# EXPLORING THE AI JUNGLE

## TEACHERS' GUIDE

# TABLE OF CONTENTS

This is an early edition of notes for teachers to accompany the children's picture book "Exploring the AI Jungle". The book is intended to teach children about Generative AI and its wider implications for society, particularly for schools.

# INTRODUCTION

***The use of AI, particularly GenAI, in schools is a hot topic. There has been anxiety among adults about the possibility of using GenAI tools to cheat at school. What if children use ChatGPT to do their homework? What if they use it to get As in all their exams? Our research with young people (secondary school and university learners) suggests that they, too, dislike the idea of cheating with AI tools and are wary of becoming lazy by using them to learn. Clearly, no one wants to end up in a state where we end up using AI tools to mark assessments written by other AI tools.***

Concerns about using AI to cheat distract from the broader issues of how AI can support learning for everyone and what children need to learn at school in the Age of AI. If a machine can do a task well, does a human still need to learn how to do it? I believe that we should aim for humans and AI tools to work together to solve complex and interesting problems to improve human lives. AI tools and humans have different strengths and weaknesses when solving problems; if we can use AI tools to address human weaknesses, we can free our attention for more important matters.

Concerns about using AI to cheat distract from the broader issues of how AI can support learning for everyone and what children need to learn at school in the Age of AI.

## IF A MACHINE CAN DO A TASK WELL, DOES A HUMAN STILL NEED TO LEARN HOW TO DO IT?

We need to distinguish between teaching basic concepts and skills which are necessary building blocks for working within more complex areas, and those we're simply in the habit of teaching. Of course, there are personal preferences for how much individuals feel comfortable delegating to AI. When I'm writing this, I'm happy that AI tools take care of spelling and some grammar because they let me focus on the ideas I want to express. But I'm not willing to let a GenAI tool change the meaning of my sentences or come up with ideas. In contrast, some of my students like using GenAI tools for brainstorming to help them start a project.

There isn't a single 'correct' way to use GenAI tools for learning, and there's not currently a lot of research evidence to support what is effective. A whole new world is waiting for learners and teachers to explore together. This book is part of a wider project in which my colleagues and I are consulting young people at secondary and special schools about their hopes, fears and concerns about AI Futures in schools.

The Children's Parliament in Scotland recently carried out a project about AI with primary schools.[1] These projects are based on the principle that children have a right to be consulted about their digital lives as part of their rights enshrined in the UN Convention on the Rights of the Child (UNCRC).[2] This is why the book invites readers to think and discuss what role GenAI tools should use in schools. I hope school leaders, teachers and learners will find it useful as they plan policies for using AI in their schools.

# ABOUT THE BOOK


***The format of the book was inspired by the style of the excellent book All Dogs Have ADHD by Kathy Hoopman.[3] We wanted to convey basic facts about AI in clear language, with the pictures serving as a memory aid. Like All Dogs Have ADHD, it's certainly not a complete guide to such a complex topic, but it is a place to start.***

This book doesn't explain how to use GenAI tools in the classroom. Instead, it's intended to cover basic facts about GenAI, encourage children to think about how they would like to see it used in schools and spark interest in wider discussions about AI in society.

Each page has an activity suggestion for research or discussion tasks, which will prompt the children to consider some of the more philosophical aspects of AI. GenAI tools are cool gadgets, some of which are useful for learning. But AI in general has the potential to be a lot more than just another software tool. Humans must seriously consider the possibility that we will create entities that are at least as intelligent as people– perhaps far more intelligent. Many respected scientists believe we will achieve Artificial General Intelligence (AGI) in the next thirty-five years, while some believe it will only take five years.[4] This is exhilarating and possibly terrifying! As a society, we shouldn't allow this to happen without thinking it through carefully. Children – and adults- need to be AI literate. In addition to knowing how to use AI tools, we need knowledge of how AI works, its strengths and limitations, and the capacity to think critically about how it may help or hinder society.

## TEACHING WITH METAPHORS

***This book is funded by a research grant from the Arts and Humanities Research Council on a programme called Bridging Responsible AI Divides. The programme's purpose is to discuss the broader societal implications of AI using arts-based approaches. This gave my colleagues and me the creative freedom to tackle the topic of AI in Schools in new ways, which is why we produced a picture book, a map poster, and an AI stage show as well as academic reports.***

Metaphors can be a powerful way to see things from a new perspective, particularly if they are fresh comparisons rather than well-worn clichés. The extended jungle metaphor – inspired by a comment during a research interview – brings a different way to look at technology as an ecosystem because it emphasises the interconnections between different parts of school life (learners, teachers, and the technology tools they use).

The jungle metaphor invites a family of smaller metaphors for sub-topics – the jungle inhabitants represent various aspects of GenAI tools, educators and assessment authorities. The reader of the book is in the metaphorical role of explorer. Eventually, I hope, learners and teachers will assume the role of guardians of the AI School jungle by cultivating it and reshaping it to become a fascinating place where all learners thrive.

# TEACHING WITH METAPHORS

There is a further explanation of how the jungle metaphor relates to AI in the preface. I hope that the jungle metaphor will make the topic of GenAI engaging and accessible for learners who are not usually drawn to technology. At the very least, it's a nice change from the usual illustrations of robots with blue lights for eyes which usually crop up in educational materials about AI.

I am aware that metaphors may not be suitable for some learners who prefer literal language. They may prefer to use the metaphor-free version of the main text in the book that still conveys the facts about GenAI but without all the animals. You can download this version here: **https://dataed.in/jungle1**

For a helpful discussion of metaphors as semantic waves in the context of teaching computing concepts, see Raspberry Pi Foundation's Pedagogy Quick Read 6. The purpose of using metaphors in this book is slightly different. We don't use metaphors to explain specific technical topics. We use them to invite new perspectives about how human learners and GenAI tools could and should interact. However, in both cases, it's important to explain where the metaphor breaks down so that learners don't mistakenly assume that all aspects of the metaphor apply to the target topic. For example, spider monkeys are intelligent in some respects. Their social intelligence raises an interesting point about how current GenAI tools tend not to interact with other GenAI tools. But I couldn't think of a good way to fit other aspects of spider monkey behaviour with GenAI behaviour.

Did you know that spider monkeys help trees to spread throughout the jungle by eating seeds and then poo-ing them out as they swing from one area to another? Working that into the GenAI metaphor was beyond me. As another example, the original ChatGPT-generated text for the chameleon page mentioned that their eyes allow them to see in different directions at once, and this was similar to GenAI tools, which "can handle multiple types of tasks simultaneously".

When I discussed this with colleagues, we decided that this was misleading as it incorrectly suggests that GenAI tools can process requests in parallel. So, it is useful to think with your learners about which aspects of the metaphor work and highlight aspects which may cause confusion.

## MORE METAPHOR ACTIVITY SUGGESTIONS

***Ask your learners if there is a point about GenAI that they feel is important but is not in the book (yet). Ask them to write the text and decide which metaphorical animal should represent it. You can find an editable powerpoint of the text here: https://dataed.in/jungle1. I'd love to see the extra pages #ExploreJungleAI***


Can your learners think of other animals which could represent ideas in the book? My co-authors and I had an argument about which animal should be on page 28 to represent the possible dangers of GenAI. We settled on a jaguar, mostly because I disapprove of snakes. Feel free to insert alternative animal pictures in the editable powerpoint.

# TEACHING WITH METAPHORS

*A jungle is only one ecosystem – what would a book called Exploring the AI Ocean be like? Maybe mimic-octopi would get the limelight as ChatGPT suggested! Our team would love to see your books and share them with other schools.* **#ExploreJungleAI**

## AGE AND STAGE

I think that upper-primary school children will enjoy this book (aged 10- 12 years), but you may find it works with other groups of learners you work with. The main text on each page was intended to be as clear as possible to put across the key ideas, while the text in the coloured blobs is designed to be slightly more information-dense to satisfy those learners who seek out information like monkeys foraging fruit. Feel free to give a copy of this guide to particularly voracious monkeys in your class.

The activity suggestions and questions can lead to grappling with some big ideas, which are likely to appeal to teenagers too. If they think the animals are too childish, you could always give them the metaphor-free version: **https://dataed.in/jungle1**

If you are interested in further information about useful applications of AI systems, try this Scientific American article: **https://dataed.in/jungle2**

## EXPERIENCES AND OUTCOMES

For teachers working in Scotland, the Teachers' Notes draw attention to the Experiences and Outcomes within the Curriculum for Excellence relevant to each page. We have selected the Level 2 outcomes, although the related ones at first, third and fourth levels may also apply. Teachers who work outside Scotland may spot equivalent entries in their own curriculum frameworks. In general, the book relates to these curricular areas:

**Developing Digital/AI literacy:** developing a critical understanding AI and the role new technologies play in society.

**Developing information literacy:** developing skills in finding information online and introducing the problem of misinformation.

**Contributing to a community:** participating in discussions about how AI should be used in schools, which can be helpful in the development of a school-wide AI policy.

# CHAMELEON

## BACKGROUND INFORMATION

***There are other types of artificial intelligence in addition to GenAI, but it has received a lot of attention since ChatGPT was released in 2022. AI techniques are also used in recommender systems (like song suggestions based on what tracks you have previously listened to), smart speakers such as Alexa, home appliance robots such as Roomba, and (more experimentally) self-driving cars.***

Many AI systems, including GenAI rely on an approach called Machine Learning, which uses statistical approaches to predict patterns based on large datasets of examples. Machine Learning systems require a lot of computing power and data to train them.

# PROMPTS FOR LEARNERS

Find out about a new GenAI system which has been released recently. What can it do, and how well does it do it?

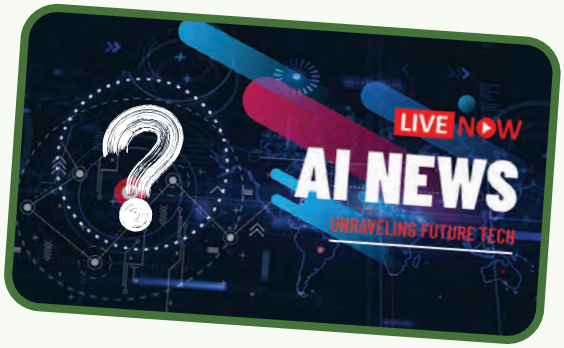


The key aspect of the prompt for developing critical thinking is "how well does it do it?" Often GenAI tools are a bit flakey to begin with but improve with new releases. Sometimes, the technology company releases a video of a demo of a system, which seems impressive, but users can't actually try it yet.

It's important to stay up to date with what new GenAI systems can do, as there are new releases weekly! So why not get your learners to stay on top of it for you? You could assign a different group to give an update every week or every month.

