



BSSS Special Interest Paper
October 2023: Students, and
Parents and Carers - Generative
AI in ACT Senior Secondary
Education.

BSSS
BOARD OF
SENIOR
SECONDARY
STUDIES

BSSS academic integrity rules apply to all teachers and students in the ACT Senior Secondary System, including the importance of natural justice.

Students have a chance to explain and justify their work

Students may show drafts or notes to help establish their authorship

Students may be asked to undertake a validation task or interview

Students have the right to query a penalty and if needed, go to an appeal.

Students will be given a letter outlining the decision of their school, and the penalty applied.

AI plagiarism is new territory for us all. Your teacher wants to be sure that students get a result that reflects the work they did. In the case of a "false positive", **work with your teacher** to help them be certain that the work is your own.

This September 2023 special interest paper has been produced by the Office of the Board of Senior Secondary Studies (OBSSS), Australian Capital Territory.

Diagrams on p.10-11 and photograph of the mushroom were created by the OBSSS. Social media tiles on p.5 and 11 were created by the OBSSS using Canva.

All other illustrations in the paper were generated using OpenAI's Dall-E which can be found here: [DALL-E: Creating images from text \(openai.com\)](https://openai.com) and DeepAI's text-to-image model, which can be found here: <https://deepai.org/machine-learning-model/text2img>

Contents

What is Generative AI?	1
How do AI plagiarism detectors work?	1
Parents and Carers: What if a student wants to appeal how the rules about AI were applied to their assignment?	2
Advice to Students: What should I do if I'm accused of AI plagiarism?	4
The ACT Senior Secondary System and AI	7
Consulting about Curriculum	7
Currency.....	7
BSSS Information Technology Courses	7
Assessment	8
What if the school down the road can use AI and my school can't?	9
What else is the BSSS doing in this area?	9
Bing's Crochet Mushroom – an AI object lesson	10
How does this relate to what students and teachers do in the classroom?.....	10
Bing's Crochet Mushroom	12
Refresher: How do Large Language models like ChatGPT, BingAI, Perplexity and Bard work?.....	13

What is Generative AI?

Generative AI is a type of computer program that creates new content like text, images, audio, and video that resembles what humans can produce. Generative AI is ‘trained’ by humans and computer programs to work out what the next most probable part of an output (a sentence, a picture, a model, etc.) should be. When asked to produce the output, it does so.

It’s very important to know that the AI does not understand what it is making. It is good at recognising patterns – even ones humans find it hard to articulate – and manipulating these patterns to make reasonable responses.

It’s also important to know that AI isn’t plagiarising from the Internet or lifting text wholesale from any particular source. This makes it very difficult to reliably detect and has been the root of academic integrity concerns in education in 2022-23.

If you’d like a more detailed explanation, please see the article in this paper titled [“How do Large Language models like ChatGPT, BingAI, Perplexity and Bard work?”](#)

How do AI plagiarism detectors work?

AI detecting software works in a variety of ways, most of which are proprietary information, but can be broadly summarised as pattern detection. This pattern is usually assumed to be an artefact of the process of generating AI text. Talk of deliberately watermarking the AI generated text by the originating AI has currently not borne fruit.

A **traditional** plagiarism detector is looking to match student work to language and/or sentences that come from text that already exists. This text might be from open access to the internet, closed databases, or depending on the model, the data of other students in the class.

AI detectors **do not work the same way as traditional detectors**. It is very important to understand that they are **not** working with text that already exists. In fact, unlike plagiarism detectors, the current AI detectors on the market do not access source texts to detect inappropriate conduct.

Some popular **misconceptions** about AI detectors include:

- Everything “stays” on the internet so text matching is possible. Unfortunately, this is not the case, and even if it were, the AI generations are usually behind some kind of password protection, so the makers of traditional plagiarism detectors cannot access them.
- ChatGPT can tell you if it wrote something. Again, this is not the case. It does not have the functions to operate as a plagiarism detector, so it will just tell you the answer it thinks you want – which is probably “yes”.
- AI detectors are 100% accurate. This is one that the creators of AI detectors are clear about – they aren’t. There is a chance of both false positives (genuine student work falsely identified as AI) and false negatives (AI falsely identified as genuine student work). Current recommendations from software providers is that it’s best for teachers to have more than one point of evidence when plagiarism is suspected. When speaking to a student about software-detected potential plagiarism, we encourage teachers to go in with an open mind and a presumption of innocence. It is also worthwhile keeping an eye out for what various AI detection platforms are publishing for advice. For example, commercial detection software Turnitin has advice sheets for both teachers and students, including [advice about the detection of false positives and how to deal with them](#).

