

Generative Artificial Intelligence (AI): Guidance for Staff

Published: 14th September 2023

Review date: 14th August 2024

Learning and Teaching Enhancement, CCCU

Generative Artificial Intelligence (AI): Guidance for Staff

Contents

AI in Higher Education	2
Generative AI	2
AI literacy	3
Using generative Artificial Intelligence for learning and teaching	4
Ethics	6
Learning to use Artificial Intelligence effectively	6
Data protection	7
Artificial Intelligence and Academic Integrity	7
Transparency.....	10
Assessment	10
Equitable access	12
Discussion, evaluation and dissemination.....	12
Further Resources	12
Selected Glossary of AI terms	13
FAQs.....	15

AI in Higher Education

The launch of Chat-GPT and DALL.E2 in late 2022 popularised the use of generative AI, which is rapidly altering the landscape of higher education. The ability of generative AI tools to produce human-like language and images has sparked an international debate about AI's influence on education, and particularly learning and the assessment of learning.

At CCCU we recognise the potential of generative AI to support learning, teaching, research and working practices. However, we are aware of the tools' limitations and the ethical complexity of their widespread use within the University. **Therefore, our approach to the adoption and use of these tools is educative rather than punitive.** In this guidance, we will explain how we recommend staff approach the learning, teaching and integration of AI into their professional teaching and research practice. As these technologies and our approach to them continue to evolve, this guidance may be subject to change.

Generative AI

What is generative AI?

Generative AI refers to a category of artificial intelligence algorithms that produce content in response to prompts. Non-generative AIs have been in use for many years, including in traditional search functions used by Google Search, Amazon, Netflix etc. Generative AI has made significant advances since 2018 and is now easily accessible through user-friendly interfaces. The most common of these are chatbots such as Chat-GPT, GPT4, Bard, Claude etc. and image generators like Jasper Art, DALL.E 2, Midjourney, Stable Diffusion etc. Generative AI is also being built into search engines to improve results from natural language prompts (e.g. Bing) and connected to plug-ins to improve capabilities in mathematical and technical problem solving (e.g. Wolfram).

I've heard about Chat-GPT. What is it?

Chat-GPT stands for Chat Generative Pre-trained Transformer. It is a chatbot built on a generative artificial intelligence system that uses natural language processing to produce human-like text responses to prompts entered by the user. The chatbot's response depends on the prompt entered, and a user can enter multiple prompts to alter the original answer given. Chat-GPT is owned by OpenAI, and trained on their large language model, GTP-3.5. OpenAI have also released a paid-for chatbot based on their more recent large language model, GTP-4. Other generative AI chatbots work in a very similar way, but may have access to different training data and use different algorithms, therefore giving different levels of accuracy to their responses.

How does it work?

Generative AI systems that focus on text, such as OpenAI's GPT-3.5, learn patterns and relationships between words and phrases in natural language use. Large language models are trained on enormous amounts of data "mined" from open-source internet content. They then reproduce these patterns in response to prompts. The GPT model has a "self-attention" mechanism that allows it to decide on the relative importance of each part of the prompt or question entered by the user and decide what information is most relevant. It can then create contextualised responses. Therefore, Chat-GPT and other text-based generative AI systems,

operate like very powerful search engines that also summarise content into human-like written responses.

What are its limitations?

1. Generative AI chatbots cannot distinguish between truth and falsehood. They generate responses based on the most likely, or most common, pattern of language available in their training data. The current free version of Chat-GPT was trained on open-access data up until 2021, and the current paid-for version, GTP-4, on open-access data up until early-2023. Any data that is published behind a paywall, such as academic or professional publications, were not included in the GTP training data.

However, internet connected chatbots, such as Microsoft's Bing and Google's Bard are able to access current information.

2. As generative AI cannot decide on, or be held accountable for, the truth of the information it reproduces, it can produce incorrect, biased and discriminatory responses.
3. It can also produce entirely fictional or nonsensical responses, called "hallucinations". These appear to occur when the training data is not sufficient to answer the prompt, so the AI makes it up in plausible-sounding language.

