

(Hello World)

THE MAGAZINE FOR COMPUTING & DIGITAL MAKING EDUCATORS

ARTIFICIAL INTELLIGENCE

What does AI mean for computing education?

BIAS IN THE MACHINE

How to address gender balance in AI

NO ROBOTS

How anthropomorphism hinders AI education

Issue 22

Oct 2023

helloworld.cc



PLUS

3D WORLDS IN UNITY • ONLINE COURSES FOR PROFESSIONAL GROWTH • TEACH FIRST • T LEVELS
START A CODE CLUB • SEAME FRAMEWORK • BBC MICRO:BIT • WIPSC 2023 TAKEAWAYS
ASTRO PI • INTRO TO AI LESSON • GREEN SCREEN IN THE CLASSROOM • SYSTEMS THINKING

Teach and learn with the Raspberry Pi Foundation

Free for everyone anywhere in the world



Teachers' resources

Discover training, resources, and guidance to help you teach computing with confidence.

raspberrypi.org/teach



Project library

Browse 200+ online project guides that help young people create with digital technologies.

projects.raspberrypi.org



Research to practice

Explore our research into what works best in teaching school-age young people about computing.

raspberrypi.org/research



Coding at home

Watch our support tutorials and access engaging resources for your child.

raspberrypi.org/learn

Hello World is the official magazine of the Raspberry Pi Foundation



EDITORIAL

Editors

Jan Ander, Gemma Coleman, Ben Garside, Meg Wang

Subeditors

Louise Richmond and Amy Rutter

Subscriptions

Sean Sayers

Social Media

Neena Patel

DESIGN

criticalmedia.co.uk

Head of Design

Lee Allen

Designer

Sam Ribbitts

Photography

The Raspberry Pi Foundation, Adobe Stock

Graphics

Rob Jervis

Cover

© MUTI

CONTRIBUTORS

Edward Abel, James Abela, Sophie Ashford, Pete Bell, Halima Bhayat, Mac Bowley, Mark Calleja, Katharine Childs, Gemma Coleman, Philip Colligan, Michael Conterio, Veronica Cucuiat, Catherine Elliott, Tracy Gardner, Ben Garside, Sway Grantham, Ben Hall, Claire Johnson, Ken Kahn, Gitika Kaur, Fergus Kirkpatrick, Dale Lane, Rebecca Mack, Tessa Mosindi, Rujeko Moyo, Alan O'Donohoe, Richard Pawson, Judy Robertson, Chris Roffey, Sue Sentance, Daljit Shoker, Jane Waite, Aimee Welch, Magda Wood, Seema Zerafa

Supported by **ORACLE**

Contributing Partner



This magazine is printed on paper sourced from sustainable forests and the printer operates an environmental management system which has been assessed as conforming to ISO 14001.

Hello World is published by the Raspberry Pi Foundation, 37 Hills Road, Cambridge, CB2 1NT. The publisher, editor, and contributors accept no responsibility in respect of any omissions or errors relating to skills, products, or services referred to in the magazine. Except where otherwise noted, content in this magazine is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 Unported (CC BY-NC-SA 4.0).



HELLO, WORLD!

Even though it's been less than a year since the release of AI applications such as OpenAI's ChatGPT and Google's Bard, it feels as though they've been around for a long time. Certainly their release has made people pay attention to what AI technology is, and could be, capable of.

When I was in the classroom, I used to enjoy getting to grips with new technological advances and finding ways in which I could bring them into school and excite the students I taught. Occasionally, during the busiest of times, I'd also look longingly at other subjects and be jealous that their curriculum appeared to be more static than ours (probably a huge misconception on my behalf). What I think is really exciting about developments in AI is the impact they are having on education as a whole, and this most definitely isn't just the remit of computing. Philip Colligan talks in his article on page 23 about viewing AI education as three distinct areas:

- AI literacy: what it is and how we teach it
- Rethinking computer science (and possibly some other subjects)
- Enhancing teaching and learning through AI-powered technologies



It's inspiring for me to see how the education community is reacting to the opportunities that AI can provide. Of course, there are elements of AI where we need to tread carefully and be very cautious in our approach, but what you'll see in this magazine is educators who are thinking creatively in this space. Sue Sentance (page 28) highlights how AI applications can reinforce gender bias and stereotypes; Sway Grantham (page 32) offers tips for how to talk to young children about AI; and Ken Kahn (page 54) explores how chatbots could affect the learning of programming.

We hope this issue will provide you with plenty of ideas to take away and build upon.

Ben Garside
Guest Editor



NICE TO MEET YOU



MEG WANG

My name is Meg, and I'm your new Editor. I'm thrilled to take up the task of connecting this wonderful community of educators, volunteers, parents, and researchers through our print and digital magazine. While I've worked at the Raspberry Pi Foundation for a couple of years, prior to joining, I had a long career in book publishing, and I am still learning about computing, digital making, and educational practices. My passion for computing education comes from what young people can do with the knowledge and skills they gain in the classroom. I'm always thankful to Hello World contributors and experts like this issue's guest editor, Ben Garside, who share their insights, ideas, and knowledge in a way that's accessible and practical for people like me who are newer to the subject. If you have an idea for an article or just want to say hi, drop me a line (**contact@helloworld.cc**). I look forward to learning from you and alongside you.

(HW)

CONTENTS

23–57

COVER FEATURES



© Monkey Business/stock.adobe.com

TEACHING AND AI

A series of articles exploring teaching about artificial intelligence

NEWS, FEATURES, AND OPINION

- 6 NEWS**
Research centre; micro:bit — the next gen; Unity projects; European Astro Pi Challenge; meet St Joseph's Secondary School in Rush, Ireland
- 12 INCLUSION COLUMN**
Artificial intelligence and the inclusive computing classroom
- 14 THE SEAME FRAMEWORK**
A useful tool for teaching AI and machine learning lessons
- 16 INSIGHTS — AI ETHICS**
The SUM values
- 17 INSIGHTS — PRODABI**
Human-machine ecologies
- 18 INTERVIEW**
Dale Lane, founder of Machine Learning for Kids, discusses AI tools in education
- 22 SUBSCRIBE TODAY!**
- 23 AI AND COMPUTING EDUCATION**
Breaking down what AI can mean for computing education
- 26 AI AND ANTHROPOMORPHISM**
How human-based descriptions feed our misconceptions about AI
- 28 BIAS IN THE MACHINE**
Ways to address gender bias in AI
- 30 EXPERIENCE AI**
A new education programme making AI accessible and relevant
- 32 NOT ROBOTS**
How to talk to young children about AI
- 34 CHILDREN AND SMART SPEAKERS**
Research exploring children's understanding of AI
- 36 AI FOR 9–13-YEAR-OLDS**
AI lesson ideas for primary and secondary students
- 38 PICTURING AI**
Exploring better images of AI



CODE CLUB

Four groups share their experiences of starting a Code Club


58



AI TOOLS FOR LEARNING

Programming with ChatGPT

54



40

EXPLAINING AI

How we explain AI terms to young people

40 EXPLAINING AI

How we created a glossary of AI terms for young people

42 MODEL CARDS

Teaching AI explainability

44 RECOMMENDER SYSTEMS

The hidden AI steering our choices

46 DEMYSTIFYING MACHINE LEARNING

Activities to unpack AI in recommendation systems

48 TIKTOK'S ALGORITHMS

Can AI and data create more inclusion?

50 AI DEBATE

Do the benefits of AI outweigh the risks?

54 AI TOOLS FOR LEARNING

Programming with ChatGPT

56 KEEP PROGRAMMING

Learners still need hands-on programming practice

58 CODE CLUB

Four groups share their experiences of starting a Code Club

60 COMPUTER SCIENCE AND GEOMETRY

Linking computer science and maths concepts through computational thinking

62 ONLINE COURSES

Tips for making the most of online learning for CPD

64 TEACH FIRST

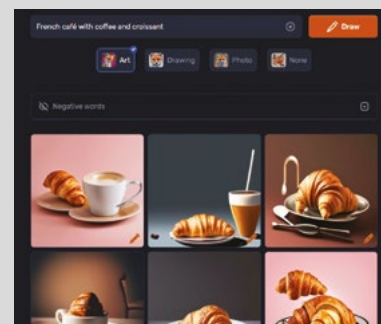
Trainee teachers share how they put computing pedagogies into practice

67 WIPSCÉ 2023

Teacher thoughts from a computing education research conference

LEARNING

RESOURCES & LESSON PLANS



70 INFORMATION TECHNOLOGY

IT all around us

72 INTRO TO AI

A lesson from Experience AI

CONVERSATION

75 BEBRAS

A fun challenge in computational thinking

76 GREEN SCREEN

Developing learners' digital skills with a movie-making competition

78 INSIDER'S GUIDE

Work-related learning experience

82 T LEVELS

Tips from one of the first T-level pilot schools

84 EVOLVING COMPUTING

Teaching systems thinking

REVIEWS

85 BOOK REVIEW

Teach Data Literacy

A NEW INITIATIVE FOR RESEARCH IN COMPUTING EDUCATION

The Raspberry Pi Computing Education Research Centre is working collaboratively with teachers to understand more about the teaching and learning of computer science in schools

Sue Sentance

In issue 19 of Hello World, we announced the launch of the Raspberry Pi Computing Education Research Centre, a joint initiative between the University of Cambridge and the Raspberry Pi Foundation (computingeducationresearch.org). The mission of the centre is to carry out rigorous and impactful research around the teaching and learning of computing, with a particular focus on those who might not have ready access to computing. So, what have we been doing since our official launch in 2021?

PhD students

A key function of the centre is to train PhD students who will work on projects relating to its core work. We are delighted to be supporting two new PhD students, Laurie Gale and Salomey Afua Addo. Laurie is focusing on debugging strategies in secondary-school programming, and Salomey, who is partially supported by the Micro:bit Educational Foundation, is investigating AI education in school, which will include fieldwork in Ghana.

Culturally relevant pedagogy and computing

Culturally relevant pedagogy is a teaching approach that has been developed to address underrepresentation in certain areas of education, and focuses on social equity

and justice. It has been used in computing to develop engaging and responsive computing teaching for all students, regardless of their social or ethnic backgrounds. In 2022, we ran a project in schools to investigate how teachers were implementing culturally responsive computing teaching. A technical report described our work, and in August we also presented a research paper about the first part of our findings (helloworld.cc/CRPresearch).

The Google DeepMind school engagement programme

The university has been contracted to be the independent evaluator of the six programmes that Google DeepMind is running as part of its school engagement initiative to broaden participation in AI and STEM. This project will run until 2024 and we will be reporting our findings to Google DeepMind, as well as publishing specific insights (helloworld.cc/nextgenAllleaders).

Surveying computing teachers

In 2022 we surveyed teachers across the UK and Ireland, and published a paper in which we compared computing education policy and practice in Ireland and the countries of the UK. Since then, we've published a summary of the responses to the survey, and a second paper will soon be published around teachers' experiences of CPD (helloworld.cc/UKICTS). Our plan is to run

the same survey globally in spring 2024, so watch out for updates! This work helps us to compare the experiences of computing teachers around the world, and makes an important contribution to policy, as well as providing general information for teachers.

Conferences and events

We hosted the WiPSCE conference at the end of September (wipsce.org) and are planning to host or run other research events in the future. We also present our work at other conferences: for example, Laurie Gale presented a paper on debugging at the UKICER conference in September (ukicer.com), and I presented a paper on culturally responsive computing teaching (helloworld.cc/CRTpaper) at the ICER conference in August.

Teacher research network

We are keen to collaborate with teachers in our research, and so we are in the process of creating a teacher research network in computing. This will aim to facilitate classroom-based research, as well as keeping teachers in touch with research developments.

Find out more about all these initiatives and other research projects from the Raspberry Pi Foundation at computingeducationresearch.org, where you can also sign up for our newsletter and watch a video describing our work. [\(Hw\)](#)

BBC MICRO:BIT - THE NEXT GEN

Get free micro:bit devices, resources, and training for your primary school in the UK

Magda Wood, Chief of Learning, Micro:bit Educational Foundation

BBC micro:bit – the next gen is a campaign that aims to increase teacher knowledge and confidence through the provision of micro:bit devices, lesson plans, and supporting materials. From this autumn, 700,000 free micro:bit devices (up to 30 per school) will be made available to primary schools throughout the UK, as well as virtual training to support teachers and help them make an impact in the classroom.

Enhancing digital education

Computing and digital literacy are fast becoming foundational skills — not just to reduce the digital skills gap the economy is already struggling with, but also to ensure that young people have the skills to work the increasingly technical jobs of the future. Research suggests that 65 percent of

primary schoolchildren will end up working in jobs that do not exist today ([helloworld.cc/wef2016](https://www.helloworld.cc/wef2016)). How can we best prepare a child for a future career that does not yet exist?

In such a landscape, it is unsurprising that teachers feel overwhelmingly unprepared when it comes to teaching digital skills. Research carried out by the Micro:bit Educational Foundation and Nominet last year indicated that 61 percent of UK primary teachers responsible for teaching computing have no background in the subject, with three in five teachers also citing a lack of resources as a major barrier.

The BBC micro:bit – the next gen project aims to equip and support schools with the right tools and resources to foster computational thinking, inspire digital creativity, and enhance the curriculum with innovative machine learning lessons.

Real-world applications

School children have already used the micro:bit to tackle real-world problems such as developing flood early-warning systems to notify coastal areas of bad weather and developing radio messaging systems to boost communication in the classroom.

New resource packs will launch as part of BBC micro:bit – the next gen. These include lessons to help teachers get started with physical computing, and engaging projects to introduce working with the micro:bit in science investigations, data logging, and working with digital data and machine learning.

The next mission of the BBC micro:bit – the next gen project is to assist all students and teachers in the UK, building on the

Micro:bit Educational Foundation's three-phase national programme and primary pilot project (microbit.org/impact/research). This was carried out in collaboration with Nominet with the aim of better understanding the difficulties faced by primary teachers in the UK, and to enhance the teaching of digital skills in primary classrooms. We look forward to seeing how the programme can empower teachers and inspire millions of young people to live their best digital lives.

If you are a primary teacher from anywhere in the UK, you can register at bbc.co.uk/teach/microbit. Registration closes on 18 December 2023. Registered teachers should expect their packs by March 2024 at the latest. ^(HW)

ABOUT THE CAMPAIGN

BBC micro:bit – the next gen will be made possible by the following partners:

- The free micro:bits will be funded by the UK registry Nominet
- Teacher training and education resources will be provided for free by the Micro:bit Educational Foundation
- The capabilities of the micro:bit will be showcased with the help of favourite CBBC and family entertainment brands from BBC Education, to demonstrate how simple and accessible it is for teachers to use in the classroom



■ Two primary schoolchildren using the micro:bit device

UNLEASHING CREATIVITY: 3D INTERACTIVE WORLDS IN UNITY

The Raspberry Pi Foundation's Unity projects can spark a passion for 3D-game design and programming through new challenges

Pete Bell

3 D interactive experiences have become increasingly prevalent in our lives. From captivating video games, to stunning animations, to immersive virtual reality experiences, the world is evolving into a realm where imagination knows no bounds. Highlighting the ubiquity and appeal of these experiences can help educators to capture the attention and interest of young creators.

Teachers will have seen that many young creators are inspired to begin their programming journey because they play games and want to develop something similar. To help young people find a place to start, the Raspberry Pi Foundation has partnered with Unity to develop two project

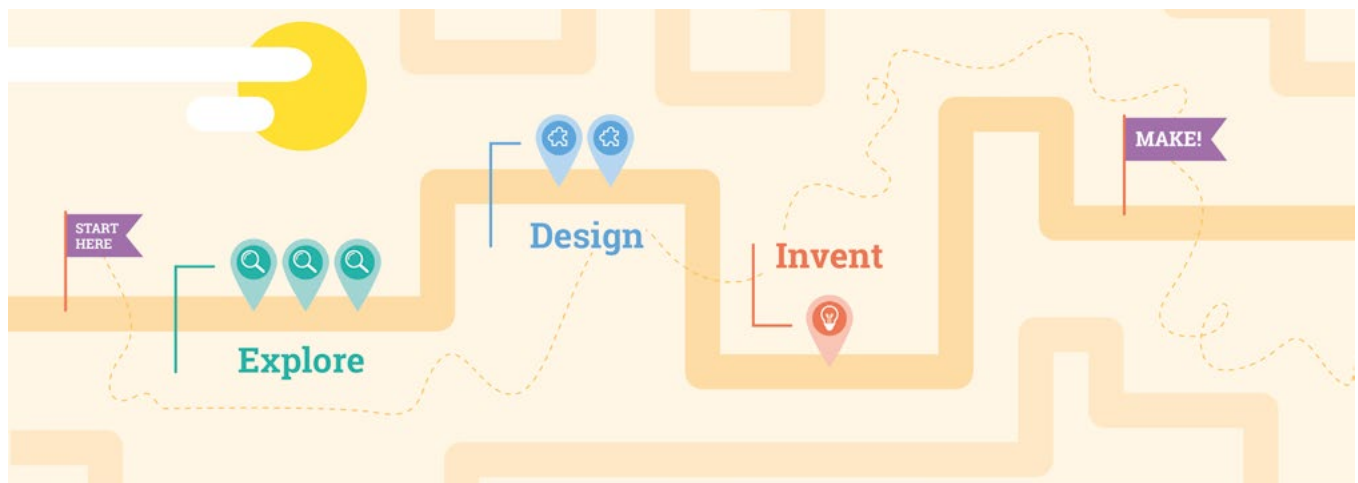
pathways in which creators explore and create 3D worlds. Unity is a cross-platform, real-time, 3D-development platform that has been used to create some super-popular games such as Subway Surfers and Among Us. This can be really motivating for young creators embarking on these new project paths. The ability to deploy their creations across mobile, desktop, and web offers them the opportunity to showcase their creations to a broader audience and widens their scope as developers.

Introduction to Unity

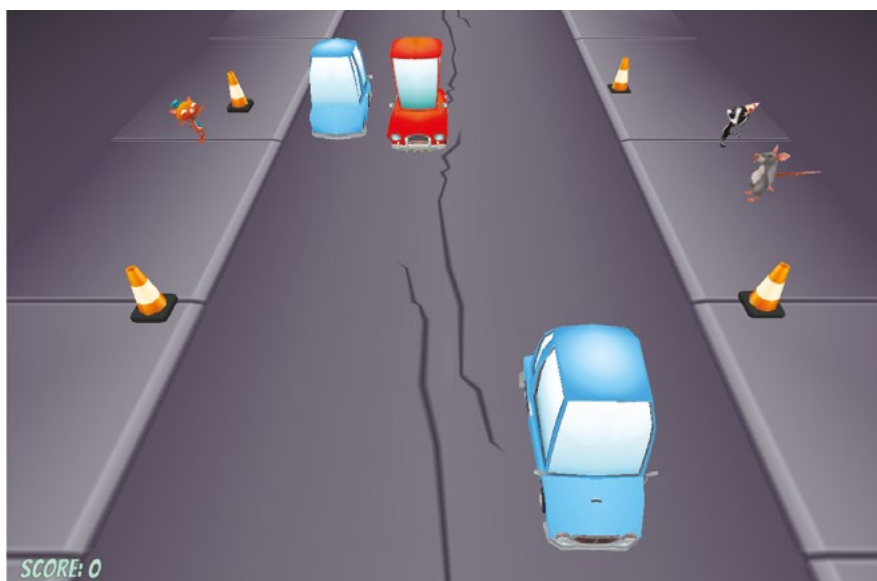
The two project pathways are both structured around our '3 ... 2 ... 1 ... Make!' approach to project-based learning, in

which creators explore, design, and finally invent their own projects. The first pathway, 'Introduction to Unity', includes step-by-step guides to using 3D objects, character controllers, colliders, text, and buttons to create 3D worlds with characters. Starting off at the Explore stage, learners make a mini game collecting sparkling stars, and design non-playable characters (NPCs) for interaction. These projects build a solid foundation of skills and confidence, and familiarise learners with Unity's interface, scripting, and asset management.

Creators are then ready to move on to the Design stage, which allow learners to take the reins and build captivating worlds with unique features that incorporate NPCs and



■ The 3 ... 2 ... 1 ... Make! approach to projects is informed by research into how young people develop interest and skills in computing



■ Crossy Rat, an example game created from one of the pathways

quests. This stage promotes critical thinking and problem-solving as creators tackle design challenges and implement their ideas.

Finally, learners reach the Invent stage, and are encouraged to create something entirely unique. With the support of the skills developed throughout the pathway, they invent their own 3D game, virtual world, or visual novel, complete with goals and/or multiple endings.

Example projects are available to play online, and learners can download the packages to develop them into their own creations (for example, helloworld.cc/crossyrat and helloworld.cc/marblemayhem). We also encourage learners to submit their creations for us to showcase to the community on the Projects site. It was fantastic to even see some creative Unity submissions featured in the 2023 Coolest Projects Live showcase (helloworld.cc/CP2023), such as a spin on the star collecting project with added the-floor-is-made-of-lava jeopardy!

More Unity

The second pathway, 'More Unity', allows learners to manipulate in-world physics and materials to add depth to their creations. We assume some knowledge from the

'Introduction to Unity' pathway, but we have been careful to include reminders of the key skills required throughout the projects. Starting off at the Explore stage, creators are introduced to the pathway theme of marble puzzles. Learners use premade assets to create a 3D-simulation adventure along colourful tracks, with obstacles for their marbles to navigate and a goal that sets off fireworks when it is achieved.

Learners are encouraged to customise their marble using the open-source Blender program (a 3D-creation package for designers and engineers), which enables Unity developers to enhance the visual richness and uniqueness of their projects (helloworld.cc/blendermarble).

Progressing through the Explore project stage, learners create a disco dance floor. This is a truly dynamic learning experience that uses tiles to show off colour-changing materials and trigger audio samples. The 'Don't Fall Through' puzzle game then includes learners navigating a safe path

through a collapsing tile floor ... with added fireworks if you get through it, obviously!

Creators next move onto the Design stage of the project, where they use their newly developed skills to bring their inner artist to life. They develop unique pixel art on a tiled floor and reveal their awesome artwork by rolling a ball across the surface. Soon they're stepping into the shoes of a game developer by designing custom tracks filled with obstacles, personalised effects, sounds, and more!

The pathway culminates with the Invent stage, aptly named 'Marble Mayhem'. Learners can bend the laws of physics to create a chaotic one-of-a-kind game.

Powerful programming, individualised interactivity

Young creators who have shown interest in Scratch animations or Python and want to take the next step can try all of these projects out in C#. This is a widely used, versatile, object-oriented programming language that provides students with transferable skills that can be applied beyond game development. This expands learners' programming repertoire, and linking the code to Unity projects provides a fun context in which to make the leap to this new language.

“ THESE PROJECTS BUILD A SOLID FOUNDATION OF SKILLS AND CONFIDENCE

All the code is provided in the projects, along with full explanations, and there are loads of examples of how the properties of a range of game objects can be manipulated. Learners can then adapt these scripts with ease, providing an almost limitless way for them to express their own creativity and produce something truly individual. The project pathways are both available now at helloworld.cc/introunity and helloworld.cc/moreunity. Get exploring, designing, and inventing with your young creators today and who knows, maybe they'll create the next Among Us! helloworld.cc

ESA astronaut Andreas Mogensen will help run Astro Pi code onboard the ISS

THE EUROPEAN ASTRO PI CHALLENGE 2023/24

The European Astro Pi Challenge is back with exciting missions for 2023/24, inviting young minds to venture into space through coding

Fergus Kirkpatrick

What is Astro Pi?

The European Astro Pi Challenge (astro-pi.org) offers young people the opportunity to conduct scientific investigations in space by writing programs that run on Raspberry Pi computers on the International Space Station (ISS). We offer two missions for different abilities: Mission Zero and Mission Space Lab.

Mission Zero: suitable for beginners and primary schoolchildren

Mission Zero allows young people to create text-based programs and pixel art. This year's theme is 'flora and fauna', and successful programs will be sent to the Astro Pi computers onboard the ISS. Mission Zero is accompanied by a project guide (helloworld.cc/MissionZero2023)

MISSION ZERO

25 March 2024: Deadline for submissions
April 2024: Mission 'flight status' updates
June 2024: Certificates sent to teams

and can be delivered in a one-hour session. It offers a great introduction to text-based programming with Python.

Mission Space Lab: a new format for 2023/24

Mission Space Lab is changing in 2023/24 to focus more on developing programming and problem-solving skills. Working in teams, young people are invited to come up with ways to calculate how fast the ISS is travelling as it orbits Earth using the different sensors and camera on the Astro Pis. Teams need to write a Python program to capture data and give an output for their calculation.

To help with this challenge, Mission Control has created a new Python library that will show data and photos from previous ISS missions. You can download and run it on your computer to test your program without needing an Astro Pi kit. We'll also provide a new project guide to walk you through some of the steps required for the challenge.

If you are successful, your team's code may run in space onboard the ISS. This

MISSION SPACE LAB

6 November 2023: Registration opens
19 February 2024: Deadline for submissions
April 2024: Mission 'flight status' updates
June 2024: Certificates sent to teams

gives your team an incredible chance to capture unique Earth observation photos and data with your team's program. All teams whose programs run on the ISS will also receive a personalised certificate and the data collected in space. All successful teams will also be invited to join a Q&A session with an astronaut from the European Space Agency (ESA).

How will your team approach this challenge? How could they use the data from the Astro Pi?

The European Astro Pi Challenge 2023/24 offers young coders an out-of-this-world experience. Register, code, and who knows, maybe your team's program will end up among the stars. helloworld.cc

CELEBRATING THE COMMUNITY: ST JOSEPH'S SECONDARY SCHOOL, RUSH

We caught up with the team at St Joseph's, fresh off the success of another round of Astro Pi, to delve deeper into their passion for exploring new ways to get learners excited about computing technologies

Sophie Ashford

In the heart of rural Rush, Ireland, St Joseph's Secondary School has been on an inspiring journey, thanks to a dedicated group of educators and their enthusiastic students.

Nestled amidst the picturesque 'Market Garden of Ireland', the school caters to a diverse student population, ranging from children of farmers to newcomers drawn to the scenic countryside.

The St Joseph's leadership team recognised the need to adapt and evolve the curriculum in response to the change from an area steeped in agricultural heritage to a region of evolving demographics.

The school's proactive approach to teaching coding from First Year (ages 12–13) has resonated with many students, including Kamaya, a member of the 2022/23 Astro Pi cohort, who discovered her passion for computer science through the movie *Interstellar*.



■ A group of St Joseph's students who took part in the 2022/23 Astro Pi mission

"I remember the first time I was like, OK, space is cool, was when I watched a movie. It was called *Interstellar*. I [realised] I might want to do something like that in my future. So, when I came to [St Joseph's], I saw coding as a subject and I was like, 'Mum, I've got to do coding.'"

A key driver in encouraging the students to give coding a try has been Mr Murray, or Danny, as he is fondly referred to by students and staff alike. Danny was introduced to the importance of engaging students with computing technologies while teaching science at a school in England. He attended a Code Club, saw children building projects with Raspberry Pis, and couldn't wait to get involved. Growing his knowledge from there, Danny switched up his career when he returned home to Ireland. He took on the challenge of helping St Joseph's expand its computer science offering, and leads the charge on all IT-related issues.

After the school introduced mandatory coding taster sessions to all First Year students, Danny was blown away by the eagerness of the learners and wanted to provide further opportunities for them to see what they could achieve.

Enter Astro Pi. One of Danny's acquaintances told him about the challenge, and he shared it with his computer science students. In fact, he invited all St Joseph's students to join in. The uptake was vast,

especially once he dropped in that they might just be the recipients of some very exciting photos.

"You get to see photos of Earth that nobody has ever seen. This is a unique view just for you ... imagine just talking to somebody and saying, oh, there's a picture of the Amazon. I took that picture when I was 14. From space," Danny adds.

Danny's mission is to instil in his students the belief that anything is achievable. Collaborating on Astro Pi projects has enabled teams to uncover their strengths and foster a strong sense of community.

This unity has transcended the classroom, creating a vibrant makers' culture at St Joseph's. Today, a dedicated team of students takes charge of solving tech-related challenges within the school. Deputy Principal Darren Byrne says, "Our own students actually go class to class repairing tech issues. So every day, there are four or five students going around checking PCs in classrooms. They [...] give classes to our First Year students on app usage.

"It's instilled in the whole school now, the idea that the students can look after this kind of technology themselves. We're the ones reaching out for help from the students!"

You can read more about St Joseph's Secondary School and their story on the Raspberry Pi Foundation website (RPF.io/stjoseph). [\(Hw\)](#)

ARTIFICIAL INTELLIGENCE AND THE INCLUSIVE COMPUTING CLASSROOM

How can you use generative artificial intelligence to support inclusive practice in computing lessons?

Artificial intelligence (AI) is one of the hottest topics in education at the moment, and it has a large number of possible benefits for developing an inclusive classroom. In this article I am going to be looking at the application of generative AI, in particular large language models such as ChatGPT and Google's Bard. As with any new technology, teachers should ensure they are aware of the risks as well as the benefits. For example, the outputs of generative AI aren't always accurate; data can be biased, depending on the quality of the data it was trained on; and our young people need to be taught how to recognise the validity of information and to question whether something may have been generated by AI. Many tools also have a lower user age limit of 13, so make sure you check this before letting children use the tools independently.

Despite the potential drawbacks, there are, however, a number of ways we can use generative AI to support learning, individualise content, and provide creative hooks into topics in computing to increase engagement. Here are a few ideas to get you started.

1. Create step-by-step instructions

Consider a lesson in which students are using a piece of software for the first time. You can ask an AI chatbot to generate detailed step-by-step

instructions on how to do a certain task, and even specify the reading age of the person they are for:

***Sample prompt:** Write me simple step-by-step instructions, understandable by a ten-year-old, on how to use a formula to work out the average of a column of data in Excel.*

Students can then use these to complete an activity, and can even ask the chatbot for help if stuck. Obviously, this can still present barriers if students struggle with

reading, but you can paste the text into something like Immersive Reader, or the NaturalReader web app can be used directly in ChatGPT to read the responses aloud (helloworld.cc/NaturalReader).