We know that students may be unused to needing to keep information, undertake validation tasks, or be queried without the hyper-accurate text matching of traditional detectors. Asking for drafts, setting validation tasks, and querying potential plagiarism are all ways that teachers are dealing with the use of AI while keeping things fair for everyone.

Parents and Carers: What if a student wants to appeal how the rules about AI were applied to their assignment?

Under *BSSS Policy and Procedure*, the right to query and appeal allows for the right of contestability to be exercised by the student. This can be a query at a class teacher or head of faculty level, or an appeal at a school or eventually Board level.

A school appeal comes after all other avenues at the school have been exhausted. All queries and appeals, including school and Board appeals, must be initiated by the student, not by a parent or carer on the student's behalf.

The student first must query with the teacher and the head of faculty before a school appeal can be raised. Not paying attention in class or not reading required course materials/student handbooks doesn't excuse students from knowing about academic integrity.

If an appeal about AI plagiarism rises to the level of a school appeal, it's important to have evidence so that the committee can see what has occurred prior to the school appeal. The below information is intended to help you to know what kind of evidence you might provide to help the panel deliberate on their decision.

Preparing for a School Appeal

Students need to gather evidence so that the appeal panel can see it. This is best done as the student goes through the process of creating the assignment. This advice has been specifically put together in the event of an appeal against an academic integrity penalty for the use of AI; not all of these factors will apply in other appeals.

Before the task was handed in

What happened in constructing the task? For example:

- Was there explicit guidance from the teacher or school about how AI should or should not be used?
- Was there reference to the school's assessment policies on the task sheet? (Note: BSSS policies identify the use of AI as plagiarism unless specified as allowed in the task sheet)
- How was the task drafted? Did the student save versions, talk with anyone about their work, take photos as they made a physical object, take notes, save articles?
- Did the student hand in a draft? Were they warned about the use of AI in response to this draft?

What evidence does the student have of this? For example:

- Cover sheet
- Assignment description/sheet
- Unit Outline
- Academic handbook
- Drafts
- Folder (digital or physical) of notes and research material
- Independent evidence of conversation – e.g., Teacher Librarian or mentor teacher able to corroborate that student discussed ideas with them
- Class work that demonstrates student's voice, evidence of complexity of language and thought, or nascent ideas that led to the draft

Once the task was deemed plagiarised

What happened in the decision about applying the penalty?

- Was the student given a chance to explain their situation?
- Which of the BSSS penalties were given, and what was the reasoning explained to the student for this penalty?
- Has a letter confirming the penalty and decision gone to the student?

What evidence should the student have? For example:

- Meeting minutes or notes
- Copy of the letter to the student outlining the plagiarism penalty and reasons for it
- If a written validation task was given, copy of the written validation task
- If a verbal validation interview was given, notes from the interview
- Evidence from an independent observer to the process (e.g., a pastoral care teacher who attended the interview)
- Head of Faculty's written feedback after the student queried them (this is an essential step prior to a school appeal)
- Evidence that was supplied to support the student's initial query of the penalty.

In a school appeal, the student will be asked to speak with the appeal team. This will be three people – two from the student's school, and one external person. Keeping good records can help the student to explain their point of view and help the appeal team make a decision that is fair and equitable for all involved.

A Board Appeal is very rare, and only occurs when there is evidence that policy and procedure has been applied incorrectly or inconsistently. If your child gets to this stage, your school will speak with them, or they can contact the Office of the BSSS at bsssenquiries@act.gov.au.



With the introduction of generative AI, academic integrity is a hot topic in ACT schools. Let's all work together to keep things fair.

It is **normal and expected** for teachers to check documents for possible plagiarism and speak with students when there's a disparity with other class work or anomalies they can't otherwise explain.

- The ways of detecting AI plagiarism are still developing, and false positives will happen.
- Students **must be given a chance to explain** their side of the story when AI plagiarism is suspected.
- Students should save drafts or notes so that they can easily show genuine authorship of their assessment work.
- If needed, a penalty can be appealed.



Advice to Students: What should I do if I'm accused of AI plagiarism?

Students need to be prepared to show genuine ownership of their work, by keeping notes or records as assignments are drafted. This is the responsibility of the student, and it will make things much easier if a student is wrongly accused of AI plagiarism. It is normal and expected for teachers to investigate possible plagiarism.