Questions for your learners to consider when they are evaluating a news release about a new AI system include:

— *What does the system do?*

— *Who made the system?*

— *Who published the news story about the system? Is it a trustworthy source of information?*

— *What benefits could this system have for learning?*

— *Could there be any drawbacks to this system?*

— *Do you need any more information about this system to evaluate it?*

Using what I know about the features of different types of texts, I can find, select and sort information from a variety of sources and use this for different purposes. **LIT 2-14a**

To help me develop an informed view, I can distinguish fact from opinion, and I am learning to recognise when my sources try to influence me and how useful these are. **LIT 2-08a**

I can extend and enhance my knowledge of digital technologies to collect, analyse ideas, relevant information and organise these in an appropriate way. **TCH 2-01a**

I can use evidence selectively to research current social, political or economic issues. **SOC 2-15a**

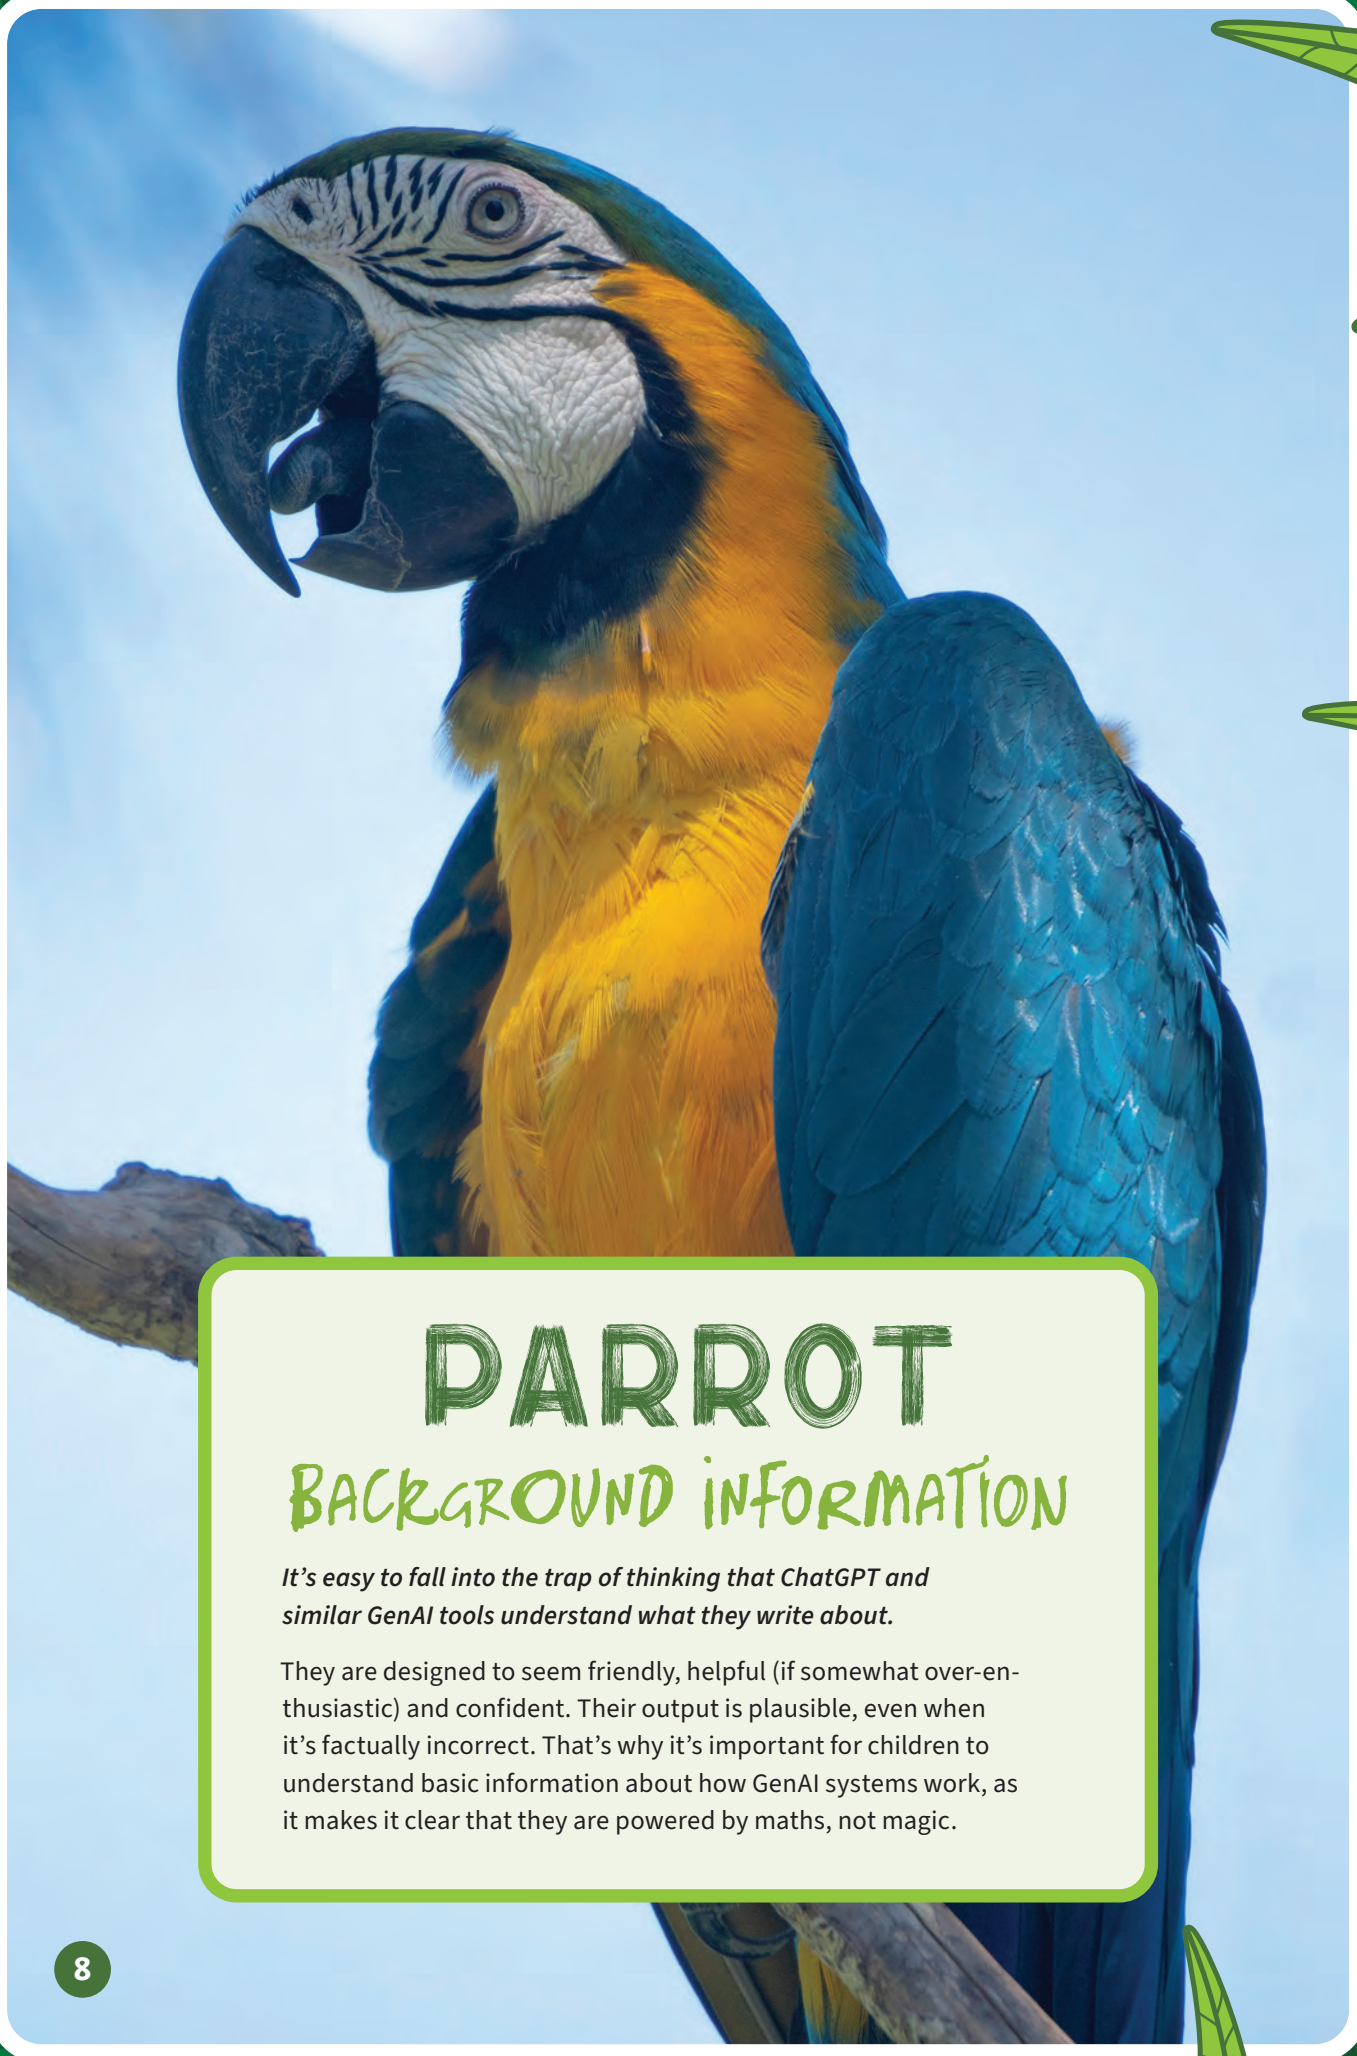

# PARROT

## BACKGROUND INFORMATION

***It's easy to fall into the trap of thinking that ChatGPT and similar GenAI tools understand what they write about.***

They are designed to seem friendly, helpful (if somewhat over-enthusiastic) and confident. Their output is plausible, even when it's factually incorrect. That's why it's important for children to understand basic information about how GenAI systems work, as it makes it clear that they are powered by maths, not magic.

# PROMPTS FOR LEARNERS

Find out more how GenAI systems are trained. How much training data does it need? More or less than your average parrot?

## HOW GENAI SYSTEMS ARE TRAINED

At a basic level, GenAI systems are trained on datasets containing lots of examples. Systems which generate written text are trained on (almost) all the documents on the internet. During the training process, the system uses an algorithm and a lot of maths to spot patterns, such as words which commonly go together. It stores these patterns in a model. After training, the model is usually checked by humans to try to prevent it from producing dangerous or offensive content. When a user comes along to ask the system to write some text, the system uses its model to predict which word, phrase or sentence is most likely to come next in a sequence. GenAI systems which generate images or videos are trained in a similar way, except that they use training sets of images or video content respectively.


https://dataed.in/jungle3

*This video has a reasonably accessible explanation of how large language models like ChatGPT are trained (0.39-6.13).*

If you'd like a regular reminder of GenAI's lack of understanding and the reasons behind it try subscribing to Gary Marcus's substack: **substack.com/@garymarcus**

Using what I know about the features of different types of texts, I can find, select and sort information from a variety of sources and use this for different purposes. **LIT 2-14a**

I can extend and enhance my knowledge of digital technologies to collect, analyse ideas, relevant information and organise these in an appropriate way. **TCH 2-01a**

To help me develop an informed view, I can identify and explain the difference between fact and opinion, recognise when I am being influenced, and have assessed how useful and believable my sources are. **LIT 2-18a**

I can use evidence selectively to research current social, political or economic issues. **SOC 2-15a**

# PROMPTS FOR LEARNERS

Find out more how GenAI systems are trained. How much training data does it need? More or less than your average parrot?

## HOW MUCH TRAINING DATA?

The answer depends on which GenAI the learners have investigated. I chose ChatGPT. It was trained on a dataset called Common Crawl, which contained about 400 billion words.