These, and other, ethical issues will be discussed throughout this guidance.

What about plug-ins?

Plug-ins, software that enhances an existing programme's performance, are now available that combine GTP4's language capabilities with their specialisms. The ability to use plug-ins is mainly available for subscription versions of generative AI chatbots. As plug-ins allow open-access generative AI chatbots to connect to the internet, they offer a work-around to access today's data. Plug-ins perform many specialist tasks, e.g. mathematics, coding, summarising PDFs, text-to-speech, diarising appointments, or even finding discount fashion, travel itineraries and restaurant reservations.

AI literacy

Non-generative AI has been in our lives for a long time already (Alexa & Google Dot, biometric passports, facial recognition on social media & law enforcement, lane assist on cars, spellcheck, autocomplete, internet search functions, etc.) and is likely to become ever-more prevalent.

One of our tasks will be to learn how to use generative AI effectively and ethically. It is important that we work with our students as they also navigate this rapidly evolving digital landscape. At the moment, the law has not caught up with generative AI, although plans for regulation in the [EU](#) and [UK](#) are underway.

We need to consider how AI affects our specific fields, how it intersects with the way we teach, how our students learn now, and how they will work in the future. We encourage academic staff to co-create with students and look for opportunities to build learning communities where insights and experiences from learners and teachers can cross-pollinate.

This may mean upskilling ourselves, revisiting current working and teaching practices and considering new ways of learning, teaching and assessing, using AI as a part of our toolkit.

Staff are encouraged to explore generative AI, including, but not limited to, Chat-GPT, so they are familiar with the technology that is available to their students. At present, we are building capacity for staff development in this area and are working with other University stakeholders, including IT and HROD, to develop a University-wide approach to supporting AI use.

Notwithstanding, if you would like to learn more about using AI within your learning, teaching, and assessment activities, please [get in touch](#) with the Technology Enhanced Learning team in Learning and Teaching Enhancement.

Using generative Artificial Intelligence for learning and teaching

Engagement and debate

As the reality of AI use in higher education is evolving rapidly, staff are actively encouraged to discuss, contribute and consult with students, each other, and central services about their experiences working, teaching, learning and researching with AI.

We encourage staff with teaching and pastoral responsibilities to discuss the pros and cons of using generative AI systems with students. Similarly, course teams may wish to facilitate a forum for discussion and debate between colleagues.

Teaching with generative AI

There are many ways generative AI can be used creatively and critically for learning and teaching. Learning with AI can incorporate simultaneously acquiring knowledge about a particular topic and learning digital literacy skills that encompass ethically aware, critical use of the tool itself. By designing teaching activities using generative AI, lecturers can integrate subject specialist knowledge, digital skills for future employment and develop traditional critical thinking skills around truth, accuracy, agency, meta-cognition and reflection.

Despite its potential for reproducing the bias that already exists in societal discourse, generative AI also offers opportunities to level the playing field for students who face barriers to producing clear written content in English, perhaps because English is an additional language, or due to learning differences. By making extremely powerful spelling, punctuation and grammar editing and proofreading openly available, assessors are free to reduce emphasis on the mechanics of

writing in their marking and feedback and focus on the quality of thinking, understanding and creativity expressed.¹

Further discussion of the permissible use of generative AI for assessments is discussed in the Assessment section below.

Information Literacy

Every subject at the University has a dedicated Learning and Research Librarian who supports staff and students. Librarians are experts in information literacy, that is finding, evaluating, organising and disseminating information. If you would like to learn more about how to teach students to find and use good quality information sources, [contact your subject librarian](#).

Foregrounding Human Intelligence

A key tenet of higher education is the development of thinking skills. Whether these are termed lifelong learning skills, critical thinking skills or meta-cognitive skills, the principle that students will learn how to think for themselves is core to our purpose. Therefore, it is important to discuss the effects of using generative AI on human agency. The decisions generative AI makes about what information to include and exclude in its responses and the language it uses to express them remain opaque. It remains vital to teach students how their specific personal, reflective, embodied and localised experiences affect the acquisition and use of knowledge, and also how their individual perspectives can lend these processes value.

