

Using generative AI in the classroom

A guide for computing
teachers



Foreword

The advent of generative AI and the explosion of interest since ChatGPT was introduced to the general public in November 2022 can hardly have escaped anyone's notice. In particular, the mixture of excitement and anxiety it has brought to schools and teachers cannot be underestimated. All over the world, schools are trying to explore how generative AI can be used in the best interests of learners and teachers, while the constant announcement of new tools and applications makes it incredibly difficult to keep up. Computing teachers are at the sharp edge of these developments as they seek to keep pupils educated in new aspects of technology, where they have a responsibility for internet safety, and where they may be seen by their schools as resident experts.

Conscious of the needs of computing teachers, including all those who may not necessarily consider themselves specialists, a small working group of interested teachers and researchers have worked together to develop guidance and examples about the area of generative AI and its use in the teaching and learning of computing in schools. We are particularly focused on the UK, but there's no reason why this document might not be useful to you wherever you are in the world. Our approach has been led by the working group's perspectives as active computing teachers so we hope that what you read in these pages is useful, accessible and relevant. The target audience for the whole document is all teachers in primary, secondary and at sixth-form level, who teach some computing at some point in the year. Other teachers may be interested in specific sections on generative AI for whole school issues and for administration.

The working group first met in December 2023, and has met online several times since. It consists of eight computing teachers (two primary and six secondary) alongside a small number of other academics and researchers. It's been a participatory and collaborative project, and we've listened to each other's perspectives and views and tried to incorporate all of these in this short document. What we've focused on is providing examples of what you might use generative AI for in the classroom, specifically for computing, and some background definitions and tips that teachers might find useful.

Some provisos. We're aware that there exists a plethora of guidance and information on the topic of generative AI; this document sits alongside that as a small contribution, and part of our work has been to signpost useful resources for further reading and research. We're also aware that this is a fast-moving field, and that documents like this will need revising sooner or later, probably sooner. That said, we hope you enjoy reading it and we look forward to hearing your comments, and welcome volunteers for the next iteration when it comes!

Sue Sentance, on behalf of the working group

Raspberry Pi Computing Education Research Centre, University of Cambridge



Contents

Introduction	4
Part one: Teaching and learning with GenAI	5
Teaching about AI and GenAI	6
GenAI and the teaching and learning of computing	8
Teaching programming with GenAI	9
GenAI technology in primary education	11
GenAI for supporting all children	13
Part two: Using GenAI	14
GenAIs for administration	15
Whole school issues relating to GenAI	16
GenAI, bias and hallucinations	17
GenAI and privacy	18
Conclusion: What lies ahead?	19
Further reading	20
Appendices	21
A. More GenAI terminology	21
B. How an LLM works	21
C. An example of prompt engineering in lesson planning	23
Working group	25



Introduction

The release of ChatGPT in November 2022 by OpenAI brought generative AI (GenAI) to the attention of educators and society. This prompted many to consider the potential role and implications for education. ChatGPT¹ is an example of a GenAI application, where the user engages with a GenAI model via a chatbot, largely in text. GenAI is a technology that can generate human-like content in text, image, audio and video as a result of the model having been trained on large publicly available datasets. It is the training on large datasets and the use of deep learning that leads to quite remarkable human-like output.

GenAI presents numerous opportunities for enhancing teaching and learning in schools. GenAI tools can assist teachers with administrative tasks so that they have more time to attend to students' learning needs. Teachers can also use GenAI to help develop educational materials. This has the potential to make it easier for teachers to tailor educational materials to suit the individual needs of their students. Moreover, GenAI tools are showing some promise in their use to support adaptive testing and feedback. For example, teachers might be able to use GenAI tools to adjust the content of assessments to match students' progress and performance.

Computing teachers at primary, secondary, and post-16 levels are likely to have started to consider issues surrounding AI and how to teach about AI and the social and ethical implications. The introduction of GenAI raises questions about how to make use of the capabilities of such tools in the teaching and learning of computing, while maintaining the integrity of assessment. Further questions arise about how to use GenAI products to assist in a range of administration functions associated with teaching including report-writing, lesson-planning, and developing teaching material and timetabling. Ethical questions are also brought to the fore over issues of biased and inaccurate output, and energy usage. There are also concerns over safety, data privacy, copyright, plagiarism and the threat to assessment and students' critical thinking. While this guidance is not aimed at answering all these questions in any detail, it serves as an introduction to the main issues while primarily aimed at giving computing teachers more confidence in taking GenAI into the classroom.

In this document, we suggest how applications based on GenAI might be used in a practical way by computing teachers, with additional pointers to some of the issues that extend beyond computing as a specific subject in the curriculum. The guidance is structured in two parts. Part one of the guidance considers teaching about GenAI, teaching computing and programming with GenAI, followed by separate sections on GenAI for primary education and on using GenAI to ensure inclusive education. In Part two, we address more generic issues around GenAI, including how it can be used for a teacher's administrative tasks and across the whole school. We then look at issues of bias and privacy that teachers will be aware of emerging from GenAI use.

¹ Other examples of text-based GenAI include Anthropic's Claude and Google Gemini



Part one: Teaching and learning with GenAI



Teaching about AI and GenAI

Before using GenAI tools in your classrooms, it may be useful for you to spend some time helping pupils understand something about AI, in an age-appropriate way. For example, helping students identify and recall systems they use in everyday life that have AI technology embedded in them is probably a good starting point. For older learners, teaching some of the foundational concepts of neural networks and machine learning could help them understand how GenAI works.