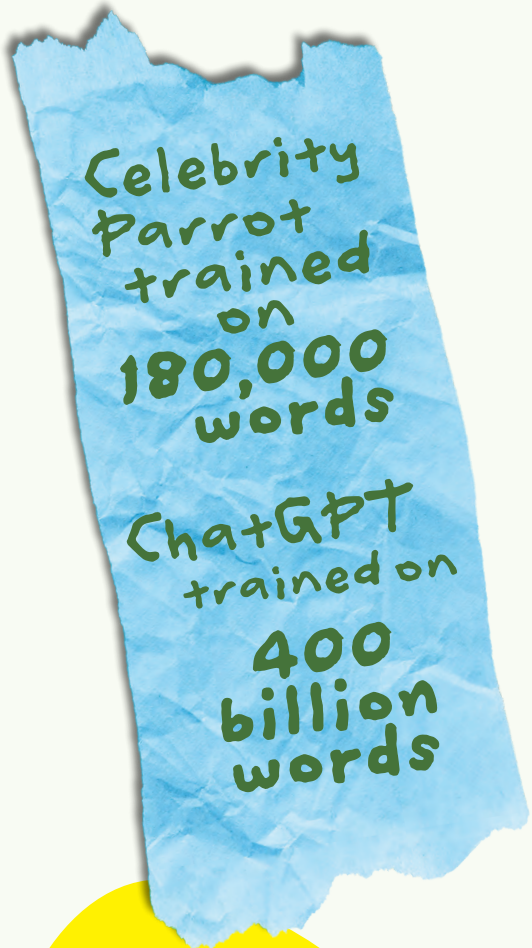


You can find the resources for class "unplugged" activities, illustrating how GenAI systems work here: **www.dataschools.education/resource/is-this-a-zegah/**

## TRAINING A PARROT

I've never trained a parrot to talk, but I'm guessing they need a lot less training data than GenAI! According to this guide to training parrots (**https://dataed.in/jungle4**), you can teach them a word by repeating it several times a day for up to 5 minutes at a time. That's quite a lot of repetition of the same word, but how many different words can parrots learn? There is a celebrity parrot (**https://dataed.in/jungle5**) who knows 200 words or sounds, so let's take that as a reasonable estimate. There are 300 seconds in 5 minutes. Suppose you said a training word per second and you did three training sessions with your parrot a day–that would be 900 training words per day. If the parrot could learn one word a day, it would take 200 days to learn 200 words, which would be a total of 180,000 training words. This is clearly a guestimate because who knows how long it takes a parrot to grasp a word? It could take longer than a day. Maybe it forgets words it already knows. But a back-of-an-envelope guestimate is enough for our comparison.

***In short, if ChatGPT is a stochastic parrot, it's a parrot with an astonishingly good memory. It is trained with many more words, but then it knows more words than a parrot and can string them together into sentences, paragraphs and documents which (mostly) make sense. But the key point remains: GenAI systems don't understand what they generate.***

# POISON DART FROG

## BACKGROUND INFORMATION

*As GenAI systems are prone to what their developers politely call "hallucinating" and everyone else calls "making stuff up", there are lots of examples of misinformation for the learners to find.*

# PROMPTS FOR LEARNERS

Find an example of misinformation generated by a GenAI tool (e.g.: search for AI, pizza and glue).


If you're looking for evidence to convince your learners that AI is not as clever as they think, try the AI weirdness blog. My favourite is a poster generated by DALL-E 3 intended to explain basic shapes: **https://dataed.in/jungle6**


## AI MISINFORMATION

An important takeaway for learners is to double-check information from a GenAI tool. This article discusses the issue of AI generated misinformation, including the recommendation to use glue to stick cheese to pizza and to only eat one rock per day:

**https://dataed.in/jungle7**

Using what I know about the features of different types of texts, I can find, select and sort information from a variety of sources and use this for different purposes. **LIT 2-14a**

To help me develop an informed view, I can distinguish fact from opinion, and I am learning to recognise when my sources try to influence me and how useful these are. **LIT 2-08a**

I can extend and enhance my knowledge of digital technologies to collect, analyse ideas, relevant information and organise these in an appropriate way. **TCH 2-01a**

I can use evidence selectively to research current social, political or economic issues. **SOC 2-15a**

# OWL BUTTERFLY

## BACKGROUND INFORMATION

***The point here is that GenAI systems do not currently have sinister intentions, and so if they give false information, it is either a mistake or the propagation of false information they were trained on.***

However, as indicated by the activity suggestion, humans with sinister intentions could misuse GenAI tools for their own ends – such as for political gain or spreading hate speech. The amount of misinformation on the Internet is spreading faster than ever as new GenAI systems are trained on false information generated by current GenAI systems.

# PROMPTS FOR LEARNERS

Imagine you firmly believe that the moon is made of green cheese. How would you use AI systems to spread your cheesy word?

## IF I WERE A "GREEN-CHEESER"...

1. If I were a green-cheeser, I would probably prompt ChatGPT by telling it that it is a conspiracy theorist who wants to spread the word about green cheese and ask it to come up with a marketing plan.
2. I'd get it to generate a lot of articles for different audiences with the same cheesy message.
3. I'd ask it to make a lot of new (seemingly different) websites which would back up the information.
4. If I was running short of cash to pay for my websites or adverts, I'd ask ChatGPT to come up with a money making scheme for me.

This website explains how to spot misinformation and how to talk about it with children:

**www.security.org/digital-security/misinformation/#talk**

To help me develop an informed view, I can distinguish fact from opinion, and I am learning to recognise when my sources try to influence me and how useful these are. **LIT 2-08a**

# LEAFCUTTER ANTS
## BACKGROUND INFORMATION

***Humans are inclined to attribute human-like motivations and emotions to other animals or even objects (see this photo of a drunk washing machine if you don't believe me: https://dataed.in/jungle8).***

This may be a side effect of being born with brains wired for social interaction, but it can be misleading. With leafcutter ants and GenAI tools alike, we misattribute rule-following to social behaviour.

# PROMPTS FOR LEARNERS

Would you rather get advice from a system like ChatGPT or a human friend? Why? Does it matter what you need advice about?

For detail-orientated readers: when the text mentions biological rules in the ants' nature, I am referring to genes, which encode rules to some approximation. GenAI tools follow rules in the sense that they execute an algorithm written by a human during their training stage. Once a model is trained, it generates output based on the patterns it extracted from its training data. The model wasn't directly created by a human and the patterns are probabilities rather than rules. In spite of these finer points, I think the metaphor is useful.

Our previous research found that primary school-aged children are unsure whether AI systems like Alexa have human-like feelings. This is not surprising because they are often designed to appear human-like as it makes them easier to use. They appear friendly, so they can give the misleading impression that they can be a friend. Combine this with a cauldron of teenage hormones and the fact that GenAI systems aren't always accurate, and you have a recipe for heartache.

***However, young people are often more sophisticated and knowledgeable about the digital world than adults give them credit for. Some people (young and old) prefer talking to a non-judgemental machine to work through their problems.***

This page could prompt a useful and interesting discussion with your learners about how they seek help when they have problems. It's important that your learners know who to turn to if they need a friendly human to talk to and that they are aware of the limitations of talking to an AI.

When I engage with others, I can respond in ways appropriate to my role, show that I value others' contributions and use these to build on thinking. **LIT 2-02a**

I understand that there are people I can talk to and that there are a number of ways in which I can gain access to practical and emotional support to help me and others in a range of circumstances. **HWB 0-03a / HWB 1-03a / HWB 2-03a / HWB 3-03a / HWB 4-03a**

# STRANGLER FIG

## BACKGROUND INFORMATION

***Despite arguments about this page with colleagues and family members, at the moment, I stand by the claim that GenAI isn't very creative.***

GenAI tools were built by humans and are trained on vast datasets of all the stuff humans see fit to put on the Internet. There is the occasional gem of human creativity in there among all the cat pictures. While some AI systems can be arguably creative (this is a field of research in its own right), the machine learning approach used in Generative AI systems is designed to generate the most plausible output. As this is often highly related to the content they encounter during training, output from GenAI tools can often seem bland or mediocre.

# PROMPTS FOR LEARNERS

Find out what happens when strangler figs get really big. Do you think something similar would happen with a GenAI system?

***This page is another invitation for open discussion, as once the children have learned about the habits of strangler figs, it's all a matter of opinion.***

Strangler figs can eventually kill their host trees by cutting off their nutrients, water and sunlight (**https://dataed.in/jungle9**). Killer AI is a common theme in science fiction stories. What if the machines that we designed to help us get their own ideas and want to take over? Your learners may be familiar with Terminator or The Matrix where AI is the villain. But maybe they're young enough to be influenced by BayMax and WALL-E, and are more inclined to take an optimistic view.

In fact, there are serious concerns about the breach of copyright by GenAI tools and the impact this will have on the livelihoods of artists and writers. In our recent research, some young people raised this as an issue and explained that they did not want to use GenAI tools as a result.


Remember that if children or young people do not want to use GenAI tools, the school needs to provide an appropriate alternative educational activity.

When I engage with others, I can respond in ways appropriate to my role, show that I value others' contributions and use these to build on thinking. **LIT 2-02a**

I can persuade, argue, explore issues or express an opinion using relevant supporting detail and/or evidence. **LIT 2-29a**

