CAMBRIDGE
√Mathematics

Engaging with
Artificial Intelligence
in
Research and Writing

Initiative led by Xinyue Li

Contents

What is Artificial Intelligence?

An introduction to Artificial Intelligence (AI)

Artificial Intelligence (AI) is an umbrella term for technology capable of performing tasks intelligently. Various definitions have been proposed by pioneers, experts, researchers, companies, organisations, the media and so forth. Some of these definitions include:

"The science of making machines do things that would require intelligence if done by men." (Minsky, 1968, p. v)

"The study of man-made computational devices and systems which can be made to act in a manner which we would be inclined to call intelligent." (Berkeley, 1997)

"It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable." (McCarthy, 2007, p. 2)

"The use of computer programs that have some of the qualities of the human mind, such as the ability to understand language, recognize pictures, and learn from experience." (Cambridge University Press, 2017)

"AI, shorthand for artificial intelligence, defines technologies emerging today that can understand, learn, and then act based on that information." (PWC, 2017, p. 2)

"AI generally refers to the ability of machines to exhibit human-like intelligence—for example, solving a problem without the use of hand-coded software containing detailed instructions." (Bughin et al., 2017, p. 7)

"A system's ability to correctly interpret external data, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation." (Kaplan & Haenlein, 2019, p. 17)

"AI is the ability of a machine to display human-like capabilities such as reasoning, learning, planning and creativity." (European Parliament, 2021)

"The ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings." (Encyclopedia Britannica, 2023)

Simply put, AI can be understood as the brain for computers, robots and other smart technologies. Just as students learn to get better grades in examinations, AI can learn to get better at various tasks. At first, it might not be very good at a specific task, again like a child first learning something new. But as AI keeps learning from data and experience, over time its performance gets better and better.

AI's definition of AI

GPT-4 was used to generate a description of artificial intelligence.

The following prompt was input into GPT-4: "Could you tell me what artificial intelligence is?"

The output obtained was:

"Artificial Intelligence (AI) is a branch of computer science that aims to create machines that mimic human intelligence. The term was first coined in 1956 by John McCarthy, and it encompasses several different approaches and techniques, but the primary goal is to create systems that can perform tasks that would normally require human intelligence.

These tasks can range from simple ones like understanding spoken or written language (Natural Language Processing), recognizing patterns or objects (Machine Vision), and making decisions based on data (Decision Making), to more complex ones like learning and adapting to new situations (Machine Learning) and exhibiting creativity (Generative Design).

AI can be classified into two main types:

Narrow or Weak AI: This type of AI is designed to perform a narrow task (e.g., facial recognition or internet searches) and it operates under a limited context. It doesn't possess true intelligence or self-awareness despite its capabilities. Examples of this kind of AI are Siri, Google Assistant, Amazon's Alexa, and many other AI-powered systems.

General or Strong AI: This type of AI, which does not yet exist as of my last training data in September 2021, is the kind of AI that has the potential to outperform humans at most economically valuable work. It can understand, learn, adapt, and implement knowledge in a broad range of tasks, demonstrating an intelligence equal to or surpassing human intelligence. It's often the type of AI portrayed in science fiction, like the HAL 9000 from *2001: A Space Odyssey* or Data from *Star Trek*.

AI technologies are increasingly being integrated into various industries and aspects of life, from healthcare and education to transportation and entertainment. While there are numerous benefits to AI, it also raises ethical and societal concerns, such as job automation, privacy issues, and the potential for misuse." (OpenAI, 2023, June 1)