Teaching the concepts related to AI and working out what pedagogy to use is likely to feel like a complex matter, particularly if you are not familiar with these technologies yourself. Pedagogical opportunities might include exploration of GenAI apps, but it's very important to consider e-safety and age-restrictions. Consider introducing research-based activities where some of the challenges and benefits can be discovered and they can link this to the pupils' developing knowledge of AI models and engines. For example, students could try to use a chatbot when carrying out a familiar task, try to spot where it was and was not accurate, and reflect on why that might be. For a general introduction, the [Experience AI resources](#), designed for 11- to 14-year-olds, can also be used for introducing the foundational concepts of AI.

GenAI tools work primarily by responding to prompts. Prompts are usually text-based inputs designed to guide the GenAI tool to produce a desired output. If students are writing prompts, they will need to learn to be precise and unambiguous and provide enough detail to get the output they require. Prompt engineering (see [key terms](#)) is the process teachers and students can use to create better prompts when interacting with various GenAI applications. In the same way as we might teach about complex search terms, prompt engineering may need to be explicitly taught; in addition, understanding how GenAI works can help students get better at prompt engineering.

Resources are available to help teachers learn more about GenAI. The [Appendix](#) of this document includes definitions of some additional terms, an explanation of how an LLM works, and an example of prompt engineering for lesson planning. The [Further reading](#) section signposts a host of other material that can help you increase your understanding of this area so that you feel confident supporting your pupils.

Example: A secondary-level enrichment session

This lesson outline could be used for a general introduction to various aspects of GenAI without becoming too technical.

1. The lesson could start with a general history of AI, and a simple explanation of GenAI and large language models as a subset.
2. Students could work in pairs to use a chatbot to prompt output to give some more details about the history of AI (for example). Some time could be spent exploring how to improve prompts to get more tailored information.
3. The teacher could suggest that students compare answers from chatbots with slightly different prompts - or pairs of students could compare whether their answers are valuable or not.
4. The students could develop guidelines for what makes a good prompt.
5. The teacher could give examples of where chatbots might hallucinate (give out false information).
6. Finally, the students could research other tools that generate images or music, and discuss what kind of inputs work well for the best results and issues of bias and misinformation.



GenAI key terms

Terminology	Explanation
Artificial intelligence	Artificial intelligence (AI) is a broad field for which many definitions exist. The Cambridge Dictionary defines AI as “the use or study of computer systems or machines that have some of the qualities that the human brain has, such as the ability to interpret and produce language in a way that seems human, recognize or create images, solve problems, and learn from data supplied to them”. ²
Anthropomorphization	The action or fact of attributing human characteristics, form, or personality to something non-human (in later use especially an animal).
Chatbots	Chatbots are systems designed to produce output that appears to be like a conversation with a person. Chatbots can be pre-programmed based on a set of rules, with predefined responses such as ELIZA. ³ However, modern chatbots such as Siri and Amazon’s Alexa are built based on the use of a data-driven approach, such as using adaptive deep learning algorithms.
Generative AI	Generative artificial intelligence (GenAI) refers to a class of technology designed to produce new content such as text, images or sound that is similar to, but not identical to, the training data used to build the underlying model. The output from GenAI models is particularly useful for building applications that require the creation of realistic images, generation of text, composition of music, and even designing new molecules.
Hallucination	Hallucination is a phenomenon whereby GenAI systems generate unreliable output. These outputs are in the form of statements that may appear reasonable but are either cognitively irrelevant or factually incorrect. Teachers are encouraged to do fact-checking with external information in the real world to either affirm or negate the statements generated by GenAI applications.
Large language model	A large language model (LLM) refers to a type of GenAI system designed to generate human-like text or speech. At its core, this technology employs sophisticated data-driven approaches to process vast amounts of data. In processing the data, representations of patterns in the data are built that can be used to produce output that mimics the patterns in the original data.
Prompts	Instructions or questions, mostly in a natural language format, used as input to an AI system to generate an output.
Prompt engineering	Prompt engineering is a practice which involves designing, refining, and implementing prompts used as input to GenAI systems. It is a relatively new area of research and it has proven to be beneficial in improving the output of GenAI systems in areas such as education, medical practice and research across various disciplines.

²<https://dictionary.cambridge.org/dictionary/english/artificial-intelligence>

³<https://en.wikipedia.org/wiki/ELIZA>



GenAI and the teaching and learning of computing

Computing is a broad subject so there may be many ways in which GenAI applications can be used to teach and assess it. As with any subject, GenAI applications can be used to create summaries of textual materials for students, and to prepare sets of questions on particular topics. Teachers have reported that GenAI applications can be very useful for developing glossaries and terminology lists for pupils at different levels of complexity. However, it's really important that teachers check any materials given to students for accuracy.

Teachers can be role models when it comes to using GenAI with students. They can demonstrate how they use it themselves, provide screenshots and other examples of the ways in which they have interacted with LLMs, and model good practice. For example, teachers could demonstrate to students how to cite image/text/code generated by GenAI and facilitate discussion around plagiarism.

Ethical issues are core to the computing curriculum and using GenAI as a case study can make these lessons come to life. For example, using a chatbot offers an opportunity to teach about real-world ethical issues relating to AI. For pupils of all ages, introducing some of the known ethical issues associated with AI would start to engage them to think critically about technology and its implications. It is a good idea to use examples of AI applications that your pupils are familiar with, such as uses of AI in Netflix or other movie recommender systems, or in what they might see using social media, about both their benefits and harms. This can lead to interesting discussions about the way that GenAI applications might be biased, fallible and hallucinate. This might help young people to think critically about these new technologies and understand their capabilities and limitations.