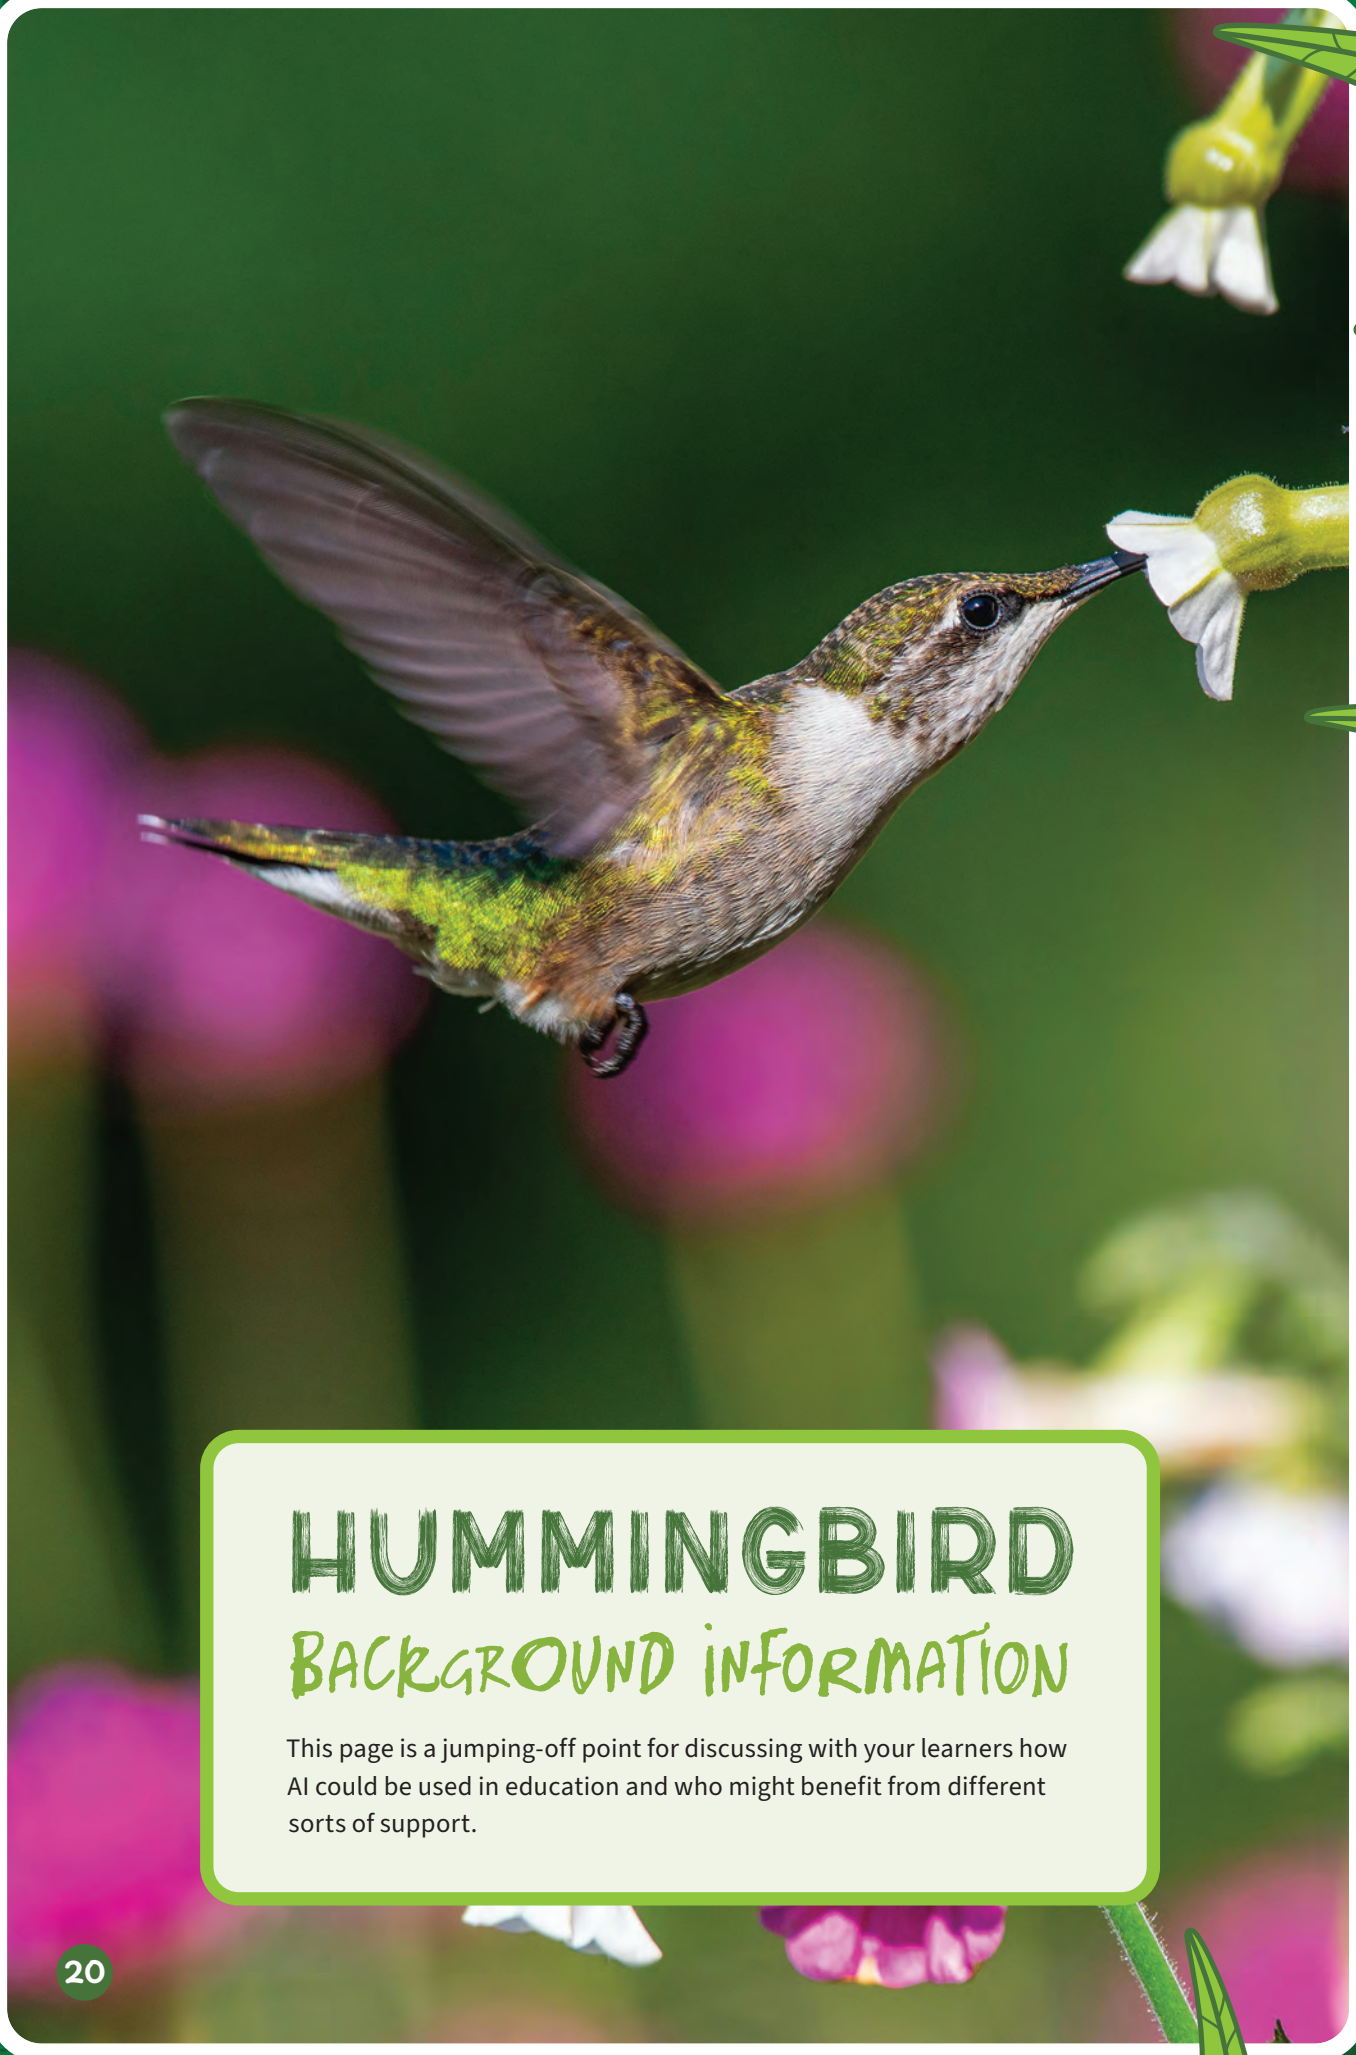

# HUMMINGBIRD

## BACKGROUND INFORMATION

This page is a jumping-off point for discussing with your learners how AI could be used in education and who might benefit from different sorts of support.

# PROMPTS FOR LEARNERS

In what ways could GenAI make learning more interesting or fun for different people?

***You could ask your class if they can think of ways in which AI is already used for learning. Their answers might include:***

— Spelling and grammar checkers use AI techniques. They are useful to almost everyone, and particularly for people with dyslexia.

— Learners could use tools like ChatGPT to summarise texts which are long or difficult to read. This could be useful for some learners with additional support needs like ADHD.

— AI tools can help learners with visual impairments by reading text aloud.

— Teachers could save time by using AI tools to plan interesting lessons.

— Image generation tools like DALL-E can be used to illustrate text to make lessons more appealing for everyone.

— Personalised AI tools could act as tutors to give learners lessons tailored to their own pace and preferences for learning.

When animals go into torpor **(https://dataed.in/jungle10)** it's like suspended animation. They drop their body temperature and their brain activity is reduced. Like the GenAI system when it is not in use, they are not thinking or learning. This is different to human sleep because our brains are active even while we sleep, making sense of the events of the day.

As I explore the rights to which I and others are entitled, I am able to exercise these rights appropriately and accept the responsibilities that go with them. I show respect for the rights of others. **HWB 0-09a / HWB 1-09a / HWB 2-09a / HWB 3-09a / HWB 4-09a**

When I engage with others, I can respond in ways appropriate to my role, show that I value others' contributions and use these to build on thinking. **LIT 2-02a**

I can use evidence selectively to research current social, political or economic issues. **SOC 2-15a**

# PROMPTS FOR LEARNERS

Find out what GenAI systems do when they are not being used? Is it similar to torpor? In what way?

***If your class is likely to enjoy a philosophical discussion, you could use the prompt to start a conversation about AI consciousness.***

When GenAI systems are not being used, they do nothing. Once the user switches away from the window or tab of the system's interface, it won't get any processor power until the user uses it again.

The GenAI system is not off plotting world domination or brushing up on its facts. Your computer will use a small amount of memory to record what the GenAI system was doing when you stopped using it so it can return to your session seamlessly when you switch back to it.

Imagine if someone programmed ChatGPT to beg not to be switched off. Could you bring yourself to do it?

I included this prompt because it introduces the topic of what it means to be conscious or alive. Humans don't usually think twice about switching between windows of their computer, or switching it off at the end of the day. They correctly infer that no permanent damage will be done and that there is no 'self' inside the computer to be offended by the rudeness of being ignored.

I am developing my understanding that people have beliefs and values based upon religious or other positions. **RME 2-09b**

When I engage with others, I can respond in ways appropriate to my role, show that I value others' contributions and use these to build on thinking. **LIT 2-02a**