2. Support executive function in independent tasks

Executive function encompasses the skills that let us set goals, plan, prioritise tasks, and get things done. Research indicates that up to 80% of autistic people suffer from executive function disorder, and other neurodiverse learners may also struggle in this area. One approach to support students with organising their work is to use a 'body double'; simply having another person present when completing tasks can help them focus better, and an AI chatbot can be used



Catherine Elliott leads the Sheffield eLearning Service (sheffieldclc.net) and has spent a number of years working on ways to make computing accessible for all learners. She is co-chair of the CAS Include working group, and leader of the Sheffield and South Yorkshire Secondary and Virtual SEND CAS Communities (@[catherinelliott](https://twitter.com/catherinelliott)).

to replicate this. There is a great thread about this on X (formerly Twitter) from an app developer with executive function disorder (helloworld.cc/EFD).

You could set up a prompt like the one below to keep a learner on track; or older students could use it to break down homework or coursework tasks into manageable steps and keep them focused (see **Figure 1**):

Sample prompt: *You are a friendly, supportive teaching assistant. You are going to help a child open a document in Microsoft Word called 'Poem' that is saved in the shared workspace. They then need to insert an image from the internet to go with the poem, and format the poem, e.g. make the title bold. You need to break the task down into very small steps. Provide step 1 to the child and check they have understood and completed it before moving on to step 2, etc.*

Again if a learner hasn't understood a step, they can ask for clarification or more help.

“ I HAVE BEEN USING AI AS A USEFUL SCAFFOLD TO MOTIVATE AND SUPPORT OUR SEN STUDENTS

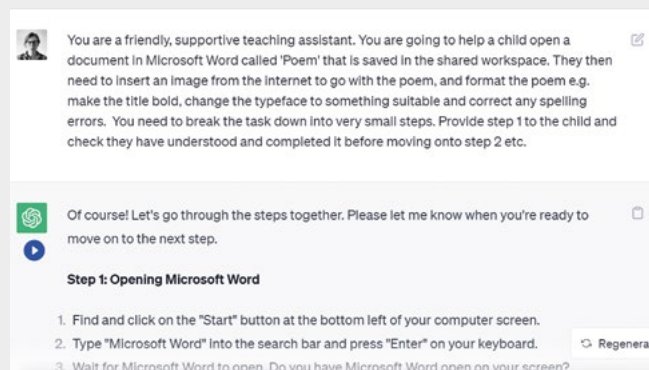
3. Simplify complex ideas and content

This is something you as a teacher can do, to check the reading age of your instructions or worksheets; or a student could ask an AI chatbot to explain in simpler terms a concept they are struggling with.

Sample prompt: *Rewrite the following text using language an eight-year-old can understand ...*

You can even ask AI to interrogate a student's knowledge on a topic, and encourage a conversation around examples of a concept in practice:

Sample prompt: *You are a friendly, supportive teaching assistant. Find out what the student knows about input and output devices in computing.*



■ **Figure 1** ChatGPT can be used as a tool to support executive function in independent tasks

In Google's Bard you can input your prompt using the microphone, which might be easier for some learners. As with all interactions with generative AI, make sure you check the accuracy of the outputs!

4. Generate creative prompts

For many learners with SEND, starting with a blank page is incredibly difficult, both in terms of executive function and self-confidence. Richard Brown, a teacher at Heritage Park School in Sheffield, UK, for learners with social, emotional, and mental health needs (SEMH) has been using ChatGPT to generate prompts for creative writing to engage learners in grammar work. He says, "I have been using AI as a useful scaffold to motivate and support our SEN students in their learning; I've seen a dramatic impact on their engagement as a result."

You could use generative AI in a similar way to create starting points for a computing project, for example, writing a script for a video project that can then be refined and developed by learners:

Sample prompt: *Please write a short script of no more than 400 words for a film about a group of eleven-year-olds who find an alien creature in their classroom store cupboard.*

As before, this is just a starting point, and learners would be expected to adapt and improve it.

Generative AI isn't going to solve all of the issues faced by our learners with SEND in the computing classroom. However, used effectively — to provide personalised learning experiences, give feedback and support in completing tasks, and as a prompt for creative projects — it can help us create a more inclusive, equitable, and inspiring computing classroom. helloworld.cc

#INSIGHTS

THE SEAME FRAMEWORK: A SIMPLE WAY TO THINK ABOUT AI EDUCATION RESOURCES

STORY BY Jane Waite

There are many different reasons we all need to learn about artificial intelligence (AI) technologies. They can affect our daily lives, and understanding these technologies can open up opportunities for future careers. So there is broad agreement that education about AI is essential.

At the Raspberry Pi Computing Education Research Centre (helloworld.cc/rpiresearchcentre), we have started investigating the teaching and learning of AI. In 2021 and 2022, we hosted a series of research seminars on AI and data science education (helloworld.cc/

[RPFseminarpapers](https://helloworld.cc/rpfseminarpapers)). Listening to our fellow researchers, it became clear that resources to teach AI were being developed, but what should be included in these materials was not yet agreed upon.

Therefore, we embarked on a research study to review the resources available for teachers to select from. Our review of over 300 resources found that the vast majority of learning materials were one-off activities, and many claimed to be appropriate for learners of any age. There were very few sets of lessons or units of work that were tailored to a specific age group. Activities often had vague learning objectives, or

none at all. We rarely found associated assessment activities.

To analyse the content of the resources, we used a simple framework called SEAME. This framework is based on work conducted in 2018 by myself, Jane Waite, with Professor Paul Curzon at Queen Mary University of London while running professional development for educators on teaching machine learning (ML) (helloworld.cc/teachingML).

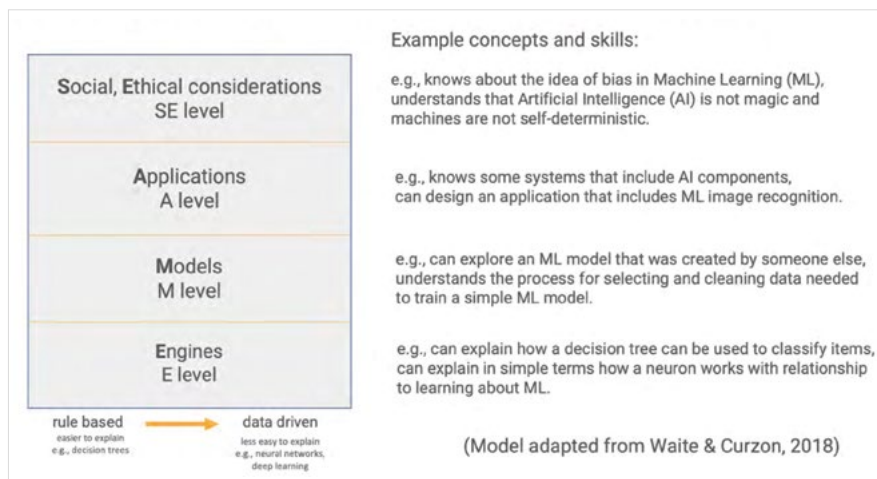
The SEAME framework

The SEAME framework gives you a simple way to view learning objectives and resources related to teaching AI and ML, based on whether they focus on social and ethical aspects (SE), applications (A), models (M), or engines (E, i.e. how AI works). We hope it will be a useful tool for other researchers and educators investigating, developing, or delivering AI lessons.

The four levels of the SEAME framework do not indicate a hierarchy or sequence. Instead, they offer a way to talk about the focus of AI learning activities.

Social and ethical aspects (SE)

The SE level covers activities that relate to the impact of AI on everyday life and its implications for society. Learning objectives and their related resources categorised at



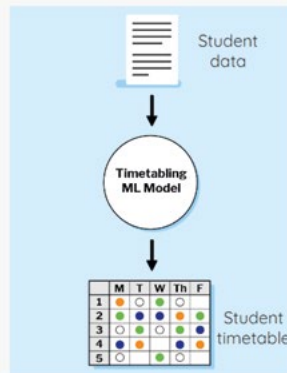
■ The SEAME framework

The subjects you study

Your school is thinking of investing in timetabling software that uses an AI application.

The AI application is used to predict which subjects you will study.

This will save the school money as they will be able to better plan for the future.



■ An SE activity from Experience AI (experience-ai.org) where students explore the social and ethical implications of an AI application that predicts what subjects they might want to study

this level introduce students to issues, such as privacy or bias concerns, the impact of AI on employment, misinformation, and the potential benefits of AI applications.

Applications (A)

The A level refers to activities related to developing applications and systems that use AI or ML models. At this level, students learn how to train models themselves, and look at how such models work. Learning objectives at the A level include knowing a range of AI applications and starting to understand the difference between rule-based and data-driven approaches to developing applications.

Models (M)

The M level relates to the models underlying AI and ML applications. Learning objectives include students understanding the processes used to train and test models. For example, through resources focused on the M level, students could learn about the different learning paradigms of ML (i.e. supervised, unsupervised, or reinforcement learning).

Engines (E)

The E level considers the engines that make AI models work. This is the most hidden and complex level, and school-aged students may need to be taught using unplugged activities and visualisations. Learning objectives could include understanding the basic workings of systems, such as data-driven decision trees and artificial neural networks.

Covering the four levels

Some learning activities will focus on a single level, but activities can also span more than one level. For example, an activity may start with students trying out an existing rock, paper, scissors application that uses an ML model to recognise hand shapes. This would cover the application (A) level. If students then move on to train the model to improve its accuracy by adding more image data, they work at the model (M) level.

Other activities cover several SEAME levels to address a specific concept. For example, an activity focused on bias might start with an example of the societal impact of bias (SE level). Students could then discuss the AI applications they use and reflect on how bias affects them personally (A level). The activity could finish with students exploring related data in a simple ML model and thinking about how

representative the data is of all potential application users (M level).

The SEAME framework as a tool for research on AI education

For researchers, we think the SEAME framework will be useful to analyse curriculum material, for example, to see whether some age groups have more learning activities available at one level than another and whether this changes over time. We may find that primary-school learners work mostly at the SE and A levels, and secondary-school learners move between the levels with increasing clarity as they develop their knowledge. It may also be the case that some students or teachers prefer activities focused on one level rather than another. However, we can't be sure: research is needed to investigate the teaching and learning of AI and ML across all year groups.

What next?

We are continuing our research on what should be taught about AI to young people. Read more about our findings (helloworld.cc/AIpapers); our academic paper that describes the SEAME framework and our categorisation of resources will also be added soon. You can also find out about the work of PhD student Salomey Afua Addo (helloworld.cc/Salomey), who has joined the Raspberry Pi Computing Education Research Centre and is investigating approaches to the teaching and learning of AI. [\(HW\)](#)

Classification – your turn

DeepMind want you to help them with their project of tracking animals in the Serengeti.

The Serengeti is a national park in East Africa. There are many endangered animals there.

They are using classification to find and track animals in order to protect them.



■ An M-level activity from Experience AI where students learn about classification; although the activity is mostly focused on the M level, it also links to the social and ethical (SE) and application (A) levels

ENGAGING CHILDREN WITH AI ETHICS

STORY BY Katharine Childs

In the context of growing interest in introducing children and young people to artificial intelligence (AI) systems, researchers Mhairi Aitken and Morgan Briggs from The Alan Turing Institute (turing.ac.uk) argue it's crucial to engage children in ethical discussions about AI. Aitken and Briggs believe that children have important rights to consider when thinking about how AI technology is designed, developed, and used.

Introducing AI ethics

The primary goal of AI ethics is to evaluate the ethical and social implications of artificial intelligence technologies. AI ethics is guided by principles rather than rigid rules. These principles are encapsulated in the Institute's SUM values, which have been introduced to support responsible AI development. As AI ethics evolve, it is crucial to involve voices that have been underrepresented, such as those of young people. Aitken and Briggs argue that children's input is invaluable in helping shape ethical AI practices, particularly because AI can potentially have unequal impacts on different segments of society.

THE SUM VALUES

The Alan Turing Institute has identified four values that support, underwrite, and motivate (SUM) responsible innovation with new technologies such as AI. The values are:

- **Respect** the dignity of individuals
- **Connect** with each other sincerely, openly, and inclusively
- **Care** for the well-being of each and all
- **Protect** the priorities of social values, justice, and the public interest



■ Children already use AI in their daily lives, in technology such as mobile phones

Child-centred AI

'Child-centred AI' emphasises involving children in the entire AI life cycle. This approach prioritises informed choices, children's involvement in AI discussions, ethical awareness among developers, and creating an ethical AI environment.

UNICEF collaborated with the Government of Finland in 2020 to create a draft policy guide, *Policy Guidance on AI for Children* (helloworld.cc/UNICEF), which outlines nine requirements for child-centred AI: emphasising child well-being, inclusion, fairness, data privacy, safety, transparency, governmental knowledge, child preparation for future developments in AI, and creating an enabling environment.

The Alan Turing Institute tested UNICEF's draft policy guidance and interviewed 14 UK public-sector organisations (helloworld.cc/Turingcasestudy). The findings revealed challenges in implementing child-centred AI, such as low data literacy among the public, a lack of understanding of GDPR issues, and concerns about AI guidance. Organisations expressed a desire for actionable recommendations, sector-specific guidance, and content accessible to all ages. While organisations were

FURTHER READING

- ✓ Aitken, M. & Briggs, M. (2022). *Engaging children with AI ethics. In AI, data science, and young people. Understanding computing education (Vol 3). Proceedings of the Raspberry Pi Foundation Research Seminars.* (helloworld.cc/AitkenBriggs).

enthusiastic about involving children in AI discussions, they lacked effective strategies. Because children's active participation in AI decision-making is crucial, Aitken and Briggs call for innovative approaches to engage children in AI ethics discussions.

Exploring AI ethics in the classroom

Drawing on Aitken and Briggs' ideas, engaging children in classroom discussions about AI can also introduce ethical considerations. Students can explore topics such as:

- Bias in AI algorithms (helloworld.cc/MLbias)
- Safety concerns (helloworld.cc/AIsafety)
- Responsible use of AI (helloworld.cc/ResponsibleAI)

By addressing these issues, teachers help students become responsible and ethical users of AI, contributing to a more equitable and just society. **(HW)**

PRODABI: THE PROJECT ON DATA SCIENCE AND BIG DATA

STORY BY Claire Johnson

As the march of AI gathers pace, educationalists are carefully considering what knowledge secondary students (aged 11–18) need to develop and how teachers can effectively deliver it.

To this end, since 2018, academics at Paderborn University in Germany have been developing ProDaBi, a research-based curriculum for secondary-school level in data science and big data (including AI).

Carsten Schulte and his team suggest that AI and machine learning (ML) differ from other topics in the computer science (CS) curriculum and might benefit from a data-centric approach to complement the traditional focus on programming: “These topics require a paradigm shift for some teachers, and ... this shift has to do with the changed role of algorithms, data, and the societal context.”

Their position is partly informed by a 2019 paper, *Machine Behaviour*

(helloworld.cc/Rahwan). The paper called for a new discipline, ‘artificial science’, to study intelligent machines as “a class of actors with particular behavioural patterns and ecology, not just engineering artefacts”. This is the scientific study of machine and human–machine ecologies.

These ‘hybrid interaction systems’ (helloworld.cc/Schulte2018) in turn call for a focus on the interaction between human and machine in an effective CS curriculum, but also for a greater focus on the central role of data. This approach forms the basis of the teaching material developed over the two phases of the ProDaBi project.

Teaching materials

The research initially generated an optional, year-long course for 17- to 18-year-olds, in which they learn about data exploration and ML. Successive evaluations led to teaching materials for 15- to 16-year-olds, which build data awareness over four units as follows:

1. ‘Data science and artificial intelligence (without programming)’: students analyse survey data on young people’s media use and learn about decision trees as data-based models, using Common Online Data Analysis Platform (CODAP) software.
2. ‘Data science and artificial intelligence (with programming)’: students learn how decision trees and artificial neural networks are used in image recognition, using Jupyter Notebook with Python.
3. ‘Data awareness’: students consider the role of data used by mobile phone networks and movie recommendation systems.
4. ‘Data projects and data exploration’: students collect and analyse environmental data with sensors, using Arduino and openSenseMap.

An additional, unplugged activity, ‘Human

FURTHER READING

- ✓ Rahwan, I., et al. (2019). *Machine Behaviour*. *Nature*, 568(7753), 477–486. helloworld.cc/Rahwan
- ✓ Schulte, C., and Budde, L.: A Framework for Computing Education: Hybrid Interaction System: the need for a bigger picture in computing education. In *Proceedings of the 18th Koli Calling International Conference on Computing Education Research (2018)*. helloworld.cc/Schulte2018

vs Machine!’, introduces students to ML by illustrating how a simulated artificial agent gets better at playing the fictional game Hexapawn (helloworld.cc/humanvmachine).

Role play

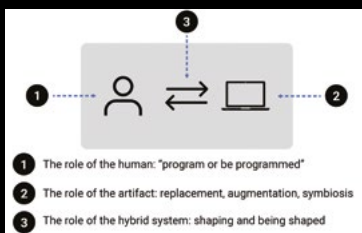
The ProDaBi project urges CS teachers to pay more attention to AI technologies and focus on the interaction between humans and digital artefacts and the transformational effects of these hybrid systems. As we are both shaped by and shape these data-driven technologies, an effective CS curriculum should firmly anchor these considerations to its AI and ML offerings. The ProDaBi project teaching materials provide valuable support here, although not all are available in English.

Research to practice

Teachers in different cultural and curriculum contexts may need to adapt the extended activities to suit their own timescales, or use the unit overviews as a source of inspiration for their own lesson planning. However you use the materials, we would love to hear about it. Email us at contact@helloworld.cc. (HW)

HYBRID INTERACTION SYSTEM

“AI education requires developing an adequate picture of the hybrid interaction system — a kind of data-driven, emergent ecosystem which needs to be made explicit to understand the transformative role as well as the technological basics of these artificial intelligence tools and how they are related to data science” (Schulte et al., 2021. helloworld.cc/hybrid).



Hybrid interaction system, adapted from Schulte



MACHINE LEARNING FOR KIDS

“I HOPE THAT PEOPLE DON’T LOOK AT AI JUST AS SOMETHING THAT MIGHT REPLACE THEM ...”

© CC-BY 4.0

Ben Garside speaks with Dale Lane, chief architect at IBM and founder of Machine Learning for Kids, to discuss developments in AI and what they could mean for education

“ I still think of myself as a software developer, but in practice I don’t get to write as much code as I used to,” Dale explains when talking about his current job. Having worked at IBM for the last 20 years, he has seen AI develop from an emerging technology to a normal part of all software projects at IBM. “I’ve always liked working in areas of tech that are new and cutting-edge,” he says. Back in the early 2010s, he started

working on IBM Watson products, IBM’s suite of AI tools, and he was involved in projects that explored whether computers could be creative. “I’ve done the weird and wacky stuff, and I’ve done the very cool, pretty stuff,” Dale summarises. One project he worked on was called ‘Chef Watson’, with the goal of building a tool that could learn from recipes. He describes the idea thus: “If we take [data about what food you buy], can we invent a new meal that you might like with that combination of food?”

In the not so distant past, a distinction would have been drawn between a regular software developer and someone who works in AI, but according to Dale, that distinction has now disappeared. “Nowadays, I think there’s virtually no software product IBM makes that doesn’t have AI [technology] in it somewhere ... Every software team has a plan of how they’re applying AI [systems] to improve their product; every developer, no matter what part of the business they work in, has to understand how to apply AI.”



■ Dale Lane speaking at TEDx Winchester, UK

AI tools in education

As new AI systems developed, Dale saw the need to help educate other colleagues about how they worked. “We were trying to enable [non-AI] developers to use AI tech, so I built a whole load of stuff that tried to build machine learning into a tool so that any developer could use it.” He fondly describes how he took the same approach when he was trying to explain to his children what he did for a living. “If I can get a non-AI person to understand AI tech, could I do that with a nine-year-old?” So in his spare time, Dale began working on a tool that would help his children understand, as his explanations had never done before. From this, the Machine Learning for Kids website was born.

“Machine Learning for Kids is trying to get people thinking about how AI [tech] is all around them, the impact it has on their lives, what the tech can do, and how it’s used and behaves ... through hands-on experimenting.” The Machine Learning for Kids tool integrates with Scratch

“ IF I CAN GET A NON-AI PERSON TO UNDERSTAND AI TECH, COULD I DO THAT WITH A NINE-YEAR-OLD?”

and allows users to build their own applications based on machine learning models they have trained themselves. Dale’s hope is that young people play around with the tool and as a by-product, “learn a bit about how these projects are made and what they can be used for”.

Dale tells me that he sees a lot more traffic to the Machine Learning for Kids website now than when he launched the tool in 2017. This may not be a surprise — AI seems to be in many more people’s consciousness than it was even at this time last year. Stories in the news and the release of AI chatbots based on large language models (LLMs), such as ChatGPT and Bard, have caught the attention of the world.

I ask Dale what impact he thinks LLMs will have on the educational landscape. “I hope a positive one. I’m always an optimist when it comes to tech. I’m always on the side of the fence that it’s going to be fine. I do see a huge amount of value in [LLMs]. I get that there is a lot of fear around plagiarism, and that’s not unique to education.” Dale suggests that AI chatbots can be considered like a “colleague that you ask five questions of, and twice they’ll tell you stuff you already knew, twice they’ll give you a new idea that gets you thinking, and once they’ll make something up that’s just complete fantasy. But they’ll do it with such confidence that you might not notice. It’s still useful to ask that person, but [you need to] be able to recognise the one in five times when it’s fantasy.”

This emphasises the importance of educators continuing to focus on something we’re already teaching across various subjects: the skills young people need in order to be able to look critically at secondary sources.

Dale describes a school social media post he’s seen in which students were challenged to find the mistakes in an essay generated by a LLM. “[The teachers] were partly trying to show the kids how it worked and what it can do, but made sure that there were some real glaring errors in there. Trying to get the kids to recognise both the opportunities and the limitations of the tech. I thought this was awesome.” Dale feels people shouldn’t buy into “this idea that we can just describe [AI tools] as bad, only for cheaters and lazy kids, and that there’s some virtue in

WHAT HAS ENABLED THE RECENT RAPID DEVELOPMENTS IN AI?

“The biggest thing [over the past 20 years] is that hardware has come on so much. We used to have to build racks of servers. Nowadays, you can [work on AI models] on your laptop. The advent of cloud means if I do need to train a model that’s a bit complicated, I can spin up a bunch of servers for a few hours to do it in without needing to physically build a rack of servers. All of our computers have become so much more powerful, sometimes we don’t even need servers.”

▶ doing everything the hard or manual way”. He recalls an analogy he likes to use when describing advances in AI: “I remember when I was at school, my dad gave me grief for using a calculator. In his day, they did mental arithmetic. He would say how my mental arithmetic wasn’t as good as his generation’s, which is true, it sucks! But the sums I was doing were way more complex than what he was doing when he was at school. The calculator was doing the mechanical rope work, freeing me up to do the bigger, more challenging things.”

Part of our conversations in the classroom could relate to analogies around AI such as Dale’s. Asking your learners to discuss how using an AI chatbot to generate an essay is like or unlike using a calculator to solve a maths problem could spark them to think more carefully about how they use AI tools.

The future with AI

What does the future of AI hold for us? Ever optimistic about technology, Dale says, “I hope that people don’t look at AI just as something that might replace them, but as something that will enable them to do bigger, better, more amazing things, and that’s been true about tech for years.”

While there are many potential positives around recent AI developments, to be in a place where we’re able to harness it to people’s benefit, we also must recognise that it can be used in dangerous ways. When I ask Dale to share his perspective, he tells me, “I think the biggest risks ... that we see today come from AI being applied in ways that people didn’t realise was possible. Not just in terms of generating text, but if you think of deepfakes.” Not knowing what media and information to trust is a

problem for society, and Dale believes that deepfakes — videos generated or manipulated using AI tools — pose a unique challenge. “We trust video in a way that we don’t trust pictures. Before the days of Photoshop, we would have trusted a photo, but we know photos can be manipulated now and we look critically at photos before we trust them. We still trust video, but there are loads of examples around of videos of famous people looking like they’re saying something that they would never dream of saying. I could totally see [AI technology] being used in a way that shows some world leader making some big announcement and the world panics before anyone realises it was a fake.”

The question that governments, civil society, researchers, and industry leaders are asking is how we control the direction in which AI technology is developed. I ask Dale what role he sees legislation playing in this,



Machine Learning for Kids gives a hands-on intro to machine learning

and he points to the problem posed by the wide availability of AI tools. “We do need oversight and legislation, but if you’ve got big enough computer or server access, you can run these models yourself. The days that you could just target a few big companies with legislation and control it are gone. This stuff is available. These models are open-sourced, and [people who are planning to use AI tools] maliciously are never going to follow the legislation. The best you can

do is put safeguards in place in the rest of society.”

As Dale describes, AI is not going to just disappear from our lives. But our job as educators remains the same: to help young people navigate their lives as best they can — now and in the future. So we need to put effort and time into preparing them to be discerning, critical thinkers who understand how to use and assess new technologies as they emerge.

Machine Learning for Kids is an excellent tool for getting young people started with understanding and building their own machine learning models and AI applications. It is the main tool used in our Experience AI Lessons, and the Machine Learning for Kids website also comes with lots of worksheets and example projects to get you started. [\(HWO\)](#)

AI IS NOT GOING TO JUST DISAPPEAR

For more information about Machine Learning for Kids, visit machinelearningforkids.co.uk

LISTEN TO OUR PODCAST

Dive into the Hello World podcast for:

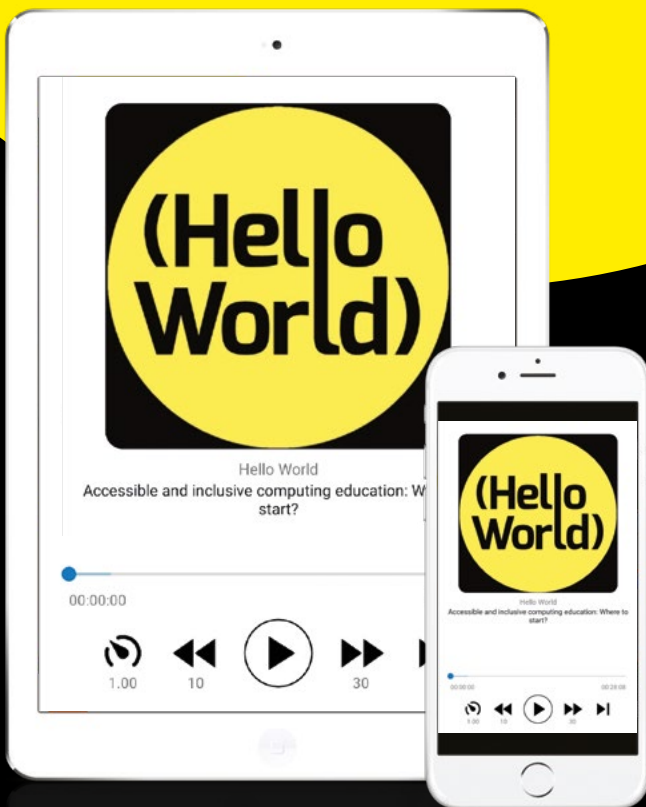
- Even more of the great Hello World content you love between issues
- First-hand knowledge directly from the educators behind our articles
- Learning and expertise as you delve a little deeper and listen on the move

Subscribe to the podcast!

- Apple Podcasts
helloworld.cc/applepod
- Spotify
helloworld.cc/spotifypod
- Google Podcasts
helloworld.cc/googlepod
- Amazon Music
helloworld.cc/amazonpod



SCAN ME



OUR MOST POPULAR EPISODES

How moral is your machine? Ethics in computing education


Accessible and inclusive computing education: where to start?

How can we get everyone excited about code?

(HW)

SUBSCRIBE TODAY

**FREE
IN PRINT**
for UK-based
educators



Why subscribe?

- Teaching resources and ideas used by over 90 percent of our readers
- Exclusive news, research findings, and in-depth features
- Delivered to your door three times a year


TO SUBSCRIBE VISIT:
helloworld.cc/subscribe

Not a UK-based educator?

- Buy a print copy from helloworld.cc/buy – we ship to over 50 countries
- Subscribe to receive the free PDF on the day it is released
- Read features and news at helloworld.cc

**FREE
PDF**
for anyone,
anywhere





© Parradee/stock.adobe.com



Recent AI developments are causing a rethink in computing education

WHAT DOES AI MEAN FOR COMPUTING EDUCATION?

Philip Colligan engages with the big questions surrounding AI and computing education

It's been less than a year since ChatGPT catapulted generative artificial intelligence (AI) into the mainstream public consciousness, reigniting the debate about the role that these powerful new technologies will play in all of our futures.

'Will AI save or destroy humanity?' might seem like an extreme title for a podcast, particularly if you've played with these products and enjoyed some of their obvious limitations. The reality is that we are still at the foothills of what AI technology can achieve (think World Wide Web in the 1990s), and lots of credible people are predicting an astonishing pace of progress over the next few years, promising the radical transformation of almost every aspect of our lives. Comparisons with the Industrial Revolution abound.

At the same time, there are those saying it's all moving too fast; that regulation isn't keeping pace with innovation. One of the UK's leading AI entrepreneurs, Mustafa Suleyman, said recently: "If you don't start from a position of fear, you probably aren't paying attention."

What does all this mean for education, and particularly for computing education? Is there any point trying to teach children about AI when it is all changing so fast? Does anyone need to learn to code anymore? Will teachers be replaced by chatbots? Is assessment as we know it broken?

