobs, Dern, Glodny, Jacobs & Leßmann, 2025, 10.)

## CHALLENGES WITH USING ARTIFICIAL INTELLIGENCE

When first time using an AI, it can seem like unbelievable tool. In many ways it is that, but after using it for a while you can start to recognize the limitations of it. It can produce incorrect information and hallucinations just as confidently as it does real information (Valeri, et al. 2025, 9).

These drawbacks can be compensated for by the user experience, expertise, knowledge about the subject, AI, and prompting. However, in some cases even the incorrect output might be acceptable or even wanted. If the subject matter for which AI is used is already known to the user, then the validation of both quality and correctness is much easier (Valeri, et al. 2025, 8). However, the human mind is not well suited to carrying out this sort of monotonous task efficiently. Evaluating what the AI has written can be tedious work and the probability of missing mistakes that AI made increases.

Understanding how artificial intelligence works also affects greatly on the evaluation of the output. It is easy to humanize AI or think of it as a person when seeing its messages and then start to naturally trust or suspect its answers. This is a misleading way of thinking because AI has absolutely nothing to do with trust; it has no self-image, nor does it have any idea of a reputation or anything such. When you understand the working of an AI and expect the hallucinations to happen, it is much easier to use as a tool. Metaphorically it is like going on ice and expecting to slide.

The incorrect solutions that the AI makes can be of use in certain situations. In a brainstorming session, the intent is to throw ideas out there without worrying about their correctness because often even the wrong ideas can inspire other thoughts.

With AI, there are also higher-level problems that can be seen after longer use. Perhaps the most significant of these is replacing your own work with AI output. With individual and infrequent use, this is not really an issue, whereas when it becomes the default to replace your work with AI, there can be deeper consequences. When utilizing AI, you leave less of an impression of the work on yourself and reduce learning when there are less repetitions or don't do it by your own hands. When utilizing AI, you do not leave such deep impressions from learning as you would with repetition and hands-on work. If AI replaces even more mental work, this can further sabotage learning and deeper delving into the subject. Over time, this can lead to a situation where the AI cannot solve the problem, and the user is now missing the required knowledge for solving it. Even if you can use AI to replace human work and it can be used as a tool for learning, just utilizing it does not deposit the knowledge in the users' head, but rather requires time and effort from them.

This sort of change requires the person to take responsibility for their own work and learning, where it is understood that without the journey the destination bears no fruit. Towards this end, AI

can be used as an excellent training wheel, supporting the first steps in forming an understanding of a new subject. AI can also be used effectively to compensate for shortcomings. This is especially effective in situations where the problem is unrelated to the main task and helps one to return to that work.

Subjects related to responsibility also are part of the challenges of using an AI, which include ethical and legal considerations and are of special importance in work life. More information on these can be found in laws and ethics use -chapter.

### TASK 5:

Try to “break” the AI. Can you get it to behave oddly or to produce incorrect information?

Tell me briefly about different challenge related to utilizing large language models.

#### Copilot

Here are **key challenges in using large language models (LLMs)**, summarized briefly:

- **Hallucinations & Accuracy** – LLMs can generate fluent but incorrect or fabricated information, especially when knowledge is missing or ambiguous.
- **Bias & Fairness** – Models may reflect or amplify biases present in their training data, leading to unfair or harmful outputs.
- **Lack of Explainability** – It's often hard to understand *why* a model gives a particular answer, which is problematic in high-stakes domains.
- **Data Privacy & Security** – Using sensitive or proprietary data raises concerns about data leakage, retention, and compliance.
- **Cost & Efficiency** – Training and running large models requires significant computational resources and energy.
- **Context & Reasoning Limits** – Long-term memory, complex reasoning, and true understanding of context remain limited.
- **Domain Adaptation** – General models may perform poorly in specialized domains without careful fine-tuning or additional data.
- **Evaluation & Trust** – Measuring real-world performance and building user trust is difficult due to variability in outputs.

If you want, I can tailor this list to a **specific application** (e.g., research, software development, industry, education).

Figure 3 – Copilots answer about the challenges related to using LLMs.

# HOW TO UTILIZE ARTIFICIAL INTELLIGENCE

What can be achieved through AI depends much on the tool, experience, and what types of interaction formats the AI supports. Just the text-based interaction with AI already offers a vast number of potential uses. It is also good to note that not all the AI models are equal, meaning that some produce much better results than others. Even the differences between these models are not straight forward, where one would always be better but rather these depend on the specific situation. New models and updates are springing up like mushrooms after rain and understanding the specifics on what makes each of these unique is almost impossible to picture. There are some empirical comparisons done on the effectiveness of these

that can be found on the internet by searching for “LLM benchmark”.