Key dates and names

| 1950 | •Alan Turing publishes _Computing Machinery and Intelligence_ and introduces the Turing Test as a way to determine a machine's ability (Turing, 1950). |

| 1956 | •The term "artificial intelligence" is coined by John McCarthy at the first AI conference at Dartmouth College. |

| 1957 | •Frank Rosenblatt introduces the Perceptron, which has been the core of Deep Learning since then. |

| 1969 | •Marvin Minsky and Seymour Papert (1969) publish _Perceptrons_, which then becomes the landmark work on neural networks. |

| 1980s | •The _backpropagation_ is popularised as a training technique for neural works (Vandeput, 2021). |

| 1997 | •IBM's Deep Blue defeats chess world champion Garry Kasparov. It is the first time a computer defeats a world champion in a complex game. |

| 2011 | •IBM's Watson defeats Ken Jennings and Brad Rutter, two _Jeopardy!_ (Amerian quiz competition game show) champions. |

| 2012 | •AI startup DeepMind develops a deep neural network that is capable of recognising cats in YouTube videos. |

| 2014 | •Facebook creates DeepFace, a deep learning facial recognition system which can recognise faces. |

| 2016 | •AlphaGo (developed by DeepMind) defeats world champion Lee Sodol in the game of Go. |

| 2020 | •OpenAI releases GPT-3, marking a significant breakthrough in natural language processing.
•AlphaFold 2 (developed by DeepMind) solves the protein folding problem, which is considered a significant achievement in computational biology. |

| 2023 | •A rise in Large Langauge Models (LLMs). |

Figure 1. A Brief History of Artificial Intelligence (AI); adapted from IBM (2023)

Applications and Limitations

Given the dynamic nature of machine learning, AI-based tools are constantly evolving and developing, leading to rapid changes in this field. Therefore, the following discussion is contextually limited to the time at which this report was written.

What is AI good for?

AI-powered tools can help us in a number of ways, for instance:

- Answering questions that have answers available on the internet or that can be accessed via a public link that the user provides; e.g., "What is the colour of an apple?"
- Reviewing and summarising written materials – this can be particularly useful when drafting a descriptive literature review or an abstract.
- Assisting in finding and filtering information for various applications, ranging from everyday life situations such as learning how to play a game to academic-related purposes such as locating relevant research literature.
- Planning and structuring written materials.
- Acting as digital personal assistants; e.g., Siri (a built-in, voice-controlled virtual assistant available for Apple users) can be used to help organise daily routines.
- Providing machine translations.
- Drafting and brainstorming.
- Generating graphs, images and visuals based on information provided.
- Helping to improve the tone and structure of writing, etc.

Examples of AI use

Table 1 below briefly lists some potential applications of AI tools in our daily work.

Table 1. Some Potential Applications of AI Tools

Role	Description	Example of Implementation
Project brainstormer	Assist in generating ideas for projects and problem-solving initiatives.	AI tools could be used to suggest ideas or approaches based on given information.
Communications aid	Assist in drafting emails and other written communications.	AI tools could help draft emails, offering suggestions for phrasing and ensuring clarity and flawless grammar.
Meetings facilitator	Help schedule and manage meetings.	AI tools could be integrated with a calendar application/software to schedule meetings.
Coding assistant	Assist with coding problems and provide explanations of programming concepts.	AI tools could provide coding suggestions, help debug issues, or explain complex programming concepts.

Professional development	Provide explanations of topics requested.	AI tools could be used to explain information, concepts and other job-specific knowledge.
Workflow coordinator	Offer suggestions for workflow optimisation and help manage tasks.	AI tools could be integrated with project management software to assign tasks, track progress, and suggest workflow improvement.
Translator	Translate text.	AI tools could translate text into the desired language.

What are the limitations?

Artificial and human intelligence are not the same; currently, AI-powered tools do not understand anything they produce.

- AI-powered tools exhibit and perpetuate stereotypes and biases. As biases can be introduced into data through how they are collected, analysed or selected for use, existing social inequalities and prejudices can be reinforced by some AI tools.
- AI-powered tools may generate well-formatted but potentially fictitious citations and references. It is crucial to double-check all references produced by the tool. In practice, this may be even more time-consuming than conducting literature research on search engines.
- Unlike search engines, which provide real-time access to the most up-to-date information, certain AI tools, depending on their training model, may not access the most current resources and could be several months out of date. With further developments in the field of AI, it is possible that this constraint will be addressed.
- Based on the training model(s), AI-powered tools perform better (e.g., with more accurate answers, richer explanations, etc.) in subjects that are widely written about, and less effectively in specialist/niche areas.
- Although the output may seem quick, rich, and well-written, AI-powered tools frequently make mistakes and should not be relied upon for factual accuracy.