For older pupils, there are many real-world examples to draw on, including how AI can be used in the criminal justice and health domains, highlighting that critical errors can occur if biased training data is used. When students are old enough to use GenAI themselves, the output of GenAI can help to introduce discussions around ethical issues relating to technology. Issues of bias, privacy and trustworthiness can be introduced and explained: it's an opportunity to think critically about the technology and its socioethical implications.

Example: Using LLMs to discuss social justice and equity issues

Culturally responsive computing teaching involves relating computing topics to students' own cultural backgrounds and life experiences. It can also highlight issues of social justice and bias relating to technology. Secondary students could watch the film 'Coded Bias'⁴ to become aware of some of the general issues around AI and bias, before focusing specifically on generative AI with respect to bias.

Secondary students can use GenAI applications to investigate social justice issues of which they

are aware. For example, teachers could design a research-based lesson, where students use GenAI and prompt engineering to investigate an issue such as the underrepresentation of women, minorities, and individuals from diverse backgrounds in technology-related fields. Students might use the GenAI application to gather information on statistics, challenges, and initiatives related to diversity and representation in the tech industry and compare the output to data from standard searches.

⁴ <https://www.codedbias.com/>



These discussions don't necessarily need to be limited to the computing classroom either: we have heard a growing interest in teaching AI issues across many different curriculum subjects. For computing teachers, this can help to highlight the interdisciplinary nature of our subject.

Teaching programming with GenAI

Programming is one aspect of computing where there has been a lot of recent attention related to the integration of GenAI technology. Tools such as GitHub Copilot⁵ can be used to generate programs, although there are some concerns that students will use them to generate code that they do not understand and can't debug. Programming IDEs are increasingly being developed with AI functionality to give augmented error messages and feedback on how to proceed when programming. This doesn't involve the teacher or student using GenAI directly but rather using tools that have that built-in functionality.

While some studies are starting to show ways in which such AI-powered tools can be used to scaffold the learning of programming, we do need more research on how effective these tools and approaches are beyond just generating the code. With programming error messages, our own research shows⁶ that teachers can be wary of tools that do not give feedback that helps the student to work out their own mistakes and learn from them.

Programming pedagogy is an area where we've increased our understanding in recent years, and the example in this section shows how a GenAI application can be used alongside the PRIMM⁷ pedagogy for structuring a lesson.

Example: Generating programming examples built around the PRIMM structure

- Let's assume you regularly use PRIMM for teaching programming but you need more examples.
- The first goal is to use a GenAI application to generate some examples for the Predict stage of PRIMM. Students could then proceed to use these for the Run and Investigate stages once they predict what the program will do.
- In order to generate programming examples for a lesson you need to use prompt engineering to be very specific about the content that you want.
- For PRIMM, you are likely to have to include a description of the PRIMM stages, although see the figure on page 10.
- The more information you can include the better, particularly the programming constructs you want to focus on, and the programming level of the learners.
- Offering examples to the GenAI application will help it to generate something that is similar to the content you usually use.
- Once you are happy with the examples (this may take a few attempts), you could try to ask the GenAI application to generate a range of Investigate questions (you would need to provide a variety of examples) for the code snippets, and some potential exercises for the Modify stage.

⁵ <https://github.com/features/copilot>

⁶ <http://rpf.io/pemstudy>

⁷ <https://computingeducationresearch.org/projects/primm/>



PROMPT: Create me a PRIMM activity for a lesson introducing for loops in Python

Absolutely! Here's a PRIMM activity tailored to introduce for loops in Python.

Title: Looping Shapes

Objectives

- Understand the concept of a for loop for repeating actions in Python.
- Utilize for loops with the Python Turtle library to draw patterns.

Materials

- Computers with Python installed (and ideally an online Python environment like Trinket [\[https://trinket.io/\]](https://trinket.io/) where they can also use the Turtle library without additional setup).
- PRIMM worksheet (provided below)

Procedure

1. Predict

- Display this code snippet:

```
Python
import turtle

for i in range(4):
    turtle.forward(100)
    turtle.left(90)
```

- Ask: "What do you think this code will do? Draw your prediction."

2. Run

- Have students run the code to see if their predictions were correct.

3. Investigate

- Explain the core elements:
 - for i in range(4): This loop repeats the code inside 4 times.
 - turtle.forward(100): Moves the turtle forward.
 - turtle.left(90): Turns the turtle.

4. Modify

- "Change the number in range() to 3. Predict what will happen, then run the code."
- "Change the angle in turtle.left(). Predict the new shape and test!"
- "Can you make an octagon?"

5. Make

- **Challenge 1:** Create a pattern of squares (Hint: You might need nested loops!)
- **Challenge 2:** Draw a colorful spiral by changing the pen color in each loop iteration (You can find color names for Turtle online).

GenAI can be used to create draft learning materials (created using Google Gemini, April 25th, 2024)



Drawing on another approach to teaching programming, research has shown that using Parson's problems is a very effective way to build programming skills.⁸ A Parson's problem is a code listing where the lines of code have been mixed up and students need to put them back into the right order. A GenAI application can help here if it is provided with an example of a program and asked to mix it up. It could be prompted to provide different versions of the mixed-up code at various difficulty levels.

GenAI technology in primary education

In primary classrooms, introducing GenAI technology as a topic is likely to be daunting for both non-specialist and specialist computing teachers alike; it is unlikely that educators will have learned about teaching these topics. Therefore, teacher professional development and resources that help demystify the content and provide pedagogical guidance will be very important.

The first use of GenAI tools in primary settings is likely to be for administrative and teacher productivity purposes (see [GenAIs for administration](#)). For example, GenAI tools might be used to help teachers create resources for classroom use, such as learning texts and questions. As primary teachers usually teach across the curriculum, it may be that GenAI tools are first used to create learning materials in other subjects, for example, for literacy, rather than computing (but this remains to be seen).