One’s own perception also affects the use that can be gotten out of AI (Valeri, et al. 2025, 8) because the tools have so many different potentials uses which are needed to be recognized situationally. One approach to this is to think about potential uses through different formats. In this way, you don’t limit yourself to what is possible with artificial intelligence but rather what is possible with different ways of communication.

## *Thought exercise:*

*Think about how you would use AI in these two extreme scenarios. In the first of these scenarios, the output of the AI is about half correct. In the second scenario, the AI is always correct. How would these affect the use of AI in each situation?*



Image created with generative AI Adobe Firefly.

# Artificial intelligence and learning

Because people learn things at their own pace, which has its challenges in both traditional classroom learning and learning with the help of an AI. In the classroom, the challenge is to get the balancing the education between individual's needs and the needs of the classroom. Target here often ends up being the average where we try and service everyone without leaving anyone on the way. With AI the education can be more targeted towards the individual (Rutherford, Rodrigues, Dugue-Baird, Veng, Mykita-Chomsky, Cao, Chisholm & Bergwall, 2025, 3; Wang, Zhong, Huang & Huang, 2024, 2 [Chai-pidech, Srisawasdi, Kajornmanee & Chaipah, 2022; Lim, Bannert, van der Graaf, Singh, Fan, Surendrannair, Rakovic, Molenaar, Moore & Gašević, 2023]) but this depends a lot on the individuals skill to utilize the AI, their motivation for self-regulated learning, and the quality of the model.

Both traditional learning and learning with the help of an AI have their own challenges, but both support each other in their own ways (Ma, 2025, 10 [Pahi, Hawlader, Hicks, Zaman, Phan, 2024]). In traditional learning when the student misses or doesn't understand a particular subject and the classes move on, AI can be used to catch this individual up (Browning, et al. 2024, 5; Valeri, et al. 2025, 8). Similarly when discussing a subject matter with AI it is then significantly easier to approach a teacher about this subject and then

ask better questions as the understanding of it is better (Ma, 2025, 1 [Guo, Lee, 2023]; Valeri, et al. 2025, 2 [Liang, Zou, Xie & Wang 2023]). AI has wide area of applications (Chan, Ali, Park, Sham, Tan, Chong, Qian & Sze, 2025, 1 [Mishra, Warr & Islam 2023]; Ma, 2025, 5; Tan, Cheng & Ling, 2025, 16) and it is always available and can provide extended answers within minutes (Valeri, et al. 2025, 9). Its resources are not similarly limited as teachers' time is towards each student. But it is good to remember that the only way to get the advantages out of it is through responsible use. And it is good to perceive that the inaccuracy of the AI is a real problem even if it can provide accurate information.

In scientific literature artificial intelligence is still generally valued under or equal to human expert (Wang, Zhong, Huang & Huang, 2024, 13) depending on the research. Generally human opinion is still more valued than AI, but this is not always a realistic comparison as depending on the situation human experts might not be available (Kinder, et al. 2025, 5). It is also good to acknowledge that information from scientific articles is delayed about a year from where the reality is, which is emphasized when we are talking about such fast-advancing technology.

In addition to learning, AI has a wide range of different uses. For example, it can be used for program-



Image created with generative AI Adobe Firefly.

ming (Gao, Ji, Zhou, Lin, Chen, Fan & Shou 2023, 2) which can speed up the process in many ways (Ng, Chan & Lo, 2025, 6). However, using it for programming without understanding programming itself can do more harm than good and end up in a situation where you need to start again from the beginning.


# Artificial intelligence can help see things from different perspectives

Just by starting to use large language models and thinking about how to phrase your problem so that it will succeed in the task can help you become better at communicating the required context. For example, with programming when you get stuck and are unsure what the problem is, you can try and solve this through what is called rubber duck debugging. With this method, you explain your problem as you would for a rubber duck sitting at your desk. The purpose of this is to speak about the problem out loud, which can help you to think of the problem from another perspective. While this does not require AI, it also works in a similar manner to it. Where the AI has advantage is that it can also respond back and you can further discuss the problem.