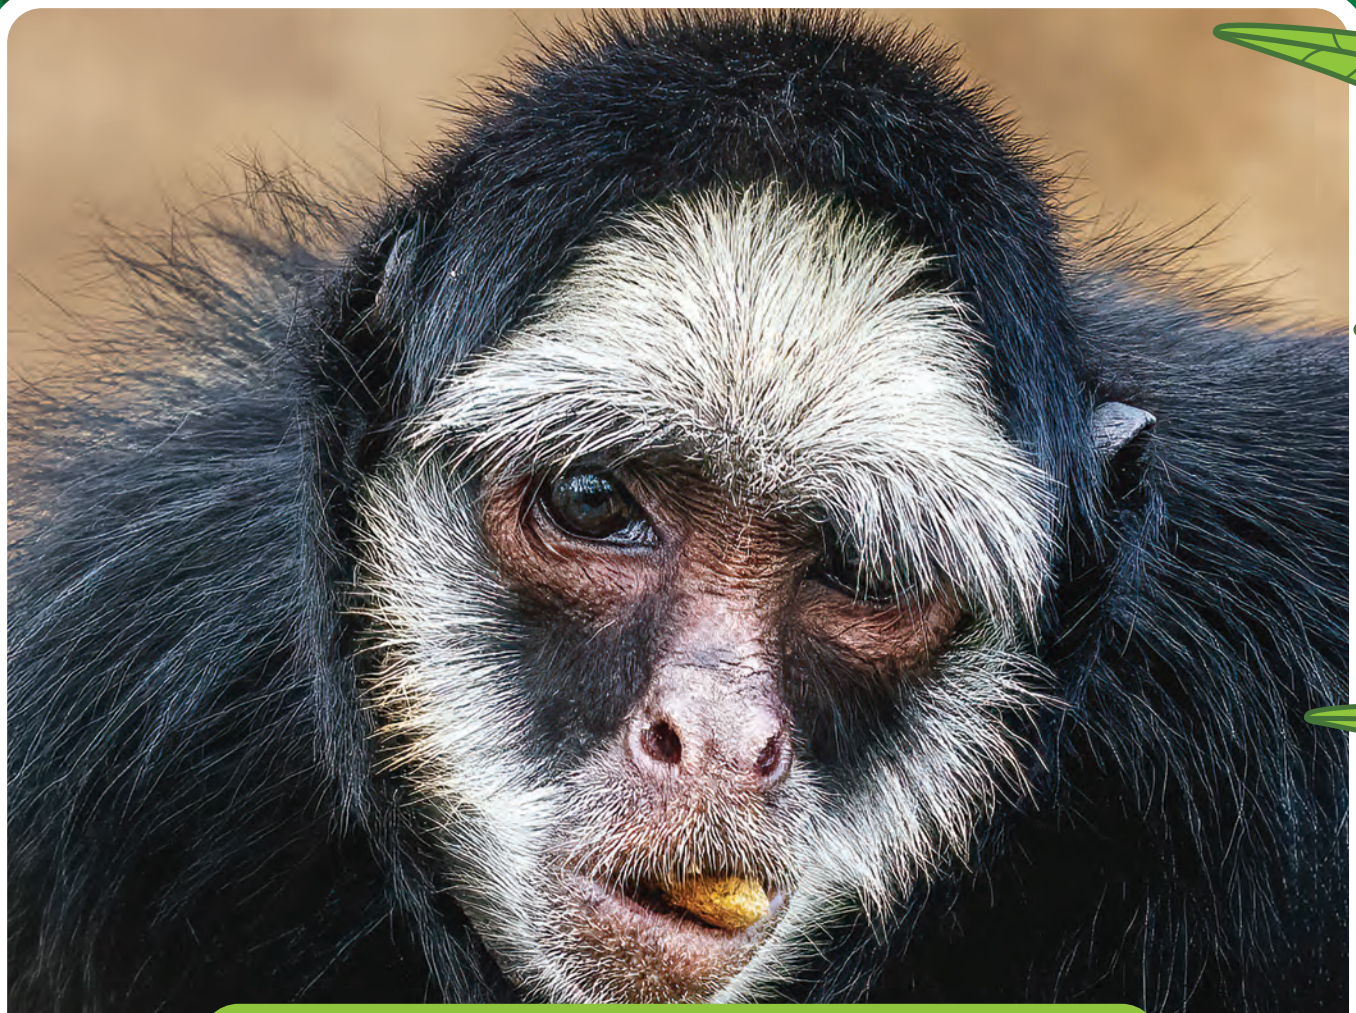

# SPIDER MONKEY
## BACKGROUND INFORMATION

***Although spider monkeys and GenAI tools could be said to show intelligent behaviour, spider monkeys have a different sort of intelligence from GenAI tools (obviously). Interestingly, some scientists think that spider monkeys, who are highly social, have evolved to have such good memories because it is useful for remembering members of the different social groups to which the monkeys belong. Today GenAI tools are not usually social – a single instance of a GenAI tool works with a single user. However, some new GenAI tools are designed to share tasks with each other. It's possible that one day you will have an AI agent which interacts with other AI agents on your behalf.***

When I say the GenAI tool "remembers ideas which it has read before" I don't mean that it stores facts like in a database. I mean that it stores associations between ideas in its model in a way conceptually similar to the way human brains store memories.

# PROMPTS FOR LEARNERS

Spider monkeys can recognize themselves in the mirror. This suggests they have a sense of who they are. Do you think AI tools will ever have a sense of self?

***This discussion prompt is an entry into a long-running philosophical debate about whether machines could ever think or be conscious. There isn't a right answer to this, and it's actually very hard to tell whether another being is conscious. But it's an interesting topic to discuss!***

Currently, there is no compelling evidence that machines including GenAI systems are conscious or have a sense of self , although a few AI scientists have made the claim that systems like GPT4 show "sparks of Artificial General Intelligence", and one former Google employee claimed that a Google AI system has become sentient. It is extremely confusing because GenAI systems like ChatGPT generate text which makes it seem as if they have a sense of self, or are sentient.

Some animals, like spider monkeys, do show an awareness of self, which is considered one ingredient of consciousness. If highly intelligent species like octopi and dolphins were conscious, they might experience it very differently from humans because they have such different sensory systems. Some philosophers have argued that only biologically-based lifeforms could ever be conscious, that it is unique to the slimy wetware of the human brain. I'm sure octopi philosophers would argue that it's unique to the slimy wetware of carbon-based brains. For other (human) philosophers, what a brain is made of doesn't matter. What matters is what the brain can do and how it behaves.


https://dataed.in/jungle11


https://dataed.in/jungle12

*These videos, made for a competition held by the Cognitive Science Society, are short introductions to the topic of whether machines can think, showing different sides to the philosophical debate.*

I am developing my understanding that people have beliefs and values based upon religious or other positions. **RME 2-09b**

When I engage with others, I can respond in ways appropriate to my role, show that I value others' contributions and use these to build on thinking. **LIT 2-02a**

I can use evidence selectively to research current social, political or economic issues. **SOC 2-15a**

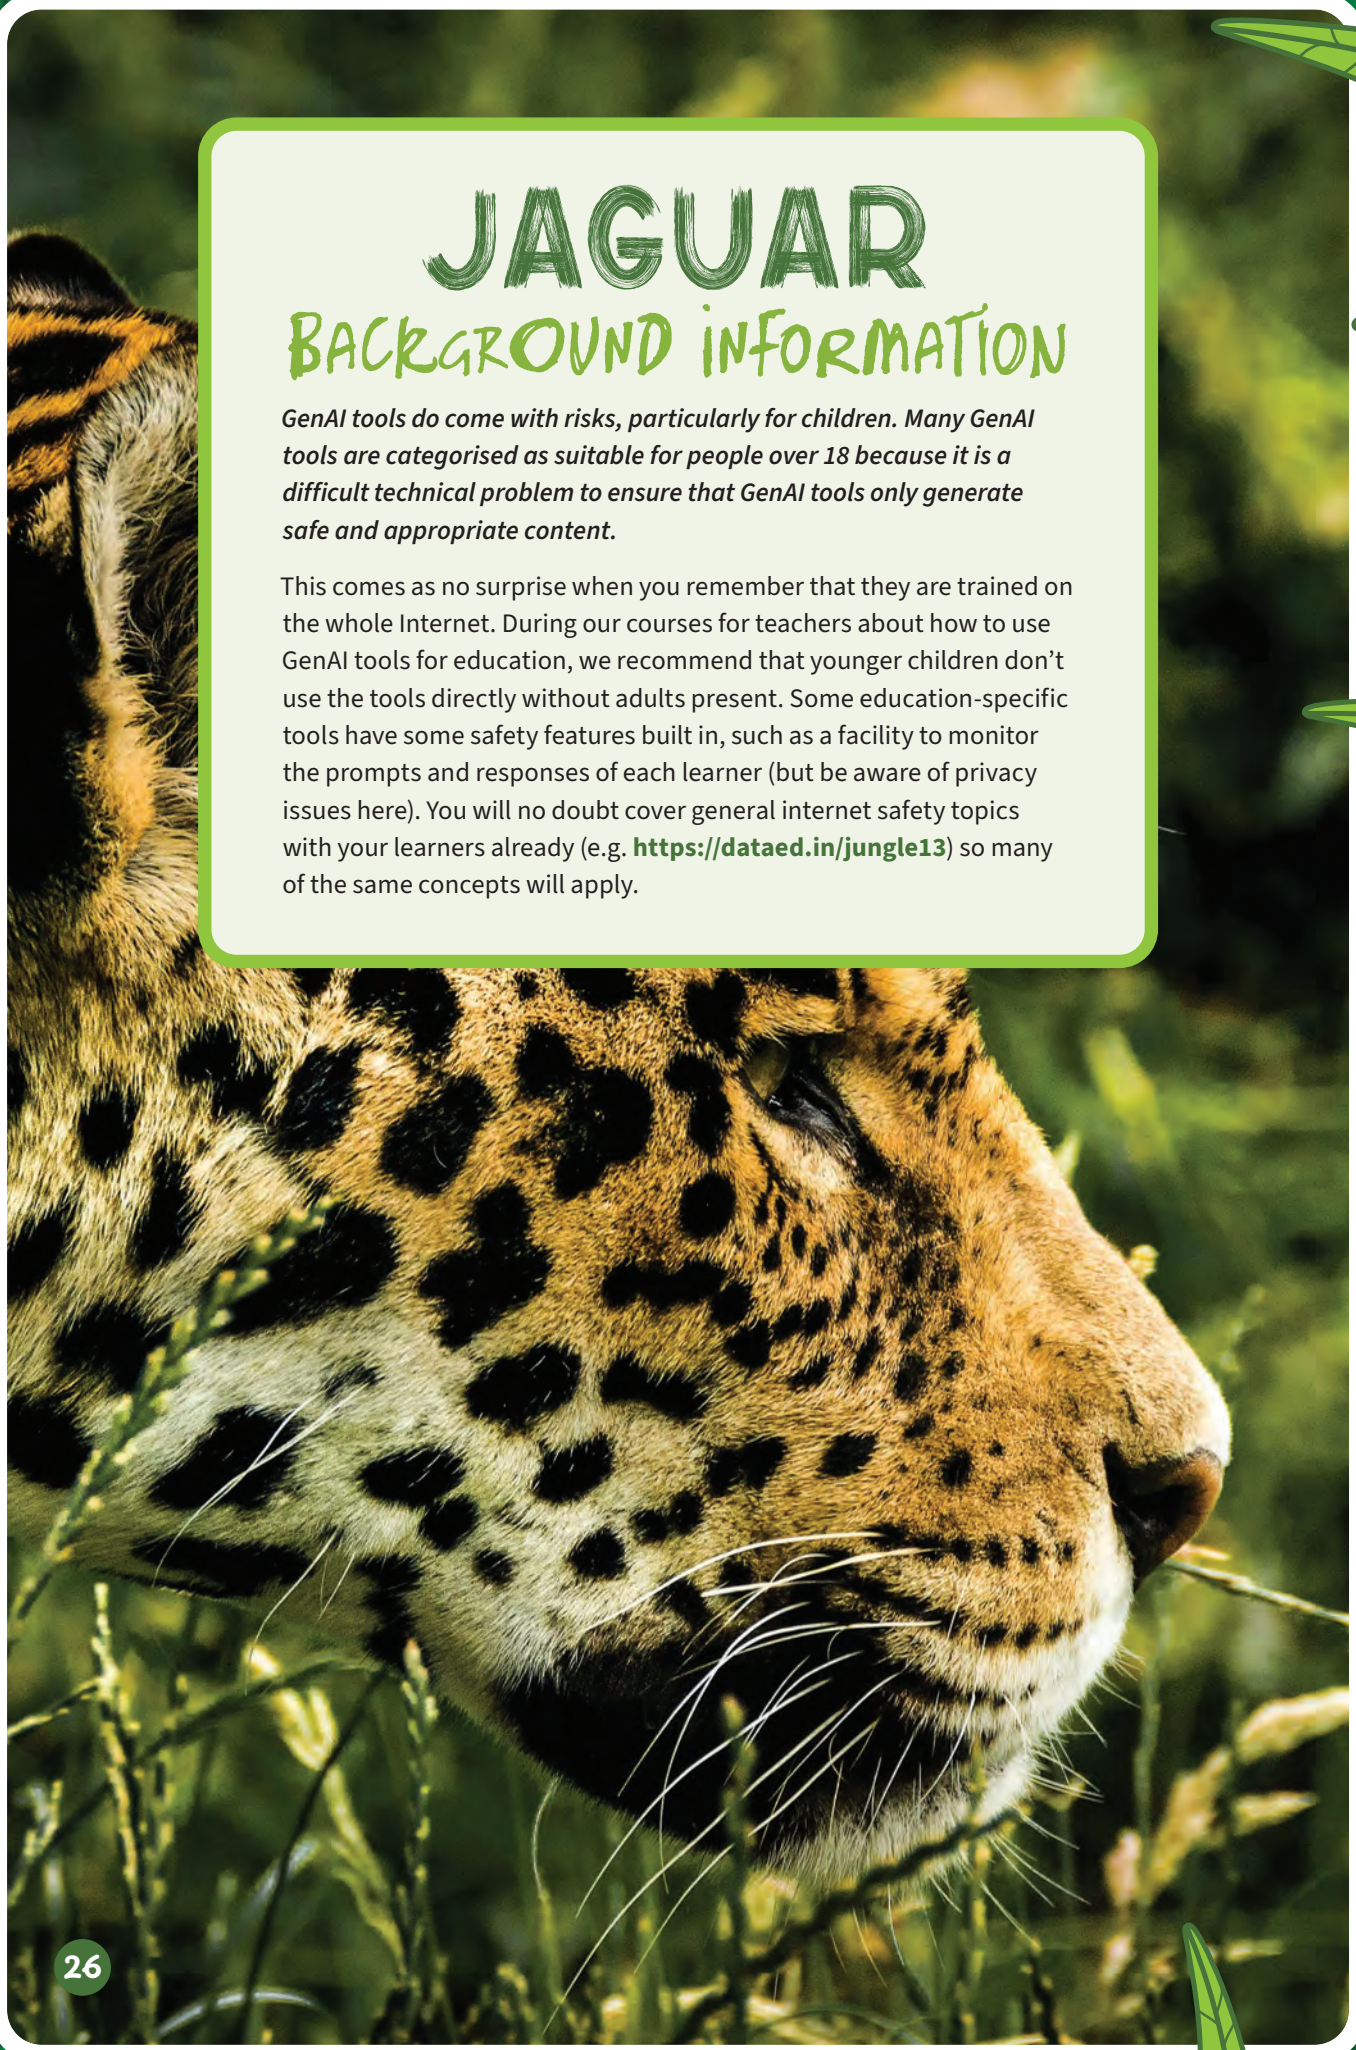

# JAGUAR

## BACKGROUND INFORMATION

***GenAI tools do come with risks, particularly for children. Many GenAI tools are categorised as suitable for people over 18 because it is a difficult technical problem to ensure that GenAI tools only generate safe and appropriate content.***