Students should be made aware explicitly that they remain wholly responsible for their outputs, regardless of whether the information and ideas contained within them are generated by AI.

¹ The quality of written English is expected to be maintained, in line with OfS requirements, however the tools to achieve this are now more powerful and widely available. Where students require subject-specific support to improve their written English, this should be provided in accordance with existing marking rubrics.

Ethics

Generative AI systems pose many ethical issues that staff and students should be aware of when using them for work, research and creating content. Many issues are still emerging through current research, but some that you may wish to **encourage discussion with your students** about, include:

- Bias in training data, including on gender, race, religion, sexuality, disability, and geo-politics, among others. This is the same discussion we would have about using quality information for academic research.
- Misinformation, including "hallucinations", where generative AI gives authoritative-sounding answers that are factually incorrect.
- The unfairness inherent in un-critically using AI to complete assessed coursework, or sensitive tasks in a work environment.
- Training AI on copyrighted images, texts, music and code, which are then remixed/reproduced without compensation to the original creator.
- The role of academic researchers in "cleaning" data, used under not-for-profit licenses which is then used by for-profit enterprises. Examples exist in the fields of facial recognition, photography and video generation.
- The use of workers in poorer economies for data-labelling tasks. Data-labelling is the process by which humans "tag" raw data to give it context for that machine-learning. This practice includes humans viewing and labelling offensive, illegal and harmful content so that generative AI systems can learn not to reproduce it in their responses to users.
- Data security, as Chat-GTP and other generative AIs, stores all input and prompts, including data and personal information shared with it, and re-uses it as training data.
- The concentration of AI technology in a handful of private companies, mainly based in the USA.

[Click here](#) for some ideas on incorporating generative AI into teaching.

Learning to use Artificial Intelligence effectively

The effectiveness of generative AI's responses relies in greater part on the prompts the user enters. The art/science of writing effective prompts has been dubbed "prompt engineering". A prompt 'sets the context for the conversation and tells the LLM what information is important and what the desired output form and content should be.'²

² White, J. et al. (2023) 'A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT'. arXiv. Available at: <https://doi.org/10.48550/arXiv.2302.11382>.

Whilst this is a very new area, here are some top tips gleaned through traditional internet searching:

1. **Be specific.** Give all key information, including the information you want, and the length and format of the output, e.g. find me a **recipe** for **vegan chocolate brownies** that uses **metric measurements**.
2. **Give context.** If you provide contextual information the response you are given will be more relevant to your specific needs, e.g. I am having a **party for 20 children** and I want to bake vegan chocolate brownies to go in their **party bags**. Find me a recipe that uses metric measurements.
3. **Provide examples.** If it's possible to offer an example of the output you would like, but with your specific data instead, you are more likely to get a suitable response, e.g. you could offer the chatbot a link to a recipe site or chef you particularly like, or provide an example of a non-vegan recipe you enjoy and ask for a similar vegan alternative.
4. **Fine tune.** If you don't get the answer you're after, try changing the vocabulary of your prompt, e.g. **yummy** vegan chocolate brownie, or **kid-friendly** chocolate brownie. You could also adopt a persona and ask it to write a response "in the style of..."

Data protection

AI companies, including OpenAI, Google, Anthropic and others, collect and store all data shared with them on sign up and during interaction with their chatbots and image creators. Companies use this information for analysis, research, development and to comply with legal obligations, e.g. detection and prevention of fraud.

Therefore, **we recommend maintaining confidentiality** for yourself and your students when using these tools by redacting personal or commercially sensitive information, or information protected as Intellectual Property (IP).

We will be working with Governance and Legal Services to ensure consistency as we develop our educative approach to AI in this area.