Example: Developing comprehension for reluctant readers

- Find out their interests (e.g. Minecraft)
- List a set of related terms (e.g characters, scenes and contexts)
- Using a GenAI tool, write a prompt, with the related terms
 - To write a draft set of stories
 - To write a draft set of retrieval questions with inference or other focus
- Check, adapt and refine the draft materials
- Share the checked and refined stories and questions with students to support teaching and learning

Beyond learning material design and administration support for teachers, primary schools may start to teach some basics **about** AI technologies. E-safety is one aspect of AI technology that is likely to be introduced first in primary schools. How this is integrated into the curriculum will vary from school to school. Some schools may add a new stand-alone topic, and others could integrate it into existing units of work, say, in computing, PSHE&C (Personal, Social, Health Education & Citizenship), literacy, art, mathematics, or cross-curricular topics. Primary-aged pupil learning outcomes might include students being aware of everyday applications that include GenAI, the potential dangers associated with such applications including bias, and issues around ethics as well as the potential benefits of such tools.

⁸ <https://dl.acm.org/doi/abs/10.1145/3230977.3231000>



There are a growing number of resources available for teachers to select from to teach about AI such as in [Mr P's online ICT resources](#) that have lesson ideas and teacher support. Many resources that teach about AI are aimed at secondary students. Some aspects of these resources could be adapted for primary school use, for example, the [Experience AI resources](#) or [Code.org AI materials](#). Beyond computing and PSHE&C, AI-produced texts and images could be introduced in literacy and art lessons and the importance of data and probability might be highlighted in maths as example contexts.

Example: Introducing e-safety and GenAI

E-safety is likely to be the first aspect of GenAI application use that will be taught in primary classrooms.

A Key Stage 2 Safer Internet Day lesson

- Start with the class brainstorming examples of applications that use GenAI
- Watch a video about GenAI applications
- Discuss the potential e-safety issues of GenAI applications
- In groups, students create a page for a class book "Watch out! GenAI apps are about!" of examples of GenAI applications and related issues (online or physical book)
- The class book is shared with another group of students such as part of a transition to secondary school or another year group in the school
- Optional homework task: The class book is shared with families and additional ideas are added

At the time of writing, GenAI applications with a prompt engineering interface such as chatbots, or those with an enhanced programming interface, have a minimum (generally 13+) user age that is above the oldest of our primary students. Therefore, primary students should not be using such tools in class. However, within existing educational software used by primary students, GenAI technologies may have already been incorporated and more are likely to start to do so. AI-based software that is available to learners will need to be reviewed to ensure that teachers are aware of potential inaccuracies, bias and inequity. Teachers will need to figure out what impact this might have on teaching and learning, including pedagogy and e-safety.

Teaching about AI in primary

"I use Code.org AI for Oceans as a starting point for my Year 6 students. This introduces real AI applications and shows machine learning and facilitates the discussion of bias." (primary teacher)

Learning objectives for primary might include (i) knowing that vast amounts of data are used to train GenAI models or (ii) being able to train simple ML models. However, exactly what concepts and skills will be covered and in what order are not yet widely agreed.



An example for introducing GenAI in the primary classroom might be to:

- (i) Share keywords (e.g. AI, pattern, data, GenAI)
- (ii) Do an unplugged activity, such as on spotting patterns in data
- (iii) Watch a video to explain concepts and showcase real-world examples
- (iv) Do a practical activity, e.g. training a model using an education software such as [Teachable Machine](#) or [Machine Learning for Kids](#)
- (v) Reflect on what students have learned and what this means in the real world
- (vi) Finish with a keyword activity

GenAI for supporting all children

GenAI, when integrated into computing education, can be used in a variety of ways that enhance learning experiences, facilitate personalised education, and promote inclusivity. The inclusion of GenAI in educational settings can be categorised into three key functions: generating explanations, offering different viewpoints, and facilitating dialogue.

Explanation generation: GenAI can generate explanations and summaries tailored to the individual. This capability is especially valuable in computing education where concepts can vary widely in complexity and abstraction. For example, a GenAI system can present the concept of algorithms in different forms. This adaptability helps in addressing the diverse needs of students, ensuring that learning materials and content are accessible.

Offering different viewpoints: In educational environments, GenAI can act as a cognitive tool that assists students in exploring different viewpoints and constructing diverse interpretations of the same content. This is crucial in subjects like computing where problem-solving often benefits from innovative and varied approaches. For instance, when a student encounters a challenging programming problem, GenAI can suggest multiple solving strategies, ranging from recursive to iterative approaches, demonstrating the strengths and weaknesses of each. Not only could this enhance the student's understanding but also encourage a deeper engagement with the subject matter.

Facilitating dialogue: GenAI can serve as an intermediary in communications between students and educators. It may be able to help students articulate their thoughts and questions more clearly, which is particularly beneficial in large classes where individual interaction with the teacher is limited. GenAI could possibly be used to reformulate a student's poorly structured question into a well-defined one that is easier for educators to resolve. The use of AI can help rephrase their experiences or questions in a manner that is likely to be better understood and more effectively addressed by peers and educators. Research is needed to understand more about the pedagogical implications of this use of GenAI, to consider the potential development of anthropomorphic thinking (see [key terms](#)) around AI and over reliance of technology, as well as the child's ability to understand the output given by GenAI.



Part two: Using GenAI



GenAIs for administration

Whatever sector you are working in, from education to manufacturing, there are likely to be opportunities to improve productivity by the use of GenAI tools. In education, there is a growing use of GenAI applications to reduce teacher workload. For example, GenAI applications can be used to help teachers with ideas for lesson plans, create potential worksheets and assessment activities, provide suggested email responses and draft student end-of-year reports. Government guidance, websites and organisations have sprung up to support teacher and school awareness of, and access to, such productivity opportunities. Some of these are listed in the [Further reading](#) section at the end of this document.