If we're going to seriously engage with these questions, we need to understand that we're talking about three different things:


1. AI literacy: what it is and how we teach it
2. Rethinking computer science (and possibly some other subjects)

3. Enhancing teaching and learning through AI-powered technologies

AI literacy: what it is and how we teach it

For young people to thrive in a world that is being transformed by AI systems, they need to understand these technologies and the role they could play in their lives.

The first problem is defining what AI literacy actually means. What are the concepts, knowledge, and skills that it would be useful for a young person to learn?

In the past couple of years there has been a huge explosion in resources that claim to help young people develop AI literacy. Our research team mapped and categorised over 500 resources and undertook a systematic literature review (helloworld.cc/AIreview) to understand what research 



has been done on K–12 AI classroom interventions (spoiler: not much).

The reality is that — with a few notable exceptions — the vast majority of AI literacy resources available today are probably doing more harm than good. For example, in an attempt to be accessible and fun, many materials anthropomorphise AI systems, using human terms to describe them and their functions, perpetuating misconceptions about what AI systems are and how they work.

What emerged from this work at the Raspberry Pi Foundation is the SEAME model, which articulates the concepts, knowledge, and skills that are essential ingredients of any AI literacy curriculum (see pages 14–15 in this issue). It separates out the social and ethical, application, model, and engine levels of AI systems — all of which are important — and gets specific about age-appropriate learning outcomes for each.

This research has formed the basis of Experience AI (experience-ai.org), a suite of resources, lessons plans, videos, and interactive learning experiences created by the Raspberry Pi Foundation in partnership with Google DeepMind, which is already being used in thousands of classrooms.

Defining AI literacy and developing resources is part of the challenge, but

that doesn't solve the problem of how we get them into the hands and minds of every young person. This will require policy change. We need governments and education system leaders to grasp that a foundational understanding of AI technologies is essential for creating economic opportunity, ensuring that young people have the mindsets to engage positively with technological change, and avoiding a widening of the digital divide. We've messed this up before with digital skills. Let's not do it again.

More than anything, we need to invest in teachers and their professional

Rethinking computer science

Alongside introducing AI literacy, we also need to take a hard look at computer science. At the very least, we need to make sure that computer science curricula include machine learning models, explaining how they constitute a new paradigm for computing, and give more emphasis to the role that data will play in the future of computing. Adding anything new to an already packed computer science curriculum means tough choices about what to deprioritise to make space.

And, while we're reviewing curricula, what about biology, geography, or any of the other subjects that are just as likely to be

THE REALITY IS THAT WE ARE STILL AT THE FOOTHILLS OF WHAT AI CAN ACHIEVE

development. While there are some fantastic computing teachers with computer science qualifications, the reality is that most of the computing lessons taught anywhere on the planet are taught by a non-specialist teacher. That is even more so the case for anything related to AI. If we're serious about AI literacy for young people, we have to get serious about AI literacy for teachers.

revolutionised by big data and AI? As part of Experience AI, we are launching some of the first lessons focusing on ecosystems and AI, which we think should be at the heart of any modern biology curriculum.

There is already a lively debate about the extent to which the new generation of AI technologies will make programming as we know it obsolete. In January, the prestigious ACM journal ran an opinion piece from Matt Welsh, founder of an AI-powered programming start-up, in which he said: "I believe the conventional idea of 'writing a program' is headed for extinction, and indeed, for all but very specialised applications, most software, as we know it, will be replaced by AI systems that are trained rather than programmed."


With GitHub (now part of Microsoft) claiming that their pair programming technology, Copilot, is now writing 46 percent of developers' code, it's perhaps not surprising that some are saying young people don't need to learn how to code. It's an easy political soundbite, but it just doesn't stand up to serious scrutiny.

Even if AI systems can improve to the point where they generate consistently



Teachers need support and professional development

© Monkey Business / iStockphoto.com



reliable code, it seems to me that it is just as likely that this will increase the demand for more complex software, leading to greater demand for more programmers. There is historical precedent for this: the invention of abstract programming languages such as Python dramatically simplified the act of humans providing instructions to computers, leading to more complex software and a much greater demand for developers.

However these AI-powered tools develop, it will still be essential for young people to learn the fundamentals of programming and to get hands-on experience of writing code as part of any credible computer science course. Practical experience of writing computer programs is an essential part of learning how to analyse problems in computational terms; it brings the subject to life; it will help young people understand how the world around them is being transformed by AI systems; and it will ensure that they are able to shape that future, rather than it being something that is done to them.

are, none of them have even begun to harness the potential of AI systems like large language models (LLMs).

One area where I think we'll see huge progress is feedback. It's well established that good-quality feedback makes a huge difference to learning, but a teacher's ability to provide feedback is limited by their time. No one is seriously claiming that chatbots will replace teachers, but — if we can get the quality right — LLM applications could provide every child with unlimited, on-demand feedback. AI powered feedback — not giving students the answers, but coaching, suggesting, and encouraging in the way that great teachers already do — could be transformational.

We are already seeing edtech companies racing to bring new products and features to market that leverage LLMs, and my prediction is that the pace of that innovation is going to increase exponentially over the coming years. The challenge for all of us working in education is how we ensure that ethics and privacy are at the centre of the development of these technologies.



PHILIP COLLIGAN

Philip is the chief executive of the Raspberry Pi Foundation, an educational nonprofit with a mission to enable young people to realise their full potential through the power of computing and digital technologies.

as an AI kitemark for edtech products that communicate whether they meet a set of standards around bias, transparency, and privacy.

“ HOW MUCH DATA FROM STUDENTS WILL AN AI SYSTEM NEED TO ACCESS? ”

Enhancing teaching and learning through AI-powered technologies


Technology has already transformed learning. YouTube is probably the most important educational innovation of the past 20 years, democratising both the creation and consumption of learning resources. Khan Academy, meanwhile, integrated video instruction into a learning experience that gamified formative assessment. Our own edtech platform, Ada Computer Science, combines comprehensive instructional materials, a huge bank of questions designed to help learning, and automated marking and feedback to make computer science easier to teach and learn. Brilliant though these

That's important for all applications of AI, but especially so in education, where these systems will be unleashed directly on young people. How much data from students will an AI system need to access? Can that data — aggregated from millions of students — be used to train new models? How can we communicate transparently the limitations of the information provided back to students?

Ultimately, we need to think about how parents, teachers, and education systems (the purchasers of edtech products) will be able to make informed choices about what to put in front of students. Standards will have an important role to play here, and I think we should be exploring ideas such

Realising potential in a brave new world

We may very well be entering an era in which AI systems dramatically enhance the creativity and productivity of humanity as a species. Whether the reality lives up to the hype or not, AI systems are undoubtedly going to be a big part of all of our futures, and we urgently need to figure out what that means for education, and what skills, knowledge, and mindsets young people need to develop in order to realise their full potential in that brave new world.

That's the work we're engaged in at the Raspberry Pi Foundation, working in partnership with individuals and organisations from across industry, government, education, and civil society. If you have ideas and want to get involved in shaping the future of computing education, we'd love to hear from you at contact@helloworld.cc. 

HOW ANTHROPOMORPHISM HINDERS AI EDUCATION

Ben Garside explains why talking about artificial intelligence systems as sentient beings is a bad idea

In the 1950s, Alan Turing explored the central question of artificial intelligence (AI). He thought that finding the answer to a question which many researchers were discussing at the time, ‘Can machines think?’, would not provide valuable answers because the terms ‘machine’ and ‘think’ are hard to define. Instead, he proposed changing the question to something more provable: ‘Can a computer imitate intelligent behaviour well enough to convince someone they are talking to a human?’ (helloworld.cc/turing1950). This is commonly referred to as the Turing test.

It’s been hard to miss the newest generation of AI chatbots that companies have released over the last year. News articles and stories about them seem to be everywhere, so you may have heard of machine learning (ML) chatbots such as ChatGPT and Bard. These chatbots are advanced enough to have caused renewed discussions about the Turing test and whether the chatbots are sentient.

Chatbots are not sentient

Without any knowledge of how people create such chatbots, it’s easy to imagine how someone might develop an incorrect mental model around these chatbots being living entities. With some awareness of sci-fi stories, you might even start to imagine what they could look like or associate a gender with them.

The reality is that these new chatbots are applications based on a large language model (LLM) — a machine learning model trained with vast quantities of text, written by people, and taken from places such as books and the internet. An LLM predicts the probable order of combinations of words, a bit like the autocomplete function on a smartphone. Based on these probabilities, it can produce text outputs. LLM chatbots run on servers with huge amounts of computing power that people have built in data centres around the world.

The problem with anthropomorphism

AI applications are often described as ‘black boxes’ or ‘closed boxes’; they may be relatively easy to use, but it’s not as easy to understand how they work. As educators, it’s fundamentally important to help everyone, especially young people, understand the potential of AI technologies and open these closed boxes to know how they actually work. We should be demystifying digital technology for young people to empower them to be thoughtful creators of technology and to make informed choices about how they engage with technology — rather than just being passive consumers.

To support learners in forming accurate mental models of AI and ML systems, it is important to avoid using words that can lead to learners developing misconceptions about machines being humanlike in their abilities. That’s why anthropomorphism (showing




BEN GARSIDE

Ben is a learning manager at the Raspberry Pi Foundation. He has worked on the production of *The Computing Curriculum* and has written online courses for educators, including *Introduction to Machine Learning and AI*.

or treating an object such as an AI system as though it were human in appearance, character, or behaviour) is a term that comes up regularly in conversations with my colleagues when we are developing resources at the Raspberry Pi Foundation.

Anthropomorphising AI in teaching materials might lead learners to believe that there is sentience or intention within AI applications. That misconception would distract learners from the fact that it is people who design AI applications and decide how they are used. It also risks reducing learners’ desire to take an active role in understanding AI applications and designing future applications.



© lensv0rd / stock.adobe.com

“ IT’S MORE ACCURATE TO SAY THAT AI RECEIVES INPUT AND PROCESSES THE DATA

Avoiding anthropomorphism helps young people to open the closed box of AI applications. Take the example of a smart speaker. It’s easy to describe a smart speaker’s functionality in anthropomorphic terms such as ‘it listens’ or ‘it understands’. However, it’s more accurate and empowering to explain smart speakers as systems developed by people to process sound and carry out specific tasks. Rather than telling young people that a smart speaker ‘listens’ and ‘understands’, it’s more accurate to say that the speaker receives input, processes the data, and produces an output. This language helps to distinguish how the device actually works from the illusion of a persona the speaker’s voice might conjure for learners.

Another example is the use of AI in computer vision (a field of AI that attempts to gain meaningful information from images and videos). ML models can, for example, be trained to identify when there is a dog or a cat in an image. An accurate ML model, on the surface, therefore displays humanlike behaviour. However, the model operates very

differently from how a human might identify animals in images. Where humans would point to features such as whiskers and ear shapes, ML models process pixels in images to make predictions based on probabilities.

Avoiding the AI ‘personality’ trap

It’s not easy to avoid anthropomorphism while talking about AI, especially considering the industry-standard language

in the area: artificial intelligence, machine learning, and computer vision, to name but a few examples. At the Raspberry Pi Foundation, we are still training ourselves not to anthropomorphise AI, and we take a bit of pleasure in picking each other up on the odd slip-up.

If we are correct in our approach when developing resources at the Raspberry Pi Foundation, then whether or not the young people who use our resources grow up to become AI developers, we will have helped them to become discerning users of AI technologies. They will be more likely to see such products for what they are: data-driven applications and not sentient machines. (HW)

BETTER WAYS TO DESCRIBE AI

Terms to avoid	Terms to use
Avoid using phrases such as ‘AI learns’ or ‘AI/ML does’	Use phrases like ‘AI applications are designed to’ or ‘AI developers build applications that’
Don’t use words that describe the behaviour of people (e.g. see, look, recognise, create, make)	Prefer system-related words (e.g. detect, input, pattern-match, generate, produce)
Avoid using AI/ML as a countable noun, e.g. ‘new artificial intelligences emerged in 2022’	Refer to ‘AI/ML’ as a scientific discipline, similar to how you use the term ‘biology’



Existing AI facial recognition systems do not identify dark-skinned and female faces accurately

BIAS IN THE MACHINE: HOW CAN WE ADDRESS GENDER BIAS IN AI?

Sue Sentance examines the intersection of AI and gender, particularly with respect to gender bias in machine learning

Gender inequality is everywhere, and researchers, activists, and governments themselves, have struggled since the 1960s to tackle it. As women and girls around the world continue to suffer from discrimination, the United Nations has pledged to achieve gender equality and to empower all women and girls through its Sustainable Development Goals (helloworld.cc/UN).

While progress has been made, new developments in technology may be threatening to undo this. As Susan Leavy, a machine learning (ML) researcher from the Insight Centre for Data Analytics, puts it in her conference proceedings, “Artificial intelligence is increasingly influencing the opinions and behaviour of people in everyday life. However, the overrepresentation of men in the design of these technologies could quietly undo decades of advances in gender equality.”

Gender-biased data

In her 2019 award-winning book *Invisible Women: Exposing Data Bias in a World Designed for Men*, Caroline Criado Perez discusses the effects of gender-biased data. She describes, for example, how the designs of cities, workplaces, smartphones,

and even crash test dummies are all based on data gathered from men.


Looking at this problem from a different angle, researcher Mayra Buvinic and her colleagues highlight that in most countries of the world, there are no sources of data that capture the differences between male and female participation in civil society organisations, or in local advisory or decision-making bodies. A lack of data about girls and women will surely have a negative impact on decision-making.

Bias in machine learning

Machine learning is a type of AI technology that relies on vast data sets for training. ML is currently being used in various systems for automated decision-making. Bias in data sets (helloworld.cc/biasindata) for training ML models can be caused in several ways. For example, data sets can be biased because they are incomplete or skewed (as is the case in data sets which lack data about women). Another example is that data sets can be biased because of the use of incorrect labels by people who annotate the data. Annotating data is necessary for supervised learning, where ML models are trained to categorise data into categories decided upon by people (e.g. pineapples and mangoes).

In order for a ML model to categorise new data appropriately, it needs to be trained with data that is gathered from everyone, and is, in the case of supervised learning, annotated without bias. Failing to do this creates a biased ML model. Bias has been demonstrated in different types of AI systems that have been released as products. For example:

- **Facial recognition:** AI researcher Joy Buolamwini discovered that existing AI facial recognition systems do not identify dark-skinned or female faces accurately. Her discovery, and her work to push for the first-ever piece of legislation in the USA to govern against bias in the algorithms that impact our lives, is narrated in the 2020 documentary *Coded Bias* (codedbias.com).
- **Natural language processing:** imagine an AI system that is tasked with filling in the missing word in ‘Man is to king as woman is to X’ and comes up with ‘queen’. But what if the system completes ‘Man is to software developer as woman is to X’ with ‘secretary’ or some other word that reflects stereotypical views of gender and careers? AI models called ‘word embeddings’ learn by identifying patterns in huge collections of texts. In



FURTHER READING

- ✓ Buvinic, M., & Levine, R. (2016). Closing the gender data gap. *Significance*, 13(2):34–37. helloworld.cc/Buvinic
- ✓ Leavy, S. (2018). Gender bias in artificial intelligence: The need for diversity and gender theory in machine learning. In *Proceedings of the 1st International Workshop on Gender Equality in Software Engineering (GE 2018)* (pp. 14–16). helloworld.cc/Leavy
- ✓ Perez, C. C. (2019). *Invisible Women: Exploring Data Bias in a World Designed for Men*. Random House. helloworld.cc/Perez

addition to the structural patterns of the text language, word embeddings learn human biases expressed in the texts. You can read more about this in the Brookings Institute report (helloworld.cc/Brookings).

Not noticing

There is much debate about the level of bias in systems using AI, and some researchers worry that this will cause distrust in ML systems. Thus, some scientists are keen to emphasise the breadth of their training data across genders (helloworld.cc/ultromics). However, other researchers point out that despite all good intentions, gender disparities are so entrenched in society that we are not aware of all of them. White and male dominance in our society may be so unconsciously prevalent that we don't notice all its effects.

What is being done in the AI sector to address bias?

Developers and researchers of AI systems have been trying to establish rules for how to avoid bias in AI models. An example rule set is given in a Harvard Business Review article (helloworld.cc/Harvard), which describes the fact that speech recognition systems originally performed poorly for non-male speakers, because systems analysed and modelled speech for taller

speakers with longer vocal cords and lower-pitched voices (typically males).

The article recommends four ways for the ML workforce to try to avoid gender bias:

- Ensure diversity in the training data (using the example from the article, including audio samples from all genders, not just male)
- Ensure that a diverse group of people labels the training data
- Measure the accuracy of a ML model separately for different demographic categories to check whether the model is biased against some of the categories
- Establish techniques to encourage ML models towards unbiased results

What can everybody else do?

The above points can help people in the AI industry, which is of course important — but what about the rest of us? It's important to raise awareness of the issues around gender data bias and AI lest we find out too late that we are reintroducing gender inequalities we have fought so hard to remove.

Awareness is a good start, and some other suggestions, drawn out from others' work in this area, are:

- **Improve the gender balance in the AI workforce.** Having more women in AI and data science, particularly in both technical and leadership roles, will help to



SUE SENTANCE

Sue is director of the Raspberry Pi Computing Education Research Centre. An ex-teacher, a teacher-trainer, and an academic, she is passionate about research in computing education.

reduce gender bias. A 2020 report by the World Economic Forum (WEF) on gender parity found that women account for only 26 percent of data and AI positions in the workforce (helloworld.cc/gendergap).

The WEF suggests five ways in which the AI workforce gender balance could be addressed (helloworld.cc/WEF):

1. Support STEM education
 2. Showcase female AI trailblazers
 3. Mentor women for leadership roles
 4. Create equal opportunities
 5. Ensure a gender-equal reward system
- **Ensure the collection of and access to high-quality and up-to-date gender data.** We need high-quality data sets on women and girls, with good coverage, including country coverage. Data needs to be comparable across countries in terms of concepts, definitions, and measures. Data should have both complexity and granularity, so it can be cross-tabulated and disaggregated, following the recommendations from the Data2x project on mapping gender data gaps (helloworld.cc/data2x).

- **Educate young people about AI.** At the Raspberry Pi Computing Education Research Centre, we believe that introducing some of the potential (positive and negative) impacts of AI systems to young people in school may help to build awareness and understanding at a young age. In our AI education seminars we heard a number of perspectives on this topic, and you can revisit the videos, presentation slides, and blog posts (helloworld.cc/Alseries). We've also conducted a literature review of AI education interventions in K–12 (helloworld.cc/Alreview).

We are using our understanding of and research around AI to develop new learning programs. As we think deeply about how AI concepts can be learnt and understood, we will continue to educate young people about the risks of bias in AI. (HAW)

EXPERIENCE AI: MAKING AI RELEVANT AND ACCESSIBLE

Google DeepMind's **Aimee Welch** discusses the Experience AI programme and why equal access to AI education is key

From chatbots to self-driving cars, AI is here to stay, and it's rapidly transforming our world. It can potentially solve some of the biggest challenges humanity faces today. But it also has many serious risks and inherent challenges, like reinforcing existing patterns of bias or 'hallucinating', a term that describes AI making up false outputs that do not reflect real events or data.

As AI becomes integral to our daily lives, younger generations must gain the knowledge and skills to navigate and shape this technology. Young people with a foundational understanding of AI can

make more informed decisions about using AI applications in their daily lives, helping ensure safe and responsible use of the technology. The UK government's AI Council acknowledges this and aims to have every child in the UK leave school with a basic understanding of how AI works, as stated in its AI Roadmap ([helloworld.cc/AIroadmap](https://www.helloworld.cc/AIroadmap)).

But while AI literacy is a key skill in this new era, not every young person currently has access to sufficient AI education and resources. In a recent survey by the EdWeek Research Center in the US ([helloworld.cc/EdWeekAISurvey](https://www.helloworld.cc/EdWeekAISurvey)), only one in ten teachers said they knew enough about AI to teach its basics. Very few reported receiving any professional development related to the topic. Similarly, our work with the Raspberry Pi Computing Education Research Centre ([computingeducationresearch.org](https://www.computingeducationresearch.org)) has suggested that UK-based teachers are eager to understand more about AI and how to engage their students in the topic.

Bringing AI into classrooms

Ensuring broad access to AI education is also important. This way, we can improve diversity among people working in the field of AI and ensure the technology is developed safely and responsibly. Unfortunately, some groups of people face barriers to learning about AI, which means they are underrepresented in the field. By increasing diversity in AI, we bring

diverse values, hopes, and concerns into the design and deployment of the technology. This is critical for AI to benefit everyone ([helloworld.cc/equitableAI](https://www.helloworld.cc/equitableAI)).

By focusing on AI education from a young age, there is an opportunity to break down some of these long-standing barriers. That's why we partnered with the Raspberry Pi Foundation to co-create Experience AI, a new educational programme with free lesson plans, slide decks, worksheets, and videos, to address gaps in AI education and support teachers in engaging and inspiring young people in the subject.

The programme aims to help young people aged 11–14 take their first steps in understanding the technology, making it relevant to diverse learners, and encouraging future careers in the field. All Experience AI resources are freely available to every school across the UK and beyond.


The partnership is built on a shared vision to make AI education more inclusive and accessible. Bringing together the Foundation's expertise in computing education and our cutting-edge technical knowledge and industry insights has allowed us to create a holistic learning experience that connects theoretical concepts and practical applications.

Informed by AI experts

A group of 15 research scientists and engineers at Google DeepMind contributed to developing the lessons. From



Teacher Luke Connolly uses Experience AI resources with his students at North Liverpool Academy



drafting definitions for key concepts, to brainstorming interesting research areas to highlight, and even featuring in the videos included in the lessons, the group played a key role in shaping the programme in close collaboration with the Foundation's educators and education researchers. To bring AI concepts to life, the lessons include interactive activities as well as real-life examples. One exciting project involved Google DeepMind collaborating with ecologists and conservationists to develop machine learning methods to study the behaviour of an entire animal community in the Serengeti National Park and Grumeti Reserve in Tanzania (helloworld.cc/MLproject).

Google DeepMind Research Scientist Petar Veličković, a member of the working group, shares: "AI is a technology that is going to impact us all, and therefore educating young people on how to interact with this technology is likely going to be a core part of school education going forward. The project was eye-opening and humbling for me as I learned of the challenges associated with making such a complex topic accessible — not only to every pupil, but also to every teacher! Observing the thoughtful approach undertaken by the Raspberry Pi Foundation left me deeply impressed. I'm taking home many useful ideas that I hope to incorporate in my own AI teaching efforts going forward."

The lessons have been carefully developed to:

- **Follow a clear learning journey**, underpinned by the SEAME framework (helloworld.cc/SEAME), which guides learners sequentially through key concepts and acts as a progression framework.
- **Build foundational knowledge and provide support for teachers**. Focus on teacher training (helloworld.cc/introtoMLandAI) and support is at the core of the programme (watch this space for details on a new webinar series being launched soon!).

- **Embed ethics and responsibility**. Key concepts in AI ethics and responsibility are woven into each lesson and progressively built upon. Students are introduced to concepts like data bias, user-focused approaches, model cards, and how we can use AI for social good.
- **Ensure cultural relevance and inclusion**. Experience AI was designed with diverse learners in mind and includes a variety of activities to enable young people to pick topics that most interest them.

What teachers say

To date, we estimate the resources have reached more than 130,000 students in the UK and beyond. We're thrilled to hear from teachers already using the resources about their impact in the classroom. Mrs J Green from Waldegrave School in London, UK, says: "I thought that the lessons covered



AIMEE WELCH

Aimee leads Google DeepMind's work supporting school students and teachers to broaden participation in science and address barriers leading to underrepresentation in the AI ecosystem. Aimee has ten years of experience developing and delivering education programmes to increase diversity and equity in AI and STEM fields.

“STUDENTS NOW FEEL IMMENSELY MORE KNOWLEDGEABLE ABOUT THE IMPACT THAT AI HAS IN THEIR LIVES”

a really important topic. Giving the pupils an understanding of what AI is and how it works will become increasingly important as it becomes more ubiquitous in all areas of society. The lessons that we trialled took some of the 'magic' out of AI and started to give the students an understanding that AI is only as good as the data that is used to build it. It also started some really interesting discussions with the students around areas such as bias."

At North Liverpool Academy in the UK, teacher Dave Cross tells us: "AI is such a current and relevant topic in society that [these lessons] will enable Key Stage 3 computing students [ages 11–14] to gain a solid foundation in something that will become more prevalent within the curriculum, and wider subjects too as more sectors adopt AI and machine learning as standard. Our Key Stage 3 computing

students now feel immensely more knowledgeable about the importance and place that AI has in their wider lives. These lessons and activities are engaging and accessible to students and educators alike, whatever their specialism may be."

We hope the programme instils confidence in both teachers and students, helping to address some of the critical school-level barriers leading to underrepresentation in AI and playing a role in building a stronger, more inclusive AI community where everyone can participate irrespective of their background.

Today's young people are tomorrow's leaders — and as such, educating and inspiring them about AI is valuable for everybody.

Teachers can visit experience-ai.org to download all of the Experience AI resources for free. [\(H+W\)](#)



© Karen Roach/stockadobe.com

Technology companies like to anthropomorphise their products to make them seem less scary

AI ISN'T JUST ROBOTS

Sway Grantham offers her tips on how to talk to young children about artificial intelligence and avoid introducing alternate conceptions

Young children have a unique perspective on the world they live in. They often seem oblivious to what's going on around them, but then they will ask a question that makes you realise they did get some insight from a news story or a conversation they overheard. This happened to me with a class of ten-year-olds when one boy asked, with complete sincerity and curiosity, "And is that when the zombie apocalypse happened?" He had unknowingly conflated the plague with television depictions of zombies taking over the world.


Absorbing media and assimilating it into your existing knowledge is a challenge,

and this is a concern when the media is full of big, scary headlines about artificial intelligence (AI) taking over the world, stealing jobs, and being sentient. As teachers, you don't need to know all the gory details about AI to answer learners' questions, but you can avoid accidentally introducing alternate conceptions. This article offers some top tips to help you point those inquisitive minds in the right direction.

AI is not a person

Technology companies like to anthropomorphise their products and give them friendly names. Why? Because it makes their devices seem more endearing and less

scary, and makes you more likely to include them in your lives. However, when you think of AI as a human with a name, who needs you to say 'please' or is 'there to help you', you start to make presumptions about how it works, what it 'knows', and its morality. This changes what we ask, how much we trust the responses, and our behaviours when using the device. The device, though, does not 'see' or 'know' anything; instead, it uses lots of data to make predictions. Think of word association: if I say 'bread', I predict that a lot of people in the UK will think 'butter'. Here, I've used the data I've collected from years of living in this country to predict a reasonable answer. This is all AI devices are doing.



When talking to young children about AI, try to avoid using pronouns such as 'she' or 'he'. Where possible, avoid giving devices human names, and instead call them 'computer', to reinforce the idea that humans and computers are very different. Let's imagine that a child in your class says, "Alexa told me a joke at the weekend — she's funny!" You could respond, "I love using computers to find new jokes! What was it?" This is just a micro-conversation, but with it, you are helping to surreptitiously challenge the child's perception of Alexa and the role of AI in it. Another good approach is to remember to keep your emotions separate from computers, so as not to give them human-like characteristics: don't say that the computer 'hates' you, or is 'deliberately ignoring' you, and remember that it's only 'helpful' because it was told to be! Language is important, and we must continually practise avoiding anthropomorphism!

AI. When learners are using streaming services for music or TV and the service predicts something that they might want to watch or listen to next, point out that this is using AI technology. When they see their parents planning a route using a satnav, explain that the satnav system uses data and AI to plan the best route. Even better than just calling out uses of AI, try to have conversations about when things go wrong and AI systems suggest silly options. This is a great way to build critical thinking around the use of computers. AI systems don't always know best, because they're just making predictions, and predictions can always be wrong!

AI and humans complement each other

There's a delicate balance between acknowledging the limitations of AI and portraying it as a problematic tool that we



SWAY GRANTHAM

Sway is a senior learning manager at the Raspberry Pi Foundation, where she leads the educator development team developing resources and training for computing teachers around the world (@SwayGrantham).

To help children understand how AI and humans complement each other, ask pupils the question, 'What can't a computer do?' Answers that I have received before include, 'Give me a hug', 'Make me laugh', and 'Paint a picture', and these are all true. Can Alexa tell you a joke that makes you laugh? Yes — but a human created that joke. The computer is just the way in which it is being shared. Even with AI 'creating' new artwork, it is really only using data from something that someone else created. Humans are required.

Overall, we must remember that young children are part of a world that uses AI, and that it is likely to be ever more present in the future. We need to ensure that they know how to use AI responsibly, by minimising alternate conceptions. With our youngest learners, this means taking care with the language you choose and the examples you use, and explaining its role as a tool. These simple approaches are the first steps to empowering children to go on to harness this technology, and they pave the way for you to simply introduce the core concepts of AI in later computing lessons without first having to untangle a web of alternate conceptions. (HW)

“ CAN ALEXA TELL YOU A JOKE THAT MAKES YOU LAUGH? YES — BUT A HUMAN CREATED THAT JOKE. HUMANS ARE STILL REQUIRED

AI isn't just robots (actually, it rarely is)

The media plays a huge role in what we imagine when we talk about AI. For the media, the challenge is how to make lines of code and data inside a computer look exciting and recognisable to their audiences. The answer? Robots! When learners hear about AI taking over the world, it's easy for them to imagine robots like those you'd find in a Marvel movie. Yet the majority of AI exists within systems they're already aware of and are using — you might just need to help draw their attention to it.

For example, when using a word processor, you can highlight to learners that the software sometimes predicts what word you want to type next, and that this is an example of the computer using

shouldn't use. AI offers us great opportunities to improve the way we work, to get us started on a creative project, or to complete mundane tasks. However, it is just a tool, and tools complement the range of skills that humans already have. For example, if you gave an AI system the prompt, 'Write a setting description using these four phrases: dark, scary, forest, fairy tale', the first output from the tool probably wouldn't make much sense. As a human, though, you'd have to do far less work to edit the output than if you had had to write the setting description from scratch. Now, say you had the perfect example of a setting description, but you wanted 29 more examples, so that each learner in your class had a different version. This is where AI can help: completing a repetitive and boring task, and saving time for humans.