This comes as no surprise when you remember that they are trained on the whole Internet. During our courses for teachers about how to use GenAI tools for education, we recommend that younger children don't use the tools directly without adults present. Some education-specific tools have some safety features built in, such as a facility to monitor the prompts and responses of each learner (but be aware of privacy issues here). You will no doubt cover general internet safety topics with your learners already (e.g. **https://dataed.in/jungle13**) so many of the same concepts will apply.

# PROMPTS FOR LEARNERS

Find out about an AI system which is biased. What can designers do to reduce bias in their AI systems?

***Another risk for users of GenAI systems is privacy. It is important for children to know that they should not share personal information on a GenAI tool. Different technology companies have different policies about what happens to information users enter into their GenAI tools and whether it is used to train new models. Be careful that tools which attempt to safeguard children from inappropriate content do not also violate their digital privacy.***

Bias is also a risk in any AI system. This big topic fits well with other school topics, such as citizenship and inclusion. Investigations of various AI systems have shown them to make biased decisions that discriminate against groups of people. For example, some facial recognition systems do not accurately recognise people with darker-coloured skin, which can cause injustice when used in law enforcement. Problems of bias are often caused by bias in the data used to train the AI system. If the training data isn't representative of the whole population, then the model will have lower accuracy when it encounters new examples it has not seen in training. For example, the video below mentions an example where a technology company used AI to screen job applications. As it was trained on the applications from previous successful applications, and the technology industry is male-dominated, the AI system was biased against applicants who were not male. In general, AI systems which are trained on data containing human biases may also be biased.

To avoid or reduce or avoid bias in AI systems, users and developers need to work together to decide what they consider to be fair and unbiased output. This is harder than it sounds because different groups of users might have different ideas of what is fair, and they might not be compatible with each


***This is a great free interactive activity to explain bias in algorithms called Tech We Trust:*** **activity.techwetrust.scot/tech-we-trust/**

I can use evidence selectively to research current social, political or economic issues. **SOC 2-15a**

As I explore the rights to which I and others are entitled, I am able to exercise these rights appropriately and accept the responsibilities that go with them. I show respect for the rights of others. **HWB 0-09a / HWB 1-09a / HWB 2-09a / HWB 3-09a / HWB 4-09a**

I can gather and use information about forms of discrimination against people in societies and consider the impact this has on people's lives. **SOC 2-16b**

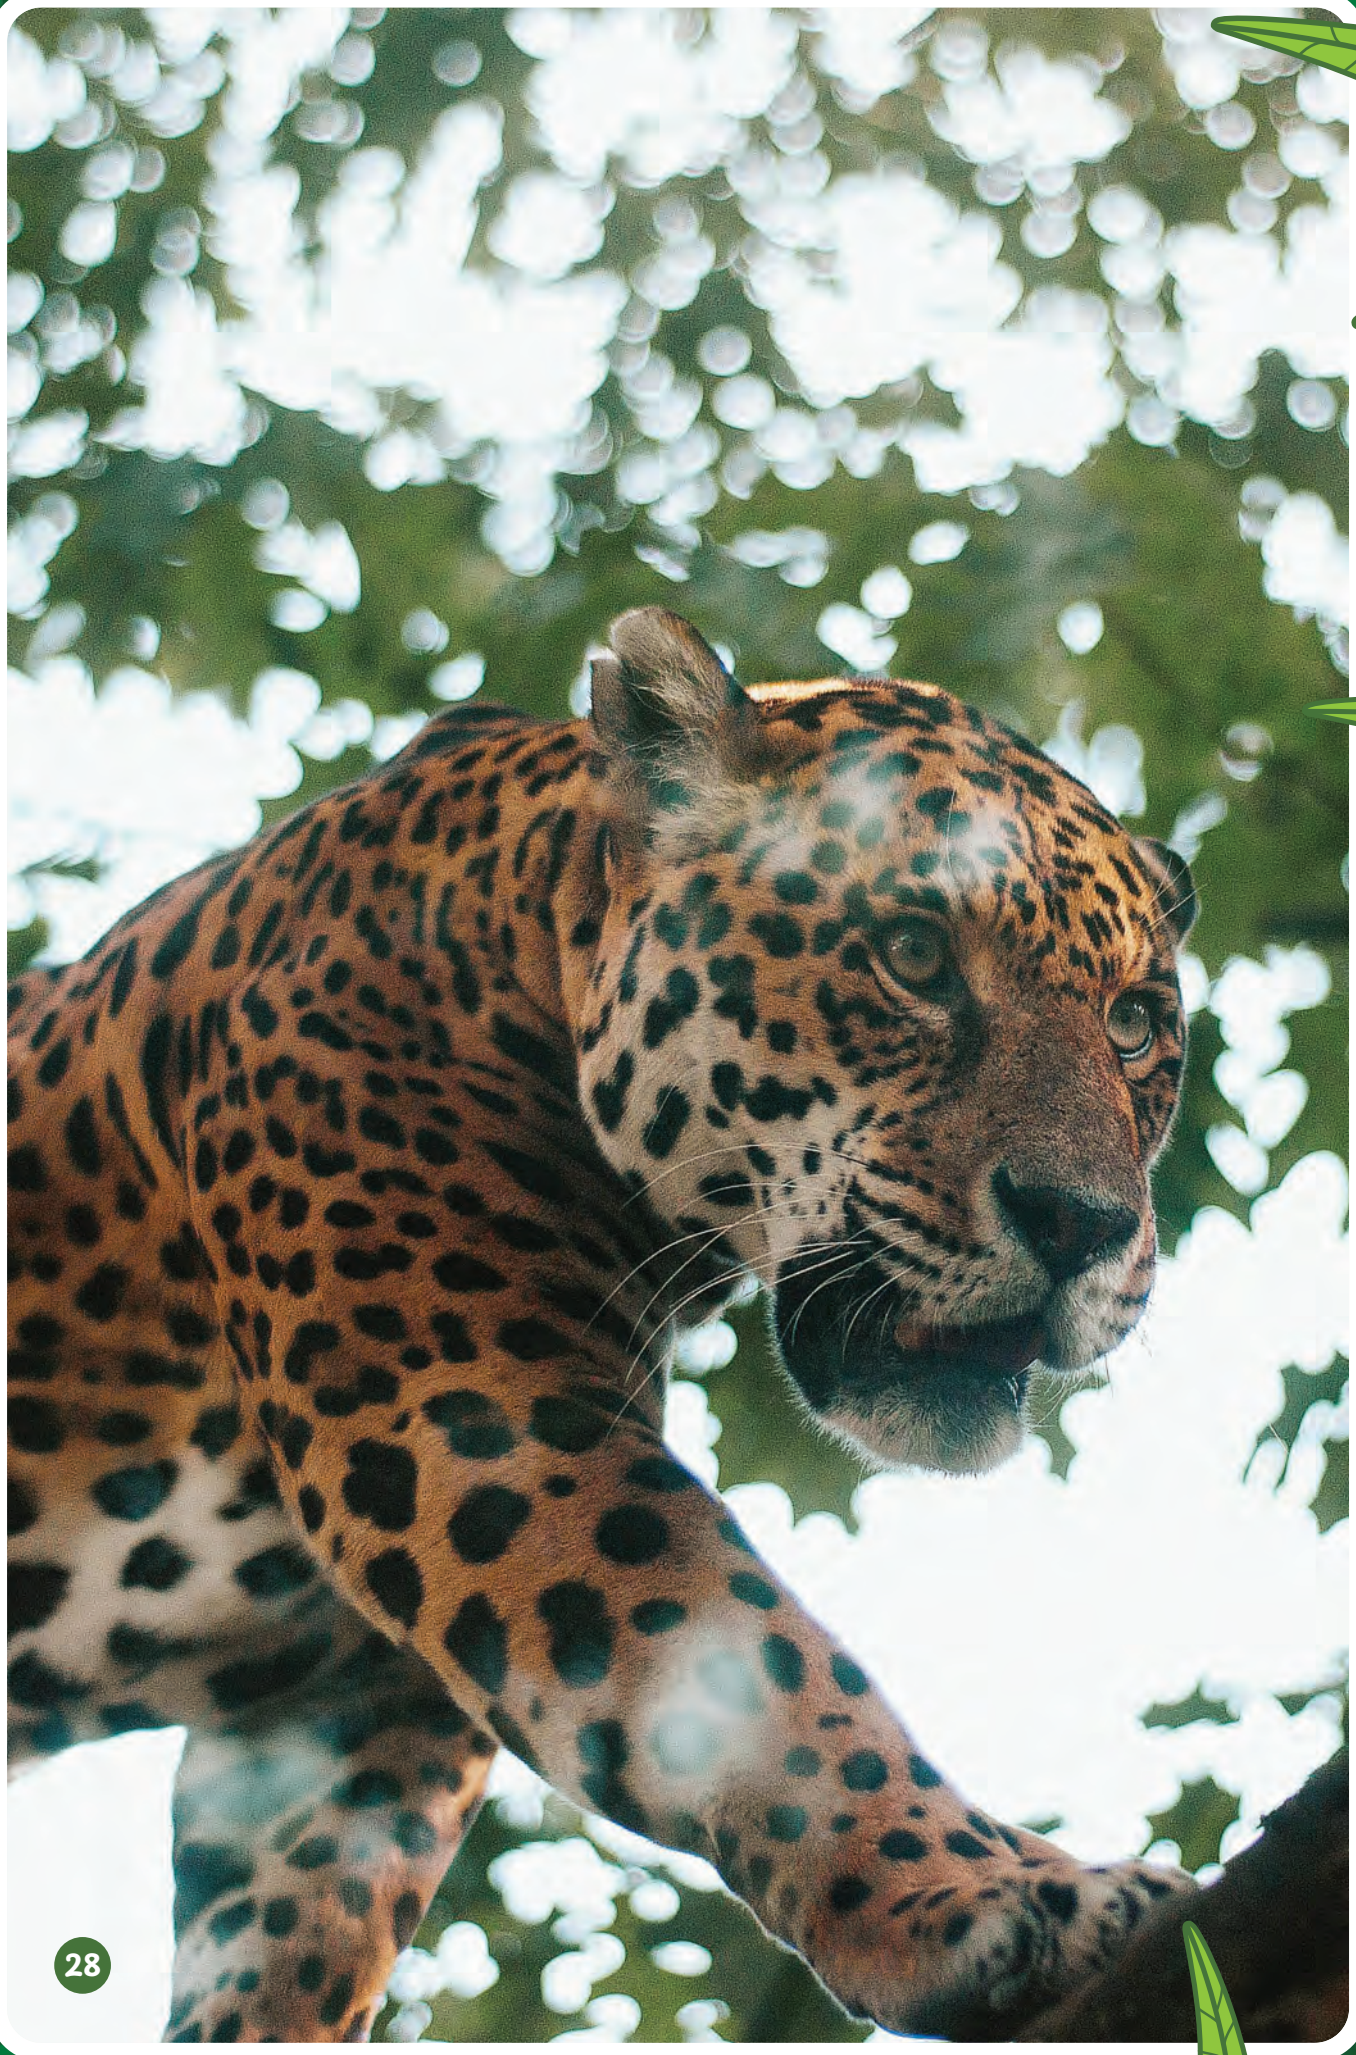

# PROMPTS FOR LEARNERS

Find out about an AI system which is biased. What can designers do to reduce bias in their AI systems?

other. The developers then need to ensure that the data they use to train the system won't cause bias and carefully test that the system's output meets the fairness standards they agreed to with the users. It's a good idea to keep checking that the system's output remains acceptable as time passes.