Despite the outputs from GenAI tools providing good starting points for administrative activities and the design of educational resources, there are issues to be aware of:

- Content generated by GenAI applications may be factually incorrect, include bias or be insufficiently personalised;
- Materials presented by teachers are trusted, so if learning resources are incorrect, this may lead to student misconceptions;
- Where communications are too generic, parents and others may lose confidence; and
- Any bias in GenAI-generated resources, if left unchecked, could have significant implications.

This does not mean that we should never use GenAI applications for administration and productivity. Instead, teachers must be 'the human in the loop' to check, adapt and refine GenAI-generated suggestions and drafts. In this way, the teacher is responsible for all materials they create and use in their practice.

For guidance on how to use GenAI tools it is likely that acceptable use policies (AUPs) will be developed by schools. Such policies, along with other resources can inform parents and families and others in the school community on the uses of AI technology for both administration and productivity as well as in teaching and learning, providing transparency.



Example: Using GenAI tools for administration and productivity

Consider using GenAI tools to:

- Support simple email responses where the content is straightforward, for example, your rushed email might be refined by a chatbot
- Create first-cut helpsheets, wordbanks, writing templates, summaries, bulleted lists, worksheets or assessments for certain groups of students e.g. with certain interests, or additional needs
- Create sample text e.g. poems, stories, instructions, introductions etc (and potential related questions) in the style of certain authors, or voices for students to use in lessons
- Create draft knowledge organisers for students to refine
- Summarise videos for learning materials
- Create answers to essay-type questions, and then ask students to mark these
- Draft student reports from a summary of achievements and keywords, but remember to check and refine
- Draft lesson plans and example activities, for a new topic
- Support supply teachers with topic information
- Develop first-cut timetables
- Format and organise work data or text

Do:

- Fact-check and refine all output, particularly for communications to families and non-fiction learning material
- Check any maths output from GenAI tools, as at present, such output is more likely to be incorrect than for literacy-based output
- Be transparent about the use of GenAI tools
- Ask your school for professional development to help you learn about prompt engineering and using GenAI tools for administration and productivity
- Consider using one of the websites that support teacher GenAI system use, they can create extra prompts (in the background) from the data you enter in tailor-made screens for specific purposes (e.g. creating worksheets)
- If using GenAI tools to mark student work, be very cautious, be transparent about the use of the tool, and double-check the marking accuracy
- Ask your school to develop an acceptable use policy for GenAI technologies for administration and productivity, teaching about GenAI technologies and for using GenAI tools to help teaching and learning in any and all subjects

Whole school issues relating to GenAI

All school stakeholders should be considered when developing school policies and responses to the use of GenAI applications in education. Stakeholders include students, teachers, families, governors, education reporting bodies, and potentially inspectorate groups, etc.

A school's acceptable use policy (AUP) will likely be a key document in informing teachers and school staff on how they should use GenAI tools for both admin and productivity purposes and in teaching and learning. The AUP or other policies should provide stakeholders with a view of the uses of GenAI applications in the school, providing transparency of uses. AUP are likely to be different for primary, secondary and other school types. Many GenAI tools have age restrictions, younger children and more vulnerable older learners will need to be considered, and the e-safety, privacy and security issues for all students should be of paramount concern.



As well as defining the acceptable use of GenAI tools for teachers and other school staff, how students might use GenAI technologies can also be defined and agreed upon with students and families. For example, schools might require students to be transparent about the use of such tools for homework.

There may be some initial nervousness in being transparent about the use of GenAI tools in schools, but openness is likely to be increasingly important as media attention increases on this topic and families start to use similar tools in their own work and lives. Developing a climate of trust in the use of GenAI tools will be essential for schools.

Teachers are likely to need professional development and guidance, beyond an AUP, to help them gain confidence and knowledge to use GenAI tools effectively and safely. There are many issues that might not be in the control of the computing teacher but are important across the school.

Specialist computing teachers may be asked to take on the role of advisor and expert on AI technologies, as they are for other software and hardware uses in school. This may be not only for potential curricula changes across subjects, but also for teacher and administration productivity opportunities. It's increasingly the case that schools will want to give one member of staff the responsibility of considering the changes that GenAI technologies will have on a whole school basis, whether this should be the computing teacher or not, requires careful consideration. It would be useful to develop case studies as to how schools are introducing GenAI.

Example: Helping families be aware of GenAI tools

Families of students might be made aware of the uses of GenAI tools in a school including the basic concepts and language used in teaching and learning through:

- Curriculum and topic summaries
- School and student-acceptable use policies
- Presentations (e.g. during e-safety events)
- Policies, posters and pamphlets
- Homework activities

GenAI, bias and hallucinations

The way in which GenAI works means that there are some important things to think about in terms of the outputs to user prompts. These include issues of hallucination and bias.

Output from a GenAI tool is data-driven rather than rule-based. This leads to the possibility of output that looks realistic and plausible but is not accurate; such an instance can be called a hallucination. Misinformation is therefore possible, which requires teacher oversight while learners become familiar with the technology; teachers should give ongoing reminders to students to check the output from GenAI tools.



Generally, GenAI tools do not search the internet when they are being used nor do they have a database of rule-based facts or information. Rather, output is based on the probability that a word or number will next occur in the output string, based on a pre-trained data-driven model, created from processing vast amounts of data. Because GenAI processing is not rule-based, there can be particular issues with the accuracy of GenAI for solving mathematics and logic problems. However, GenAI tools are capable of generating explanations of mathematical ideas and logical principles.