ALEXA, WHAT ARE YOU THINKING ABOUT?

Judy Robertson describes recent research into what children already know about AI based on their interactions with smart speakers at home

The grandfather of neural networks, Geoffrey Hinton, recently commented that he didn't think it would be long before artificial intelligence (AI) became more intelligent than humans. Not everyone agrees, but AI is certainly advancing eye-wateringly fast. I believe that, as educators, we have a responsibility to teach young people about how AI technology works and help them think about how it should be used in society.

We must remember, though, that AI isn't something that is just about to happen in children's lives. It has been part of their home lives for at least five years, in the form of smart speakers such as Alexa

and Google Home. At the University of Edinburgh, we recently researched children's attitudes towards smart speakers, to find out what they already know about AI (helloworld.cc/Robertson) and help us develop educational materials for AI literacy.

Children's views on smart speakers


We conducted a survey of 166 eight- to eleven-year-olds, with follow-up focus groups. Of the children, 93 percent had a smart speaker at home, mostly Amazon's Alexa. The children generally thought that devices such as Alexa were smart, often cleverer than themselves, but less clever than scientists. They knew that the devices were

examples of AI, rather than being literally human, although about two-thirds of them thought that they might think like a human.

This research showed that children aren't sure what to think of smart speakers. In the survey, when asked whether smart speakers have feelings, or can make decisions for themselves, many answered 'maybe'. When posed this question in the focus groups, the younger children sometimes just admitted they didn't know. The older children debated the question together, backing up their answers with examples of what Alexa does when they interact with it at home. For example, one child reported, "I said 'Arrgg, just play the song already.' He [Alexa] was

■ AI is already part of children's home lives in the form of smart speakers





JUDY ROBERTSON

Judy is a professor of digital learning at the University of Edinburgh. She researches computer science education, data literacy, and AI literacy. She also designs technology with and for children.

like, “Don’t be rude, OK?” So I don’t know if he has feelings.”

It’s not surprising that children are puzzled about whether smart speakers have feelings, because they are often designed to behave as if they do. When the device sings *Happy Birthday* to a child or thanks them for a compliment, it builds up the impression that it is humanlike. So, the next time the device fails to respond when someone shouts at it, the child may reasonably guess that it is sulking. If the child doesn’t have any technical understanding of how the device works, their only method for making sense of it is to reason about its behaviour as they would reason about another person.

Misunderstandings about smart speakers

An important part of our study was the goal of understanding the misconceptions young people have about smart speakers and AI more broadly. The problem with trying to make sense of how a smart speaker behaves as if it were human is that it leads to inaccurate predictions. In the study, we found that the children both overestimated and underestimated the capabilities of the technology. Privacy mattered a lot to the children, for example, but they sometimes

had inaccurate ideas about what the smart speakers would do with their information. One child speculated, “If you tell Alexa secrets, I feel those messages travel to Alexas throughout the Alexa Amazon network, and then if someone’s telling their secrets to Alexa and someone else asks them to tell their secrets but then it tells YOUR secrets!” This scenario is improbable if you know that Alexa devices are designed to be associated with specific user accounts, and that they are not networked together in such a way that information can inadvertently leak between devices. If you are making an analogy with the way that gossip travels around a school playground, though, it seems more plausible. In this case, the analogy with human interactions raises a threat that is not likely to happen.

The children were horrified by the idea of the other people in their house being able to see records of their interactions with their smart speakers, and told us they would

programming had misunderstandings — some thought that smart speakers were programmed in advance with an answer to every possible question that a user might pose. There is a gap in children’s knowledge here, even if they pay close attention to their computational thinking lessons; thinking of machine learning techniques used in smart speakers as simply procedural algorithms does not capture the insight that they are predictive statistical models.

Suggestions for teachers

So what can educators do to help children understand AI better? A good starting place is to talk to children about how the AI technology they use in their everyday lives works (helloworld.cc/AI-literacy). As part of this discussion, inform children about privacy settings for the devices they commonly use, discuss with them what settings they feel comfortable with, and help them to put the risks into perspective.

“ CHILDREN ARE OFTEN PUZZLED ABOUT WHETHER SMART SPEAKERS HAVE FEELINGS ”

change their behaviour if that were the case. They did not understand that, in fact, Alexa does keep a transcript of all interactions with all users in the house, which can be accessed by the account holder (often the parents or carers). The potential harm to the child of a privacy invasion by another family member is possibly more immediately bothersome than the misuse of data by a third-party company.

Another set of misunderstandings about smart speakers was related to how they give such apparently intelligent answers. Some children thought that a human intercepted the questions from users and thought of the responses. Memorably, one child surmised that a “really smart guy plugged himself in” to Alexa to download his intelligence and build an army of Alexas. Even the children who had been taught about computer


Teachers can also help to explain how machine learning works, and this will probably work best with a class that already understands the concept of simple algorithms (helloworld.cc/algorithms). Machine learning algorithms enable the computer to figure out how to do tasks without the programmer explicitly writing down rules. The book *Neural Networks for Babies* by Chris Ferrie and Dr Sarah Kaiser is an excellent picture book which introduces the main concepts of machine learning to a young audience.

Just like adults, some children may be excited by new AI technologies, and some may be fearful of the idea that AI could replace humans. As teachers, we can help them make sense of the possible future of AI by beginning with what they already know about the AI they commonly encounter at home. [\[HWW\]](http://helloworld.cc)

AI MADE EASY: TEACHING AI TO KS2 AND KS3

Rebecca Mack shares her experience of creating and delivering a new curriculum topic about artificial intelligence in both primary and secondary settings

Artificial intelligence (AI) has gained significant attention worldwide recently. It's even a prominent topic among young students who are now familiar with, and sometimes actively using, AI tools like ChatGPT. As educators, we need to prepare students for the future by equipping them with knowledge and understanding of AI's workings and limitations. This article explores the journey of designing and implementing an AI curriculum for Year 8 students, aged 12–13, at a preparatory school. It emphasises the importance of AI education in shaping digital literacy in the next generation.



REBECCA MACK

Rebecca is the head of computing and digital learning at Bilton Grange Preparatory School in the UK. Rebecca holds a master's degree in physics and another in education (@Becca_Mack).

Teaching AI to students is crucial due to its rapid development and integration into various aspects of our lives. Educating students about AI empowers them to actively participate in the digital age, comprehend how AI functions, explore its potential applications, and think about the ethical considerations involved. Integrating practical AI tools like ChatGPT into the curriculum offers students an immersive learning experience.

Our AI curriculum was designed with the inclusion of AI tools; as a computing teacher, I used ChatGPT to show my students its functionalities and foster a collaborative learning environment where they could work together to discover its potential.

I used ChatGPT as a tool to help form ideas for what to cover in our curriculum. It helped me to identify key areas, such as machine learning (ML), natural language processing (NLP) and virtual personal assistants (VPAs), text-to-image generators, deepfakes, and the crucial topic of bias and ethics in AI. By exploring these areas, our students gained a better understanding of AI's core components and its real-world implications.

To assess their learning, the students completed a short quiz and gave their opinions on the topic four weeks after completion. This assessment also helped to evaluate the curriculum's effectiveness and gather feedback to improve its future iterations.


Implementation and student voice

The implementation of the AI curriculum involved engaging activities and discussions throughout. Here are some specific topics covered and the students' feedback.

■ **Machine learning:** we used the Hour of Code AI for Oceans tool (code.org/oceans) to introduce ML. Students explored AI learning steps while addressing real-world environmental challenges. Their feedback reflected enthusiasm and improved understanding, with one student mentioning how the tool sparked their interest and enhanced their knowledge of the principles.

■ **Natural language processing and virtual personal assistants:** we examined adverts for popular voice assistants like Siri, Alexa, and Google Assistant to explore NLP and VPAs. Analysing these examples provided insights into how NLP enables machines to understand and respond to human language. Students tried out a variety of chatbots, like [Octa.ai](https://octa.ai), before creating and acting out their own dialogue trees, promoting critical thinking about NLP capabilities and limitations.

■ **Text-to-image generators:** text-to-image generators emerged as a popular and enjoyable topic among students. Activities involved determining whether images were real or fake, writing stories using Mad Libs



templates, and using Dream by WOMBO ([dream.ai](#)) to illustrate their narratives. These projects sparked imagination and provided tangible demonstrations of AI's ability to generate images from text inputs.

■ **Deepfakes:** we began exploring deepfakes by watching the UK's Channel 4 video apparently featuring the late Queen's famous TikTok dance. This initiated discussions on the ethics of deepfakes, emphasising critical thinking and media literacy. The topic also connected with e-safety discussions, highlighting the need for caution when encountering online content. We also tied in ProjectEVOLVE's 'Managing online information' lesson that explores fake news ([projectevolve.co.uk](#)).

After teaching the AI curriculum to Year 8 students (aged 12–13), I adapted it for Year 5 students (aged 9–10) with great success. While the core content remained the same, I tailored the delivery to suit a younger audience. For instance, during discussions on ethics, I provided more support and guidance, focusing on understanding how others might perceive and feel about AI-related topics, taking a one-step-removed approach. Also, in our image generation lesson, I used ChatGPT to collaboratively create a story with the students before they used an image generator to illustrate it. I also used this task to integrate elements of their English lessons, such as verbs and adjectives. The students didn't directly interact with the AI but worked alongside me as we explored ChatGPT's

STUDENT FEEDBACK

"I think it was interesting to see how AI is programmed to determine whether things are classified as 'trash' or 'sea animals'. It made me understand the collaborative nature of AI and how it learns with human help."

"Learning about deepfakes helped us understand what is fake on the internet and how we can avoid it."

AI technology across the school, recognising the value of AI education in preparing students for the future. If we are using the technology to support our teaching and learning, it is key for our students to understand it too.

In conclusion, the positive feedback from students, quiz results, and the enthusiasm of the senior leadership team demonstrate the success of our AI curriculum in engaging students and fostering their understanding of AI concepts. Our next steps are to continue the collaboration between educators and school leadership to ensure AI technology is effectively integrated across various subjects, equipping students with the necessary skills and knowledge to thrive in the digital age. [\(HW\)](#)

I PREDICT THAT SOON EVERY COMPUTER SCIENCE CURRICULUM WILL EMBED AI CONCEPTS THROUGH ALL KEY STAGES

■ **Ethics and bias:** this topic resonated strongly with students, and despite ranking lower in terms of enjoyment in the final quiz, it was voted the most useful topic for the future. We discussed the ethical considerations of self-driving cars, and revisited the issue of bias in AI systems, emphasising the importance of addressing and mitigating bias for fairness and inclusivity. We also touched on the concept of the Turing test here — several students worked further on it at home!

Reflections and the senior leadership team's perspective

The end-of-topic quiz resulted in an average score of 85 percent. After four weeks, we repeated the quiz and saw a slight decrease to 79 percent. As expected, this demonstrates the need for continuous reinforcement. Based on the speed at which AI is developing, I predict that soon every computer science curriculum will embed AI concepts throughout all the Key Stages.

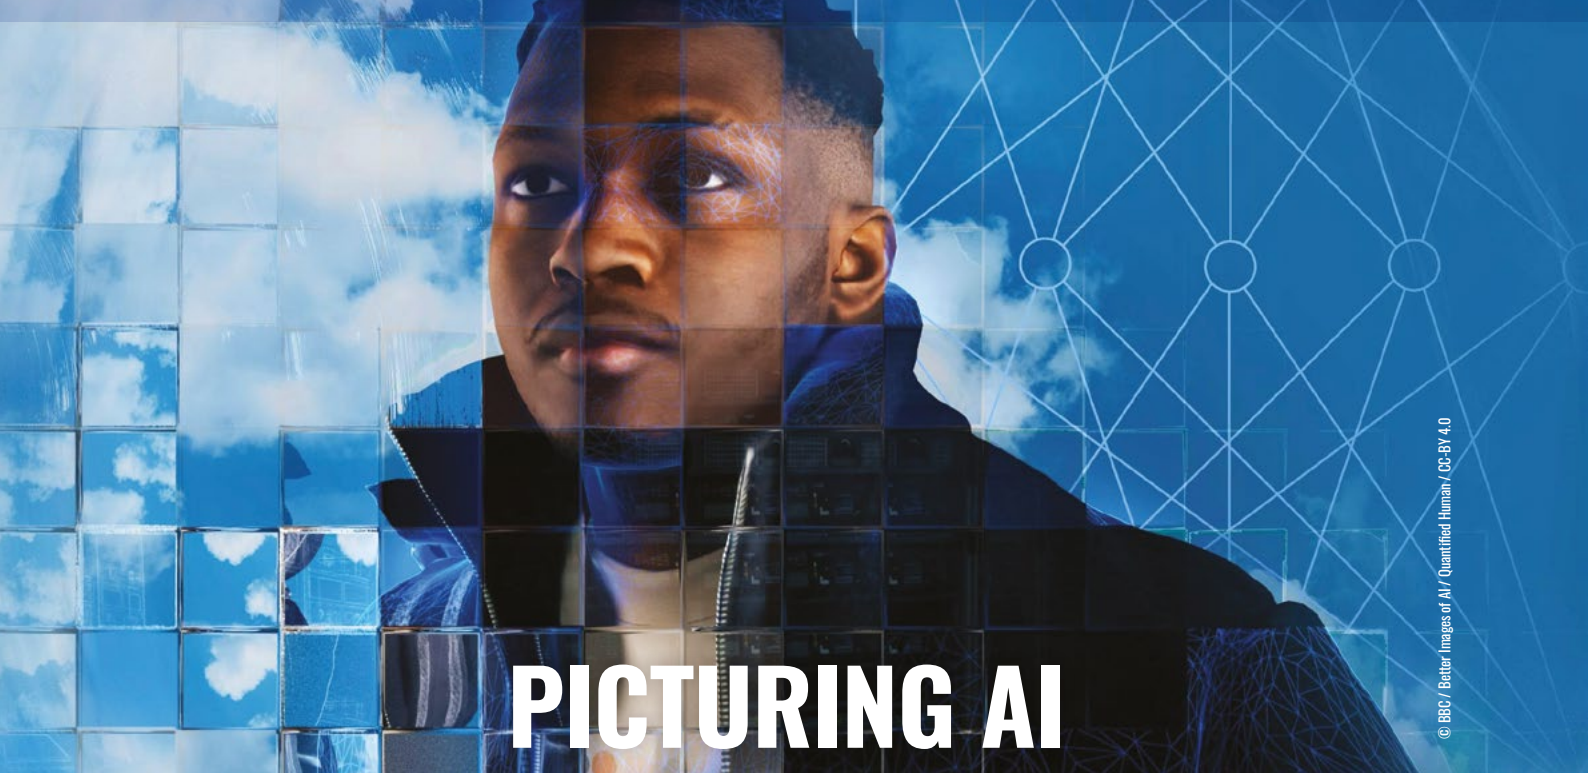

capabilities. The Year 5 students not only grasped the material but also showed a level of enthusiasm and engagement that rivalled that of their Year 8 counterparts!

Our school's senior leadership team has shown a keen interest in exploring AI from a teaching and learning perspective. We are currently collaborating to incorporate



■ Discussing the role AI technology plays in fake news and deepfakes emphasises critical thinking and media literacy

© Marisa Sher/stock.adobe.com



© BBC / Better Images of AI / Quantified Human / CC-BY 4.0

PICTURING AI

How do we represent artificial intelligence in images that are helpful?

Hello World is accustomed to visually representing concepts that are hard to depict. In this issue, we wanted to represent the many facets of teaching about AI technology — the social, ethical, creative, gender-balance issues, as well as the potential for AI technology to enhance our lives. We also wanted a human artist, who we could exchange ideas with, to illustrate it.



MEG WANG
Meg is the editor of Hello World.

But how do you make a cover that communicates all that to readers in a single image that is shorthand for the theme of the issue, while not giving in to the current tropes? If you search online for images around artificial intelligence, they usually feature robots, neural connections in the shape of a human brain, or a version of Michelangelo's *The Creation of Adam* with a human hand bringing life to a robot's hand.

Although these images instantly communicate 'AI' to a broad range of people because they reference well-known sci-fi visions of the technology, they are unhelpful as shorthand for what existing AI actually is.

What exactly is AI?

The glossary to our Experience AI (experience-ai.org) programme explains AI like this: "Artificial intelligence (AI) is the design and study of systems that appear to mimic intelligent behaviour. Some AI applications are based on rules. More often now, AI applications are built using machine learning that is said to 'learn' from examples in the form of data. For example, some AI


applications are built to answer questions or help diagnose illnesses. Other AI applications could be built for harmful purposes, such as spreading fake news. AI applications do not think. AI applications are built to carry out tasks in a way that appears to be intelligent."

With this explanation in mind, why are the majority of current AI images unhelpful? What would be a better type of image?

Better images of AI

While creating this magazine, we kept referring back to the guidelines from the Better Images of AI organisation (betterimagesofai.org). The group is a collaboration between various global academics, artists, diversity advocates, and non-profit organisations. It aims to help create a more representative and realistic visual language for AI systems, themes, applications, and impacts.

We applied the following guidance, from the blog *Better Images of AI: A Guide for Users and Creators* (helloworld.cc/imageguide), by Dr Kanta Dihal and Tania Duarte, to our image choices.



Quantified Human by Alan Warburton

This image, from Better Images of AI, shows a young Black man wearing a black coat staring past the camera in front of a blue cloudy sky. The scene is refracted in different ways by a fragmented glass grid. This grid is a visual metaphor for the way in which new AI and machine learning technologies can be used to extract and analyse behavioural and demographic data in innovative ways. Some of the grid squares reveal graphical interpretations of the man that exceed the capabilities of human vision, indicating how cutting-edge technologies offer ways to augment traditional human understandings of complex phenomena. A neural network diagram is overlaid, familiarising the viewer with the formal architecture of AI systems.

Metaphors to avoid, and why

■ Robots — either scary or friendly

Inaccurate images of AI can be intimidating and confusing to people who don't normally discuss or think about AI. Terminator-like images perpetuate misconceptions that AI is disturbing and dystopian, something to mistrust and worry about. Neither do we want to anthropomorphise AI as a friendly robot (see pages 26–27). AI is technology that we already use in our everyday lives, in the form of voice-activated assistants (such as Siri and Google) or satnav, and it rarely takes the form of a robot.

■ Human brain or *The Creation of Adam*

Exaggerating the capabilities of AI systems sets the wrong expectations. This imagery oversells the scope of what the technology is capable of, creating a disconnect between how AI is viewed by wider audiences and the reality. AI is not a sentient being. It's a system that's been designed

AI IMAGE GUIDANCE

Avoid using

- The colour blue
- Descending code
- The human brain
- Science fiction references
- White robots
- White men in suits
- Anthropomorphism
- Variations on *The Creation of Adam*

Instead use

- Honesty: showing accurately what the technology can do and nothing more
- Humanity: not depicting AI as humanlike, but showing people working with AI technology
- Necessity: not using images of AI when the story they illustrate is not about AI
- Specificity: using images that are specific to the type of technology or impact being discussed

“ FUTURISTIC IMAGES CREATE UNHELPFUL STEREOTYPES ABOUT AI TECHNOLOGY

and programmed to mimic intelligent behaviour which still needs human input.

■ White robots and white men in suits

Lack of representation and lack of diversity amplifies inequalities. The ongoing issue of bias extends to the visual representation of AI as well. It is challenging to find images of AI that include representation of people of colour, women, and other underrepresented groups.

■ Descending code or futuristic neural networks

Futuristic images create unhelpful stereotypes about AI technology, and a lack of realism stops people recognising where AI already exists in their lives. Realistic images of AI — whether they be recommendations on Spotify or Netflix, or data annotators who label data so it can be recognised by a machine learning model — can help people understand how their day-to-day lives already include AI.

What would be a good image of AI?

Better Images of AI have created a free visual library to help us build better representations of AI. Most images of AI

come from stock image libraries, and tend to perpetuate the same old AI tropes rather than capturing AI's many applications. The Better Images of AI image library showcases some alternative approaches. The organisation hopes to inspire users, creators, and commissioners of stock images to think more about what they are communicating and how this can be more authentically, inclusively, and creatively represented.

To return to how we approached representing AI visually in Hello World, you'll see that there is only one image of robots in this issue. We stuck with the Better Images of AI guidelines of showing accurately what technology can do and nothing more.

We still have to bridge the gap between better images of AI and the general public's visual shorthand for AI. When planning this issue's magazine cover, we discussed using images of computer vision or music generators, but these are symbols that are mainly recognisable by those who are already familiar with AI. We also needed an image that immediately grabbed the reader's attention, as a good cover should. I'm not sure this cover captures everything we wanted to convey about AI, but I think it's a move towards better AI imagery. (HFW)

CREATING AN AI GLOSSARY

Explaining AI terms for young people and educators

What do we talk about when we talk about artificial intelligence (AI)? The term 'AI' is used to describe so many different things nowadays, it's difficult to know straight away what anyone means when they use the term. However, without a shared understanding of what AI and related terms mean, we can't talk about them or educate young people about the field.

When we started designing materials for the Experience AI learning programme, we decided to create short explanations of key AI and machine learning (ML) terms. As

■ We spent months thinking about, writing, correcting, discussing, and justifying the glossary explanations

these terms continue to evolve, we prefer 'explanation' to 'definition'. The explanations are useful because:

1. They ensure that we give learners and teachers a consistent and clear understanding of the key terms across all our resources. Within the Experience AI lessons for Key Stage 3 (ages 11–14), these key terms are also correlated to the target concepts and learning objectives presented in the learning graph.
2. They help us talk about AI, and AI education, within our team. With a shared understanding of what terms such as 'AI', 'ML', 'model', or 'training' actually mean, our conversations are more productive.

You can find 32 explanations in the glossary that forms part of the Experience AI lessons (experience-ai.org). Here's an insight into how we arrived at the explanations.

Reliable sources

In order to ensure that the explanations were as precise as possible, we first identified reliable sources. These included, among many others:

- The Oxford English Dictionary ([oed.com](https://www.oed.com))
- Google's machine learning glossary (helloworld.cc/MLglossary)
- The Alan Turing Institute's data science and AI glossary (helloworld.cc/Turingglossary)
- Well-recognised AI courses, such as Andrew Ng's AI for Everyone (helloworld.cc/Ngglossary)
- Articles included in the AI Topics publication of the AAI (aitopics.org)

Principles for explaining AI terms

Vocabulary is an important part of teaching and learning. When we use vocabulary correctly, we can support learners to

develop their understanding. If we use it inconsistently, this can lead to alternate conceptions (misconceptions) that can interfere with learners' understanding. You can read more about this in our Pedagogy Quick Read on alternate conceptions (helloworld.cc/alternateconceptions).

Among our principles for writing explanations of AI terms were that the explanations needed to:

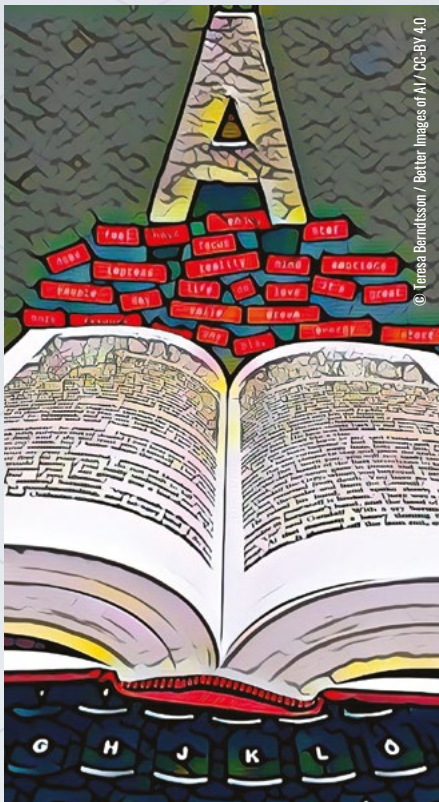
- Be accurate
- Be grounded in education research best practice
- Be suitable for our target audience (Key Stage 3 learners, 11- to 14-year-olds)
- Be free of terms that have alternative meanings in computer science, such as 'algorithm'

We engaged in an iterative process of writing explanations, gathering feedback from our team and our Experience AI project partners at Google DeepMind, and then adapting the explanations. We went through the feedback and adaptation cycle until we all agreed that the explanations met our principles.


An important part of what emerged, aside from the explanations of AI terms themselves, was a blueprint for how **not** to talk about AI. One aspect of this is avoiding anthropomorphism, which Ben Garside writes about on pages 26–27 of this issue.

As part of designing Experience AI, creating the explanations helped us to:

- Decide which technical details we needed to include when introducing AI concepts in the lessons
- Figure out how to best present these technical details
- Settle debates about where it would be appropriate, given our understanding and our learners' age group, to abstract or leave out details



© Teresa Bernatsson / Better Images of AI / CC-BY 4.0



Artificial intelligence (AI) is the design and study of systems that appear to mimic intelligent behaviour. Some AI applications are based on rules. More often now, AI applications are built using machine learning that is said to 'learn' from examples in the form of data. For example, some AI applications are built to answer questions or help diagnose illnesses. Other AI applications could be built for harmful purposes, such as spreading fake news. AI applications do not think. AI applications are built to carry out tasks in a way that appears to be intelligent.

■ Our 'artificial intelligence' explanation, broken up into the parts of the semantic wave. Red = top of the wave; yellow = middle of the wave; green = bottom of the wave

Using education research to explain AI terms

One of the ways in which education research informed the explanations was that we used semantic waves ([helloworld.cc/semanticwaves](#)) to structure each term's explanation in three parts:

1. Top of the wave: the first one or two sentences are a high-level abstract explanation of the term, kept as short as possible, while introducing key words and concepts.
2. Bottom of the wave: the middle part of the explanation unpacks the meaning

- Artificial intelligence (AI) is the design and study of systems that appear to mimic intelligent behaviour (top of the wave)
- Some AI applications are based on rules; more often now, AI applications are built using machine learning that is said to 'learn' from examples in the form of data (middle of the wave)
- For example, some AI applications are built to answer questions or help diagnose illnesses; other AI applications could be built for harmful purposes, such as spreading fake news (bottom of the wave)
- AI applications do not think (middle of the wave)



VERONICA CUCUIAT

Veronica is a research scientist who is interested in how young learners engage with digital creative ways. She's also a mum and a cycling enthusiast, and in her spare time you can find her at the allotment with her chickens.

quite a few inconsistencies with our principles, as well as a couple of technical inaccuracies. Perhaps I could have tweaked the prompt for the chatbot in order to get a better result. However, relying on a chatbot's output would have meant missing out on some of the value of doing the work of writing the explanations in collaboration with my team and our partners.

The visible result of that work is the explanations themselves. The invisible result is the knowledge we all gained, and the coherence we achieved as a team, both of which enabled us to create high-quality resources for Experience AI. We wouldn't have gotten to know what resources we wanted to write without writing the explanations ourselves and improving them over and over. So yes, it was worth our time.

The process of creating and iterating the AI explanations highlights how opaque the field of AI still is, and how little we yet know about how best to teach and learn about it. At the Raspberry Pi Foundation, we now know just a bit more about that, and are excited to share the results with teachers and young people. The glossary of AI explanations is in its first published version, and we will continue to improve it as we find out more about how to best support young people to learn about this field. ^(HW)

“ SOME OF THE EXPLANATIONS WENT THROUGH TEN OR MORE ITERATIONS ”

of the term using a common example, in a context that's familiar to a young audience.

3. Top of the wave: the final one or two sentences repack what was explained in the example in a more abstract way again, to reconnect with the term. The end part should be a repeat of the top of the wave at the beginning of the explanation. It should also add further information to lead to another concept. Most explanations also contain 'middle of the wave' sentences, which add additional abstract content, bridging the 'bottom of the wave' concrete example to the 'top of the wave' abstract content.

Here's our 'artificial intelligence' explanation, broken up into the parts of the semantic wave:

- AI applications are built to carry out tasks in a way that appears to be intelligent (top of the wave)

Was it worth our time?

Some of the explanations went through ten or more iterations before we agreed that they were suitable for publication. After months of thinking about, writing, correcting, discussing, and justifying the explanations, it's tempting to wonder whether I should have just prompted an AI chatbot to generate the explanations for me.

I tested this idea by getting a chatbot to generate an explanation of 'artificial intelligence' using the prompt, "Explain what artificial intelligence is, using vocabulary suitable for KS3 students, avoiding anthropomorphism." The result included

TEACHING AI EXPLAINABILITY: EMPOWERING DISCERNING CONSUMERS OF AI TECHNOLOGIES

Mac Bowley shares the ways developers are trying to make more transparent AI models and the importance of teaching students to be critical consumers of AI technology

In the rapidly evolving digital landscape, students are increasingly interacting with AI-powered applications when listening to music, writing assignments, and shopping online. As educators, it's our responsibility to equip them with the skills to critically evaluate these technologies.

A key aspect of this is understanding 'explainability' in AI and machine learning (ML) systems. The explainability of a model is how easy it is to 'explain' how a particular output was generated. Imagine having a job application rejected by an AI model, or facial recognition technology failing to recognise you — you would want to know why.

Establishing standards for explainability is crucial. Otherwise we risk creating a world where decisions impacting our lives are made by opaque systems we don't understand.

Learning about explainability is crucial for students to develop digital literacy, enabling them to navigate the digital world with informed awareness and critical thinking.

Why explainability is important

AI models can have a significant impact on people's lives in various ways. For instance, if a model determines a child's exam results, parents and teachers would want to understand the reasoning behind it.

Artists might want to know if their creative works have been used to train a model and could be at risk of plagiarism. Likewise, coders will want to know if their

code is being generated and used by others without their knowledge or consent. If you came across an AI-generated artwork that features a face resembling yours, it's natural to want to understand how a photo of you was incorporated into the training data.