Current government [guidance for Civil Service employees](#) offers an indication of the information security concerns surrounding the use of generative AI in the workplace.

Artificial Intelligence and Academic Integrity

Detection

Controlling the use of AI with surveillance and detection software is not feasible. Companies such as Turnitin have developed AI detection software that provides a percentage value indicating how much of a document may have been written by AI writing tools. Early studies have shown these tools are easily circumvented by changing just a few words in a paragraph.

There is also a significant risk of false positives, with original work that has been translated into English more likely to be flagged as AI generated.

We think that altering our teaching and assessment practices is a more pedagogically sound alternative to relying on detection and punitive arrangements to manage the arrival of open access AI writing tools. As with all digital innovations in learning and teaching, we continue to investigate, research, and monitor any developments in the field so that our approach remains appropriate and up-to-date.

Academic Integrity Policy

Our existing [Academic Integrity Policy](#) places the responsibility on lecturers and course teams, with support of professional services, proactively to teach all students about academic integrity at all levels of study (1.4; 3.2). We encourage all staff to familiarize themselves with the policy, however the key points for explaining what constitutes the acceptable and unacceptable use of AI to your students are:

1.3.1. take responsibility for your own work

1.3.2. refrain from any actions that would give you an unfair advantage over other students

1.3.3. be honest in presenting your work for assessment and ensure it is your own work. You must not get others (including generative AI) to complete the work for you, in whole or in part

1.3.4. acknowledge the work of others (including generative AI), by including a complete reference, where it contributes to your work

and

3.4. Students may choose to ask someone to proof-read their work before submitting it. Any editorial help must be acknowledged at the beginning of the work, and follow the guidance in the Academic Misconduct Procedures.

Use of AI in assessment is treated as any other third-party source and its outputs/use should be fully referenced and acknowledged in a reference list (see next sub-section).

Generative AI tools highlight the urgency and importance of integrating conversations about Academic Integrity, including plagiarism, copyright, bias and inaccuracy, into teaching at regular points throughout a student's journey through their course.

If you decide to allow the use of AI within your assessments, ensure that students are clear about the parameters of acceptable use and the reason for any restrictions, how you will check their fair use, and that they are able to reference it correctly.

We recommend that Module Leads place a statement in the assessment section of their relevant student handbooks and Blackboard sites outlining any restrictions or guidelines related to the use of AI on that module. This will ensure that students are aware of any limitations on the use of AI and can help to promote consistency across different assessments within a module.

If you suspect a student has gone beyond the boundaries of acceptable use of generative AI, as defined by your assessment criteria in conjunction with the [Academic Integrity Policy](#), then the [Academic Misconduct Procedures](#) apply as normal.

Acknowledging generative AI in assessed work

All use of generative AI in assessed work should be acknowledged and fully referenced.

1. We recommend that course teams provide students with a choice of generic statements that can be used verbatim, or adapted, to add at the top of their written work:
 - a. No content created by generative AI has been included in this work.
 - b. I acknowledge the use of outputs from [insert the name of generative AI tool used] in the learning, preparation, planning or proofreading of this work.
 - c. I acknowledge the inclusion of outputs from [insert the name of generative AI tool used] in this work, in modified form.
2. We recommend that students include all prompts used and outputs generated by generative AI as an appendix in their written work.

A majority of generative AI chatbots have a facility to copy/email your conversation, so this should not add substantively to students' assessment preparation time.

3. All use of generative AI should be referenced. See Cite Them Right for advice on how to reference following your subject's conventions. The guide for Harvard is available [here](#).

Acknowledging generative AI in published work

Researchers are discovering the potential for generative AI to speed up a number of research-related tasks, particularly around literature reviews and data analysis. There are tools, such as Research Rabbit and Elicit, that are explicitly designed and pitched at the academic community for this purpose. Things to consider when using generative AI to conduct research for publication include:

- The proposed journal's rules on AI use. Springer Nature has led the way with [a policy](#) and [ground rules](#) that explicitly prohibit generative AI as a co-author on grounds of accountability and responsibility, but permit its use when fully acknowledged.
- The uploading of published academic work to generative AI tools for summary, as this may breach the University's copyright arrangements with publishers.