Similarly, when prompts are written to elicit information, the output may be biased toward particular groups and ideas. This is because during GenAI model training any biases that already existed in the training dataset, such as from large volumes of publicly available text on the internet, will be replicated in the model.

Some have suggested that hallucinations, biases, and misinformation are limitations of GenAI. Others suggest that it may be better to see this as a feature with opportunities of translating ideas and text into other forms, like summarising text, suggesting alternative interpretations, explaining ideas, rephrasing text, or assisting in developing and evaluating computer code.

Understanding and using GenAI in this way means that hallucinations, biases and misinformation are less of a problem. It is important in introducing learners to GenAI to emphasise this approach in order that learners can use it effectively and safely.

Example: Exploring the concept of 'trust' in AI

A lesson could be designed around the concept of 'trust' in AI. After an exercise using GenAI to research a particular topic, ask students how much they trust the information they have been given. This can lead to a discussion about why and how AI can/cannot be trustworthy and how students should make sure they can trust the information they are presented with. Follow-up questions can revolve around the fact that bias in GenAI models can stem from biases present in the training data, and the potential reinforcement of stereotypes. The teacher can alert students to different types of bias including gender bias, ethnic and racial bias, socioeconomic bias, age bias and cultural bias.

GenAI and privacy

The General Data Protection Regulation (GDPR) imposes stringent requirements on data privacy and security relating to GenAI. There are two aspects of this, the first is in relation to the training of GenAI and the second is when using it, i.e., inputting prompts.

Training: It is indeed possible, and often likely, that personal data is used in the training data of GenAI systems. This data can come from various sources, including publicly available datasets, user-generated content, or data from the internet. Often, this training data can include personal information, either intentionally as part of the design to enhance the AI's understanding of human contexts, or unintentionally through the inclusion of data that has not been properly anonymized. However, in training, personal data is not stored as we understand it in the traditional sense. Despite not storing personal data in the traditional sense, the trained model can still reflect patterns, biases, or sensitive information inferred from the training data. This can lead to concerns about privacy and the potential for re-identifying individuals

indirectly through model outputs or behaviour. Because of the novelty of GenAI technology, it remains a contested issue about how and even if GDPR can be applied to GenAI as the regulation's inception was based on more orthodox digital and archival data storage.

Data privacy and prompts: Most of the popular GenAI models do not use user prompts for training their models. Although they might use user prompts to improve the model's interface. It is worth checking how user prompt data is used with the model that you want to use in the classroom. However, the most widely used models do claim that the interaction with a GenAI model is GDPR compliant. The overall advice to teachers is to take reasonable precautions and not use personal data in a prompt unless it is anonymised.

Conclusion: What lies ahead?

This guide has highlighted ways in which computing teachers can use GenAI in their teaching and other aspects of their work. The insights and examples provided here are based on the real-world experiences and suggestions of our working group, which have been tailored to meet the diverse needs of educators at the primary, secondary and sixth-form level. As we explore the potential of GenAI in education, it's crucial to prioritise internet safety, digital agency and awareness, and collaboration among stakeholders. By working together and embracing innovative approaches, we can ensure that both educators and learners are well-prepared to thrive in an AI-enhanced educational environment.

GenAI is here to stay but what might be next? Looking ahead, the future of AI in computing education presents exciting opportunities to improve teaching and learning practices. With AI-powered tools, educators can customise learning experiences, enhance teaching methods, and broaden access to computing education for all students. However, as we embrace these advancements, it's important to uphold ethical standards and promote responsible AI use among learners. Collaboration among educators, researchers, and industry professionals will be key to driving innovation and ensuring that AI enhances the teacher-student relationship rather than replacing it.

As computing teachers, we have a unique opportunity to shape the future of education in an AI-enhanced world. The rapid advances in technology may at times seem overwhelming and daunting, but collaboration between schools and between teachers to address the opportunities and challenges can help. Hopefully, we all share the goal of empowering our students with the skills and knowledge they need to succeed in the AI age while fostering inclusive, ethical, and effective learning environments.

If you've found this guidance useful, please share it with your colleagues!



Further reading

1. Guidance

Department for Education, England (2023). [Generative Artificial Intelligence in Education](#).

Department for Education, England (2023). [Artificial intelligence in schools – everything you need to know](#). The Education Hub.

Joint Council for Qualifications (2023). [Artificial Intelligence \(AI\) Use in Assessments: Protecting the Integrity of Qualifications](#).

OECD (2023). [Opportunities, guidelines and guardrails for effective and equitable use of AI in education](#).

UNESCO (2023). [Guidance for generative AI in education and research](#).

2. How GenAI technologies work

Elastic (2024). [What is a large language model \(LLM\)?](#) [Website]

Henrik Kniberg (2024). [Generative AI in a Nutshell - how to survive and thrive in the age of AI](#). [YouTube video]

Stephen Wolfram (2023). [What is ChatGPT doing...and why does it work?](#) [Blog post]

UNESCO (2023). [Education in the age of artificial intelligence](#). [Periodical issue]

Visual Storytelling Team, & Madhumita Murgia (2023). [Generative AI exists because of the transformer](#). Financial Times. [Website]

YouTube (various)

- [Machine Learning Explained in 5 Minutes](#). [YouTube video]
- [How AI is used in image recognition](#). [YouTube video]
- [What is Machine Learning?](#) [YouTube video]

3. Resources for teachers

AI for Education. [Free AI Resources for your School or Classroom](#). [Website]

Beverly Clarke. [Exploring Computer Science AI Alternate Curriculum Unit](#). University of Oregon. [Curriculum resource]

Code.org. Learn about AI. <https://code.org/ai>. [Website]