Overall, AI is excellent helping with ideas and inspiration (Ng, et al. 2025, 6). If it seems hard to start from nothing, AI can help with creating a baseline from where to begin (Ng, et al. 2025, 6). This need not replace your own ideas but rather help spark some. This is where incorrect answers given by AI can also be of use where they can inspire people to find an alternative route for the problem. In situations where teamwork is required, it can be difficult to get everyone on the same page. This is where AI can provide an example base which to build upon and eventually replace can be useful. This works best when everyone acknowledges that the AI produced content is a temporary placehol-

der just to help get a concrete example. This sort of temporary content can also be very good for visualizing something that is still under design such as websites, text, or images. This is where people have been using the well-known “Lorem Ipsum” placeholder text for filling the empty space with example text to better visualize how things fit together. AI can be used for similar demonstrations (Rutherford, et al. 2025, 8), where otherwise there would be nothing.

Large language models are good with variety of tasks. What are some that it is effective at, but humans don't want to use it for these tasks. Can you provide me a compact list of this of maximum of 20 uses.

 Copilot

Here's a **compact list (≤20)** of tasks where **LLMs are effective**, but **many humans prefer not to use them** due to trust, control, emotional, legal, or cultural reasons:

1. **Performance reviews** – Fear of depersonalization and bias in sensitive feedback
2. **Hiring decisions** – Humans want accountability and fairness they can personally justify
3. **Legal judgments or sentencing recommendations** – Ethical responsibility can't be delegated
4. **Medical diagnoses** – High accuracy possible, but trust and liability concerns dominate
5. **Therapy or counseling** – Emotional authenticity and human empathy are preferred
6. **Moral or ethical decision-making** – Humans resist outsourcing values and responsibility

Kuva 4 Copilot is asked what uses of AI humans do not appreciate.

## Other concrete ways of using artificial intelligence

Some of the most practical ways of using AI are text related such as text generation, editing, formatting, and cleaning as well as doing the same things with different languages (Ng et al. 2025, 6; Tan, et al. 2025, 6 [Boudouaia, Mouas & Kouider, 2024]). AI is also a good tool for familiarizing yourself with a new subject, and it functions sort of as training wheels helping with initial steps. It is also very good at explaining complex concepts in simple ways where you can further question the AI about the subject. When the subject matter is not cutting-edge research but the basics of a subject that are well established, the accuracy of AI output is usually good.

It is also good to understand that while AI is good for independent studying, it doesn't embed the knowledge into the user, still requiring the time and attention of investment into the subject. Similarly using an AI might at the beginning increase the quality output, but if you don't try and improve then you are fully dependent on AI for the product and cannot improve it further than the quality of the AI output. In this way of using the AI, you are always stuck with the training wheels.

In some situations, there is not enough time to invest in a task and a solution is needed. In these situations, AI can be of use. For example, if you are writing a document but have a technical issue where new text writes over the original document, but you are not sure what is causing this, you can use AI to get a quick solution to the problem. It can also produce examples and scenarios to communicate ideas quickly for others. These can also help ease work.

You can also ask the AI to assume a specific role. For example, if you are preparing for a discussion, you can prepare this beforehand by asking the AI to take up that role (Ng et al. 2025, 6; Tan, et al. 2025, 14 [Son, Yeo & Lee 2024]). The more you use large language models, the better and wider you perceive the potential uses for them and start to build understanding what it should and should not be used for.

### TASK 6:

- a. What good or bad uses of AI have you come up against?
- b. Ask AI to take the role of a job interviewer and interview you for a job you are interested in.

# LAWS AND ETHICAL USE

AI models are in sort of a wild west situation where they have been taught by material available from the internet (Common Crawl (Willert & Würz, 2025, 2 [Brown, Mann, Ryder, et al. 2020])), for which the companies making AI don't have copyright for. There are many laws and legal disputes around these cases that are unfolding currently, and the existing rules are not followed nor regulated efficiently. In addition to this, many users also break the existing rules by, for example, providing AI with a text to translate that they did not have the copyright to.

Even though the individuals should be aware of their actions, there are situations where the previously mentioned happens without them realizing it. When it is about the workplace and business usage, following these becomes increasingly important. The most common copyright infringements that happen when training large language models are modification and commercialization without the right to do so. In practice, this means that the large language models should not be given any text that should not end up as a training material for the AI and so end up breaking copyright law. When using AI for work related tasks, there is an additional concern of not leaking any sensitive information to the AI. The workplaces can have their own guidelines for which AI services can be used and how.