See also: [Academic Integrity Toolkit](#)

[Academic Misconduct](#) student webpages

Transparency

AI writing tools have the potential to be extremely useful to both staff and students in many areas of work and study life. We must use AI responsibly ourselves and teach our students to do the same.

Course teams should be direct and transparent about the use of AI in developing teaching and assessment activities and provide a clear rationale for any restrictions on its use by students. This can help to promote academic integrity and prevent any misunderstandings or ethical issues.

See also: [Course Directors' Handbook](#)

[Module Leaders' Handbook](#)

Assessment

Assessment Design

Modifying the design and structure of assessments can greatly diminish the probability of students engaging in cheating while also improving their learning experience.

Studies have revealed that when students have a personal connection to the questions they tackle, comprehend how the assessment contributes to their long-term learning objectives, complete preliminary work beforehand, and engage in discussions about their initial efforts with peers, they are less inclined to cheat.

These qualities can be incorporated into authentic assessments, which are assessments that mirror real-life scenarios, demand cognitive rigor, and necessitate genuine evaluation.

Tasks which are more difficult for AI to do well include:

- Reports on independent research activities,
- Progressive/reflective portfolio-style assignments that are built up over time,
- Briefs for real clients, or simulations of real-world professionals,
- Assessments based on location, needing in-place knowledge (e.g. field trips),
- Interactive oral assessments,
- Programme-level or synoptic assessment,
- Analysis of images or videos,
- Video-based assessments,
- Interpretation or creation of real-world artefacts.

Formative assessment

Generative AI synthesises its answers from training data, namely the internet, without necessarily discerning between high- and poor-quality information. It is a useful tool for forming ideas into sentences, if you already know the information contained is correct. However, core

critical thinking skills remain vital to discern between correct and incorrect knowledge.

Educating students about the advantages and limitations of generative AI through taught synchronous or a-synchronous activities is an important safeguard lecturers can build into their usual critical thinking development activities.

Opportunities to test students' ability to do the things AI does reasonably well, e.g. re-produce knowledge, apply theory or summarise a research paper, are best done formatively. Concept checks of core knowledge are best scaffolded into formative assessment. Alternatively, the weighting given to these tasks in summative assessments could be reduced, with more weighting given to the elements of learning that generative AI cannot easily replicate, such as evaluation, explaining meaning in context and reflection on personal learning.

"At risk" assessments

Assessment type changes require minor modifications. This is typically a lengthy process which negates a reactive approach to assessment re-design. If you are concerned that your assessment is at risk from generative artificial intelligence, there are a few steps you can take to mitigate this now:

1. Change the assessment question to ensure reasoning skills are required. Generative AI does not have capacity to reason, yet.
2. Re-visit your marking rubrics to ensure that to pass, a student must exhibit analytical, evaluative and critical skills, appropriate to their level of study. This must be kept in line with learning outcomes.
3. Introduce a reflective element to written tasks, particularly reflecting on learning, or events in synchronous teaching time.
4. Introduce annotated bibliographies, and tailor these to texts that are on the core reading list and you know well.
5. Require information or sources that are under 2 years old.
6. Exclude the use of lists, as generative AI is very adept at producing these.

Equitable access

Currently there are many open-access generative AI writing tools, however paid versions are beginning to be brought to market. The newer versions are trained on more up-to-date data. Free versions often restrict the number of requests a user can make per day, whereas paying users have unlimited access. It is likely that future iterations of generative AI tools will require payment or subscription.

Therefore, there is an emerging equality issue surrounding the use of generative AI for teaching and assessment, as students with lower incomes or less internet connectivity may be less able to afford access to emerging technology. Equitable access must be considered when designing learning, teaching and assessment activities using generative AI.