AI has changed a lot during 2023. With the introduction of generative AI, some people have been using programs to plagiarise their work, knowing that traditional text-matching plagiarism checkers will find it hard to pick up. Some companies claim to be able to pick up AI generated content, but these checkers are not perfect; they miss AI content, or unfairly accuse innocent people.

You need to be prepared for false positives¹, and prepared to discuss your work in detail, even if you did not use AI. If you refuse to participate in these processes and/or do not provide any evidence to demonstrate the work is your own, a plagiarism penalty may be applied as per BSSS Policy and Procedure. This is because your teacher needs to be certain that the work is your own, because they are certifying officially that you have skills and knowledge from your courses, and that those skills and knowledge are of a particular level.

The bottom line is that AI has made plagiarism a challenge for everyone – students and teachers. We all need to act together to help prevent plagiarism. It is normal and expected for teachers to check work for plagiarism, and to query work when it shows anomalies that can't be explained by other factors. The teachers must do it, so that they can be confident that everyone is being fairly marked on the work they've done!

Before you hand in the assignment

When you do the assignment, keep records of what you've done to put it together. We know that is time-consuming, and we hope that future developments in technology mean we can go back to easy, more accurate plagiarism checking, but until that happens, we all need to work together to keep things fair.

You might:

- Save different versions of the document if you're using a word processor that doesn't save history
- Set up "draftback" or a similar browser extension to record your process
- Put all your printouts and notes into an envelope or folder that you keep in your study space until the assessment period is over
- Take photos of your assignment each time you finish doing some work on it and put them in the same album on your phone
- Put all your digital documents relating to that assignment into the same folder on your computer or cloud storage system (e.g., Google drive, iPhone files)

Having a designated "spot" to put things while you work will mean you don't need to stress if you get asked for your notes. Even if it's just a tub that everything goes into, or a keyword that you use in filenames so you can use the search function, it's much easier than trying to work out where everything is. It is advisable to keep these notes at least until the end of the assessment period when you have your scores and grades and you know you will not be seeking an appeal or picked up for possible plagiarism.

¹ False positive: a result where someone who hasn't used AI has been wrongly flagged as using it.

If you are accused of AI plagiarism

You will have a meeting with your teacher where your teacher explains what it is about your work that is indicating that it might be AI plagiarised. These might be things like:

- A mismatch between the language in different sections of your work
- A mismatch between the quality of your usual work, and the quality of this particular piece of work
- Significant similarities between your work and someone else's
- An AI detector says there's a possibility of plagiarism

You will be given the chance to explain your side of the story. If you need to get things to help explain (e.g., your notes), then you should bring them with you, or organise another meeting so you can have what you need with you.

When you have the meeting...

Don't:

- Don't go into a meeting about plagiarism angry; if you need to, ask to reschedule, and explain that you'd like to pause to get your feelings in order
- Don't threaten - it's never okay to threaten other people
- Don't have this conversation in a corridor, or when you/your teacher is on the way to another class
- Don't bring a belligerent support person: if you bring a friend, that's fine, but they need to be quiet support (i.e., not acting like a "bush lawyer" or getting angry/threatening on your behalf)



Do:

- Do save your notes and drafts and bring them to show evidence of your originality
- Do talk with people in pastoral care or student services if you're worried or nervous about speaking to the teacher: they can help you prepare, and even come with you to a meeting
- Do be prepared to explain how you came up with your ideas, and your process of creating the piece
- Do be prepared to do a validation task or interview
- Do make some notes about what happened once you've talked to the teacher
- If you brought a friend or support person, debrief with them, and ask them to make some notes
- If you did use AI, admit it: this will save a lot of heartache, time, and stress for everyone involved!



What happens then?