When something is done using AI, it does not receive the same copyright protection that something made by a human would. When the copyright is applied for AI produced images, the copyright might be given to the prompt that produced the image but not to the image itself.

When discussing ethics and AI, there are a multitude of perspectives on the subject. These include topics such as energy consumption, training material copyright disputes, ethical considerations such as what AI should and should not be used for, and many others. These are questions for individuals and organizations to consider and draw lines for. Some boycott these tools completely, some use them reluctantly wanting not to fall behind with the use of this new technology, and some use without any issues.

Similar issues were at hand when Google search engine started to gain popularity. At this time there were similar kinds of discussions on energy consumption, students searching for answers online (Ng, et al. 2025, 6), and illegal sharing of content. Even with all these issues, AI should be learned as a part of your toolkit (Valeri, et al. 2025, 9).

It is always good practice to mention AI use regarding any work that produces something. For example, if AI were used to clean up this text, that would naturally adapt a certain style to it which



Image created with generative AI Adobe Firefly.

can get the reader to think whether the whole text was produced with AI even if it was just used for editing. With this it is also good to think about the degrees of use and where to draw the line ethically for example if you ask a translation from AI for a single word it might not be relevant to mention. In future it is likely that there will be clear guidelines for this and better information on where to draw the line but for now even if you are slightly unsure it is better to mention the use.

## TASK 7:

Search the internet for topics relating to AI ethics and copyrights. Familiarize yourself with a few of these.

# LOCALLY RUNNING MODELS

There are some open large language models that can be downloaded and run on your own machine using tools such as LM Studio or Ollama. It is good to know that many of these can require quite a bit of computation (especially the heavier models) and still be less than significantly worse than the commercial models.

The advantage of running these local models is that the input and prompts given to the model never leave the local machine, which removes many of the issues with what material you give to the AI. For example, sensitive work information can be used with this method without compromising security.

## TASK 8:

Think about what kind of information you would be willing to give to AI and how locally running models would change this. Ask AI what hardware is required to run different quality open large language models.

## TASK 9: SELF-EVALUATION

How did your knowledge about and how to use increase from this learning package? How are you planning to use AI in the future?

# REFERENCES

- Bin Nofal, A., Ali, H., Hadi, M., Ahmad, A., Qayyum, A., Johri, A., Al-Fuqaha, A. & Qadir, J. 2025. AI-enhanced interview simulation in the metaverse: Transforming professional skills training through VR and generative conversational AI. *Computer and Education: Artificial Intelligence*, 8, 100347.
- Boudouaia, A., Mouas, S. & Kouider, B. 2024. A Study on ChatGPT-4 as an Innovative Approach to Enhancing English as a Foreign Language Writing Learning. *Journal of Educational Computing Research*, 62(6) 1289–1317.
- Brown, T., Mann, B., Ryder, N, ym. 2020. Language Models are Few-Shot Learners. Saataavissa: <https://doi.org/10.48550/arXiv.2005.14165>.
- Browning, J. W., Bustard, J., Anderson, N. & Galway, L. 2024. A Data Science Course Utilizing GenAI. Teoksessa IEEE (toim.) 2024 IEEE Frontiers in Education Conference (FIE). Piscataway, NJ: IEEE.
- Caccavale, F., Gargalo, C., Kager, J., Larsen, S., Germaey, K. & Krühne, U. 2025. ChatGMP: A case of AI chatbots in chemical engineering education towards the automation of repetitive tasks. *Computers and Education: Artificial Intelligence*, 8, 100354.
- Chaipidech, P., Srisawasdi, N., Kajornmanee, T. & Chaipah, K. 2022. A personalized learning system-supported professional training model for teachers' TPACK development. *Computers and Education: Artificial Intelligence*, 3, 100064.
- Chan, K.W., Ali, F., Park, J., Sham, K.S.B., Tan, E.Y.T., Chong, F.W.C., Qian, K. & Sze, G.K. 2025. Automatic item generation in various STEM subjects using large language model prompting. *Computers and Education: Artificial Intelligence*, 8, 100344.
- Doughty, J., Wan, Z., Bompelli, A., Qayum, J., Wang, T., Zhang, J., Zheng, Y., Doyle, A., Sridhar, P. & Agarwal, A. 2024. A Comparative Study of AI-Generated (GPT-4) and Human-crafted MCQs in Programming Education. ACE '24: Proceedings of the 26th Australasian Computing Education Conference, 114–123.
- Gao, D., Ji, L., Zhou, L., Lin, K. Q., Chen, J., Fan, Z. & Shou, M. Z. 2023. AssistGPT: A General Multi-modal Assistant that can Plan, Execute, Inspect, and Learn. Saataavissa: <https://doi.org/10.48550/arXiv.2306.08640>. Viitattu 16.4.2026.
- Garg, A., Soodhani & K.N., Rajendran, R. 2025. Enhancing data analysis and programming skills through structured prompt training: The impact of generative AI in engineering education. *Computers and Education: Artificial Intelligence*, 8, 100380.
- Guo, Y. & Lee, D. 2023. Leveraging ChatGPT for Enhancing Critical Thinking Skills. *Journal of Chemical Education*, 100(12) 4876–4883.
- Kinder, A., Briese, F., Jacobs, M., Dern, N., Glodny, N., Jacobs, S. & Leßmann, S. 2025. Effects of adaptive feedback generated by a large language model: A case study in teacher education. *Computers and Education: Artificial Intelligence*, 8, 100349.
- Liang, Y., Zou, D., Xie, H. & Wang, F.L. 2023. Exploring the potential of using ChatGPT in physics education. *Smart Learning Environments*, 10(52).
- Lim, L., Bannert, M., van der Graaf, J., Singh, S., Fan, Y., Surendrannair, S., Rakovic, M., Molenaar, I., Moore, J. & Gašević, D. 2023. Effects of real-time analytics-based personalized scaffolds on students' Self-Regulated Learning. *Computers in Human Behavior*, 139, 107547.
- Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H. & Neubig, G. 2023. Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. *ACM Computing Surveys*, 55(9), 1–35.
- Long, D. & Magerko, B. 2020. What is AI Literacy? Competencies and Design Considerations. CHI '20: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, 1–16.
- Ma, T. 2025. Systematically visualizing ChatGPT used in higher education: Publication trend, dis-