Discussion, evaluation and dissemination

1. Talk with students and colleagues about it!
2. Course teams are encouraged to create local groups to discuss the use of generative AI in teaching and assessment.
3. Colleagues are encouraged to engage with generative AI in their discipline, sector and/or industry, for research and teaching and consider its impact on students' subject learning.
4. Case studies are encouraged – please submit your preferred format to LTE for publication on [Prism](#). Guidelines for submissions can be found [here](#).
5. The University has many interdisciplinary networks through which to share research and best practice. Course Directors are encouraged to facilitate and disseminate learning through these channels. Co-creation across disciplines will become increasingly necessary as generative AI advances, and we are all learning together.

Further Resources

[Ada Lovelace Institute](#) (Accessed: 7 Sept 23)

Atwell, S. (2023) "[Generative AI and how students are using it](#)" JISC: National Centre for AI. (Accessed: 7 Sept 23)

Department for Science, Innovation and Technology (2023) "[A pro-innovation approach to AI regulation](#)" (Cmnd. 815). (Accessed: 7 Sept 23)

Department for Science, Innovation and Technology (2023) [Office for Artificial Intelligence](#). (Accessed: 7 Sept 23)

Equality and Human Rights Commission (2023) "[Artificial Intelligence in public services](#)" (Accessed: 7 Sept 23)

- JISC (2023) [National Centre for AI](#) (Accessed: 7 Sept 23)
- JISC (2023) "[A Generative AI Primer](#)" (Accessed: 7 Sept 23)
- Nerantzi, C., Abegglen, S., Karatsiori, M. and Martínez-Arboleda, A. (eds.) (2023) *101 Creative ideas to use AI in education, A crowdsourced collection*. Zenodo. doi: [10.5281/zenodo.8072950](https://doi.org/10.5281/zenodo.8072950). (Accessed: 7 Sept 23)
- NHS (2023) [AI Dictionary](#) (Accessed: 7 Sept 23)
- Russell Group (2023) "[New principles on the use of AI in education](#)" (Accessed: 7 Sept 23)
- Sabzalieva, E. and Valentini, A. (2023) "[ChatGPT and Artificial Intelligence in higher education](#)" United Nations Educational, Scientific and Cultural Organization (UNESCO). (Accessed: 7 Sept 23)
- Taylor, S. (2023) "[\(If you\) USEME-AI...a draft structure to support conversations and actions.](#)" (Accessed: 7 Sept 23)
- The Alan Turing Institute (2023) [AI for Public Services](#) (Accessed: 7 Sept 23)
- Tobin, J. (2023) "[Artificial intelligence: Development, risks and regulation](#)", *In Focus* (House of Lords Library). (Accessed: 7 Sept 23)
- UCL (2023) "[How UCL is redesigning assessment for the AI age](#)" (Accessed: 7 Sept 23)
- University of Kent (2023) [Learning Technologies Blog](#) (Accessed: 7 Sept 23)
- UKRI (2023) "[Transforming our world with AI: UKRI's role in embracing the opportunity](#)" (Accessed: 7 Sept 23)
- UNESCO (2023) "[Artificial intelligence](#)" (Accessed: 7 Sept 23)

Selected Glossary of AI terms

Algorithm – a defined procedure, or set of rules, used in mathematical or computational problem solving to complete a particular task.

Artificial Intelligence (AI) – a branch of computer science dealing with the ability of computers to simulate intelligent human behaviour.

Artificial General Intelligence (AGI) – the theoretical ability for computers to simulate general human intelligence, including the ability to think, reason, perceive, infer, hypothesise, but with fewer errors. This is not yet reality, but is the goal of many large AI research enterprises.

Artificial Narrow Intelligence (ANI) – the ability of a computer to simulate human intelligence when focused on a particular task. A well-known example of this is the development of chess-playing computers, which beat the best human players in the early 2000s and are now accepted as training tools available on any smartphone.

Bard – A generative AI chatbot developed by Google.