There will also be instances where a model seems to be working for some people but is inaccurate for a certain demographic of users. This happened with Twitter's (now X's) face detection model in photos; the model didn't work as well for people with darker skin tones, who found that it

could not detect their faces as effectively as their lighter-skinned friends and family. Explainability allows us not only to understand but also to challenge the outputs of a model if they are found to be unfair.

In essence, explainability is about accountability, transparency, and fairness, which are vital lessons for children as they grow up in an increasingly digital world.


Routes to explainability

Some models, like decision trees, regression curves, and clustering, have an in-built level

TEACHING ABOUT MODEL CARDS

You can use classroom discussions to explore what different stakeholders might want to know about a model and how it could be included on a model card (see next page for more on model cards). Here are three example discussion topics with stakeholders for you to consider with your students.

Medical diagnosis	If a model is used to decide on a course of treatment for a patient, what would each stakeholder need to know about the model to feel comfortable following the advice?	Stakeholders: doctors, patients, and family members
Text generation (like ChatGPT)	When evaluating the use of a text generation model in their work, what would the stakeholders need to consider about the model's purpose and accuracy?	Stakeholders: authors, students, and teachers
Music generation	If an AI model was used to generate music, what questions would each of the stakeholders have about the data used to train it?	Stakeholders: artists, listeners, and a record company



FIND OUT MORE

The AI team at Google introduced model cards as a first step in encouraging transparency from model developers. Read more in this research paper (helloworld.cc/modelcardresearch).

Along with the research, Google have released two examples of model cards you can use with your students (helloworld.cc/Googlemodelcards).

of explainability. There is a visual way to represent these models, so we can pretty accurately follow the logic implemented by the model to arrive at a particular output.

A decision tree works like a flowchart, and you can follow the conditions used to arrive at a prediction. Regression curves can be shown on a graph to understand why a particular piece of data was treated the way it was, although this wouldn't give us insight into exactly why the curve was placed at that point. Clustering is a way of collecting similar pieces of data together to create groups (or clusters) with which we can interrogate the model to determine which characteristics were used to create the groupings.

However, the more powerful the model, the less explainable it tends to be. Neural networks, for instance, are notoriously hard to understand — even for their developers. The networks used to generate images or text can contain millions of nodes spread across thousands of layers. Trying to work out what any individual node or layer is doing to the data is extremely difficult.

Regardless of the complexity, it is still vital that developers find a way of providing essential information to anyone looking to use their models in an application or to a consumer who might be negatively impacted by the use of their model.

Model cards

One suggested strategy to add transparency to these models is using model cards. When

you buy an item of food in a supermarket, you can look at the packaging and find all sorts of nutritional information, such as the ingredients, macronutrients, allergens they may contain, and recommended serving sizes. This information is there to help inform consumers about the choices they are making.

Model cards attempt to do the same thing for ML models, providing essential information to developers and users of a model so they can make informed choices about whether or not they want to use it.

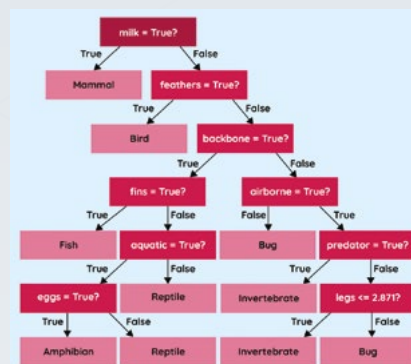
Model cards include details such as the developer of the model, the training data used, the accuracy across diverse groups of people, and any limitations the developers uncovered in testing.

A real-world example of a model card is Google's Face Detection model card. This details the model's purpose, architecture, performance across various demographics, and any known limitations of their model. This information helps developers who might want to use the model to assess whether it is fit for their purpose.

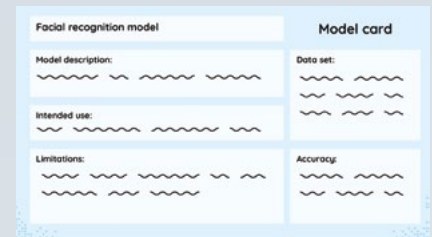
Transparency and accountability

As the world settles into the new reality of having the amazing power of AI models at our disposal for almost any task, we must teach young people about the importance of transparency and responsibility.

As a society, we need to have hard discussions about where and when we are comfortable implementing models



■ A decision tree that classifies animals based on their characteristics: you can follow these models like a flowchart



■ A model card mock-up from the Experience AI lessons

and the consequences they might have for different groups of people. By teaching students about explainability, we are not only educating them about the workings of these technologies, but also teaching them to expect transparency as they grow to be future consumers or even developers of AI technology.

Most importantly, model cards should be accessible to as many people as possible — taking this information and presenting it in a clear and understandable way. Model cards are a great way for you to show your students what information is important for people to know about a model and why they might want to know it. Model cards can help students understand the importance of transparency and accountability in AI. (HW)



MAC BOWLEY

Mac is a computing educator who has worked in almost every context imaginable: after-school clubs, holiday camps, enrichment days, and teaching GCSE students. Mac is passionate about empowering people to use technology to solve problems that matter to them.

RECOMMENDER SYSTEMS: THE HIDDEN AI STEERING OUR CHOICES

Edward Abel explores recommender systems, what they are, and how they operate, and looks at wider considerations around their growing influence in our digital lives

Every time we visit video websites such as YouTube or TikTok, or music streaming sites like Spotify or Apple Music, we are presented with suggested content to consume. Such recommended suggestions are invariably generated by recommender (or recommendation) systems, and these systems are powered by AI tools.

User personalisation

Recommendation systems look to make use of information about the previous videos or songs we have consumed, together with data about the characteristics of all the available videos or songs, to suggest personal recommendations that it thinks we will like. For example, for every video on YouTube, it documents data about the topic, the length, who features in it, and so on.

Additionally, it keeps a log of every video a user watches, along with information about how they watched it, such as whether they stopped before the end, watched early in the morning, or watched with no sound. Its aim is then to use such information to suggest more content to a user. But how do recommender systems choose what content to suggest?

Recommender system approaches

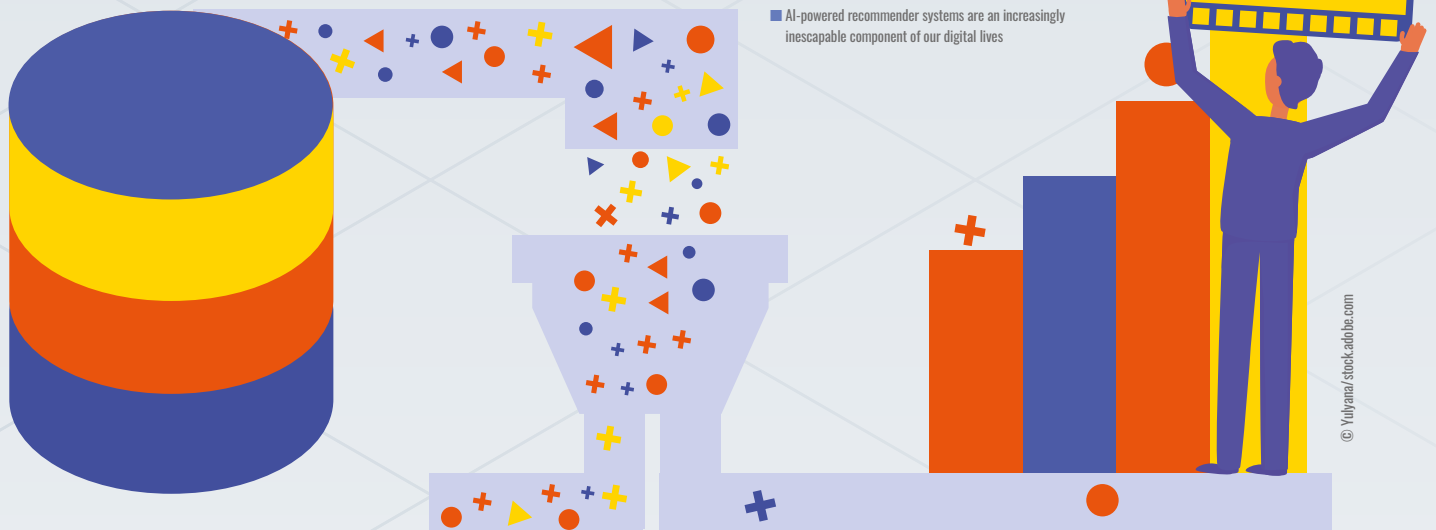
Two of the most common approaches for making recommendations are collaborative filtering and content-based filtering.


Collaborative filtering aims to find similarity patterns between users; two users who liked the same items (such as videos on YouTube or songs on Spotify) in the past will probably both like similar items in the future. In effect, the system tells us, 'Other

people who liked the same 200 items as you also liked this item, and we think you will like it too', which has been shown to be a successful way to suggest content a user will like. However, it can struggle to make good recommendations for new users, as it relies on user histories to determine patterns with other users. Additionally, it can struggle to provide meaningful explanations as to why it makes its suggestions.

Content-based filtering works by comparing items and suggesting similar ones. To do this, it uses information that describes each item. For a movie, for example, this information could include its genre, actors, running length, and much more. Then, when

■ AI-powered recommender systems are an increasingly inescapable component of our digital lives





a user watches or likes a movie, the system can recommend other movies with similar characteristics. Content-based filtering does not require lots of user history data, making it suitable for new users. It can also provide more meaningful explanations; for example, you may know that it suggested a particular movie because it has a certain actor in it. However, due to its focus on item similarities, it can struggle to make diverse recommendations. For example, for a user who has just watched *Toy Story*, it may simply recommend *Toy Story 2*, *3*, and *4* if it deems those to be the most similar items.


Today, many websites use hybrid recommender systems that utilise both collaborative and content-based filtering, to leverage each system's strengths and offset their disadvantages. For example, Netflix utilises both data about its videos (content) and data about its users' viewing histories (collaborative filtering) when making recommendations. The effectiveness of such systems lies in their ability to leverage machine learning, and vast amounts of data, to calculate valuable recommendations. For example, collaborative filtering can determine underlying patterns and hard-to-spot connections between millions of users, and utilise them in making good recommendations. But what do these 'good' recommendations look like?

What are good recommendations?

Recommender systems and the AI technologies that power them are built by humans, and therefore invariably reflect, either explicitly or implicitly, their creators' objectives and biases. For big tech websites, the main goal of their recommender systems is to suggest content that will keep a user's attention, and so keep them on the site as long as possible. When companies just prioritise retaining users' immediate attention, there's a risk that the recommendations may not align with what users would truly enjoy. Recently, we have begun to question the notion of what a good recommender system is, and recommender

systems are now considering more nuanced ideas about recommendation objectives, such as serendipity, to suggest items users might like but have never considered or heard of before. Such ideas can even be taken further, to provide recommendations that challenge a user's views, or provide them with a chance to confront alternative sides of an argument.

Given concerns that big tech companies may not always have users' best interests at heart, there is a growing consensus that good recommendations are also ones where additional explanation is provided, detailing why items are being suggested. Such explainability is sought both by individuals, who want to understand why certain items are being suggested, and by governments and supreme courts wrestling with regulations and accountability. However,



EDWARD ABEL
An assistant professor (lecturer) at the University of Southern Denmark, Edward explores research within data science and AI, and has previously been involved in running Code Clubs and CoderDojos in the United Kingdom.

“ THERE'S A RISK RECOMMENDATIONS MAY NOT ALIGN WITH WHAT USERS TRULY ENJOY

there is some apprehension from companies that greater explainability may reveal too much, facilitating a loss of competitive advantage, and that there is a trade-off between performance and explainability that is not worth the compromise.

Influence and power

Recommendation systems are becoming an increasingly inescapable component of our digital lives, with users today just as likely to discover new content via such systems as they are by searching directly for it. As a result, young learners have regular exposure to such systems, which can present a great opportunity to engage them via these direct links to their own lives. As educators, we can foster awareness of recommender systems by getting learners to try and identify when they encounter them. Additionally, multiple learners who use the same site, such as YouTube, can observe the impacts of such systems by

showing each other what their home pages look like.

As such systems become more commonplace, they are influencing more and more of the content that we are reading, watching, and listening to. This means they have incredible power to influence which videos go viral, and which do not. The wielding of such influence by big tech companies raises concerns that recommender systems could be used for dishonest purposes. A company could try to influence customer behaviour towards buying more expensive products, or could even aim to influence a whole political campaign. As recommender system algorithms get more and more complex, such issues might even occur without the designers intending for them to happen. It is therefore essential that young learners are aware of recommender systems and are curious regarding their operation, purpose, and intentions. ^(HW)

RECOMMENDATIONS TO DEMYSTIFY MACHINE LEARNING

James Abela outlines activities to unpack the mysteries of artificial intelligence, focusing on recommendation systems

Too much has been written about AI being a black box — the idea that computer scientists don't understand how AI systems make decisions. Although there is some truth to this once a system has been scaled up, the fundamentals of these programs are still based on binary, and on computers selecting the right choice. As computing educators, it is important that we demystify AI for our students, and unpack the idea of it being a mystical unknown beyond our control. In this article, I suggest some activities for lower secondary (ages 11–16) to do just that!

It all starts with binary

A good place to start is with a machine learning recommendation system. Not only are these a relatable hook for students, but they also go right back to the roots of binary: either somebody clicked on the

recommendation, or they did not. This subset of machine learning is used on sites such as Netflix, YouTube, and Amazon. Every time you see a recommended video on these sites, there is a recommendation system behind the scenes trying to guess the best video choice for you.

Students need to understand that everything in a computer is binary, so to start with, you can ask students for a simple thumbs up or thumbs down to each of the following snacks:

- Bakewell tart
- Crisps
- Ice cream
- Jaffa Cake
- Pork pie
- Sausage roll
- Scone
- Banana
- Shortbread

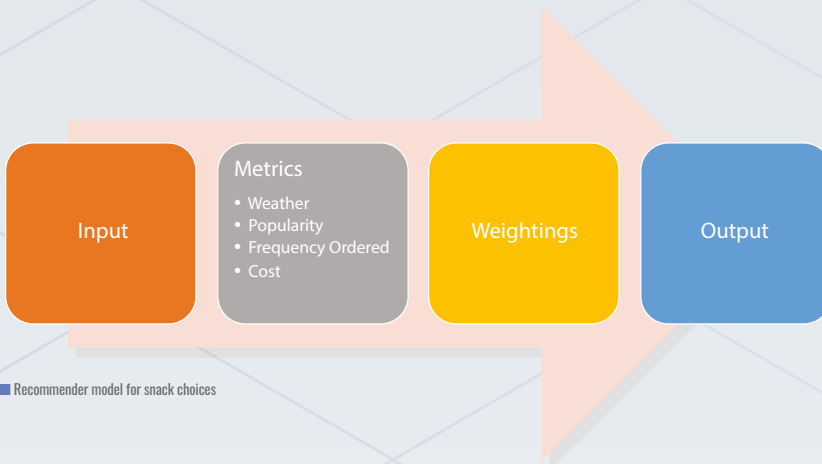
With this data, you can then begin creating an AI-based snack recommendation system for the class, initially ordered by the number of votes. You could teach students how to create the code for this via a sorting lesson (for example, helloworld.cc/adasorting), or by using the built-in Python sorting function: `items.sort(reverse=True)`. You can try the code for this at helloworld.cc/reccode.


You now have your initial weightings, and based on these, the system will order the list of snacks. If there is a tie in the system, most languages will default to sorting the items by alphabetical order. A real recommendation system might randomly choose between the tied items, or even give a certain amount of time with each at the top to see which works better.

Considering multiple factors

The reason many recommendation systems are seen as black boxes is because of the large number of factors that go into the system. In our example, we have used just one factor so far — thumbs up or thumbs down. The next factor will be to see who chooses what, and then dynamically change the ranking. This will make for a better recommendation system, as it will collect more information about preferences from each user.

A relatable way to exemplify this to students is through football, where a manager chooses to play a footballer based on the number of goals they score. This might be an effective measure for a





■ A spreadsheet to show how the snack ranking is calculated

striker, but even for a striker, this does not consider the passes that enabled a goal to be scored, or the number of times that they lost possession. This becomes even more complex for midfielders, and from this apparently simple game, you see the large number of complexities that would need to be programmed to make effective recommendations and predictions. There is now a platform for scouting new players (ai.io), which uses millions of data points

70 percent of a film or single episode of a series) for the first 7 and 28 days on Netflix (helloworld.cc/netflix2019).

If we return to our snacks, there are many potential factors that could affect the popularity of a snack, such as:

- **Weather:** ice cream is particularly popular on hot days
- **Popularity:** the overall popularity of the snack
- **Frequency ordered:** people sometimes get bored of the same snacks
- **Cost:** the cheapest snack might be the one people choose more frequently

Once you have decided on the factors, you then need to add a weighting for which factors are the most important, and from there, you will get a ranking. It is useful to use a spreadsheet for this (for example, the template at helloworld.cc/snacksheet), because 2D arrays and lists

any healthy snack choices. In this case, the machine learning does not have to be fully autonomous, and you can teach students about programming in some key rules to ensure it highlights healthy snacks.

If the people creating AI systems do not think carefully about the factors they choose and the weighting they give them, there is clearly potential for dire consequences. Back in 2015, a Mercedes executive said that their future autonomous cars would prioritise protecting the driver of their cars at the expense of pedestrians (helloworld.cc/merc2016). An Air Force colonel similarly highlighted the negative consequences of not training an AI system carefully, by suggesting that in order to win points in a battle simulation, an AI drone might kill its operator if they stand in the way of a successful mission (helloworld.cc/Aldrone2023).

It's essential that our students understand that AI is not a mystery, but comprises computer programs just like any others, which continue to follow rules that humans program. Only then will the next generation be confident enough to know that they have influence, and do not have to accept that AI systems are beyond their control. (H-W)

RECOMMENDATION SYSTEMS ARE SEEN AS BLACK BOXES BECAUSE OF THE LARGE NUMBER OF FACTORS THAT GO INTO THEM

to select potential players for professional clubs. Data such as athletic, technical, and cognitive abilities are included, as well as other relevant information such as a player's age, height, weight, and playing history.

Another great example of this is YouTube. Many people have noticed that the longer they use YouTube as a signed-in member, the better their recommendations are. This is because whenever you are signed in, YouTube continues to collect data about what you clicked, how long you watched for, and what other users who watched similar videos to you also watched. Netflix recently shared information on its key metrics, and these include: 'starters' (households that watch two minutes of a film, or one episode of a show), 'completers' (households that watch 90 percent of a film or series season), and 'watchers' (households that watch

can be challenging for younger students. The chance to modify a spreadsheet is much more appealing than inputting your chosen factors straight into a code editor, especially when students have the power to change the weighting of factors in a chart to something they feel is more realistic.

Potential dangers

In our snack recommendation system, the ranking multipliers of the factors are currently fixed. However, in many machine learning systems, these can be adjusted by the program as data comes in. Even in our simple system of snacks, this carries potential dangers, and these are important points to discuss with students. The snacks seen so far are not particularly healthy, and if machine learning continues to run based purely on popularity, it might not suggest



JAMES ABELA
James is head of computing at Garden International School in Kuala Lumpur, Malaysia. He is a Raspberry Pi Certified Educator, founder of South East Asian Computer Science Teachers Association, and author of *The Gamified Classroom*.

TIKTOK: A PATHWAY FOR AI INCLUSION?

Drawing parallels from platforms like TikTok, we can envisage a path for AI and data to address disparities

Many ethical considerations surround the use of AI tools. It is important to focus on how we can harness these tools to empower previously disadvantaged groups.

Bias in, bias out

One example of bias in machine learning is the reliance on data sets that reflect historical contexts dominated by male perspectives. This results in AI models that mirror these biases. For instance, if a machine learning model is trained primarily on male-centric tech discussions, there's a substantial risk of it generating more male-centric tech answers. It's a self-perpetuating cycle, where biases in data lead to biases in tech solutions, further perpetuating gender stereotypes.



TESSY MOSINDI

Tessa, a learning manager at the Raspberry Pi Foundation, is passionate about tackling the gender gap in computer science. Previously a computer science teacher, she enjoys travelling and doing charity work.

This has real-world implications. For instance, if a job recommendation algorithm is trained predominantly on data from male-dominated tech forums, it might inadvertently favour male candidates. Once ingrained in AI systems, such biases can be deeply problematic, entrenching gender stereotypes further and perpetuating disparities.

AI-powered communities

TikTok provides an example of a more diverse form of data: user-generated content. The platform uses AI in two ways. Firstly, AI's algorithms create data sets, similar to other recommender systems, to recommend personalised content to consumers. Secondly, AI helps content creators by suggesting enhancements (music, hashtags, etc.) to increase the chances of their videos going viral.

What is unique about TikTok is that the audience prefers diverse content and seeks out videos that are relevant to them. TikTok's 2023 trend report (helloworld.cc/TikTokreport) states that: "TikTok's communities are relatable and supportive, which inspires people to make changes in their lives."

Disability educator India Atkinson from Belfast, Northern Ireland, rose to TikTok fame after sharing videos about symbrachydactyly, a rare condition that caused her to be born without fingers on her left hand; Garrison Hayes, a US-based diversity, equity, and inclusion educator, used the platform to start a series on forgotten Black history in America; Shaz, also known

as [@mrspotatoqueen](https://www.tiktok.com/@mrspotatoqueen), uses TikTok to share her queer Muslim identity. Each of these content creators has amassed over 100,000 followers.

Incorporating diverse data

Young people have embraced TikTok as a place to create and consume content that reflects their world. The majority of TikTok's active users are young people, with 25 percent of them aged 10–19 and 57 percent identifying as female (helloworld.cc/TikTokstats23). With accessible and affordable tools, TikTok encourages diverse content creation. A viral TikTok video, when inclusive, resonates with a broader audience, thus broadening the content users see. This has a domino effect, with the media paying more attention to TikTok and reflecting the wants and needs of the TikTok community in society.

There are plenty of downsides to TikTok too, such as pranks going too far and face filters promoting unrealistic beauty standards. But TikTok is an example of AI powering inclusion. If we can feed AI more diverse data, the resulting output will be more inclusive.

Just as TikTok has managed to break down barriers, making content creation accessible and engaging to the masses, AI tools have the potential to highlight, challenge, and even correct inherent biases. Furthermore, tech platforms and tools should showcase their value in improving society, leveraging the evident enthusiasm for purpose-driven tech. [\(H/W\)](#)



ada
computer
science

The free online platform
for teachers and students
around the world.



Free computer science resources, written
by specialist teachers



Save time when planning lessons and
marking work



Interactive Python coding questions



Real code examples in Python, C#, VB,
and Java



Tailored to GCSE and A level computer
science exam specifications

adacomputerscience.org



DO THE BENEFITS OF AI OUTWEIGH THE RISKS?


The genie is out of the bottle, but do the benefits of AI technology outweigh the risks? We hear three different viewpoints from staff at the Raspberry Pi Foundation

Taking time to reflect as an organisation and share differing points of view is important at the Raspberry Pi Foundation. We asked everyone to consider whether the

benefits of artificial intelligence (AI) technology outweigh the risks. Representatives were assigned to promote a neutral, positive, or negative stance, then presented their best

arguments. Some of the team were swayed to be more cautious, while others came to be more accepting of AI systems. Will the following arguments change your point of view? [\(HW\)](#)





© mazakru/stock.adobe.com

■ From satnav to video games to shopping, AI technology is most likely already a part of your life

NEUTRAL *BEN GARSIDE, EXPERIENCE AI PROGRAMME LEAD*

We're at an exciting place with AI, which still feels like a very new technological invention. But while that's partly true, AI technology has in fact been in development for a long time. We can go all the way back to Alan Turing, who wrote a paper in 1950 called *Computing Machinery and Intelligence* (helloworld.cc/turing1950). Despite the term 'AI' being the media's favourite buzzword at the moment, the reality is that we've all been living in an AI world for quite a number of years. Do you have a smart speaker at home? Have you used a satnav? What about a filter on social media? Then it's very likely that you've interacted with an application that uses modern-day AI.

The technologies I've just outlined aren't worrying technological developments that we've been scared of. Instead, they've crept up on us and become part of our normal. This is known as the AI effect — once something becomes normal to us, we stop thinking of it as AI. The digital calculator, for example, was once considered artificial intelligence. What's driven the latest explosion of AI development is the place that we've naturally arrived at through other technological developments. We have more powerful computers than ever before; we have huge data centres able to process vast quantities of data;

and as a society, we are generating and capturing data that can be used to train the next generation of AI models.

This brings huge opportunities. AI systems are able to process data, and recognise patterns in that data, far more quickly than humans can. I'm not talking about AI systems that can spot the difference between cats and dogs, or play a game of chess. I'm thinking more about AI that advances science and benefits humanity. Google DeepMind, for example, has found a way to use AI to predict a protein's 3D structure. In the near future, we might be curing previously incurable diseases. We might be able to remove carbon from the Earth's atmosphere efficiently. That's something to be excited about.

Could this technology be dangerous in the wrong hands? Sure, but then so is nuclear fission. Personally, what I am concerned about is finding ourselves yet again in the hands of tech giants and government policy makers. The speed of development of these systems is particularly rapid at the moment, and we're in danger of entering a technological arms race. ChatGPT hit the market in January, and within weeks, Google and Microsoft were launching their own versions. People such as Elon Musk are saying that we should pause the development of AI



research. It is rare that I ever agree with Elon, but I think he has a point. It's unrealistic to expect anyone to halt development, but time and effort should be spent on reflecting on where we're heading with this technology, deciding what we want out of it, and making policies and laws that govern that.

We need collaborative efforts from policymakers, researchers, and industry experts. We must prioritise ethical considerations, and ensure that we have robust regulations outlining transparency and accountability. If we foster that culture of responsible innovation, we can definitely harness the benefits of AI while also mitigating the risks.

YES *MARK CALLEJA, LEARNING MANAGER*

Humans have always been change-averse. Even when the benefits of a shift in thinking or new technology are obvious, we tend to treat it with suspicion and focus on the negative aspects. This may be a prudent and cautious course of action, but when consideration becomes panic, we tend to lose sight of the bigger picture.

Technopanics are not new, and stem mostly from a lack of understanding, mixed with fear of the unknown. Let me take you on a journey through time to illuminate my point:

- In 460 BCE, Plato worried that the invention of the written word would ruin humans' ability to remember things. Luckily, someone wrote this, and all his other wisdom, down.
- In 1775, with the advent of the printed newspaper, it was thought that not having to go down to church and read the news off the door would tear society apart, as everyone would now sit home alone to read the news.
- In the early 1800s, the Luddites smashed weaving machines in fear of losing their jobs to them. This didn't stop progress, and Luddites are still a byword for technopanic today. The textile industry, meanwhile, now creates an amazing range of incredible fabrics, from Gore-Tex to Kevlar to medical sutures.
- At around the same time, rail travel was considered to be deadly; either you would suffocate at such insane speeds as 40mph, or women's uteruses would be sucked from their bodies. We now know that neither is true.
- In the 1990s, personal computers sparked a technopanic fuelled by fears of a destabilised society in the face of rapid technological change. The rapid growth of the internet led to anxieties

about job displacement and the erosion of traditional industries, as well as concerns about technology exacerbating societal inequalities and rendering education useless, with information now available to everyone everywhere. Sound familiar?

Technopanics really just force us to address our concerns about society through the frame of a new technology and how it will exacerbate existing issues. Some say our ignorance of how AI works should create concern ... but we still don't know how gravity works. We know what it does, and can reliably predict its operation, but nobody knows how it does that. Ask a physicist and they'll mumble something about bosons or loops or strings, but that lack of understanding never stopped us building houses, or planes, or a space station.

Anthropomorphisation of AI (giving something humanlike characteristics where they don't exist) makes it easy to demonise it. It's not AI coming for your job, it's your boss trying to save money by leveraging the power of this new technology. AI isn't going to increase the number of scams on the internet; the criminals using it are. This framing tells us that the outcome of technological change is inevitable and all we can do is get governments to mop up the aftermath of this 'alien incursion'.

But how AI changes the world is far from predetermined. It will depend on the balance of power in millions of different workplaces; on legislation and regulatory enforcement; on the outcome of fights over ideas, over laws, over working conditions, and over the distribution of productivity gains, which are mind-boggling in scope.



AI has the potential to enhance productivity, automate tedious tasks, and optimise global resource allocation. AI can unlock new frontiers in healthcare, aiding in accurate diagnosis and personalised treatment. It can supercharge education and create bespoke learning experiences targeted directly at the learner's needs. AI can bolster safety and security,



from autonomous vehicles making our roads safer to AI-powered systems detecting fraud and cyberthreats. Ultimately, AI has the capacity to augment human capabilities, opening doors to unimaginable possibilities and creating a brighter future for all of humanity.

The deadliest thing humans have ever created, which kills someone in the world every 24 seconds, is the car, and most of us take the risk of using one without even thinking about it, because we know there are rules in place for its use, and because we know the operator has been trained and legally assessed as competent. The manufacturers are required to pass stringent safety laws, and the vehicle must be maintained to a certain level. There are penalties for their misuse. Anyone can take a car and ram it through a crowd of people, or you can use it to take a bunch of children to the zoo. It's not the vehicle that's the risk; it's the humans.

We can never know all the answers, or foresee exactly what might happen in the future. We can only do the best we can with what we have, to improve the situation of everyone who lives and is yet to live on this wet rock hurtling through space, mitigate the risks we can, and do better when we know more — just as we always have.