In general, overreliance on AI-powered tools can significantly impact our critical thinking and digital literacy skills.

The safe use of AI tools

General use

The following flowchart is for general informational purposes only; it does not, and is not intended to, constitute legal advice. One thing that needs to be emphasised is that the user of AI tools is responsible for content, including input (content provided by the user) and output (content generated and returned by AI tools based on the input). The user is responsible for ensuring that the content does not violate any applicable law or the Terms of Use of a specific AI tool. It is important to note that AI tools developed for a specific company or project (i.e., uploaded data are not externally accessible) are not included in the discussion here.
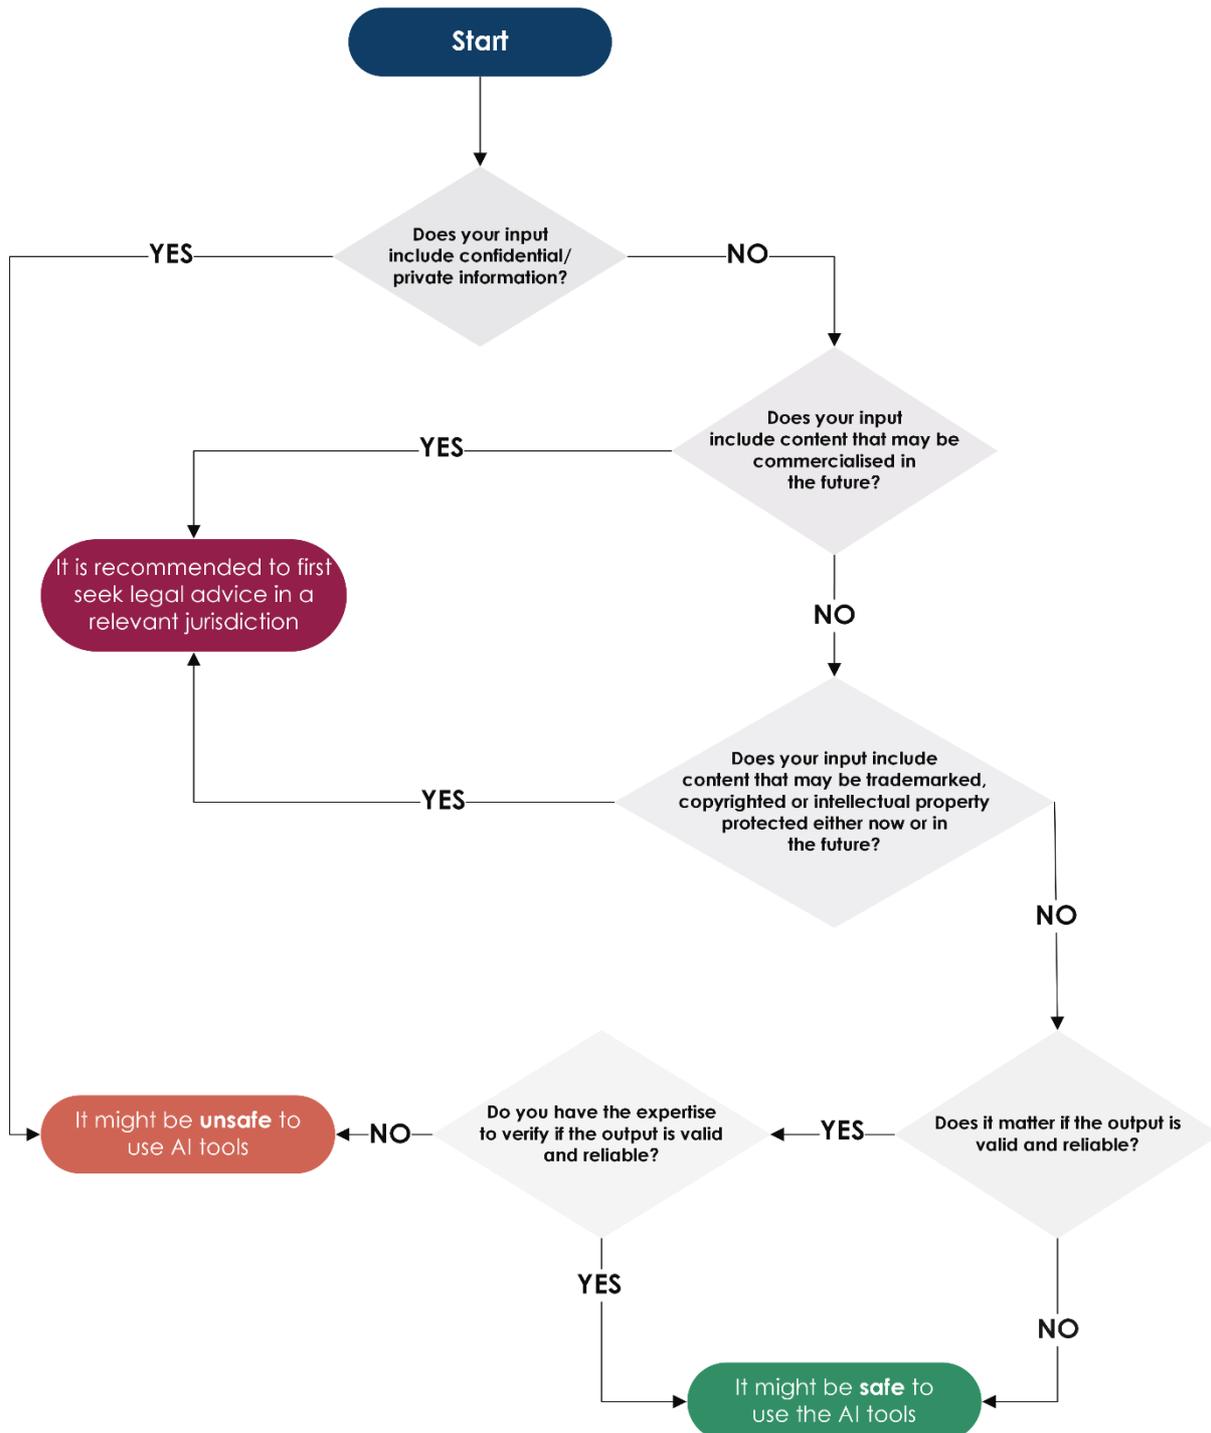
Figure 2. Safe and Unsafe Use of AI Tools