If you **did** use AI inappropriately in your assignment:

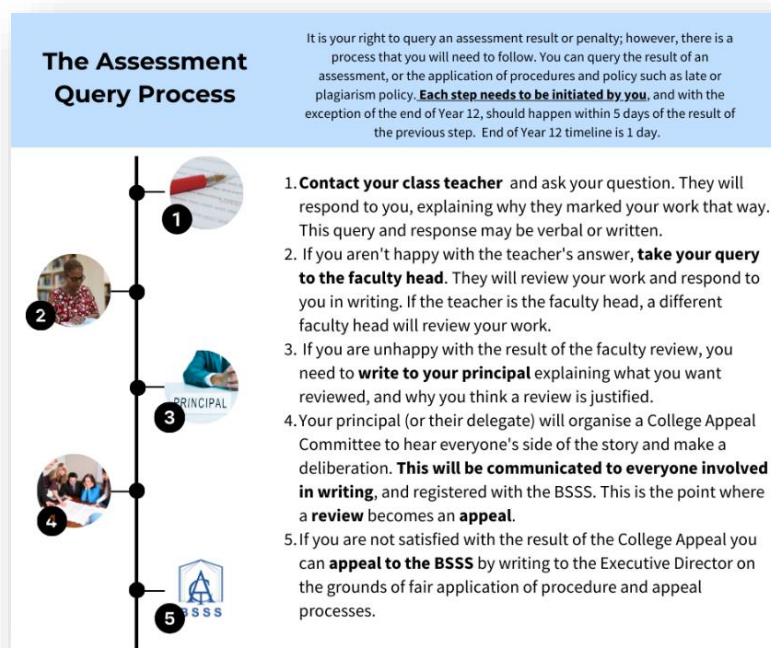
- Academic Integrity breaches don't happen in a vacuum. Usually there's a reason; e.g., mental wellbeing, prioritisation, carers responsibilities, procrastination, misunderstanding of the rules.
- Talk with your teacher about the reasons why you felt you needed to use AI and seek help from your teacher or another trusted adult to address these reasons. This might be applying for special consideration, or putting together some strategies to help you establish better study habits.
- You will probably get a penalty on the work, which will be drawn from the penalties in the Academic Integrity section of BSSS Policy and Procedure.
- One plagiarism penalty will not ruin your life and career. You are not a bad person; you just made an inappropriate choice. You'll do better next time.

If you **didn't** use AI inappropriately in your assignment:

- You will be okay; you get to explain your side of the story, and if a penalty is incorrectly applied, you have the right to query and appeal. This can go all the way to the Board if conditions apply.
- Remember: your teacher is doing their job, and something about your assignment has led them to make this query. It could be a result from an automated plagiarism checker, or your work might be very different to anything they've seen from you before. Maybe it is that you don't put much effort in when doing class work, and you put in a lot of effort this time. This doesn't mean the teacher dislikes you or has it out for you – it just means they're doing their job.
- Be prepared to demonstrate genuine authorship of your work. This might be explaining your thought process that led to choices in your work; it might be showing drafts or notes; it might be drawing on a support person at your school, like your Teacher Librarian or mentor teacher if you showed them a draft.
- Be prepared to do a validation task or undertake an interview where you explain your work in more detail.

What if I don't agree?

You can appeal the penalty. You will need to provide evidence as to why the penalty was unjustly applied – the same kind of evidence you would have produced to demonstrate your authorship initially.



The ACT Senior Secondary System and AI

Consulting about Curriculum

A student from the BSSS Student Forum asked a question as an eloquent summation of the paradox of AI:

“How will things be the same but also different?”

Technology is not, and is highly unlikely to ever be, the primary reason why students are at school. As Schleicher (2023) notes:

*“Smart education is not about technology, but about radical reimagination of what teaching and learning can be powered by technology. From learning **technology** to learning **activity**. Education is not a transactional business but a relational enterprise.”*

Curriculum and relationships remain the basis of what we do in school; they set out what skills, paradigms, ideologies, and methodologies are valued in and valuable to our society. The students we are teaching now will “...live in a world of many portals: physically, emotionally, spiritually” (Inayatullah, 2023). The ACT Senior Secondary System gives teachers the flexibility to personalise to context, to students and to innovation that will help prepare the next generation of young people for a rapidly changing world. All ACT Senior Secondary courses contain critical thinking in their course design and achievement standards.

Currency

ACT Senior Secondary courses are designed so that teachers can make use of topics, texts, ideas, and technologies that are emergent as well as those that are established. Andreas Schleicher noted at the recent Australasian Curriculum, Assessment and Certification Authorities (ACACA) conference that the kinds of tasks and employment that are growing in the world are those which are technology-intensive and non-routine. These are tasks that some students find very challenging; we encourage parents to keep supporting students in persevering through challenge and finding solutions to problems.

Where these courses really shine is the opportunity for teachers to play to the students’ interests and to emerging important issues: whether that is a focus on climate disasters in Interdisciplinary Science; AI, consciousness, and cognition in Philosophy; digital currencies in Business and Economics; copyright and ‘copywrong’ in Art and Design; or changes in assistive technology in Social and Community work.