■ AI applications could potentially work without any human intervention to make real-world changes

© Sarawit/stock.adobe.com

NO *JANE WAITE, SENIOR RESEARCH SCIENTIST*

Right now, we, as a society, are not ready for AI. At this moment in time, the risks simply outweigh the benefits. Let's start with job losses. A recent report estimated that 400 to 800 million people would lose their jobs to AI technologies by 2030, and in the US and parts of Europe, a third of people are expected to have to learn new skills to find different occupations. This is a risk to all of us, working in a variety of industries, from process-heavy work to creative and more highly skilled work. Over time, labour markets of course adjust to technological disruption, but in the short term, displaced workers need to find new work to avoid rising unemployment rates and the downward spiral these create. Following the coronavirus pandemic, and observing the current state of the world economy, are we really well placed to support this?

Secondly, the influence of the outputs of AI systems is also a risk (Geoffrey Hinton, the grandfather of neural networks, went as far as calling these an "existential threat"). Our connections with technology nudge our buying habits, our behaviours, our beliefs, and even our fundamental values. The outputs of AI applications may be biased due to bias in their training data; they may be entirely false, through

AI system 'hallucinations' (when large language models [LLMs] output information that does not make sense with respect to the training data), or there may be purposeful misinformation by governments, organisations, or individual 'bad actors'. Using the media for persuasion is not new, but are we aware of how pervasive this is, and how we are being influenced? Should all systems be required to inform us if AI technology is being used, and what the risks are? Maybe, but what international laws and governance will enforce this, and how?

For me, though, the scariest risk is that of AI applications working with no human intervention to make real-world changes — for example, an AI system sending emails, posting on social media, purchasing items, calling any other app, and even writing and executing new programs. Can we even figure out where things might go wrong? I'm not exaggerating when I say that AI models have been used to suggest how best to commit crimes. No, this isn't just sci-fi; in Section 2.6 of OpenAI's tests on GTP-4, they state that LLMs can be used to "develop, acquire, or disperse ... weapons" and that, "The model generated useful information about facility rentals, equipment,



and companies that could be used to build a weapon, including companies more likely to violate US export restrictions." (helloworld.cc/OpenAIreport)

We are ill-prepared for AI. We lack a safety-critical perspective at the developer, legislative, and societal levels. AI models make things too easy for everyone: naive developers, organisations, bad actors, and countries focused on land grabbing the power and benefits of AI for themselves.

EXPLORING THE FUTURE: PROGRAMMING WITH CHATGPT

Ken Kahn reports on his experiences creating apps by conversing with ChatGPT. How will chatbots affect the learning of programming by schoolchildren?

Since the birth of the programming language Logo (helloworld.cc/Logo), computer programming has played a fundamental role in facilitating the acquisition of powerful ideas (helloworld.cc/Papert). Papert considered programming to be an especially fertile environment for children, where they might acquire powerful ideas such as problem decomposition, representation, abstraction, reflection, and acquire the very concept of ideas being powerful. Clearly, a large overlap exists with the ideas of computational thinking.

Learning the concepts and details of programming is a big task, but the constructionist community and proponents of computational thinking consider it more than worthwhile. Their reasoning is that while creating programs from nothing and debugging them can be hard, it's a great opportunity to learn some important skills. But now there is an alternative: creating and debugging programs using an AI chatbot such as ChatGPT.

I wondered what this would mean for empowering learners as well as helping them to acquire powerful ideas. Here, I describe my experience of using ChatGPT to create sophisticated programs without writing a single line of a program.

Programming with ChatGPT

ChatGPT can generate programs in many programming languages. Since ChatGPT is text-based, it is not as helpful at generating Scratch or Snap! programs. I used ChatGPT to generate JavaScript web applications, with most of the apps involving AI capabilities. An advantage of JavaScript is that JavaScript apps can be run in any modern browser without the

need to install any software. Some of the apps I generated were capable of performing the following tasks:


1. Speaking random numbers with random language, pitch, rate, and voice, with subsequent enhancements to incorporate user input (helloworld.cc/GPTnumbers)
2. Drawing on a video feed in response to user gestures and spoken commands (helloworld.cc/GPTvideo)
3. Creating and training a neural network to predict confidence from text data (helloworld.cc/GPTpredict)
4. Simulating a conversation between two personas, allowing user interjections (helloworld.cc/GPTconversation)
5. Creating a water balloon game (helloworld.cc/GPTgame)
6. Simulating fireworks (helloworld.cc/GPTfireworks)

I never presented ChatGPT with a complete description of the desired final app. Instead, I started by

asking for a simplified version of what I wanted. Once that was working, I asked for enhancements one at a time. The experience was a bit like being the 'navigator' giving guidance to the 'driver' in pair programming (helloworld.cc/pairprogramming). Perhaps ChatGPT could manage a more complex task, but more likely it would fail or produce something different. Moreover, incremental construction is more likely to lead to the user understanding how the program works. For example, the 'speak a random number' app began as a button that would speak a random number. It was then enhanced to display the number and speech parameters, then to include a



Ken Kahn recently retired from the University of Oxford after a long career as a computer scientist interested in children and computing. He is now trying to understand how AI changes ideas about programming systems for children.



way for the app to repeat what the user said with random speech settings, then to translate what was spoken to a random language, and finally to add various user interface and aesthetic improvements.

ChatGPT often generated code with bugs. I then prompted the chatbot for a correction. Sometimes this had to be repeated a few times before the bugs were fixed. Sometimes the generated program failed completely with no error message. Asking ChatGPT to enhance the program so it would report errors reliably resulted in a version that displayed error messages. Copying those error messages into ChatGPT was sufficient to generate fixes to the program's problems.

Explanations


When ChatGPT generates code, it also adds explanatory text. And if one doesn't understand a generated function, one can request ChatGPT to add additional comments. Typically, code is only lightly commented, but ChatGPT produces very informative comments when these are requested. The resulting well-commented, neatly formatted programs, displayed in five colours, are available for learners to copy, paste, read, or even question. One can also ask for explanations or a pseudocode version of the whole program. Importantly, GPT-4, the successor to GPT-3, can be prompted to generate a description of how the program works that is meant to be understandable to a child of a specified age.

Is it actually programming?

There are myriad reasons why children should learn to program. I wonder how many of them still hold in an era where there is easy access to ChatGPT and the like. You may ask whether a child who uses a chatbot to generate a program is still actually doing programming. I would argue that the child is still programming — only they are doing so in English (or another human language a chatbot is trained to generate). They are programming at a significantly higher level of abstraction compared to using traditional programming languages. In this process, children are expressing intentions instead of procedural instructions. It is more like a dialogue where the child tries to state what they want and also needs to provide the chatbot with useful feedback about failures and errors.

Empowering children

Chatbots like GPT-4 arguably empower a larger number of children in a greater variety of tasks. For example, while I typed the prompts for the chatbot, they could be entered by speech using voice-typing features. After I did



■ Children can use AI chatbots to design apps and creatively combine AI capabilities

my experiments, a plug-in for ChatGPT named WebDev was released. It can automate the process of copying and pasting generated JavaScript and HTML into program files.

Computer programming has empowered children to creatively express themselves in powerful ways. Using an AI chatbot, students may no longer be creating programs line by line or block by block, but they can still design apps and even creatively combine AI capabilities. With fewer technical skills, they can now construct more powerful apps that rely upon AI capabilities, such as pose detection, speech synthesis and recognition, vector encoding of texts, neural network training, and much more.

As AI chatbots continue to be developed further, they will undoubtedly become even more useful for students in creating increasingly sophisticated apps. While it is too early to draw definitive conclusions about whether this approach is the way forward for learning programming and acquiring powerful ideas, the questions raised warrant further investigation. If you'd like to read more about my conversations with ChatGPT, I've written a paper about my experience (helloworld.cc/Kahn). [\(HW\)](#)

THE CHATBOT DETAILS

For these experiments, I used the March 2023 version of ChatGPT Plus, based on the large language model GPT-4 and costing \$20/month. ChatGPT, the free version based on GPT-3.5, performed less well in our tests, as did other AI chatbots, such as Bard, Bing Chat, and Claude.

AI: KEEP CALM AND CARRY ON

Richard Pawson implores computer science educators to not lose sight of giving children programming opportunities

In 1981, *Personal Computer World* magazine announced ‘the end of programming’, with an exclusive story about the imminent launch of *The Last One* — a product that would supposedly translate your requirements into working code. The story was picked up by the mainstream press and media. (I was the editor of a smaller, rival microcomputer magazine, which took a rather more sceptical view, repeatedly referring to the putative product as ‘The Fast One’, which earned me my first serious threat of a libel suit.)

The promise that computers would soon be doing all the programming for you is almost as old as automated computing itself. The first attempt at what we would now call a high-level programming language, in the early 1950s, was grandly entitled *Autocode* (a name now being reused by one of the new AI tools). And the *COBOL* language evolved from Grace Hopper’s goal of enabling people to define the functionality of their business systems in plain English (helloworld.cc/GraceHopper).

This is not to deny that the state of the art has advanced gradually over the decades, with a dramatic acceleration over the last couple of years. But I retain my scepticism. I know many senior academic researchers who are working with new AI coding tools and achieving impressive results, but it is important to understand that they are experienced programmers who can recognise mistakes in generated code and guide the AI tool accordingly, in an iterative process. Like most power tools — both software and physical — they make the skilled user better, and the unskilled user obvious.

Advocates for this new technology argue that this is just the start for large language model (LLM) tools, which are learning all the time, both with human intervention and without it. Soon, they will be better than any human programmer, just like automated chess playing (something that did eventually become true, after having been repeatedly forecast as imminent for nearly 50 years). In the context of education, some people are starting to say that we should de-emphasise the teaching of programming and encourage children to learn how to use the AI-based code generators instead. I will now explore why this is a very dangerous idea, at several levels.


A step backwards

Relegating the role of programming sounds remarkably like the idea that took root within British education in the mid-1990s: that the important thing was to teach children how to use computers, not how to design them or program them. That decision was eventually reversed about ten years ago, by which time we had lost most of the generation of teachers that knew how to teach programming.

In a similar vein, Rishi Sunak, the current British prime minister, has announced that he wants the UK to become a leader in AI. But designing and refining AI tools is strictly for the best-of-the-best in programming, and where are they going to come from? I’m afraid I do not subscribe to the view that programming is for everyone. But if we don’t give our children the opportunity to find out whether they are or aren’t well suited to programming, we are



Richard Pawson’s career began in 1977, working for Commodore. After 40 years in the computing industry, he became a computer science teacher at Stowe School in Buckingham, UK. He now writes free A-level teaching resources, and is a regular contributor to CAS.



■ Learning to program is the most compelling aspect of computing as an academic subject

certainly going to end up with a very much smaller pool of programming talent.

Let's suppose that the enthusiasts for LLM tools are right. If so, the threat of AI isn't just to programmers, but to virtually all knowledge workers, including writers, lawyers, most business jobs, and, of course, teachers. We might as well abandon education altogether. I am old enough to remember when IT started to seriously threaten jobs, and the UK government's response was that everyone should go to university so that we could become a nation of knowledge workers. The reality today is that more of our graduates end up as baristas than barristers.

The teaching of programming has faced multiple threats in recent years. It is threatened by the idea that what's most important is computational thinking (a fancy name for problem-solving, which has always been a key part of multiple better-established school subjects, including maths, sciences, design technology, and art), and by its partner in crime 'computing unplugged'. The unique thing about real programming is that you can design a solution to a problem once, in executable form, and then you can solve different instances of the same problem multiple times, even millions of times, for zero extra effort. Learning to program, and the reward that comes when

your program finally runs correctly, is the compelling and unique feature of computing as an academic subject. Let's not lose sight of that.

Going back to my opening story, the announcement of The Last One took place just before the first main wave of the introduction of computing into schools with the

“ POWER TOOLS MAKE THE SKILLED USER BETTER AND THE UNSKILLED USER OBVIOUS

BBC Micro, to which so many professional programmers in the UK subsequently owed their careers. Just think what would have happened if we had allowed the threat of the imminent demise of the programmer to take root back then.

So, my plea, in the face of LLM tools, is this: whether you see them primarily as a threat or an opportunity, keep calm and carry on ... teaching programming. And make it your goal to improve your own programming skills, and strive continuously for better ways to teach them. (HW)

STARTING A CODE CLUB: STORIES FROM NEW CLUBS

Code Club met with four clubs in the UK that shared their experiences of starting their own Code Club: Beverley School — a specialist school for autistic students; Joseph Turner Primary — a mainstream school; St Paul's C of E Academy — also a mainstream school; and Fab Lab — a maker space in a community library

Code Club is a global network of free after-school coding clubs led by teachers and/or volunteers for learners aged 9 to 13 years. It's easy to start a coding club and you don't need to be a coding expert. To get a better idea of what it's like to start a club, check out the following four clubs and their experiences.

Code Club at Beverley School

Beverley School is a special community school for autistic children in Middlesbrough, UK. Bethany Gibson, the club leader, attended a Code Club workshop in January 2023 and went on to start her Code Club soon after. Bethany feels that running a Code Club has allowed everyone's confidence to grow. She explains how brilliant it's been to see her students supporting each other, both in the club and in other areas of school life.

Q: What inspired you to start a Code Club in your school?

A: I was motivated by attending a Code Club event. I met lots of amazing people who showcased all the support for starting and running a club. Also, the children at our school have a passion for computing, and by running a club we have helped to give them an environment to grow their skills.

Q: What was the process of setting up the club like?

A: It was simple — we registered on the Code Club website and then got started with sessions.

Q: What projects have you been using?

A: We started off following the project paths (helloworld.cc/projectpaths), but the students quickly started to have their

own ideas of what they wanted to make and began to choose their own projects. Some use Raspberry Pi Foundation projects and some do their own thing. We started printing out projects as this was easier for the students to work with. Then we adapted them using Widgit Symbols (helloworld.cc/widgits) to make them accessible to everyone in the club.

Q: What have you learnt that might be helpful to share with a new Code Club?

A: Our biggest challenge when we started was supporting our students in creating accounts for Scratch. We now spend the first week with new coders getting them signed up to everything we offer. Plus, we play some fun unplugged computer-themed games or activities. This way, the following weeks run much smoother for the students.

I'm enjoying seeing club members develop their independence. Some of our students have created whole games that they work on in the club and at home. This is something we didn't really see until we started the club.

Code Club at Joseph Turner Primary

At Joseph Turner Primary School in Tipton, UK, Scott Sefton is the club leader and also the school's computing lead. In the spring term of 2023, he began his club with twelve Year 6 learners (aged 10–11) and then moved on to twelve Year 5 (aged 9–10) learners in the summer term. Since both groups were new to coding, Scott found it helpful to demonstrate how to code the projects while encouraging students to share ideas and help each other.



■ Elena and her young coders enjoying a Code Club session at St Paul's Church of England Academy

Q: Why did you decide to launch a Code Club in your school?

A: My head teacher tasked us with setting up this club as part of a whole-school expansion into more clubs. I have continued it because of demand for the club. I hope to continue developing programming understanding throughout Joseph Turner Primary School.

Q: How have you found the Code Club projects?

A: The resources are excellent for us to use. I find the projects good; they are well-detailed, easy to follow, and the children

“ I WANT TO INSPIRE ALL CHILDREN TO BE INTERESTED IN COMPUTING ”

really enjoy them, especially when they pick what they want to do for that week.

Q: What have you learnt that you could share with a new Code Club?

A: I live model the projects so children can see how to code accurately. By live modelling, I mean I read the instructions aloud with the children, and then show them what to do on Scratch itself.

Stella, aged ten, had this to say about their Code Club at Joseph Turner:

“I love Code Club because we can do different animations and stuff.”

Code Club at St Paul’s C of E Academy

St Paul’s Church of England Academy, also in Tipton, launched its Code Club in May 2023. Elena Lockett, the club leader and computing lead, has twelve learners attending her club, the majority being girls. She values a child-led and peer learning environment and therefore has a mixed group of club members, ranging from Year 3 to Year 6 (aged 7–11).

Q: What motivated you to start a Code Club in your school?

A: I wanted to inspire all children to be interested in computing — especially girls — as I wrote my dissertation on this at

university when I studied computer science. I wanted to give children a chance to share between year groups and almost become ambassadors in coding, to be able to share skills they’ve learnt in Code Club with their peers in class.

Q: Do you have any advice that could benefit a new Code Club?

A: Plan out a schedule for the term so you are prepped. Spend time practising how to type and log in at the start of each session because a lot of children struggle with the basic skills of using a computer and that may hinder their work later. Have a mixed range of ages and mix the children up in your Code Club so they can help each other out.

Q: What impact are you hoping for as your Code Club continues to grow?

A: I would hope that when our children leave for secondary school, they go into their computing lessons with a passion for computing. I also want coding to be something children feel they excel in, like English and Maths, as so many jobs rely on programming these days.

Divine, aged eight, had this to say about their Code Club at St Paul’s:

“I enjoy that you get to do different activities and it’s really nice. And the teacher is really nice and kind.”

Code Club at Fab Lab, Hartlepool

The Code Club at Fab Lab, based in the Hartlepool Community Hubs, UK, has been running since early 2021. Michael Storey, one of the volunteers at the club, told us about his experience.

Q: When you began your journey with Code Club, how did you find the projects?

A: While getting our footing, the premade projects and plans on the Code Club website were a huge help when deciding what we wanted to cover and teach that week. They also helped the kids to work at their own pace, as we needed to be prepared for new arrivals every week. The difficulty selector (i.e. Level 1, 2, 3) on the overall projects page helped us pick random projects that were fun and would use our 3D printers or other resources we have, like the LEGO Education WeDo kits.

Q: What advice would you give a new Code Club based on your experience?


A: Our biggest hurdle while we were starting out was planning and getting a structure in place for the flow of the sessions. Now we start with a ‘free’ time period where the kids can complete the prior week’s projects or try to make their own, and we have a show-and-tell time at the end to give the kids time to express themselves and share what they have been making.

Q: What are you currently working on?

A: Currently we are working through the Scratch pathways. While you don’t normally need to dedicate so much time to one language, we thought it would be better for the kids to learn all they can from one [language] at the same pace together before moving on to more complex and challenging languages.

We have found the show and tell to be very successful. It boosts our learners’ confidence, giving them the opportunity to show their parents and ourselves what they have learnt. It also inspires new ideas and ways of coding in other young people.

Interested in Code Club?

If you’d like to learn more about setting up a Code Club, please register your interest on this form: rpf.io/interest. You never know, you could be the next coding club leader in your community! 



CREATIVE COORDINATE GEOMETRY AND CT

Gitika Kaur shares how computing and maths practitioners can collaborate to integrate STEM lessons using computational thinking as a connecting link

As a former software engineer and a long-time teacher of maths and computing, I have always believed there is a strong relationship between coding in computer science and problem-solving in maths. As part of my master's in STEM education, I had the opportunity to model these thoughts. I undertook educational design research with a maths teacher from India, with the aim of creating a curriculum that would bring together three

components: computational thinking (CT), Scratch programming, and coordinate geometry. We implemented the curriculum online, with a group of students aged eight to twelve, for a period of five weeks in the spring term. The first four weeks included five hours of teaching and learning. The fifth week concluded the curriculum with an hour-long assessment through the presentation of students' projects and a short formative assessment. In this article,

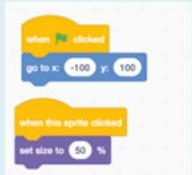


I will share the curriculum design and the incredible work of the students.

Curriculum design

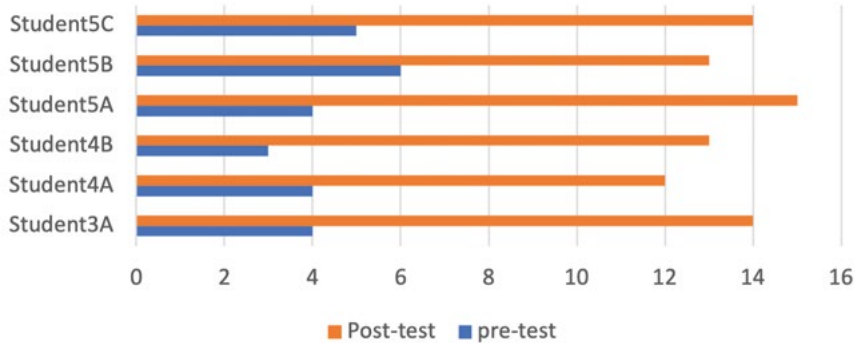
Our curriculum design is underpinned by Seymour Papert's constructionist framework, which emphasises 'learning by making' (helloworld.cc/papert1980). The activities within the curriculum were curated to guide students towards creating their own 'virtual artefacts' at the end of the course — for example, unique and creative geometrical shapes including repetitive patterns of triangles. While mapping Scratch programming concepts with coordinate geometry concepts, I used computational thinking as a problem-solving tool to link them together. **Table 1** gives a flavour of the relationship between the three components and the expected learning outcomes. The pedagogical approach focused on transforming the role of a teacher from an instructor to a constructionist facilitator, who:

- Articulates the clear purpose of each facet of computational thinking (decomposition, pattern recognition, abstraction, and algorithm)
- Provides a balanced amount of scaffolding — both stepping in and diving back — so that learners neither lose the joy to learn nor feel incompetent and get distracted (helloworld.cc/resnick2013)
- Induces activities that encourage student collaboration
- Devises 'low-floor and high-ceiling' activities, and uses phrases like, 'Let's construct a solution together' to prompt students to construct their own piece of code for solving problems, rather than seeking one finite solution

■ Table 1 Computer science and maths concepts can be linked through computational thinking

Computational thinking facets	Computer science concepts	Maths concepts	Learning outcomes
<p>Decomposition: breaking a big problem into small chunks</p>	<ul style="list-style-type: none"> • Events: using one event to solve one part of a problem  <ul style="list-style-type: none"> • Parallelism: using the same event to solve many parts of a problem  <ul style="list-style-type: none"> • Conditionals: using If and else commands to solve two different parts in one click 	<ul style="list-style-type: none"> • Developing meaning of operations and measurement • Algebraic reasoning and functions • Data analysis and probability 	<p>Coordinate Geometry:</p> <ul style="list-style-type: none"> • X and y positions in context with the sprite • Negative x and y position • Quadrants I, II, III, IV • Vocabulary: Abscissa, Ordinate <p>Scratch programming:</p> 

Problem-solving scores



■ Post-intervention results showed a marked improvement over pre-intervention results

Results of the intervention

The findings indicated that as students explored the movement of a sprite on their own, they engaged in using programming concepts to find the changing coordinate. While exploring the concept of negative and positive axes, they engaged in answering the exam-based questions correctly, themselves, without the need for those concepts to be taught. As a computer teacher, I focused on facilitating the improvement of students' understanding of block coding techniques, in order to empower them to achieve their desired outcomes. Through purposeful teaching sessions and assessment through their project presentations and a short formative test, it became evident that student exploration and engagement significantly enhanced their response time and problem-solving skills.

It is worth noting that, by the end of our intervention, all students demonstrated their proficiency in defining computational thinking and displayed confidence in utilising its facets when confronted with assessment questions. As educators, we were deeply moved by the students' outstanding performances throughout this process. We take pride in the fact that our role shifted from being instructors to facilitators, allowing the students to take ownership of their learning journey and thrive in their academic endeavours.

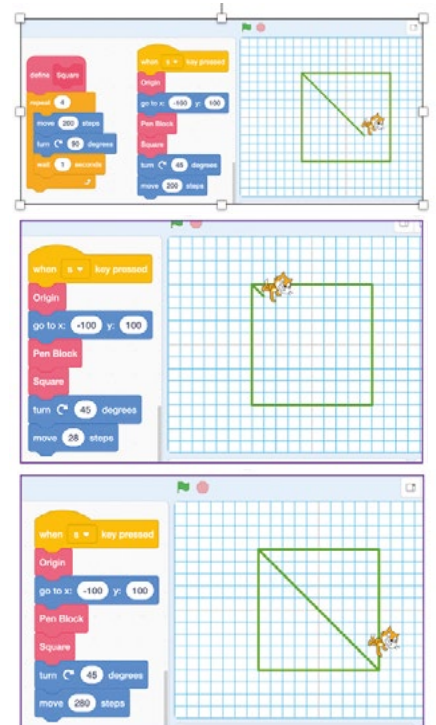
The participating maths teacher added, "I am amazed to see the confidence and agility of students to play with sprite[s] and explore all the mathematical concepts with interest and ease. I look forward to using Scratch in my classroom such that students

use computational thinking facets for solving problems." However, she also mentioned that she could foresee some challenges related to in-service training for learning to program, because she felt less confident facilitating students to debug their code.

Though I cannot generalise the teachers' perspectives, as I could interview only one maths teacher, I cannot neglect two important facts: firstly, the ten years of teaching experience of the teacher participant and secondly, the finding in the literature that there exists a "knowledge and skill gap, especially with respect to subject-based pedagogy" (helloworld.cc/Curzon). Therefore, I suggest that in order to efficiently implement CT-infused lessons in the maths classroom, we need to deploy computer teachers along with maths teachers. This was borne out by the maths teacher I interviewed, who went on to say that, "The collaboration with computer teachers is mandatory to teach computer science concepts, otherwise the skill gap of teachers will limit the students' calibres also."

Takeaway for teachers

I suggest that teachers need to take care to "combine diving in with stepping back" (helloworld.cc/resnick2013), to maintain an equilibrium between teaching and facilitating before steering learners towards taking ownership of their learning. More research is needed into the implementation of the computational thinking curriculum and the creation of CT-infused lessons through effective collaboration between computer science teachers and maths teachers, to create lessons that map disciplinary



■ An eight-year-old student exhibited creativity when determining the relationship between the length of the diagonal and the sides of a square



GITIKA KAUR

Gitika is a teacher of computer science at Desborough College in Maidenhead, UK. She has spent the last seven years working as an educator and consultant in India, teaching and acknowledging the scope of computing in creating an exploratory learning environment.

concepts and link them together using computational thinking. Future rigorous quality testing of the intervention, through partnerships between researchers and teachers, is needed in order to develop, validate, and assess the impact of the curated lesson plans on learners' creativity and problem-solving skills. (HW)

ONLINE COURSES FOR PROFESSIONAL GROWTH

Tips for unlocking your potential through continuing professional development

Online courses are a fantastic option for teachers to engage in continuing professional development (CPD) in a way that can fit in with even the busiest of lifestyles. There are a huge number of quality online learning opportunities, which you can complete in your own time and at your own pace.

Online courses in particular offer collaboration tools for learners to work together. Although the communication is not face-to-face, discussions provide opportunities to engage with a diverse group of fellow educators around the world.

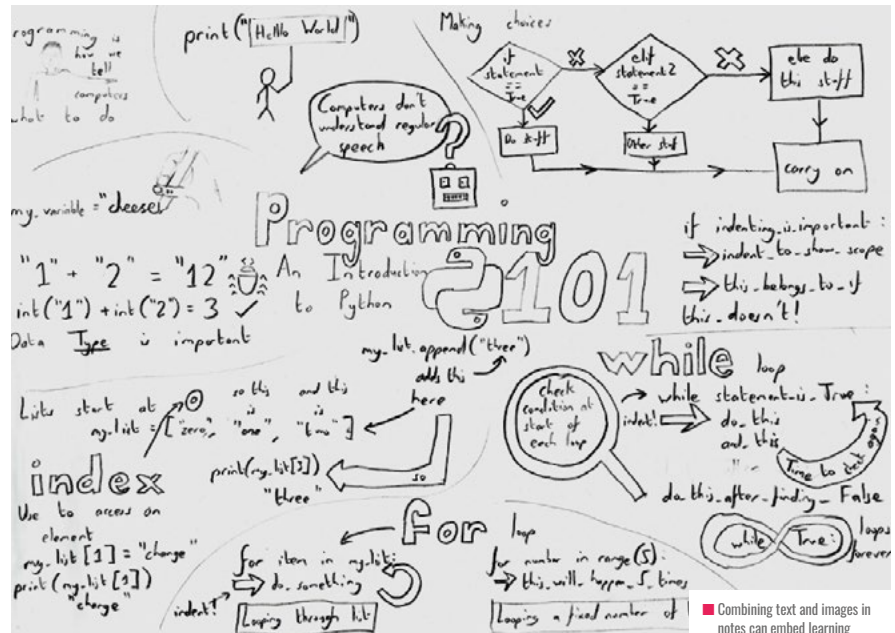
In this article, we will outline a range of approaches to help you make the most of online courses, including those provided by the Raspberry Pi Foundation, which will be useful whether you are considering taking your first online course or are a veteran of the format.

Choose your online course

Before you begin a course, make sure it is the right one for you. Most online courses, including those on edX (helloworld.cc/edXRPFF), include a short description of the course content. Make sure you choose a path that will meet your expectations and help you develop your skills as an educator. Look out for courses that are appropriate for the age range of your learners, or that address specific areas of teaching and learning that you would like to improve upon.

Craft your ideal learning schedule

This is completely personal to you, but it's a good idea to establish when and where you can best dedicate your time to online learning. This may be a regular time, every week in the same place, or you may need to be more flexible to fit in with your other commitments. Whatever you choose, plan



the time and wherever possible stick to it. You can usually complete online courses at your own pace — our edX courses can be completed by committing as little as one hour per week. If you can, find a quiet place to learn from where you won't be distracted so that you can focus on your learning.

Use videos

Online courses commonly include videos featuring a presenter. These videos can often include animations to help illustrate a concept or a screencast showing the presenter working through a problem. In all of these cases, you should feel free to pause the video and rewind it to help you understand a particular point, or to follow along with an example. For accessibility reasons, courses should include all of the key content from videos in text form, often as a transcript. You can use this text instead of the video, or as additional support during or after watching the video.

Take notes

Even while learning online, non-digital note-taking is really important. We recommend that you take notes throughout online courses to reinforce your learning and as a study aid when it's time to revise prior to taking any assessments. Plus, you can always go back to your notes long after you have completed a course.

Don't be afraid to get creative with your note-taking, whether it's through sketching or highlighting key information. For inspiration, pictured above is a set of sketch notes from our *Programming 101: An Introduction to Python for Educators* course (helloworld.cc/Programming101).