Using AI ethically

An "Ethics guidelines for trustworthy AI" document was published in 2019 by the High-Level Expert Group on Artificial Intelligence set up by the European Commission. The guidelines put forward four ethical principles, including "respect for human autonomy," "prevention of harm," "fairness" and "explicability," and seven key requirements, including "human agency and oversight," "technical robustness and safety," "privacy and data governance," "transparency," "diversity, non-discrimination and fairness," "societal and environmental wellbeing" and "accountability" (2019, p. 8). In addition, Cambridge University Press has launched the first AI research ethics policy (2023) that can be applied to research papers, books and other scholarly works. More recently, on 4th July 2023, the Russell Group (an association consisting of 24 world-class and research-intensive universities including the University of Cambridge, the University of Oxford and UCL) published their principles on the use of generative AI tools in education, emphasising the importance of ethical and responsible use of generative AI in education.

Using AI in research

This section was adapted from the online article "AI contributions to research content" published by Cambridge University Press (2023).

- AI DOES NOT meet the Cambridge requirements for authorship, given the need for accountability. Therefore, AI and LLM (Large Language Model) tools may not be listed as co-authors in our work.

- The use of AI must be declared and clearly explained in research publications, just as we expect scholars to declare and explain the use of other software, tools, and methodologies.

- Authors are accountable for the accuracy, integrity, and originality of their research publications.

- Any use of AI must adhere to Cambridge's plagiarism and academic misconduct policy (University of Cambridge, 2019). Scholarly works must be the author's own and not present others' ideas, data, words, or other material without adequate citation and transparent referencing.

Acknowledging AI sources

We must acknowledge the use of AI sources following these guidelines, which were adapted from the "Engaging with AI in your education and assessment" briefing, published by UCL (2023).

Acknowledgement

We must include the name of the tool, and how it was used in the research and/or writing. The following style should be used:

- "No content generated by AI technologies has been presented as my own work.
- I acknowledge the use of <insert the name of AI tool(s) and link> to generate materials for background research and self-study in the drafting of this paper/document.
- I acknowledge the use of <insert the name of AI tool(s) and link> to generate materials that were included within my writing/research in modified form" (UCL, 2023).

For example:

I acknowledge the use of GPT-4 (https://chat.openai.com/) to generate some initial ideas which I used in background search and self-study in the drafting of this paper.

Description

We must describe how the information or material was generated, including the prompts we used, what the output was, and any alterations we made to this output. For instance:

- "The following prompts were input into <name of AI tool(s)>: <List prompt(s)>
- The output obtained was: <paste the output generated by the AI tool(s)>
- The output was changed by me in the following ways: <explain the actions taken>" (UCL, 2023)

For example:
I utilised GPT-4 to establish a high-level structure, which I then adapted in my writing, choosing to focus on <topics>. I employed an additional GPT-4 prompt to generate high-level ideas about <topics>.

Example of a description:
I used GPT-4 to generate a description of artificial intelligence which I then cited in this document. The following prompt was input into GPT-4: "What is artificial intelligence?"

The output obtained was:

"Artificial Intelligence (AI) refers to the simulation of human intelligence in machines that are programmed to think like humans and mimic their actions. The term may also be applied to any machine that exhibits traits associated with a human mind such as learning and problem-solving.

The ideal characteristic of AI is its ability to rationalize and take actions that have the best chance of achieving a specific goal. A subset of AI, Machine Learning (ML), uses statistical techniques to enable machines to improve with experience.

AI can be categorized as either weak or strong. Weak AI, also known as Narrow AI, is an AI system that is designed and trained for a particular task. Virtual personal assistants, like Apple's Siri, are a form of weak AI. Strong AI, also known as Artificial General Intelligence (AGI), involves machines that possess the ability to understand, learn, adapt, and implement knowledge from different domains, much like a human being.

As of my last training cut-off in September 2021, we are yet to build a truly strong AI. Most of the current applications of AI are narrow, designed to perform specific tasks, like voice recognition, image recognition, recommendation systems, etc.

AI research and development are ongoing, with efforts focused on improving the capabilities of existing AI technologies, developing new AI techniques, and, to a lesser extent, working towards the long-term goal of strong AI." (OpenAI, 2023, May 12)

Reference

We must include the use of AI tool(s) in our reference list.

For example (in APA 7 style):
OpenAI. (2023, May 12). *GPT-4* [Large language model]. https://chat.openai.com/

How to cite images generated by an AI tool

When inserting an AI-generated image, we start with a figure number followed by an image title. Beneath the image, we need to include a note with details about the prompt and the sources. The full citation should be included in the References section.