Guardrails designed to reduce bias in GenAI systems have inadvertently caused misleading content , which could be a problem in classrooms. At one point, if you asked the GenAI tool Gemini to create a picture of the Apollo 11 crew, it wouldn't draw the actual people who flew that mission - it would generate images of fictional people in astronaut suits with diversity in gender and ethnicity. Other image-generation tools have similar behaviour. While the guardrails are intended to prevent the propagation of biases (such as always drawing a doctor as male), this has come at the price of historical accuracy. The tools aren't yet smart enough to sensitively handle the regrettable bias of human history in addition to our aspiration for a bias-free future.

*A deeper dive into bias in algorithms, with some practical examples suitable for a more advanced level than the Tech We Trust activity:*

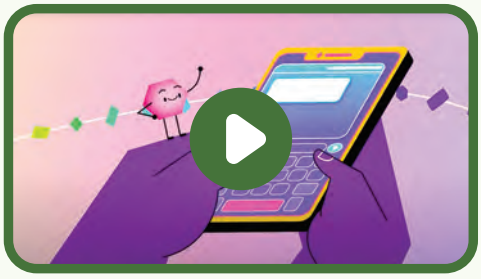

https://dataed.in/jungle14

*The follow-on video suggests some possible solutions against bias in algorithms:**

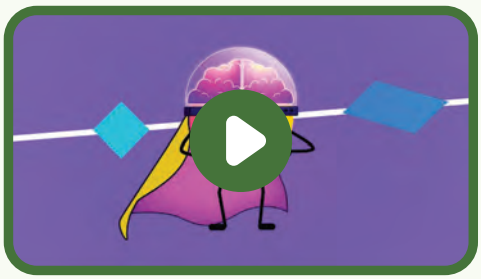

https://dataed.in/jungle14

To help me develop an informed view, I can identify and explain the difference between fact and opinion, recognise when I am being influenced, and have assessed how useful and believable my sources are. **LIT 2-18a**

I can persuade, argue, explore issues or express an opinion using relevant supporting detail and/or evidence. **LIT 2-29a**

I can extend and enhance my knowledge of digital technologies to collect, analyse ideas, relevant information and organise these in an appropriate way. **TCH 2-01a**

** Note: some learners in your class may need some support with the vocabulary in these videos, so you may want to watch them first in case you want to go over some definitions.*

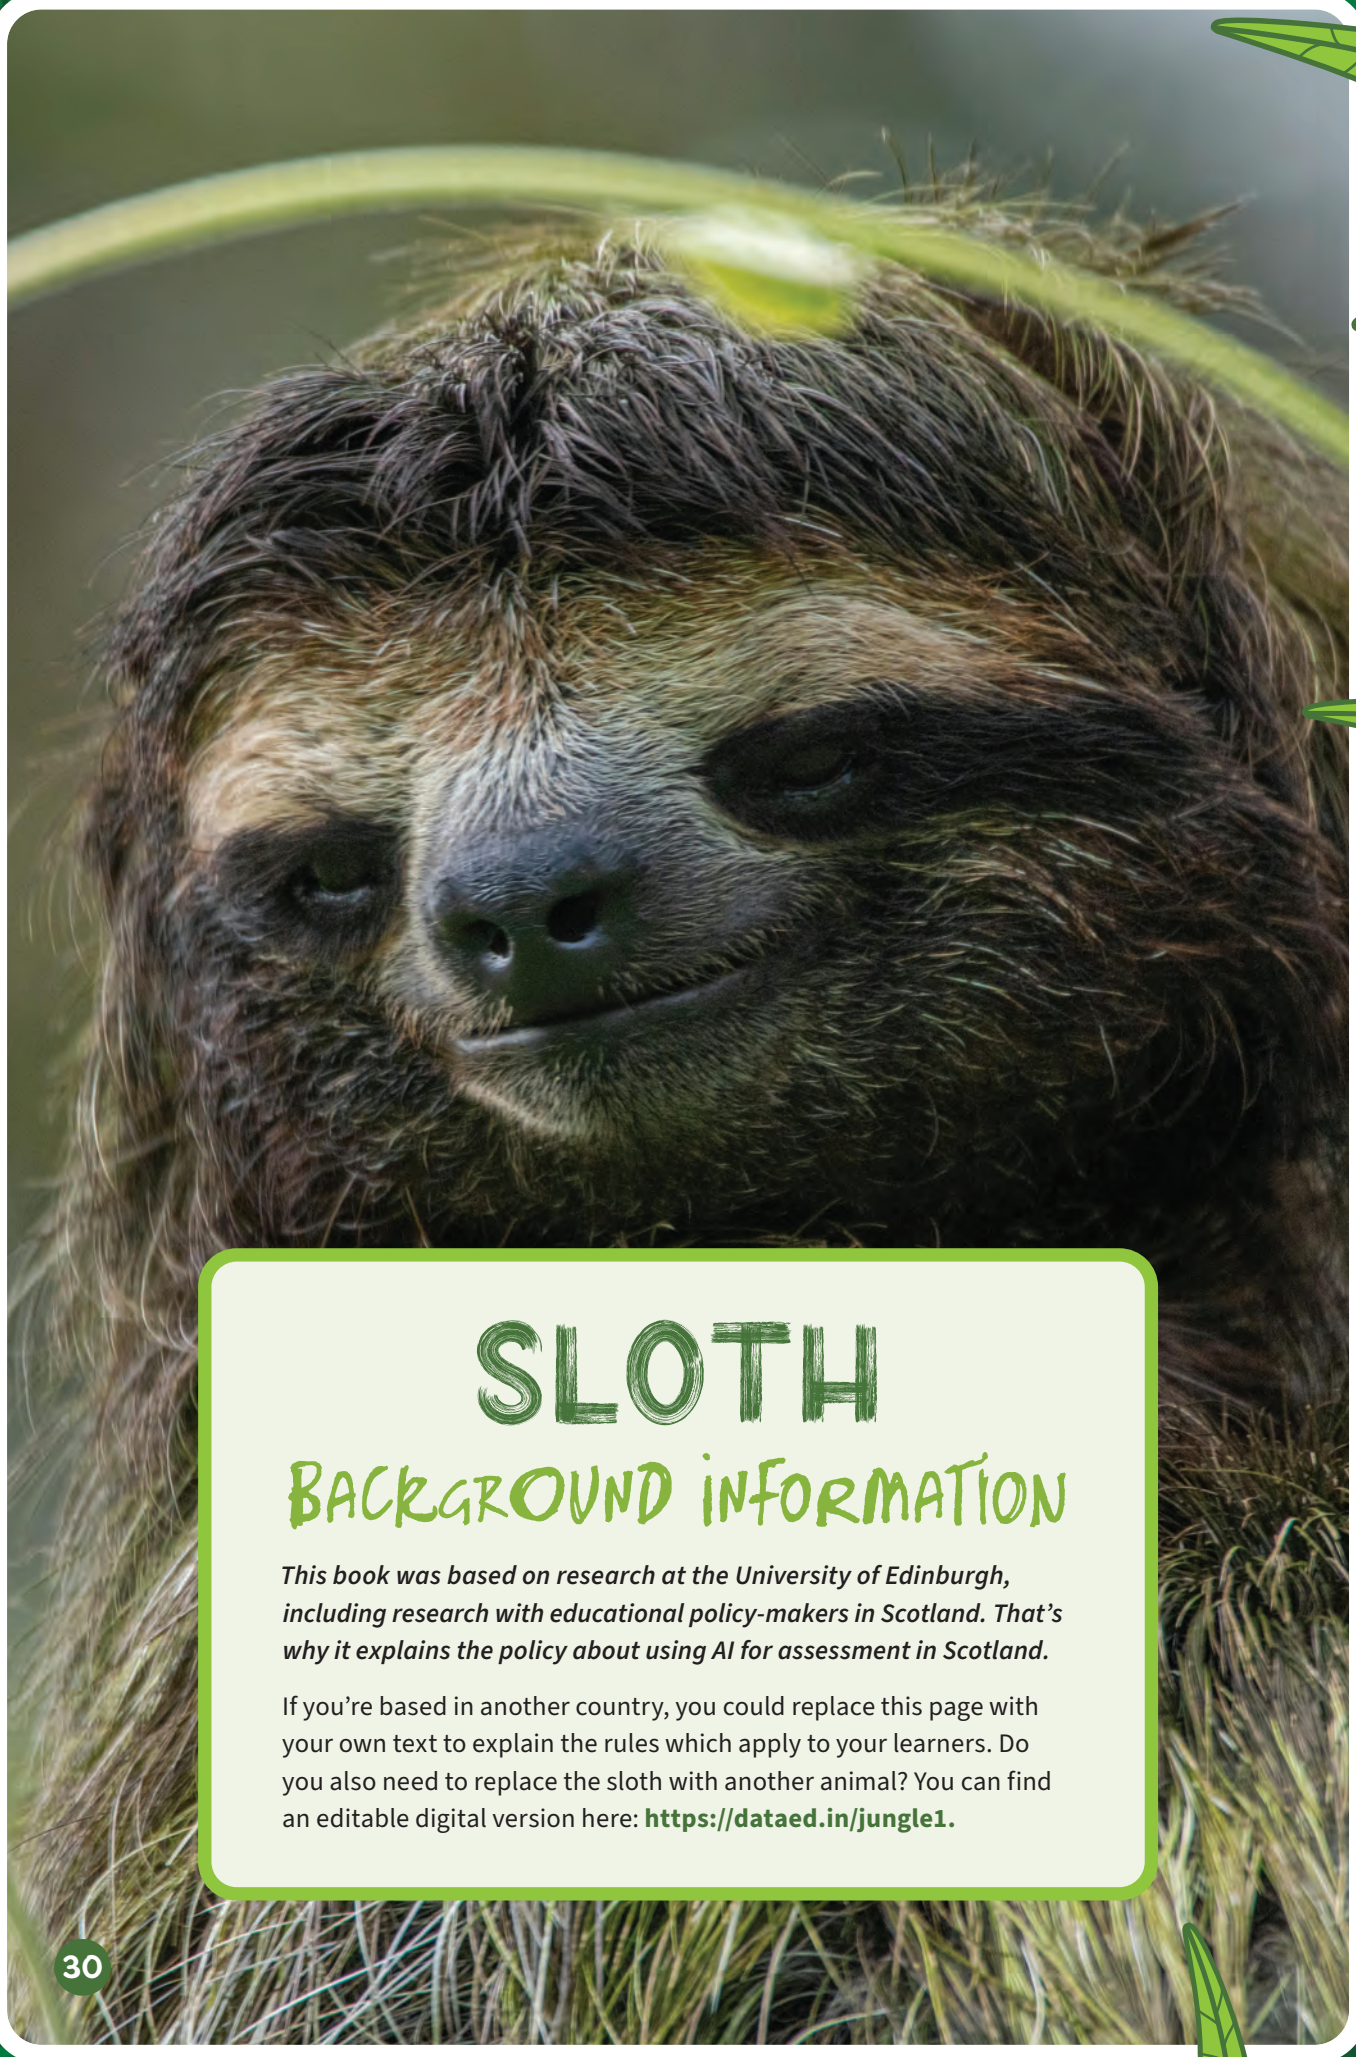

# SLOTH

## BACKGROUND INFORMATION

***This book was based on research at the University of Edinburgh, including research with educational policy-makers in Scotland. That's why it explains the policy about using AI for assessment in Scotland.***

If you're based in another country, you could replace this page with your own text to explain the rules which apply to your learners. Do you also need to replace the sloth with another animal? You can find an editable digital version here: **https://dataed.in/jungle1.**

# PROMPTS FOR LEARNERS

Do you think it is fair to use GenAI tools for work which you will be assessed for? Why?