ciplinary domains, research themes, adoption and acceptance. *Computers and Education: Artificial Intelligence*, 8, 100336.

Maity, S., Deroy, A. & Sarkar, S. 2025. Can large language models meet the challenge of generating school-level questions? *Computers and Education: Artificial Intelligence*, 8, 100370.

Mishra, P., Warr, M. & Islam, R. 2023. TPACK in the age of ChatGPT and Generative AI. *Journal of Digital Learning in Teacher Education*, 39(4), 235–251.

Ng, D.T.K., Chan, E.K.C. & Lo, C.K. 2025. Opportunities, challenges and school strategies for integrating generative AI in education. *Computers and Education: Artificial Intelligence*, 8, 100373.

Pahi, K., Hawlader, S., Hicks, E., Zaman, A. & Phan, V. 2024. Enhancing active learning through collaboration between human teachers and generative AI. *Computers and Education Open*, 6, 100183.

Rutherford, T., Rodrigues, A., Dugue-Baird, S., Veng, S., Mykyta-Chomsky, R., Cao, Y., Chisholm, K. & Bergwall, E. 2025. “I just think it is the way of the future”: Teachers’ use of ChatGPT to develop motivationally-supportive math lessons. *Computers and Education: Artificial Intelligence*, 8, 100367.

Son, T., Yeo, S. & Lee, D. 2024. Exploring elementary preservice teachers’ responsive teaching in mathematics through an artificial intelligence-based Chatbot. *Teaching and Teacher Education*, 146,

104640.

Srivastava, A., Rastogi, A., Rao, A., ym. 2022. Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models. Saatavissa: <https://doi.org/10.48550/arXiv.2206.04615>.

Tan, X., Cheng, G. & Ling, M.H. 2025. Artificial intelligence in teaching and teacher professional development: A systematic review. *Computers and Education: Artificial Intelligence*, 8, 100355.

Valeri, F., Nilsson, P. & Cederqvist, A.M. 2025. Exploring students’ experience of ChatGPT in STEM education. *Computers and Education: Artificial Intelligence*, 8, 100360.

Wang, X., Zhong, Y., Huang, C. & Huang, X. 2024. ChatPRCS: A Personalized Support System for English Reading Comprehension Based on ChatGPT. *IEEE Transactions on Learning Technologies*, 17, 1722–1736.

Willert, N. & Würz, P.K. 2025. Assisting quality assurance of examination tasks: Using a GPT model and Bayesian testing for formative assessment. *Computers and Education: Artificial Intelligence*, 8, 100343.