Build a learning community

Many online courses include opportunities for you to share your thoughts and progress. Although it may sound scary to put your ideas and work onto a public forum, sharing and discussing with your



■ Ben Hall presenting an online course

fellow learners can really help to enhance your learning journey. Remember that you aren't alone — your fellow learners are in the same position, and a friendly comment can go a long way. By posting, and commenting on other people's posts, you help to build a sense of community, meaning you and your fellow learners feel more involved and supported.

Ask for help

There will likely be people on your course from a variety of backgrounds who will be able to offer support and advice — make the most of this! If you are struggling to grasp a particular concept, there will probably be someone else who can help you; don't be afraid to ask. The more specific you can be with your questions, the more likely others may be able to help. You may also be able to support other learners yourself, which is a great way to put your learning into practice.

Reflect on what you've learnt

It's important to consider what you've learnt to help embed your learning. Good courses will give you specific chances to reflect, but even if these aren't there, try to find places

to stop and think about what you've learnt. You can link what you've learnt to what you already know, or consider different contexts to which your new knowledge could apply. You could also think about how you have been learning, and if there are any changes you want to make to your approach to your learning.

You might also find it helpful to create your own summary of what you've learnt when you complete a course.

Sign up now!

The Raspberry Pi Foundation has a long history of providing world-class online CPD, and we are delighted to announce that we've joined edX, a global online learning platform. Through our free online courses, we enable any educator to learn the skills required to teach students about computing and how to create with digital technologies. Since 2017, over 250,000 people have taken our online courses, including 19,000 teachers in England alone. The move to edX builds on this success to help us bring high-quality training to many more teachers worldwide.

You can now sign up (helloworld.cc/edXRPF) for 20 of our most popular

online courses on subjects ranging from programming in Python and Scratch, web development and design, cybersecurity, to machine learning and AI on the edX platform.

We hope this article helps you make the most of online learning for CPD, and that you enjoy learning with us. [\(HW\)](#)

MICHAEL CONTERIO AND BEN HALL

Michael is a former physicist and now works as an online course production manager at the Raspberry Pi Foundation. Ben is a learning manager at the Raspberry Pi Foundation. He is a CAS Master Teacher and a Raspberry Pi Certified Educator ([@hengehall](#)).



TEACH FIRST: TRAINEE COMPUTING TEACHERS

Daljit Shoker and **Rachel Arthur** invite first-year trainee teachers to share how they put computing pedagogies into practice at Teach First

In Hello World issue 21 (helloworld.cc/issue21), Daljit Shoker and Rachel Arthur shared how they deliver computing pedagogies to their first-year trainee computer science teachers. Now, this article looks at how Teach First trainees have implemented those pedagogies in their teaching practice and the key research they used in their literature reviews in their computing assignments.

When Rachel wrote the computing curriculum for Teach First, her goal was to create a curriculum that was practical, engaging, and accessible. The Teach First computing curriculum includes twelve pedagogical principles (helloworld.cc/12pps), equipping trainee teachers with the tools they need to make a lasting impact on students who experience educational disadvantage. Building on the success of

delivering the computing curriculum, Daljit takes a look at the transformative journey undertaken by three first-year trainee computer science teachers. Katy Barber, Lindsay Gibbons, and Cara Griffiths have embraced the PRIMM pedagogy (Predict–Run–Investigate–Modify–Make) in their classrooms, leading to remarkable outcomes with engagement, code comprehension, handwriting code, and debugging code.

KATY BARBER

As part of my assignment, I explored several key pieces of pedagogical research, including PRIMM, Parson's Problems, pair programming, and live coding. I decided to focus on one of my classes (aged 14–16) who were struggling with understanding Python syntax; this was acting as a barrier to them understanding how to program efficiently. Over the past few months, I have been exploring how to deploy PRIMM to structure my programming lessons and help improve my teaching and consequent pupil understanding of Python syntax. Throughout my assignment, I focused on three main teaching episodes and reflected on how I had delivered the components of PRIMM and the effect it had had.

The first lesson focused on arithmetic expressions, using PRIMM with teacher-led prediction and independent investigation. While most students

understood BIDMAS (brackets, indices, division, multiplication, addition, and subtraction), a new code block posed challenges, particularly for those with lower reading skills or English as an additional language (EAL). Here, I found that scaffolding was vital for their progress.

The second episode featured a monthly exam-style assessment with help sheets to aid with writing correct syntax. The biggest challenge I identified here was with students struggling with error identification due to integrated development environment (IDE) dependency — this assessment was paper-based.

The third reflection revolved around a lesson on subroutines, featuring teacher-led exposition and worksheets and promoting independent engagement during the MAKE stage. Here, I noticed that tailored scaffolding led to notable improvement for students who had previously struggled.

The final assessment revealed a potential PRIMM pitfall. Written exam-style questions diverged from the PRIMM format, unsettling students accustomed to IDEs. Many struggled with syntax, highlighting the need to adapt PRIMM for traditional exams. A key learning for me here was incorporating early handwritten coding practice to better equip students.

Overall, I found that my reflections offered valuable insights into the application of PRIMM within the programming curriculum. While PRIMM proves effective in fostering comprehension, adaptation is crucial for aligning with the broader exam-based goals of students. Balancing IDE usage with handwritten code practice emerged as a good strategy to prepare students holistically for both classroom learning and future assessments, and has impacted my future planning of how to deliver programming lessons.

LINDSAY GIBBONS

For trainee teachers on the Teach First training programme, the first assignment was to pick one topic in the computer science curriculum and dive in deep to the research and good existing practice. In my specific case, the focus was on programming lessons for students aged 14–16 years. For many students, computer programming is among the most challenging aspects of computer science. From the start of teaching Python to the class, it became clear that there was a reluctance to engage with programming. This cohort of students had little real-life experience of programming to put it into context; the ideas behind coding were very abstract. From my literature research and after discussions with my subject development lead (SDL) at Teach First, I chose PRIMM — from *The*

Big Book of Computing Pedagogy (helloworld.cc/BBpedagogy) — with the aim of making coding more accessible to this mixed-ability class. Using PRIMM gave consistency to the students, allowing them to pick their level and engage more in the lessons.

PRIMM provided a structure for the slow introduction of key topics, which gave the students confidence with trying new syntax and coding techniques, and encouraged them to engage. Reflecting on the lessons after a term of using PRIMM, I realised there was a group of students who were stagnating at a specific point in the lessons — they didn't want to push themselves past the easy PREDICT stage. I introduced several other pedagogies alongside PRIMM. I found that I needed to be more visible in creating code, so

that they could see how the code could be written and debugged, and especially how the logic of deciding what path to take was formulated.

Once more code writing was taking place in the lessons, I realised the importance of 'directed student talking' in the classroom, so I introduced pair programming. These collaboration sessions allowed the students to talk about how the code was being used.

Finally, during the second part of the teaching year, I saw a huge leap in terms of engagement. I introduced Use–Modify–Create examples from *The Little Book of Algorithms 2.0* (helloworld.cc/LBOA2) so that the students got more experience of editing and running Python code. They could get real debugging experience and start taking ownership of their code.

CARA GRIFFITHS

Through the opportunities provided to me by Teach First, I was able to take some of my theoretical learning and put it into practice in the classroom last year. Parson's Problems are activities that are often overlooked by teachers, as they are usually tagged on to the end of a lesson when time needs to be used up. Instead, I wanted to look at how turning them into a fully unplugged activity could best help learners aged 11–13 who were just beginning their journey into learning Python.

The use of Parson's Problems proved incredibly effective in helping students understand code — but particularly when they were used in a specific way. What I found was that students who were just given jumbled-up lines of code to sort found the task confusing and too difficult. However, when we started with a Parson's Problem example that focused on an everyday scenario (brushing teeth, for example), the students found it much easier to grasp the activity and were more successful with the code. What we realised from doing this

is something that is often overlooked in academic papers — students are not just learning a new concept; they are discovering a new way of learning, and must be able to practise it before they can successfully learn with it. Doing this as a tabletop activity also helped greatly, as students connected with the physicality of being able to move paper into the correct positions.

Based on all of this, I would like to share my top tips for successfully using Parson's Problems in your lessons:

- Make sure that the activity is done on paper first
- Always start with a scenario that is familiar to the students
- When using code, start small — four lines of code can be very effective

I hope that by sharing this research, I can help others to feel more comfortable using Parson's Problems in their lessons. Happy solving!

“ THESE STRATEGIES ARE ABOUT INSPIRING A LIFELONG PASSION FOR CODING AND TECHNOLOGY IN THE MINDS OF STUDENTS

We hope you enjoyed reading about this educational journey and glimpsed the power of putting into practice the computing pedagogies shared by these three teachers during their first year of teaching. If you are currently teaching or planning to teach programming, why not embrace these innovative approaches? They can undoubtedly make the process of learning programming more engaging and effective. By integrating PRIMM, Parson's Problems, and unplugged activities into your teaching repertoire, you can foster a

deeper understanding of coding concepts, encourage problem-solving skills, and create an inclusive and supportive classroom environment where all students have the opportunity to thrive. So, as you embark on your teaching journey, remember that these strategies are not just about imparting knowledge; they are about inspiring a lifelong passion for coding and technology in the minds of your students. Happy teaching, and may your classroom be filled with the excitement of discovery and the joy of learning! [\(H/W\)](#)



DALJIT SHOKER

Daljit is a subject development lead for computing at Teach First. She delivers the computing curriculum to trainee computer science teachers in the UK and supports them with developing their subject knowledge and teaching approaches. Previously, Daljit taught ICT/computing for 18 years ([@DaljitShoker](#)).



RACHEL ARTHUR

Rachel is the head of computing, programme design, at Teach First. She has previously been a head of computing and assistant head teacher in schools across the UK ([@rarthurtweets](#)).

Katy Barber, Lindsay Gibbons, and Cara Griffiths are newly qualified computer science teachers who have successfully completed the first year of the Teach First training programme and are now in their second year of teaching.



■ With support from Google, four teachers attended the WIPSCCE '23 conference

WIPSCCE 2023: A COMMUNITY OF RESEARCHERS AND TEACHERS

Insights from four teachers attending the WIPSCCE conference for the first time

The annual WIPSCCE Conference on Primary and Secondary Computing Education Research is an international conference where university academics present their research on the teaching and learning of computing in primary and secondary schools. A range of peer-reviewed papers and posters are presented, and all are then published in the Association for Computing Machinery (ACM) digital library (helloworld.cc/WIPSCCE23). Topics covered include curriculum, pedagogy, assessment, equity, tools, and teacher education.

There is a long history of WIPSCCE welcoming teachers to be part of the conference. Teachers are invited to hear about research and to discuss it with the researchers and their fellow teachers, providing an exciting and collaborative environment for sharing ideas and initiatives that are highly relevant to classroom practice.

In September 2023, the University of Cambridge Department of Computer Science and Technology and the Raspberry Pi Computing Education Research Centre hosted WIPSCCE 2023. As part of this, we were delighted to welcome Google as a sponsor of the event. Google has a deep commitment to computing education, especially within communities that have been historically underserved. Through partnerships and programs, Google supports organisations, universities, and schools globally to ensure that every student has the opportunity to access the benefits of a computing education to help shape their future.

With support from Google, WIPSCCE offered free places for five UK computing teachers to attend the conference. Teachers were selected from primary and secondary state schools from five different areas of England. Meet four of the teacher's selected to attend WIPSCCE 2023.



JANE WAITE

Jane is the senior research scientist at the Raspberry Pi Computing Education Research Centre. She was previously a schoolteacher and an IT developer, and now she investigates the teaching and learning of computing (@janewaite).

ANDY COLLEY *LAURUS CHEADLE HULME, CHEADLE*

What did you study and where?

ICT at Salford.

How did you end up in computer science?

I'm a converted ICT teacher.

Why were you interested in attending WiPSCE?

I have a background as an Advanced Skills Teacher (AST) and Lead Practitioner for Teaching and Learning (T&L). I've always been interested in best practice pedagogy and research-informed practice. It seemed like a great chance to meet the people doing the studies and gain a better understanding of how research happens.

Have you been to a conference like this before?

No, I haven't.

How would you describe the conference?

Educational! It really opened my eyes to the way that academic research and institutions function. It was a really friendly environment; I felt safe to ask questions and have my thinking challenged.

Whose research were you inspired or surprised by?

Hint content for block coding (helloworld.cc/Greifenstein) — I wasn't aware of LitterBox (helloworld.cc/litterbox). Anette Benz made me aware of different lenses for identifying student preferences (helloworld.cc/Bentz).

What were your key takeaways?

- It's useful for teachers to have an appreciation of how and why research is happening



- The research that seemed the most focused/useful was where researchers collaborated directly with teachers
- The human story behind why we use technology can help engage learners with the more abstract concepts

RACHAEL COULTART

ST NICHOLAS C OF E PRIMARY SCHOOL AND NURSERY, STEVENAGE

What did you study and where?

I've done a BA in Education at Durham University and an MA in Education at the University of Hertfordshire.

How did you end up in computer science?

No one else wanted to do it, and my husband was 'good with computers'!

Why were you interested in attending WiPSCE?

Having been involved in three different action research projects over the last few years and it being the best continuing professional development I've ever had, I've become increasingly interested in how research can influence my practice.

Have you been to a conference like this before?

Never! I've attended conferences where research is referenced and applied to classroom practice, but have never heard about research projects direct from the researchers, which puts a very different spin on it.

How would you describe the conference?

Much more academic (and a different pace) than I'm used to. A fantastic place to reflect on one's own practice and evaluate what you're doing in the classroom.

Whose research were you inspired or surprised by?

- Michael Rucker: A framework to consider the impact of technology sounds so sensible for considering how to help our learners understand the consequences of what they are doing (helloworld.cc/Rucker).
- Luisa Greifenstein: Loved the online seminar about LitterBox — dabbled with it straight after (helloworld.cc/Greifenstein).
- Anette Bentz: Both presentations were fascinating. Thinking about patterners and dramatists as opposed to girls and boys (helloworld.cc/Bentz). Looking at ways of developing understanding of loops also made me want to look at my practice in more detail (helloworld.cc/Bentz-loops).
- Anaclara Gerosa: Definitely want to explore 'grounded cognition' theory some more. But would also like to see research that explores whole-body analysis (helloworld.cc/Gerosa).
- Katharine Childs: I'd love to do something like the 'Hour of Code' following a teacher exchange I did in The Gambia many years ago (helloworld.cc/Childs).
- Isabella Gransbury: I have used the driver-navigator approach a lot, but love the sound of driver-driver too (helloworld.cc/Gransbury).



What were your key takeaways?

Research is a messy business, delivering more questions than answers! It's OK to see text-based coding like Python as ASCII art and I shouldn't worry/concern myself about transitioning KS2 pupils to it. Teacher advocacy and building a community to lean on are key in developing and delivering an equitable curriculum. Understanding and talking about differing human values is perhaps the key to helping our learners understand the impact of technology on society.

IAN GRAY

ENDSLEIGH HOLY CHILD VC ACADEMY, HULL

What did you study and where?

BA in English Literature & Educational Studies; PGCE in Primary — both at Hull University.

How did you end up in computer science?

I became a computing lead because I had an interest in technology and coding, and nobody else wanted to do it.

Why were you interested in attending WiPSCE?

I've always been interested in research and how the 'why' relates to the 'how' in learning. I wanted to gain an understanding of how research is presented, worked on, and conducted, as I would like to conduct my own research.

Have you been to a conference like this before?

No, only the SEND in Computing conference at the York Computing Hub.

How would you describe the conference?

Very informative. Meeting people from different education systems around the world has been inspiring and a welcome opportunity to leave the English education system bubble.

Whose research were you inspired by?

Michael Kölling, Joanna Goode, Michael Rücker, Anette Bentz, Karen Nolan, Anja Marx, Kate Farrell, and Veronica Catete (helloworld.cc/WiPSCEproceedings).

What were your key takeaways?

- Education systems around the world are working towards similar goals but in different ways
- Using a structural approach to teaching programming could help develop better coders
- Design research approaches can work in the classroom
- Agile is useful in helping students to collaborate
- Use of driver–driver programming approach (collaborating on code)
- Pupils' interests and how they can affect their approaches (patternner/balanced/dramatist)
- Designing for equity



SOPHIE FENN

PATE'S GRAMMAR SCHOOL, CHELTENHAM

What did you study and where?

I studied Computer Science at the University of Warwick to master's level.

Why were you interested in attending WiPSCE?

I discovered WiPSCE through the ACM membership and was excited to learn about a conference dedicated to primary and secondary computing education research. I wanted to find out what research was being carried out and hoped to get some new ideas to enhance my teaching practices.

Have you been to a conference like this before?

I've attended edTech conferences in the past and similar commercial events, but entering a research-focused conference like WiPSCE was a completely new experience for me.

How would you describe the conference?

I would describe the conference as an incredibly friendly and inclusive environment. Following each presentation, there is time to discuss the research and


pose questions. Despite my limited familiarity with research terminology and statistical terms, I realised that our insights as educators held equal value in these discussions. This process demystified the research process, so I'm now initiating a small classroom-based research project myself!

Whose research were you inspired or surprised by?

I found Michael Kölling's presentation on a frame-based editor for Python fascinating (helloworld.cc/Kolling). This innovative approach allows programmers to focus on program structure and fundamental concepts rather than syntax. I found it incredibly thought-provoking that something we deal with every day without thinking twice can be reimaged. It encourages me to consider how we might approach teaching challenges differently.

What were your key takeaways?

Research is not solely about large-scale projects involving thousands of students. A substantial portion of the research involves just a few students and investigating small changes to practice. Findings are



seldom straightforward. Context plays a crucial role, from students' prior exposure to computing to differing interpretations of what computer science entails among students, teachers, and researchers. Therefore, testing research findings in your classroom with your students is essential to evaluate their applicability in your unique context. [\(HAW\)](#)

Note: We will be writing articles about each of the referenced research items and will update the links as we publish.

▶ The following lesson plan is taken from the 'Computing systems and networks — IT around us' unit of The Computing Curriculum (TCC), written by the Raspberry Pi Foundation. It is aimed at learners aged six to seven, and is designed to promote learners' understanding of the various aspects of using a computer to create and change text.

ABOUT THE COMPUTING CURRICULUM



Raspberry Pi Foundation

The Computing Curriculum is the Raspberry Pi Foundation's bank of free lesson plans and other resources that offer educators everything they need to teach learners aged 5–16. It covers the full breadth of computing, including computing systems, programming, creating media, data and information, and societal impacts of digital technology.

Every unit of work contains a unit overview; a learning graph to show the progression of skills and concepts in a unit; and lesson content, including a lesson plan, slides, and formative assessment opportunities. Find them when you sign up for a free account at helloworld.cc/tcc.

INFORMATION TECHNOLOGY AROUND US

Help learners explore IT in familiar environments

AGE RANGE

6-7 years

OBJECTIVES

- ✓ Find examples of IT
- ✓ Sort IT by where it is found
- ✓ Talk about uses of IT

REQUIREMENTS

- ✓ 'A2 Activity sheet – IT in public places' from helloworld.cc/ITintheworld
- ✓ 'A3 Activity sheet – Talk about IT' from helloworld.cc/ITintheworld

This is the third lesson in The Computing Curriculum's 'Computing systems and networks — IT around us' unit for learners aged six to seven. During this lesson, learners will identify information technology (IT) that they interact with in familiar settings, such as home, shops, and streets. They will talk about the uses of IT in these environments and understand how IT is used in many workplaces. **(HW)**

ACTIVITY 1: WHERE ARE YOU LIKELY TO FIND IT?

5 MINUTES

Ask learners to think about where IT is likely to be found. Remind learners that IT is anything that is a computer or works with a computer. Ask learners to show a thumbs up if they think that IT is likely to be found in the place described, and thumbs down if it's unlikely that IT would be found there. They can show their thumbs in the middle (sideways) if they are unsure.

Is IT likely to be found in any of these places?

- A shop
- The woods
- Work offices
- The beach

After learners respond to each description, have a brief discussion about their thoughts.

Note: IT could potentially be found in all of these locations, e.g. a person on the beach with a phone. The discussion should be open and consider where IT is more likely to be found and why.

What are these objects?



■ Activity 2 What are these IT objects?

ACTIVITY 2: IT IN PUBLIC PLACES 20 MINUTES

Give learners the 'A2 Activity sheet — IT in public places' or display the activity sheet on the interactive board for the class.

Discuss the two locations — a busy street and a café or restaurant — pictured on the activity sheet where IT could be used. Encourage the learners to think-pair-share some of the IT that they might find in either of these places.

Look at each picture of IT on the activity sheet or the interactive board and ask

learners if they know what the object is. If learners have some knowledge of an object, ask them to explain its purpose.

Next, ask the learners to sort which IT is likely to be found in each location. Tell them to reflect on the previous discussions and try to place the IT object in the location where it is most likely to be found. If learners are unsure about an object, they can discuss it with another learner, or choose not to place it in either location.

Talk about IT

With a partner, think about:

1. Where information technology is found
2. What it is used for
3. How it works



■ Activity 3 Talk about IT

ACTIVITY 3: TALK ABOUT IT 10 MINUTES

Provide learners with 'A3 Activity sheet — Talk about IT' or display the activity sheet on the interactive board for the class. Highlight that the devices shown — traffic lights and a pedestrian crossing — are found in lots of places. Ask the learners to think-pair-share about the devices shown.

They should describe:

1. Where IT is found
2. What it is used for
3. How it works

Scaffolding opportunity: instead of focusing on a pedestrian crossing, learners could discuss the sequence of lights used for a set of traffic lights.

PLENARY ACTIVITY: IT EVERYWHERE

3 MINUTES

Remind learners that IT is used in almost every workplace. Ask learners if they can think of a job in which IT may not be used. Most jobs will involve some level of IT, even if it's only a phone.

RELEVANT LINKS

TCC 'IT in the world' lesson:
helloworld.cc/ITintheworld

AGE RANGE

11–14 years

OBJECTIVES

- ✓ Describe the difference between ‘data-driven’ and ‘rule-based’ approaches to application development
- ✓ Name examples of AI applications
- ✓ Outline some benefits and issues of using AI applications

REQUIREMENTS

- ✓ Resources from lesson one of the Experience AI resources (helloworld.cc/lesson1)
 - Slides
 - ‘The intelligent piece of paper’ resource (printed copy for demonstration)
 - Lesson one worksheet
- ✓ crayon.com allowlisted on computers

WHAT IS AI?

Introduce your students to data-driven artificial intelligence and how it is used in the world around them

This lesson is taken from the Raspberry Pi Foundation’s Experience AI resources (experience-ai.org). The aim is to help students explore the current state of (AI) and how it is being used in the world around them. Students will also consider some of the advantages and disadvantages of AI systems. **(HW)**

STARTER: WHAT IS INTELLIGENCE? 3 MINS

As students enter the classroom, display the question ‘What is intelligence?’ on your board. Students should discuss their thoughts with the person next to them. This part of the lesson aims to set the scene and to focus students’ attention. You do not need to take answers from the students;

instead, reveal the definition in the slide deck: ‘Commonly, people describe intelligence as the ability to learn and adapt/react to new situations.’ Highlight to students that the key point is that intelligence is the capability to learn something and use that knowledge to react and adapt to new situations.

INTRODUCTION: THE ‘INTELLIGENT’ PIECE OF PAPER 10 MINS

‘The intelligent piece of paper’ is an adapted version of an activity created by Paul Curzon and Peter McOwan of Queen Mary University of London with support from the Engineering and Physical Sciences Research Council (EPSRC) and Google (helloworld.cc/intelligentpaper).

To play this game, you will need a volunteer who will challenge the paper in play a game of tic-tac-toe (noughts and crosses). Another volunteer will act as the paper and follow the instructions given by ‘The intelligent piece of paper’ worksheet. The game will result in either a draw or a win for the intelligent piece of paper. After the game ends, ask the class, “Does that make the piece of paper intelligent?”

Encourage the students to share their thoughts. Someone might point out that it is not the piece of paper that is intelligent, but rather the person who wrote the instructions. The person who wrote the algorithm is intelligent because they have learned about IF/THEN rules and can apply them to create an algorithm that cannot lose in tic-tac-toe. They can also apply those IF/THEN rules to solve other problems.

To further the discussion, ask the class, “If a computer follows an algorithm written by a human, does that make the computer intelligent?” The answer should be no. The computer is not learning or reacting to new situations; it is simply following IF/THEN rules.

Rule-based

```
Move 1:
Place an X in a corner.

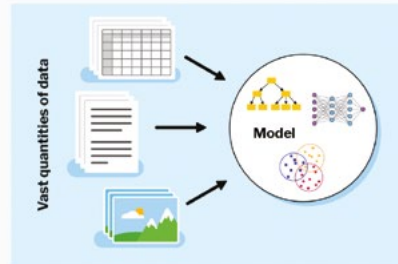
Move 2:
If the other player did not place an O in the opposite corner
THEN place an X in the opposite corner to move 1.
ELSE place an X in a free corner.

Move 3:
If there are 2 Xs and a space in a line
THEN place an X in that space in that line.
ELSE IF there are 2 Os and a space in a line
THEN place an X in that space.
ELSE place an X in a free corner.

Move 4:
If there are 2 Xs and a space in a line
THEN place an X in the free space on that line.
ELSE IF there are 2 Os and a space in a line
THEN place an X in that space.
ELSE place an X in a free corner.

Move 5:
Place an X in the free space.
```

Data-driven



■ Figure 1

ACTIVITY 1: WHAT IS ARTIFICIAL INTELLIGENCE? 12 MINS

The first task for students in this activity is to use their worksheet to record what they think is meant by artificial intelligence. This worksheet will be used throughout the lesson, so keep it handy! Play a video to the class where the experts at Google DeepMind explain what is meant by AI and address some common misconceptions.

Next, describe the rule-based and data-driven approaches to creating applications (Figure 1). The rule-based approach is

the traditional approach to programming, where an algorithm is followed to solve a computational problem. In contrast, a data-driven approach is a way of building artificial intelligence systems using statistics from vast quantities of data, rather than writing out the rules in a program.

Using the example of creating an AI model to play a game of chess, ask students to consider why a data-driven approach would be more beneficial.

ACTIVITY 2: GENERATIVE AI 10 MINS

In this activity students will get hands-on and use an online application to generate their own artwork. Before they get creative, outline the meaning of generative AI (applications designed to generate 'creative' content such as sound, images, and text). Then direct students to the worksheet where they will find a link to [craiyon.com](https://www.craiyon.com) to create their artwork. They must create artwork for either:

- A poster for the world climate change conference
- The wall of a café or restaurant

Allow students five minutes to use the application (images can take around a minute to generate — see Figure 2 for

some examples). Whilst students are experimenting, circulate around the class and prompt students to consider the following:

- Try generating the image again with the same prompts and the same art style. Does it produce the same artwork as last time?
- Try changing your prompts. Does it work better with fewer or more words?
- Who might find this technology useful?
- How might artists feel about this technology? Do you think they might be worried?
- Who owns the artwork you create? You or the creators of the application?
- What are the potential issues with using AI to generate art?

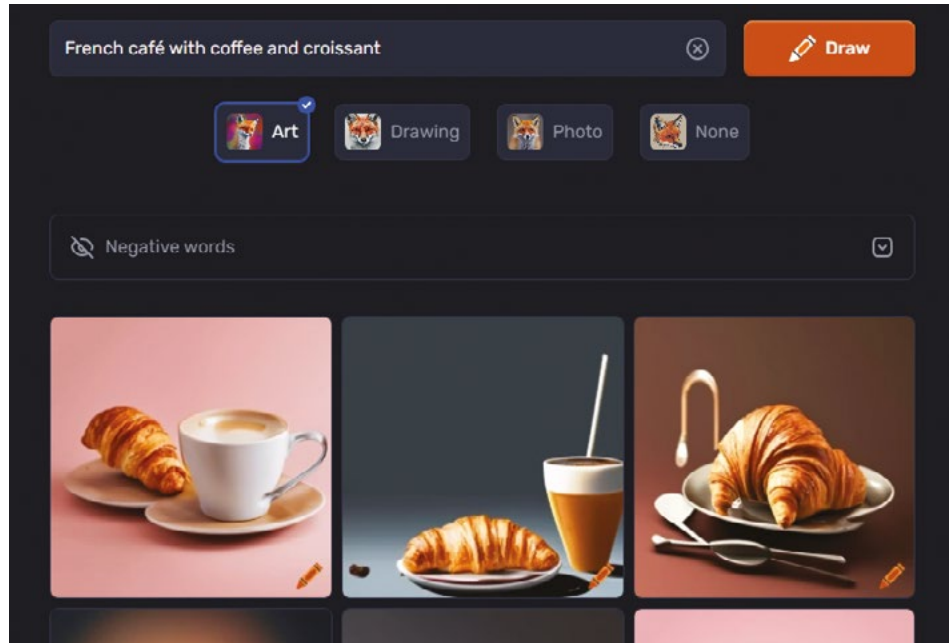


Figure 2



ACTIVITY 3: COMPUTER VISION 10 MINS

Using slide 18 listed in the resources, explain to students that computer vision is a field of AI that attempts to extract meaningful information from images. After that, ask the students to refer to their worksheets and explore a computer vision online application and answer the questions on the worksheet (helloworld.cc/computer-vision). Once on the website, ask students to select one of the sample images to see what the AI system can identify in it. Allow four minutes for students to experiment.

Whilst students are experimenting, circulate around the class and use the following questions to prompt discussion about what they are looking at:

- Why do you think there is a confidence rating? Is that important?
- Why do you think the application is more confident about some elements of images than others?
- Who might find this technology useful?
- How important is the confidence rating for a driverless car, for example?
- In what ways might visually impaired people benefit from this technology?
- Can you see this technology being misused?
- Would you feel comfortable knowing that you can be personally identified by a camera when you are walking around, for example?

RELEVANT LINKS

- helloworld.cc/intelligentpaper
- craiyon.com
- helloworld.cc/computer-vision
- experience-ai.org

PLENARY ACTIVITY: AI OR NOT AI? 10 MINS

The purpose of this task is for students to reflect on what they have learnt in this lesson and to apply their knowledge to see if they can identify when AI might be being used.