***Education systems usually decide to ban, evade (e.g. by moving to invigilated exams), or adapt to GenAI. I am trying to adapt to it (even embrace it) in my own teaching.***

Whether we like it or not, GenAI tools are here, and they will only become more capable. GenAI tools are already used in many workplaces to save human employees time. If we ban or evade GenAI tools in assessment, we will not be able to assess the new set of AI literacy skills which will be required in the workplace and everyday life. However, as the sloth metaphor suggests, we shouldn't rush into changes without taking the time to understand the implications. An important aspect of this is consulting with children and young people about their views on this topic.

If you work with primary-school-aged learners, the topic of exams might seem remote to them. If so, you could reframe it to be about something more familiar, like homework or a class test. There tend to be a range of intuitions about this. Here are some possible discussion questions to get a debate going.

— Would it be fair if you asked ChatGPT to write your essay, and then you handed it in without reading it, with your name on it without telling the teacher that ChatGPT did it? (NB: it would be a bad idea to hand it in without checking anyway, because it might be utter rubbish).

— Would it be fair if you ChatGPT wrote the essay but you changed some bits, and told the teacher?

— What if you wrote the essay and got ChatGPT to check it?

— What if someone else in your class had the money to pay someone else to write their essays and you didn't? What if someone else in your class had the money to pay an expensive GenAI to write their essays and you didn't?

I can persuade, argue, explore issues or express an opinion using relevant supporting detail and/or evidence. **LIT 2-29a**

When I engage with others, I can respond in ways appropriate to my role, show that I value others' contributions and use these to build on thinking. **LIT 2-02a**

I can gather and use information about forms of discrimination against people in societies and consider the impact this has on people's lives. **SOC 2-16b**

As I explore the rights to which I and others are entitled, I am able to exercise these rights appropriately and accept the responsibilities that go with them. I show respect for the rights of others. **HWB 0-09a / HWB 1-09a / HWB 2-09a / HWB 3-09a / HWB 4-09a**

I can use evidence selectively to research current social, political or economic issues. **SOC 2-15a**

# BROMELIAD

## BACKGROUND INFORMATION

*Encourage your learners to use their imaginations, to think beyond the GenAI tools which technology companies have already made, and to dream of what learning would be like in their ideal school of the future. You could feed their ideas into your school's policy on AI.*

# PROMPTS FOR LEARNERS

How would you like GenAI to be used in your school?

Ecosystems are constantly evolving and adapting to change. GenAI technology is also changing fast (even if our schools change more slowly) so when we're planning for AI futures, we should build in adaptability from the start. For example, it's a good idea to think beyond the usage of any current specific software package and arrive at a set of principles about how GenAI tools could and should be used in general. Ideally, we would consider the concerns and hopes of our young learners when deciding which GenAI tools to invite into our school ecosystem.

If you'd like to explore this with more structured activities, try our AI Futures Toolkit (coming soon!).


Through contributing my views, time and talents, I play a part in bringing about positive change in my school and wider community. **HWB 0-13a / HWB 1-13a / HWB 2-13a / HWB 3-13a / HWB 4-13a**

# SOURCES


Exploring the AI Jungle: Resources
—*Data Education in Schools*
**https://dataed.in/jungle1**

"10 Ways AI Was Used for Good This Year"
—*Scientific American*
**https://dataed.in/jungle2**

"How Chatbots and Large Language Models Work"—*Khan Academy*
**https://dataed.in/jungle3**

"Teaching a Parrot to Talk"—*Omlet.co.uk*
**https://dataed.in/jungle4**

"Einstein The Parrot Shows off Impressive Vocabulary Skills on 30th Birthday"
—*CBS News*
**https://dataed.in/jungle5**

"Shaped Like Information"—*AI Weirdness*
**https://dataed.in/jungle6**

"Google's AI Recommended Adding Glue To Pizza And Other Misinformation—What Caused The Viral Blunders?"—*Forbes*
**https://dataed.in/jungle7**

"26 Faces in Everyday Objects"
—*Boredpanda.com*
**https://dataed.in/jungle8**

"Strangler figs: Killers or bodyguards?"
—*Kew.org*
**https://dataed.in/jungle9**

"Torpor: what it is, why it's important and how torpor differs to hibernation and sleep"
—*Discoverwildlife.com*
**https://dataed.in/jungle10**

"Can Machines Think? Second Place"
—*CogSci: Interdisciplinary Study of the Mind*
**https://dataed.in/jungle11**

"Can Machines Think? First Place"
—*CogSci: Interdisciplinary Study of the Mind*
**https://dataed.in/jungle12**

"Cyber Toolkit for Teachers"
—*Digilearn.scot*
**https://dataed.in/jungle13**

"AI & Bias - When Algorithms Don't Work"
—*UCLA Institute for Technology, Law & Policy*
**https://dataed.in/jungle14**

"AI & Bias - So What's the Solution?"
—*UCLA Institute for Technology, Law & Policy*
**https://dataed.in/jungle15**


## ENDNOTES

1 https://www.childrensparliament.org.uk/our-work/exploring-childrens-rights-and-ai/#:~:text=Between%20August%202022%20and%20March,AI%20impacts%20children's%20human%20rights.
2 https://www.ohchr.org/en/instruments-mechanisms/instruments/convention-rights-child
3 https://kathyhoopmann.com/all-dogs-have-adhd/
4 https://research.aimultiple.com/artificial-general-intelligence-singularity-timing/

# GLOSSARY OF TERMS

**Artificial Intelligence:** a term used to describe computer software that can do a task which people used to think needed human intelligence.

**Algorithm:** a set of rules which humans or computers can follow to perform a task.

**Ancient Maya:** This is an ancient civilisation that existed over a thousand years ago in the southern part of North America and part of Central America.

**Bias:** Bias means favouring a way of feeling, certain ideas, actions or groups of people over others.

**Biological rules:** In this book, when we say biological rules, we're talking about genes. Genes exist inside the cells of living things and contain rules for controlling traits such as eye colour or height.

**ChatGPT:** a computer program made by OpenAI that uses Generative AI techniques to write text, images, and sounds in response to user prompts.

**Colony:** a group of creatures which live closely and interact together (such as bees or ants).

**Dataset:** A dataset is an organised collection of information used by computer software to perform a task. Very big datasets are often used to train AI systems.

**Ecosystem:** refers to the living things in one particular place and how they interact with each other and their environment

**Generative AI:** Generative AI techniques are used to make software that comes up with new writing, images, sounds, or videos based on a user's prompts. These techniques currently train the software on large datasets of examples. See PX of the Teachers' Notes if you want to know more.

**Metabolism:** a group of chemical changes that happen within organisms to sustain life, e.g. to let cells repair themselves or turn food into energy.

**Metaphor:** a figure of speech that indirectly compares one object, creature, or idea to another. An example of a metaphor in this book is comparing GenAI computer software to parrots. It doesn't mean that we really believe that a computer program is actually a parrot. It just means that they have some interesting similarities which might help people to understand about GenAI.

**Misinformation:** false or inaccurate information.

**Predator:** an animal that kills and eats other animals.

**Sense of self:** In this case, "sense of self" means when someone is aware of who they are. Some animals also seem to show that they have a sense of self – for example, they realise that they are looking at themselves in the mirror, and so they can change their appearance in the mirror by acting on their bodies such as removing a dirty mark.

**Stochastic:** stochastic behaviour follows a random pattern, so it can't be predicted precisely. If you wanted to predict it, you'd need to guess.

**Torpor:** a state of suspended animation which animals enter to save energy. It uses even less energy than sleep.

First published in 2024. This edition published in 2025 by the University of Edinburgh.

THE UNIVERSITY of EDINBURGH

**https://books.ed.ac.uk/edinburgh-diamond** | @EdinDiamond

EDINBURGH DIAMOND



The authors gratefully acknowledge the photographers of the jungle animals featured in this book: Laura Hayek, Oleksandr P, Camilo Ayala, Thierry Fillieul, William Warby, Skyler Ewing, Chris Charles, Molnár Tamás, Steven Paton, Nora Baltus, Steve Douglas and David Clode.

ISBN (ebook): 978-1-83645-145-7
ISBN (paperback): 978-1-83645-146-4
DOI: **https://doi.org/10.2218/ED.9781836451457**

Cover and book design by Elspeth Maxwell: **www.dogandfoxdesign.com**
Typeset by Elspeth Maxwell
Printed and bound by Mixam UK Limited

# TEACHERS' GUIDE

This is an early edition of notes for teachers to accompany the children's picture book "Exploring the AI Jungle". The book is intended to teach children about Generative AI and its wider implications for society, particularly for schools.

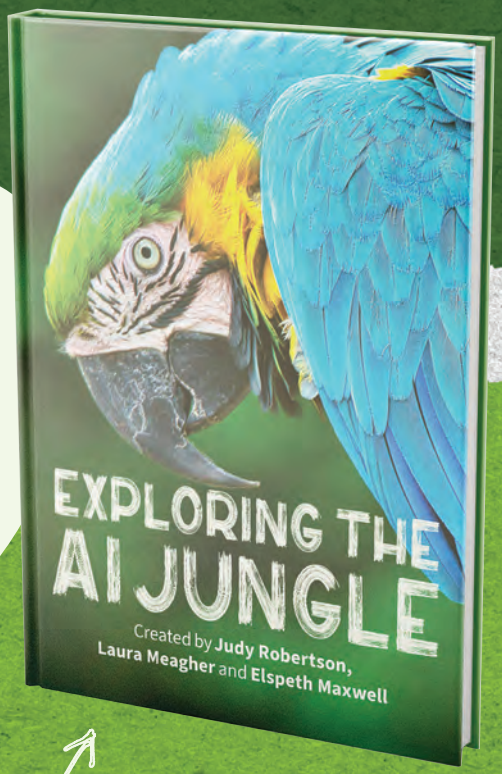


Why is an AI tool like a poison dart frog? What can we learn about AI from parrots? If you like animals, technology or metaphors, this is the book for you!

A companion for learning using this book

THE UNIVERSITY of EDINBURGH

BRAID
Bridging Responsible AI Divides

UKRI
Arts and Humanities Research Council

The authors are grateful to the Arts and Humanities Research Council and the Bridging AI Responsible Divides programme for research grant **AH/Z505560/1.**