Computing at School. [CAS AI Community](#). [Website]

Fahad Tamton (2022). [Ethical AI Frameworks Education](#). [TED Talk]

Microsoft. [AI for education: Resources and learning opportunities](#). [Website]

MIT Sloan Teaching & Learning Technologies. [AI Hub: Generative AI for Teaching & Learning](#). [Website]

Mr P ICT Online CPD. [AI resources](#). [Website]

OpenAI. Teaching with AI. <https://openai.com/blog/teaching-with-ai>. [Blog post]

Raspberry Pi Foundation. Hello World Magazine Issue 22. [Teaching and AI](#). [Periodical issue]

Raspberry Pi Foundation and Google DeepMind. Experience AI resources. <https://experience-ai.org/en/>. [Curriculum resources]

Sal Khan (2023). [How AI could save \(not destroy\) education](#). [TED Talk]

Teaching with ChatGPT. [Transform your teaching with ChatGPT](#). [Website]

We Are Teachers. [ChatGPT for Teachers: 20+ Tips, Ideas, Prompts, and More](#). [Blog post]

Appendices

A. More GenAI terminology

Terminology	Explanation
Diffusion models	Diffusion models are a type of generative AI model designed to generate images of excellent realism and quality. Diffusion models can process text, audio, images and video data.
Natural language processing	Natural language processing (NLP) is a discipline that evolved from computational linguistics. It involves building systems that can manipulate human language or data that resembles human language in the way that is written, spoken, and organised.
Self-attention mechanism	Self-attention mechanism is a technique that enables GenAI to dynamically determine the relative importance of various words in a sequence, improving the ability to capture long-range dependencies.
Tokens	Tokens are small units derived from breaking down texts. A token may be a word, a part of a word or characters.
Transformer	A transformer model is a neural network that uses context-specific data to track relationships in sequential data such as words in sentences.
Word embeddings	<p>Word embeddings are dense numerical representations of the individual words in a text taking into account the context and other surrounding words that that individual word occurs with.</p> <p>In GenAI models, words are first transformed into word embeddings before they are used as input to train models.</p>

B. How an LLM works

LLM has numerous applications across various domains, including natural language processing, content generation, conversational agents, language translation, and more. Most LLMs use transformer architecture to process sequential data such as text using self-attention. In sequential data, the interpretation of individual elements of the sequence is dependent on their position relative to other elements and the context of the data sequence. For example, 'bank' could mean a financial institution or the side of a river. In the context of the sentence 'The bank of the river was flooded', 'river' and 'flooded' indicate that 'bank' refers to the side of a river.

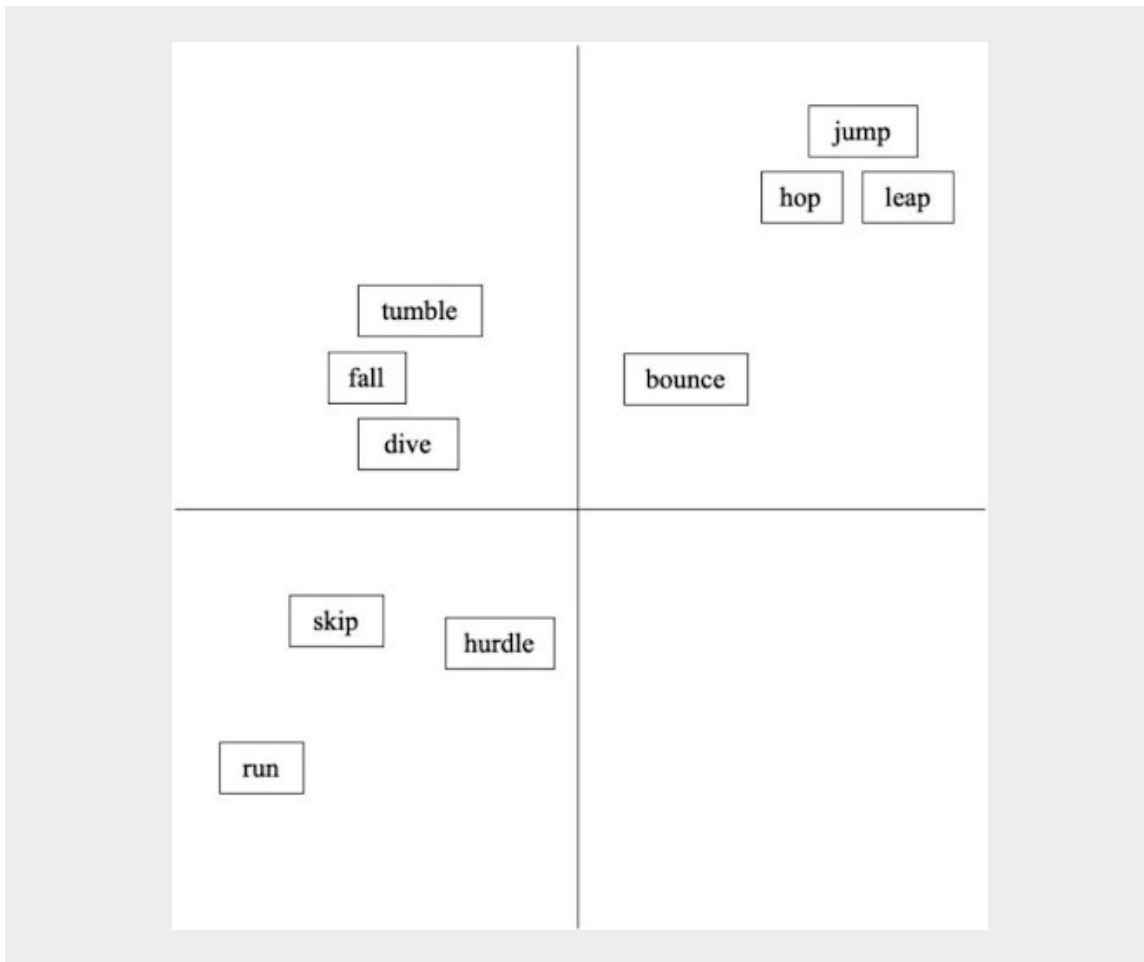
There are similarities and differences in the ways that LLM transformer architecture and diffusion models work and generate output. This explanation focuses on text-based LLMs only.