For example:

Figure 3. Example of an AI-Generated Image <insert the name of the image here>

Followed by: *Note*. Image generated using the prompt "describe your prompt here", by *<insert the developer's name>*, *<insert the name of the AI tool>*, <insert the versional year> (URL of the AI tool)

Example for Figure 4 above:
Note. Image generated using the prompt "A baby seal resting on the beach on a sunny day", by *Open AI, DALL-E*, 2023 (https://labs.openai.com)

In the References section:
Developer's Name. (Versional Year). *Name of AI tool* (Versional number) [Type of AI model]. URL of the AI tool

Example for Figure 4 above (in APA 7 style):
OpenAI. (2023). *DALL-E* (Version 2) [Large language model]. https://labs.openai.com

Ownership of the copyright and intellectual property

The rapid development of AI is posing many challenges to the current copyright field; it has raised many concerns about various legal implications, including uncertainty about ownership, copyright and intellectual property issues. There are many questions which remain unanswered, for example, can AI-generated content be qualified as a "work" protected under copyright law? Who owns the copyright of the content (e.g., text, image, music, etc.) generated by an AI tool? Are AI-generated logos trademarkable? This section is not intended to answer these questions or to provide legal advice. Instead, its aim is to briefly present some relevant information on these complicated issues.

AI-generated content may come with legal risks, including copyright, patent and trademark infringement, direct or unintentional. As it is relatively clear that the user is responsible for the input and AI-generated content (output), "including for ensuring that it does not violate any applicable law or these Terms" (OpenAI, 2023, March 14), the user needs to be aware that training data of AI tools might include unlicensed works, and AI-generated content might be unauthorised derivative works that are not covered by "fair dealing", which "is a legal term used to establish whether a use of copyright material is lawful or whether it infringes copyright. There is no statutory definition of fair dealing – it will always be a matter of fact, degree and impression in each case" (Intellectual Property Office, 2021). Some exceptions allowed under fair dealing include the use of small samples of the work for research and private study, criticism, news reporting, etc.

UK legislation has a definition for computer-generated works. Under the Copyright, Designs and Patents Act 1988, these works are considered as "generated by computer in circumstances such that there is no human author of the work" (Intellectual Property Office, 2018, section 178). The law suggests that content generated by AI can be protected by copyright. However, as AI cannot be treated as an author or owner under the Copyright, Designs, and Patents Act 1988, it is clear that AI itself cannot own copyright. But this does not mean that the user of AI owns the copyright.

Article 2(1) of the Berne Convention (Ricketson & Ginsburg, 2005) extends copyright protection to "every production in the literary, scientific and artistic domain, whatever may be the mode and form of its expression". In addition, in the Guide to the Berne Convention for the Protection of Literary and Artistic Works (World Intellectual Property Organisation, 1978, p.13), it is stated that "The expression 'literary and artistic works' must be taken as including all works capable of being protected."

The US Copyright Office (USCO, 2023) issued a statement of policy to clarify registration guidance on works that contain materials generated by AI technology, which became effective on 16 March 2023. According to USCO, AI-generated images are not eligible for copyright protection as the work is "not the product of human authorship" (p. 16192). For example, AI-generated logos do not fall under copyright protection and cannot be trademarked.

Given the information above, it is clear that the laws relevant to copyright vary from one jurisdiction to another. Furthermore, these laws are evolving continuously to cope with the technological advances related to AI, hence it is important to seek legal advice from a counsel in a relevant jurisdiction if you have any concerns, questions or uncertainties regarding these issues.

References

Berkeley, I. S. N. (1997). *What is Artificial Intelligence?* The University of Louisiana at Lafayette. https://userweb.ucs.louisiana.edu/~isb9112/dept/phil341/wisai/WhatisAI.html

Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlström, P., Henke, N., & Trench, M. (2017). *Artificial intelligence: The next digital frontier?* McKinsey Global Institute. https://www.mckinsey.com/~/media/McKinsey/Industries/Advanced Electronics/Our Insights/How artificial intelligence can deliver real value to companies/MGI-Artificial-Intelligence-Discussion-paper.ashx

Cambridge University Press. (2017). Artificial intelligence. In *Cambridge academic content dictionary*. https://dictionary.cambridge.org/dictionary/english/artificial-intelligence

Cambridge University Press. (2023). AI contributions to research content. In *Authorship and contributorship*. https://www.cambridge.org/core/services/authors/publishing-ethics/research-publishing-ethics-guidelines-for-journals/authorship-and-contributorship#:~:text=AI%20Contributions%20to%20Research%20Content&text=Any%20use%20of%20AI%20must,adequate%20citation%20and%20transparent%20referencing.