Bias – Any pre-learned attitude or preference that affects a person's response to another person, thing or idea. In the context of AI it commonly refers to a chatbot's reflection of bias present in its training data (namely, the internet) in its responses to users' queries.

Big data – an accumulation of a data set so large that it cannot be processed or manipulated by traditional data management tools. It can also refer to the systems and tools developed to manage such large data sets.

Chat-GPT – The first publicly available generative AI chatbot. Developed by OpenAI and released in November 2022.

Claude 2 – A generative AI chatbot developed by Anthropic.

DALL.E – A generative AI tool that produces images in response to text prompts. Developed by OpenAI.

Data mining – the practice and process of analysing large amounts of data to find new information, such as patterns or trends.

Data tagging – a process in data classification and categorisation, in which digital 'tags' are added to data containing metadata. In the context of generative AI, training data for Large Language Models is tagged by humans so the AI can learn whether to include or exclude it from its responses. This may be to comply with legal requirements, or ethical and moral codes.

Deep learning – a subset of machine learning that works with unstructured data and, through a process of self-correction, adjusts its outputs to increase its accuracy at a given task. In the context of AI, this process is closely related to reinforcement learning.

Generative artificial Intelligence – Generates new content from existing data in response to prompts entered by a user. It doesn't copy from an original source, but rather paraphrases text or remixes images and produces new content. It learns via unsupervised training on big data sets, but does not reason or think for itself.

GPT4 – The newest, most powerful chatbot from OpenAI. It is a subscription service.

Hallucination – Answers from generative AI chatbots that sound plausible, but are untrue, or based on unsound reasoning. It is thought that hallucinations occur due to inconsistencies and bias in training data, ambiguity in natural language prompts, an inability to verify information, or lack of contextual understanding.

Large Language Model (LLM) – A specific type of generative AI that specialises in producing natural language. Chatbots such as Bard, Claude and GPT-4 are examples of LLMs.

MidJourney – A generative AI tool that produces images in response to text prompts. Developed by MidJourney, Inc.

Machine Learning – The capacity of computers to adapt and improve their performance at a task, without explicit instructions, by incorporating new data into an existing statistical model.

Natural Language Processing (NLP) – The analysis of natural language by a computer. It is a field of computer science that takes computational linguistics and combines it with machine and deep learning models to allow computers to understand text and speech, and respond accordingly. Voice activated “smart” technologies and translation software are two of many everyday uses for NLP.

Neural networks – a computing architecture, or machine learning technique, where a number of processors are interconnected in a manner inspired by the arrangement of human neurons; they can learn through a process of trial and error.

Pattern recognition – a data analysis method that uses machine learning to recognize patterns, or regularities, in big data.

Plug-in – a small piece of software that enhances the ability a larger system or application to fulfil a specific task, e.g. a referencing plug-in on a web browser can pull the metadata from a webpage into a reference management system to create a bibliographic entry.

Recurrent neural networks (RNN) – a neural network that is trained to “remember” past data to predict what should come next. It is used for ordinal tasks, such as language translation and natural language processing, as language is sequential arrangement of letters to create meaning.

Reinforcement learning – a process whereby a deep learning model learns to become more accurate at a specific task based on feedback. The process by which it improves its accuracy is called backpropagation.

Supervised learning – a machine learning method where the model is trained using data that has been labelled by a human, i.e. training using examples. This is useful for predicting future outcomes based on past trends where data already exists.

Turing Test – a ‘test’ devised by British computer scientist Alan Turing to distinguish if a computer was “intelligent”. He posited that a human interrogator must ask questions over a fixed period of time to both a computer and a human, and distinguish which was which based on their replies. A computer would be deemed to have passed the Turing Test when a human could not distinguish between its responses and a human’s.

Unsupervised Learning – a machine learning method where the model identifies patterns in unlabelled data and makes inferences for use on future, unseen data. This is useful for looking at raw data and seeing if there are patterns occurring.

FAQs

To allow for regular updates the FAQs are hosted on [this Padlet](#).