The training of an LLM can be explained in a simplified way using the idea of word embeddings in text data (see the figure on page 22). Each word is represented as a point in space. Words that mean similar things are placed close together, and words that are very different are farther apart. The 'space' where these words are located can have hundreds or thousands of dimensions, each dimension capturing something about a word's meaning,



usage or relationship to other words. LLMs are not limited to whole words but analyse text by breaking it down into smaller units called tokens. Tokens can represent whole words, parts of words, symbols, or even punctuation marks. This allows LLMs to be able to deal with multiple languages and computer code. The high dimensional vector representation of tokens is fundamental to weighting a GenAI model's parameter. Natural language processing (NLP) is a consequence of LLMs and transformers. NLP allows interaction between computers and humans through natural language, enabling computers to understand, interpret, and generate human languages. This means that when the LLM is prompted in natural language (text, speech-to-text, or image-to-text) it can produce, for example, computer code.

For example, when prompted to write a Python function for the Fibonacci sequence up to a certain number "n", an LLM discerns the task as coding, focusing on "Fibonacci sequence" and "up to n" as key elements. Through self-attention and word embeddings, it 'understands' the Fibonacci sequence as numbers formed by the sum of the two preceding ones, with "n" marking the generation limit. Recognising the need for an input parameter and sequence generation method within the limit, the LLM generally constructs syntactically and logically correct code to fulfil the request, although it has no database of syntax or logic capability.



Simplified illustration of word embeddings in 2-dimensions using proximity to the word 'jump' based on meaning



C. An example of prompt engineering in lesson planning

1 Using ChatGPT4 to create a PRIMM activity (25th April 2024)



You

Create me a PRIMM activity for a lesson introducing for loops in Python



ChatGPT

Creating a PRIMM activity for a lesson introducing "for" loops in Python is a fantastic way to guide students through a structured learning process. PRIMM stands for Predict, Run, Investigate, Modify, and Make, providing a comprehensive approach to learning coding concepts. Here's a detailed PRIMM activity centered around understanding and using "for" loops in Python.

Topic: Introduction to "For" Loops in Python

1. Predict

- **Task:** Show the students a simple Python code snippet using a "for" loop without explaining how it works. Ask them to predict what the output will be and discuss their predictions.
- **Example Code:**

```
python Copy code  
  
for i in range(5):  
    print(i)
```

2. Run

- **Task:** Have students run the code in a Python environment (like an online compiler or local IDE). They can then compare the actual output with their predictions.
- **Expected Output:**

```
Copy code  
  
0  
1  
2  
3  
4
```

3. Investigate

- **Task:** Break down the code together with the students. Discuss the role of `range()` and how the `for` loop iterates over the sequence it generates. Ask guided questions such as:
 - What does `range(5)` mean in this context?
 - How many times does the loop run and why?
 - What would happen if we change `range(5)` to `range(1, 6)`?

4. Modify



- **Task:** Encourage students to modify the original code to see different outcomes. Offer a few suggestions:
 - Print numbers from 1 to 5 instead of 0 to 4.
 - Use the loop to print each number doubled.
 - Create a loop that counts backwards from 10 to 1.
- **Example Modification:**

```
python Copy code
for i in range(1, 6):
    print(i * 2)
```

5. Make

- **Task:** Have the students create their own "for" loop for a new purpose. They can choose one of the following tasks or come up with their own idea:
 - Print all even numbers up to 20.
 - Display a list of square numbers up to 100.
 - Create a list of words and print each word in uppercase.
- **Guidance:** Provide help with syntax and logic as needed. Encourage them to think creatively about how loops can be used to solve repetitive tasks.



Additional Tips:

- Incorporate peer discussions after each step to encourage collaborative learning.
- Use visual aids or live coding to demonstrate changes in loop behavior with different `range()` settings.
- Include debugging sessions to help them understand common errors like off-by-one errors or syntax mistakes.

This structure not only introduces the concept of loops but also encourages active participation, problem-solving, and creativity in learning Python.



These answers will help the teacher explain the mechanics of the "for" loop and the `range()` function effectively, ensuring students have a solid foundational understanding as they proceed to more complex loop usage.





Working group

Teachers and advisors

Philip Arthur, St Peter's School, York

Brett Becker, University College Dublin

Sonya Coakley-Hanan, Three Bridges Primary School

Martyn Colliver, The National Mathematics and Science College

Veronica Cucuiat, Raspberry Pi Foundation

Lyndsay Hope, Monmouth School

Emma Staves, St. Francis' College

Chris West, Lace Hill Academy

Aidan Weston, Long Road Sixth Form College

Eliot Williams, Merchant Taylors' School

Researchers

Salomey Afua Addo, Raspberry Pi Computing Education Research Centre,
University of Cambridge

Sue Sentance, Raspberry Pi Computing Education Research Centre,
University of Cambridge

Shengpeng Shi, Faculty of Education, University of Cambridge

Jane Waite, Raspberry Pi Foundation

Steve Watson, Faculty of Education, University of Cambridge

Bo Yu, Faculty of Education, University of Cambridge