Encyclopedia Britannica. (2023). *Artificial intelligence.* https://www.britannica.com/technology/artificial-intelligence

European Parliament. (2021, March 29). *What is artificial intelligence and how is it used?* https://www.europarl.europa.eu/news/en/headlines/society/20200827STO85804/what-is-artificial-intelligence-and-how-is-it-used

High-Level Expert Group on AI. (2019). *Ethics guidelines for trustworthy AI.* European Commission. https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai

IBM. (2023). *What is artificial intelligence (AI)?* https://www.ibm.com/topics/artificial-intelligence

Intellectual Property Office. (2018). *Copyright, Designs and Patents Act 1988.* https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/957583/Copyright-designs-and-patents-act-1988.pdf

Intellectual Property Office. (2021, January 4). *Guidance. Exceptions to copyright.* https://www.gov.uk/guidance/exceptions-to-copyright

Kaplan, A., & Haenlein, M. (2019). Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons, 62*(1), 15–25. https://doi.org/10.1016/j.bushor.2018.08.004

McCarthy, J. (2007). *What is artificial intelligence?* https://www-formal.stanford.edu/jmc/whatisai.pdf

Minsky, M. (Ed). (1968). *Semantic information processing.* The MIT Press.

Minsky, M., & Papert, S. A. (1969). *Perceptrons.* The MIT Press.

OpenAI. (2023). *DALL-E* (Version 2) [Large language model]. https://labs.openai.com

OpenAI. (2023, May 12). *GPT-4* [Large language model]. https://chat.openai.com/

OpenAI. (2023, June 1). *GPT-4* [Large language model]. https://chat.openai.com/

OpenAI. (2023, March 14). *Terms of use*. https://openai.com/policies/terms-of-use

PWC. (2017). *Bot.Me: A revolutionary partnership. How AI is pushing man and machine closer together.* https://www.pwc.in/assets/pdfs/consulting/digital-enablement-advisory1/pwc-botme-booklet.pdf

Ricketson, S., & Ginsburg, J. (2005). Appendix 1: Latest text of Berne Convention (1971 Paris Act plus Appendix). In *International copyright and neighbouring rights: The Berne Convention and beyond* (2nd ed.). https://global.oup.com/booksites/content/9780198259466/15550001

Russell Group. (2023). *Russell Group principles on the use of generative AI tools in education.* https://russellgroup.ac.uk/media/6137/rg_ai_principles-final.pdf

Turing, A. M. (1950) Computing machinery and intelligence. *Mind, 49,* 433–460. https://redirect.cs.umbc.edu/courses/471/papers/turing.pdf

UCL. (2023, February). *Engaging with AI in your education and assessment.* https://www.ucl.ac.uk/students/exams-and-assessments/assessment-success-guide/engaging-ai-your-education-and-assessment

University of Cambridge. (2019, November 27). *Plagiarism and academic misconduct.* https://www.plagiarism.admin.cam.ac.uk/definition

USCO. (2023, March 16). Copyright registration guidance: Works containing material generated by artificial intelligence. *Federal Register, 88*(51), 16190–16194. https://www.govinfo.gov/content/pkg/FR-2023-03-16/pdf/2023-05321.pdf

Vandeput, N. (2021, July 20). *A brief history of neural networks.* Medium. https://medium.com/analytics-vidhya/a-brief-history-of-neural-networks-c234639a43f1

World Intellectual Property Organisation. (1978). *Guide to the Berne Convention for the protection of literary and artistic works (Paris Act, 1971).* https://www.wipo.int/edocs/pubdocs/en/copyright/615/wipo_pub_615.pdf