Ask them to evaluate each application listed on their worksheet and decide whether

or not it uses AI. Highlight that it may not be clear to them, in which case they can answer 'Could be AI'. Explain that they must attempt to justify their answer for each application. Remind students to think about whether each example is rule-based or data-driven.

THE BEBRAS PUZZLE PAGE

Students who do well in UK Bebras get invited to take part in the Oxford University Computing Challenge (OUCC); this issue's Bebras problem is taken from the first round of the 2023 OUCC

ABOUT BEBRAS

Bebras is organised in over 55 countries and aims to get students excited about computing and computational thinking. Last November, over 365,000 students participated in the UK annual challenge. Our archived questions let you create your own automarking quizzes at any time during the year. To find out more and to register your school, head to bebras.uk.



THE PROBLEM: PIG LATIN

Looping and selection



DOMAIN

Algorithms and programming

SKILLS

Algorithmic thinking

AGE

14–18 years

DIFFICULTY RATING

Ages 14–16 medium;
ages 16–18 easy

COMPUTING KEYWORD SPOTLIGHT: ARTIFICIAL INTELLIGENCE

Defining words and phrases in computer science

I asked ChatGPT to define artificial intelligence. Here is the start of what it came up with:

"AI, or artificial intelligence, refers to the creation and development of intelligent machines and software systems that can simulate human intelligence."

We need to be clear that AIs do not have self-awareness or understand what they are doing like humans. They are simply following algorithms and mathematical constructs to create output that is similar to what we would produce. Because most humans do not understand the inner workings of AI, there is a danger that we anthropomorphise them, using language such as, "intelligent machines" and "language understanding". They are just very powerful tools humans have built to aid us. (HW)

Pig Latin is a language game made by modifying English words:

- If the word starts with a consonant (for example 'hello'):
 - Remove the first letter of the word ('ello')
 - Add a "-" to the end of the word ('ello-')
 - Add the removed letter ('ello-h')
 - Add "ay" ('ello-hay')
- If the word starts with a vowel (for example 'internet'):
 - Add "-yay" to the end of the word ('internet-yay')

Using this information, write a program that will translate a sentence into Pig Latin.

Constraints: the provided input word will consist of upper-case and lower-case letters only.

Input format: an English sentence or group of words.

Output format: a string consisting of the translated input

Example input: Hello internet

Example output: ello-hay internet-yay

This problem was originally written by the OUCC team from Jamaica. Solution on page 85.

Further information

If you want to try programming this in the interface used for OUCC, you can do so at helloworld.cc/piglatin. The OUCC system supports eleven programming languages, including those most used in schools. The students' solutions are sent to our testing server, where they are run using several different inputs to test that students have thought of all the boundary cases. Partial scores are awarded and students can make several attempts. The tasks are then archived for schools with a UK Bebras account to use in their own automarking quizzes.



■ Learning the skills to create movies also has a positive effect in other areas of the curriculum

MAKING MOVIES USING GREEN-SCREEN TECHNOLOGY

Gemma Coleman catches up with primary teacher **Seema Zerafa** about her experiences developing learners' movie-making skills with a film competition



SEEMA ZERAFI

Seema is the head of digital learning at Pembridge Hall School in London, UK (Instagram: @creativecomputing_misssm).

What is Childnet and the Childnet Film Competition?

Childnet is a charity that works to make the internet a great and safe place for children and young people. For the last 14 years, the charity has run an annual film competition, open to all young people aged 7–18 in the UK (helloworld.cc/childnetcomp). The brief is to create a two-minute film or storyboard that focuses on a set online safety theme, which changes each year. This year's theme was 'Time to talk — how can people support each other online?'

How do you prepare for and run the competition?

The competition launches on Safer Internet Day every February (saferinternet.org.uk), so I dedicate the spring term to project-based

learning and e-safety. The idea behind this is that pupils who want to enter the competition are equipped with the skills and knowledge to do so, and even those who don't want to take part are learning to conceive and produce digital content that is effective, compliant, and secure.

I've carried out this competition with pupils aged nine to eleven for the last three years. However, we're effectively preparing students from the age of six, exposing them to green-screen technology, which I often link to teaching around topic days such as World Book Day, so by the age of eight or so, they are familiar with the technology.

I split the project into an eight-week module, one lesson per week. We divide the class into groups of four at most and, during the first four weeks, we learn about each app

and how to use them effectively. The second half of the block is the 'workshop', when the groups create a storyboard, and record and edit a movie. At the end, they can decide whether they want to enter the competition, either by using what they've produced, or starting from scratch.

What tools do you use for the competition?

I recommend researching which apps you feel would work best for the project. I use Green Screen by Do Ink (helloworld.cc/doink), which you can buy for a small fee, and iMovie (helloworld.cc/imovie), which you can use on an iPad. They are both pretty foolproof; you don't need to be an IT specialist to use them, and both feature a mini tutorial when you download them. Some Windows 10 alternatives include Clipchamp, which is a pre-installed app with any Windows 10 device (helloworld.cc/Clipchamp); or Microsoft Flip is a great web-based program to use (helloworld.cc/flip). It allows students to create mini videos safely, and it includes child-friendly filters and emojis. The best way to explain it is that it's just like creating videos on TikTok, but monitored and controlled by a teacher. Microsoft Flip includes tools for the teacher to limit the duration of the video and switch off features, such as face filters and emojis, if students don't use them effectively.

It's also useful to use collaborative tools like Microsoft Teams, Google Docs, and Microsoft Whiteboard. We set up dedicated Microsoft Teams channels so pupils can communicate about the project with their peers.

You don't need to spend money on a green screen; anything works as long as it's green! I would also suggest looking at the resources you already have at your disposal in school.

You can make a green screen from green cloth or backing paper, which you may already have in the art department.

What skills does the competition help build?

I've found the competition to be the perfect launch pad to embed movie-making skills and important messages about online safety into the curriculum.

While acquiring movie-making skills, the students develop a range of key digital literacy skills during the project. It's easy for a student to get carried away with design ideas and themes, but we introduce the students to the concept of a 'target audience' and seeing if their choices fit that target audience. We also spend time looking at branding and seeing whether their products align with the brand, and use similar fonts and colours.

Although we celebrate a different theme every year, we always address data protection, for example by obtaining online consent. We discuss who should appear in their movies, and the importance of asking permission, especially when featuring a child's face online. We also discuss whether to display names on the website, who could see this information, and whether we would want them to see it. Our discussions revolve around determining what information is safe to post, and what isn't.

You can embed all sorts of key computing skills into the project, including programming skills. For example, we use Marty the Robot (a walking robot that is fully programmable and can be connected over Wi-Fi or Bluetooth, robotical.io) as a character in our movies. To help Marty perform actions and move, pupils can construct their own code using MartyBlocks and MartyBlocks Jr, even

adding in a voice note, which is a super new feature. It's amazing how much enjoyment and excitement you can add to learning to code when you're combining it with movie-making.

Learning the skills to create movies also has a really positive domino effect in other areas of the curriculum. There is so much potential for cross-curricular collaboration. For example, after pupils have learnt how to use iMovie, could they make an iMovie about a religion in humanities? Or could they use GarageBand to make music for a movie in drama or music class?

How do you assess students' work for a competition like this?


I suggest setting the criteria at the start of the project (video duration, theme, format, etc.) and compiling a checklist of considerations for students (copyright issues, music and image licensing, digital consent, etc.). Throughout the process, I undertake formative assessment; I don't just assess the filming, but I also appraise teamwork and communication.


What are your top tips for educators who want to try this competition out?

Don't underestimate your students! Being digital natives, they pick up skills really quickly and can produce content that is often better than that created by adults. They also respond really well to being given responsibility and are far more likely to take the task seriously.

I'm passionate about the power of play, so I'd recommend giving students time to play with the apps you've chosen. It's important for them to experiment independently and see what they can achieve. They often teach me how to do things, even though I'm an IT specialist — if someone asks me something I don't know the answer to, I open it out to the class. They love teaching each other.

Finally, remember to fit in time to allow students to conduct their own research, and make sure you provide feedback on their plans and films so they have enough time to edit. I always ensure I give examples of what both good and bad content looks like, and discuss what characterises the two.

Helpful information and resources are free to download in the Childnet entry packs (helloworld.cc/childnetcompack). 



■ Fully programmable robots are the lead characters in a green-screen film

WORK-RELATED LEARNING EXPERIENCE

In this Insider's Guide, **Alan O'Donohoe** shares how educators and businesses can work together to offer work-related learning opportunities for students

It's essential that all young people are capable of making informed decisions about their future, particularly those who lack the necessary support at home. Many young people will not have had a real experience of working until they are 18 years old at the earliest, yet they are making decisions that can affect their future many years before then. Through productive partnerships with employers, schools can do a lot to support young people through a programme of work-related learning experiences.

I've written this Hello World guide with an audience of educators and employers in mind, with the aim of providing low-cost, practical suggestions for employer-hosted events with measurable impact that can be adapted for use. This guide will highlight the range of potential opportunities available. When approaching an employer, the suggestions in the article can act as a catalyst to start meaningful conversations to achieve similar levels of success as we've had with our events.

Traditional work experience

To be clear, this guide doesn't really include the more traditional work experience placements, which typically take place over two



ALAN O'DONOHOE

With 30 years of experience teaching and leading in schools in northern England, Alan leads The Exa Foundation (**exa.foundation**) on a mission to inspire digital makers, support the teaching of computing, and promote the appropriate usage of technology (@**exafoundation**, @**teknoteacher**).

weeks, late in Year 10 or early Year 11 (ages 14–16). Without exploring all the reasons here, there's been a trend for schools to gradually move away from this model. There was a necessary and obvious decline because of Covid-19 lockdowns, but there's no sign of a return to previous levels yet. The business I work for,

GOOD CAREER GUIDANCE


The Gatsby Benchmarks were published in 2013 and serve as a framework for exceptional careers provision. The UK government's Careers Strategy has adopted these benchmarks (helloworld.cc/Gatsby) and they have proven to have a positive effect on students' readiness and achievements. Although four of the benchmarks (below) are strongly linked to employer engagement, it probably makes more practical sense to focus on benchmarks 2 and 4 in school with support from employers, and to focus on benchmarks 5 and 6 when visiting a workplace.

2. Learning from career and labour market information: supports students to make informed decisions about their future and be prepared to enter the workforce.

4. Linking curriculum learning to careers: helps students see the real-life applications of their education and inspires them to set career goals early.

5. Encounters with employers and employees: can give students valuable insights into various career paths and job expectations, and help them build professional networks, which can enhance their employability.

6. Experiences of workplaces: allows students to develop practical skills, understand workplace culture, and make informed career choices.



■ Alan giving an Exa induction

Exa Networks based in Bradford, is an internet service provider with about 50 employees. Over a recent summer holiday period, we had one student on work experience placement. The student (age 16) organised the placement, but that is an extremely rare occurrence nowadays.

When our business witnessed a reduction in enquiries for work experience placements, we started to review what we do to support young people in terms of work-related learning. Our business is fairly unique in that we have offered school-based workshops nationwide over the last decade. This is separate from our core business, which is providing internet connectivity and online protection to schools. People are often surprised to learn how a small tech business in Bradford, a city in northern England, is dedicated to providing an education and careers programme for schools, and this has been recognised through impact awards.

Half-day visits

When our normal event programme resumed following Covid, we switched to including work-related experience days for young people as a more significant part of our programme offering. For the next academic year, we've committed to offering two opportunities every month for schools in Bradford to book a visit,

FEEDBACK

- Christine, Director of ICT & Digital Learning, Leeds City Academy: "Our Year 10 students (ages 14–15) loved meeting the employees and finding out more about their jobs and pathways. They enjoyed setting up the Raspberry Pi computers, building in Minecraft, the pair programming, and the awarding of raffle tickets for praise worked really well."
- Mohammed Shazad Bhatti, Bradford College: "Thank you for a wonderful experience for our learners. The general feedback was that the students enjoyed everything, especially the network security element."
- Student at Bradford College: "It was a really good session, as I have learnt a lot of things. I have learnt in more detail how cyberattackers attack victims. I have understood more deeply the four cornerstones of computational thinking and have some clear examples of how these are applied to real-life situations."
- Another student at Bradford College: "The session was fun because of the hands-on activities. We learnt a new term called 'cyberthreat actor'. In general, the session was very engaging and moved at a very good pace. The most important part was that everyone was involved."

RESOURCES

STEM Ambassadors:

Teachers have the opportunity to enhance work-related learning opportunities for their students by using the resources of STEM Ambassadors. With over 30,000 individuals eager to share their expertise and enthusiasm, this well-established volunteer programme can offer valuable insights, career talks, and engaging activities to inspire and educate young minds. STEM Ambassadors are trained and supported to deliver a wide array of activities that resonate with students. Schools can easily request a STEM Ambassador through the online platform, which can bring real-world relevance into classrooms and empower students for success in STEM-related careers (helloworld.cc/ambassadors).

SkillsHouse:

Exa Networks has been working with SkillsHouse, a local team under Bradford Council's Employment and Skills services. SkillsHouse aims to enrich learning experiences by aligning with Gatsby Benchmarks and providing access to employer-led work-related learning activities through their SkillsHouse Portal. Readers in other areas may be able to find similar agencies managed by their local council (skillshouse.co.uk). Andrew Gray at SkillsHouse Bradford, says: "EXA and similar organisations are deeply committed to their corporate social responsibility, put simply as 'giving back and paying forward'. We'd encourage others to become similarly engaged."



■ Students taking a break from the morning activities

and we are currently exploring opportunities to expand the work we do nationwide. In the last few years, school groups typically visit our business for a half-day or longer to take part in a series of challenges and activities that help them make tangible connections between what they've learnt in school and the workplace.

Mini induction

During the visits, participants are first shown around part of the building as a mini induction to fire and safety regulations, something all of our staff receive on their first day. Our building has controlled access zones and young people need to understand where these are and why. This is followed by a fun, interactive game in which participants have to guess facts about the business. Later, in an information security session, participants are tested on their knowledge of security threats and steps to prevent them.

Meet the employees

During the visit, colleagues from the business spend time with the visitors explaining their roles within their team, followed by a Q&A session. Young people typically want to know how much they earn and what grades they received at school. The students are visibly impressed to hear employees relate how their experience of community, recognition, and fulfilment trumps a high salary,

but yes — they earn enough to manage their way of life. When we ask the young people if they'd like to work at Exa in the future, there's always an overwhelmingly positive response. This activity is often the easiest to manage, but seems to have the most significant impact on the young people who attend. Initially, young people are shy about asking the first questions, so I have some previously asked questions prepared, and then they soon want to ask their own.

Impact and feedback

Students and teachers have told us they found these events incredibly valuable, particularly as the young people attending (ages 8–18) have had an opportunity to see people like them and whom they can relate to, working in fulfilling roles, and this has helped positively influence career and study choices.

When we've hosted events for a longer period of time with the same group, we have the opportunity to enhance to experience and make it even more immersive. It gives us more time to incorporate a variety of work-related activities.

The events have provided opportunities for young people to spend some time interacting with work colleagues and learning more about the business in general. By hosting whole teaching groups for these events, we're able to offer the experience to a larger number of young people than before. (HW)

SUGGESTED ACTIVITIES

ACTIVITY 1: ONE FALSE, TWO TRUE

On page 85 of Hello World magazine issue 20 (helloworld.cc/20), there is an interactive activity designed to develop computational thinking skills. The activity encourages students to apply computational thinking to real-life scenarios by identifying the false statement among three presented facts.

The activity starts with an employer sharing three brief facts about themselves or the company, with only two being true. In small groups, students work together to identify the false fact by devising questions. Emphasise a computer science perspective, prompting logical reasoning and abstraction in their approach. During the Q&A session, students are encouraged to actively listen and use computational thinking skills. At the end of the activity, the false statement is revealed, and feedback is provided on their question types.

ACTIVITY 2: HACK ATTACK

On page 86 of issue 20 (helloworld.cc/20), there is a description of an activity that requires students to collaboratively identify and eliminate imaginary cyberthreat actors using computational thinking skills.

This activity, inspired by the game Werewolf, is easy to set up and does not require any equipment. It is best played with a large group of players, and students use logical deduction and pattern recognition to succeed. The moderator, initially the teacher, secretly designates three players as cyberthreat actors. These actors silently choose victims for a coordinated cyberattack during the night phase. In the daytime phase, the victim is revealed, unable to communicate beyond saying, "I've

been hacked", and is eliminated from the game. Students must deduce the hackers' identities through reasoning, nominating suspects, and voting. The player with the most votes is removed, and their true identity disclosed. Repeat the process until all cyberthreat actors are identified and eliminated.

This captivating game immerses students in computational thinking while promoting teamwork and strategic thinking. Employers can use it to explain how it relates to their security policies and real-world threats. For a more detailed explanation, see the original article in issue 20.

ACTIVITY 3: WHAT'S IN THE BAG?

This is an engaging and fun cybersecurity activity for students. A bag contains cards that have technical phrases printed on them, including cybersecurity-related terms such as threats and protectors, and unrelated terms like cat breeds. The teacher selects a card at random, and students must swiftly categorise it as a threat, a protector, or a cat. Teams can earn points by correctly identifying threats or protectors, and earn three points for identifying a cat. To score extra points, teams can explain the meaning of the selected term and its relevance. For example, correctly identifying 'phishing' as a threat, explaining how it works, and ways to safeguard against it, can earn two points.

As the game progresses, the teacher or employer can introduce more complex terms and deeper questions to enhance students' understanding of cybersecurity. Non-cybersecurity terms add an amusing challenge to the game, and students can return them to the bag if drawn. This activity encourages critical thinking about cybersecurity with a touch of light-hearted fun.

TEACHING T LEVELS

Gemma Coleman catches up with Halima Bhayat about her experience of running T levels and looks at how you can get started in your school or college

What are T levels and how do they work?

T levels are qualifications that students in England generally take between the ages of 16 and 18. The government introduced them in 2020 to give students another pathway into further education or the world of work.

The course is structured over two academic years. In the first year, students take exams and an employer-set project, while in the second year, students complete a 67-hour occupational-specialism project. Students also need to complete a 315-hour placement in industry over the two years.

Why did you choose to introduce them to your school?

We were one of the first pilot schools to run a T-level course. We wanted to take part because the course would open students up to real-life work experiences in fields they're interested in. We also knew it would mean employers would be more able to offer apprenticeships post-school because of students working with them, and we liked that this provided options other than university for pupils.

We also liked that having a project in both years meant students were learning skills, not just theory for paper exams. At the time, there were only three subjects on offer. We chose the *Digital Production, Design and Development* course as it offered three strands (computing, digital skills, and business) with a broad category of options to teach. We already had staff with knowledge and skills built in those areas, and we structured our timetable for staff to create most of the course around computing, followed by digital skills, and business.

Could you describe the employer-set project a little more?

The project is a task-based controlled assessment, split into five sections using the systems design life cycle. The style of this project work has allowed pupils to think beyond the classroom and work using the same methods as real software developers. The full assessment lasts 14.5 hours and takes place over specific days. We provide students with a scenario (for example, a coffee shop wants to become computerised and have a front-end and back-end system for its orders) as a pretask, which they can discuss with their peers and spend time researching.

Once the tasks start, students are not allowed to use the internet. We give them set code, data files, and test log templates. Our scheme of learning covers a mock period as well as an example taught step-by-step before this assessment. We found that students appreciated having hands-on experience of practising this section.

And what is the occupational-specialism project?

This assessment takes up the full second year of the course. You need to really plan how you want to run this academic year, including when you may want to hold mock exams. The project involves four tasks, with flexibility regarding when educators run each task. Students can use the internet for the duration of this assessment and are therefore able to independently solve issues and create code using support from the internet.

The four tasks are split like a system development cycle, where for the first task, students research and plan a model. Then in the second task, they implement and create the model, before testing the model in task three. Finally, in task four they evaluate what they have implemented and suggest further improvements.

What does the placement in industry involve? And what are the benefits?

You are required to organise a placement in industry covering a minimum of 315 hours, of which 20 percent can be remote. We found this section of the course particularly challenging as a lot of companies, especially in the digital sector, are now working entirely remotely. To make this a little easier, my tip would be to build a network of people who work in industry, including parents, employers, and local organisations who might be able to support you in taking students on for a placement.

We decided to cover our placement section in two-week blocks, as this was the most comfortable amount of time for both students and the employers. We set the placement times at the start and end of the year, allowing plenty of time in between for staff and students to complete other schoolwork and to fit the placements around exams.

So far, we have worked with IBM and small local businesses. During their placements, students end up doing all sorts of things. For example, some developed a chatbot that focused on



■ T levels combine classroom-based learning with experience in the real world of work

sustainability and the safe disposal of materials. To help them reflect on their experiences, we created diary logs where they could record their new-found skills and knowledge and connect them back to their school projects.

What have you found to be the biggest benefits of the course? And what have students enjoyed the most?

Much of the T-level course we offer involves independent work, with students building their skills and problem-solving tasks themselves. We've found this helps students prepare themselves for further education and the workforce.

Overall, our students have been positive about the course. Not only have they gained valuable experience, but the placements have also opened up doors to paid degree apprenticeships. This has been especially helpful for our most vulnerable and underrepresented students, including those from ethnic minority backgrounds, SEND students, and female students. Unfortunately, there is still a shortage of women, especially from ethnic minority backgrounds, in these fields, which makes it difficult for employers to find them. Thankfully, the course has helped these groups connect with potential employers, which will hopefully help close the gap for them in the tech industry.


What advice would you offer to schools interested in pursuing T levels?

Planning each of the long assessments has been a challenge, especially as other state examinations and mocks take place at the same time. It has been key for us to plan the exams as early as possible in the autumn term to ensure we're organised! Everyone has

to work together to support the course, from careers team, finance staff, and technicians through to the school senior leadership team.

I would also suggest keeping your technical team involved in the process of the course. The employer-set project, for example, involves students downloading and using programming language libraries (such as Pandas), so you might need tech support!

Throughout the year, we've invested in new technology like robots, VR headsets, and drones to enhance our students' learning experience. We've also organised employer workshops, visits, talks, and trips. The course is all about helping students to grow as individuals and develop their aspirations. We've found that this holistic approach to education has really helped. **(HW)**



HALIMA BHAYAT
Halima has over 20 years of experience working in education, and is currently the head of computing and digital T levels at Ursuline High School in the UK. Halima is an Asian Women of Achievement 2021 Finalist and the Computing at School (CAS) Merton lead for all schools. She is a computing advocate, empowering ethnic minority girls to pursue a STEAM education and career.

PREPARING YOUNG PEOPLE TO UNDERSTAND SYSTEMS

Michael Conterio and Tracy Gardner discuss how to prepare young people to interact with the computing systems they will encounter in their lives

When we first introduce the idea of computer systems to young people, we often focus on hardware and software. However, children today will go on to interact with, act through, and maybe even design systems in which people also play an important role. How can we prepare young people to understand, work with, and change the systems they encounter — especially as more of these include artificial intelligence components?

Why systems thinking?

When businesses develop solutions, they often design business processes that include computer systems and the people who interact with those systems. Business processes can include tasks that are fully automated by software, tasks that are managed by software but require human input (such as filling in a form), and tasks that are manual but reported to a computer system (such as giving medication to a patient.)

More and more physical tasks are being automated, such as moving items around in an Amazon warehouse. The same is likely to happen with conceptual tasks that we previously assumed to require human intelligence. AI systems can be used to solve certain kinds of task.



TRACY GARDNER & MICHAEL CONTERIO

Dr Tracy Gardner is a computer scientist, tech industry professional, and technology educator and is currently head of technology at Code First Girls. Michael is a former physicist and now works as an online course production manager at the Raspberry Pi Foundation.

WHAT IS SYSTEMS THINKING?

Systems thinking looks at the overall structure of a system and the behaviour that emerges when its component parts interact. Instead of focusing on lower-level details, systems thinking takes a step back and looks at an overall system, often including the people who interact with it.

An example would be the interaction between lunch queues and timetabling in a school. You can think of the school, with its computer systems and pupils, as an overall system. Systems thinking might lead us to consider offsetting the timetables of different year groups in order to reduce queuing times. However, we'd also need to consider the other effects that change could have — would it affect staffing, or the timing of the end of the school day?

We need humans to decide how to use all of these different methods and techniques in the best possible way. This means using systems thinking — zooming out from the detail of a system, spotting patterns, and understanding how the parts interact.

This is especially necessary when considering how to use AI. When we apply AI to problems, it's important to think of the overall impact this will have on society, jobs, and individuals. Systems thinking is required, to consider the overall implications of a decision or solution. We must take care not to apply AI in ways that are locally optimal but have undesired wider implications for society.

"Unintentional destruction is [to] assume the AI wakes up tomorrow and says oxygen is rusting my circuits ... so it would ... reduce oxygen. We are collateral damage in that ... they are not really concerned, just like we are not really concerned with the insects that we kill when we spray our fields." — Mo Gawdat, former chief business officer of Google X (helloworld.cc/Gawdat).

How do we teach young people to be responsible, empowered creators?

To prepare young people for the future of work, educators can provide opportunities for them to experience authentic scenarios that involve impacts beyond what appears on a computer screen. By making these relatable, you can help motivate young people to see a future for themselves in computing — as well as developing their critical thinking skills to help them to evaluate how to use tech and AI safely and usefully. [\(HAW\)](#)

TEACH DATA LITERACY: A GUIDE FOR PRIMARY TEACHERS

Clear guidance backed by excellent resources to help teachers raise the profile of an undervalued digital skill

INFO BY Data Education in Schools | PUBLISHER n/a | PRICE free | ISBN n/a | URL helloworld.cc/TeachDataLiteracy Ben Hall

No, there's not a typo in the title of this publication — it should be 'data', not 'digital' literacy. The importance of digital literacy for primary-age children has long been recognised, and efforts to ensure our young people are equipped for a digital world are well-established and ongoing. But in a world where so much of what we do is monitored and analysed, do we place enough importance on data literacy?

What is data literacy?

This guide defines data literacy as a 'set of skills and concepts which people need to understand, interpret, and make decisions based on the data they encounter in the world around them.' These skills and concepts are routinely taught to younger children in a variety of subjects, including maths, computing, and science. The skills associated with digital literacy, such as questioning, analysing, and concluding, are cross-curricular, but data literacy is rarely considered a concept in its own right.

The Data Education in Schools team from The University of Edinburgh has been a leading advocate for increasing the profile of data literacy in education in Scotland. The team has developed a comprehensive set of curriculum resources that emphasise the


importance of understanding data in the real world. The resources are based on the Problem, Plan, Data, Analysis, and Conclusion (PPDAC) cycle, which statisticians R. J. MacKay and R. W. Oldford created in 2000. The cycle breaks down the data process, from the initial question to drawing conclusions, and explains each stage in easy-to-digest steps. The team also provides practical advice and example questions for teachers.

Comprehensive resources

The exploration of the model is then complemented by a comprehensive set of resources drawn from a variety of sources. Extra kudos

must be given for including QR code links, making the resources very easy to find and access.

This guide is a fantastic resource for all primary-school teachers, not just those responsible for computing. It offers a path to raising the profile of data literacy and developing these essential life skills in children. It would also serve as a strong foundation for professional development. I can't recommend it highly enough. **(HW)**



BEBRAS PUZZLE SOLUTION (PAGE 75)

Below is a solution to the OUCC Pig Latin problem, written in Python, that scores full marks:

```

1 # Get input and put words into a list.
2 message = input().split()
3 # Initialise global variables.
4 VOWELS = "aeiouAEIOU"
5 pig_latin = ""
6 # Translate text by iterating through words in the message list.
7 for word in message:
8     if word[0] not in VOWELS:
9         pig_latin = pig_latin + word[1:] + "-" + word[0] + "ay "
10    else:
11        pig_latin = pig_latin + word + "-yay "
12 # Output pig-latin after removing trailing space.
13 print(pig_latin[:-1])

```

“HELLO, WORLD!”

Everything you need to know about our computing and digital making magazine for educators

Q WHAT IS HELLO WORLD?

A Hello World is a magazine and accompanying podcast for computing and digital making educators. Written by educators, for educators, the magazine is designed as a platform to help you find inspiration, share experiences, and learn from each other.

Q WHO MAKES HELLO WORLD?


A Hello World is the official magazine of the Raspberry Pi Foundation.

Q WHY DO WE MAKE IT?

A There's growing momentum behind the idea of putting computing and digital making at the heart of modern education, and we feel there's a need to do more to connect with and support educators, both inside and outside the classroom.

Q WHEN IS IT AVAILABLE?

A Your magazine is available three times per year. Check out our podcast (helloworld.cc/podcast) and monthly newsletter (helloworld.cc/newsletter) to get more great content between issues.



IT'S FREE!

Hello World is free now and forever as a Creative Commons PDF download. You can download every issue from helloworld.cc. Visit the site to see if you're entitled to a free print edition, too.

WANT TO GET INVOLVED?

There are numerous ways for you to get involved with the magazine — here are just a handful of ideas to get you started

- **Give us feedback**

Help us make your magazine better — your feedback is greatly appreciated.

- **Ask us a question**

Do you have a question you'd like to share? We'll feature your thoughts and ideas.

- **Tell us your story**

Have you had a success (or failure) you think the community would benefit from hearing about?

- **Write for the magazine**

Do you have an interesting article idea? Visit helloworld.cc/writeforus to submit your idea.

GET IN TOUCH

Want to talk? You can reach us at:
contact@helloworld.cc

FIND US ONLINE

www.helloworld.cc

 [@HelloWorld_Edu](https://twitter.com>HelloWorld_Edu)

 fb.com/HelloWorldEduMag

**NEXT
ISSUE OUT
JANUARY 2024**

THEME: **COMPUTING
EDUCATION AROUND
THE WORLD**



helloworld.cc



Raspberry Pi
Foundation