BSSS Information Technology Courses

The Office of the BSSS established a panel review of the suite of Information Technology Courses (i.e., Data Science A/T/V, Networking and Security A/T/V, Robotics and Mechatronics A/T/N/V, Digital Technologies A/T/M/V and Digital Products A/V) in light of generative AI. This panel included members from university, industry, curriculum experts in IT, and experienced IT teachers representing the three sectors.

The Panel put forward the case that students need to understand IT ethics, concepts, and processes to use AI successfully and effectively. They concluded that AI tools are tools like other IT tools that students will learn about in their courses, and that teachers will teach how and when to use them to support learning and complete projects.

The panel noted that generative AI tools may provide opportunities for more cycles of iterative improvements or heightened complexity of projects, but that this can only be successful with a deep understanding of the tools.

The recommendations of the panel included:

- Teachers should teach how to use these tools and the knowledge and skills necessary to use these tools effectively and encompass AI as they would any significant technological development. This includes making students aware of the strengths and limitations of generative AI.

- Teachers will need to adjust work and assessment regularly to keep up with the rapidly evolving field.
- Teachers should make use of the existing content strand on ethical use of IT to explore the current and emerging ethical implications of these technologies.
- Ethical questions should have some prominence in assessment to ensure significance is apparent to students.
- BSSS Courses should not be restructured or subject to minor variations.
- Due attention should be paid to academic integrity in formulating assessment tasks to ensure that AI or coding repositories do not take the place of active work by students.
- That students have access to variety of digital tools including large language models to investigate ethical and appropriate use in the context of these courses, but also for all BSSS courses.
- The OBSSS should monitor IT results and return to this review if results become bimodal² (i.e., two distinct spikes instead of a curve).

Assessment

Developing quality assessment tasks is important for the integrity of the ACT Senior Secondary System. To support teachers and ensure quality, the ACT Board of Senior Secondary Studies has released a set of [Quality Assessment Guidelines](#), based on contemporary research and designed to develop a common understanding and language of how to develop assessment to meet the needs of students.

The six areas of the Quality Assessment Guidelines are:

- coverage of the curriculum
- reliability
- bias
- provision for a range of thinking levels
- student engagement
- academic integrity.

The Office of the BSSS encourages schools to consider a broad range of assessment types; however, the current situation with AI has led schools to walk a tricky line as they try to deal with the changes that AI has wrought on assessment. When we assess students, we are looking to provide a window into their understanding and knowledge with the intent to both measure their progress and help them improve, and to report on their achievements. With those goals in mind, teachers want to make sure that the work being assessed *is* the student's work.



While some schools are doing more assessment in class, there is no mandate to do so. However, there may be factors at play indicating that the fairest way to assess students is through supervised in-class, work.

In the past 30 years, assessment has changed significantly – as the World Wide Web opened new doors, the needs of a changing society were reflected in assessment that privileged transference of skills and knowledge to novel contexts, higher order and critical thinking, and reflection.

These trends are likely to continue into the future; Schleicher comments that assessment in the current world is “...not whether you can remember answers, but ask the right questions, triangulate sources, and see what’s coming around the corner.” (Schleicher, 2023).

² See the explanation of point 4 in the crochet mushroom discussion for why this might be of interest.

What if the school down the road can use AI and my school can't?

The ACT BSSS does not have oversight or control over the provision of AI – or, indeed, any technology -- in schools. The short answer is, it's not a problem if different schools have different conditions, as long as they're applying equivalent conditions within their own units.

The ACT has a school-based assessment system that gives schools flexibility to meet the needs of their students and address AI as appropriate in the context of each course and assessment item.

In terms of differences between schools and courses in the use of AI in assessment, two points should be noted.

First, when grading assessment and when that assessment is quality assured at Moderation Day, the specific assessment conditions are taken into account. As a teacher would consider the length of time for a test, in-class vs. take-home assessment, or the degree of scaffolding provided in the assessment, the use of AI in an unsupervised assessment will also be considered when evaluating student assessment as the evidence of their learning and skills.

Second, for marking purposes, such as determining unit and course scores, the marks awarded in an assessment item are only compared to the marks of other students completing the same assessment in the same course in their school only or markbook course grouping in their school only.

Teachers use a range of measures to ensure assessment item results are comparable and have the equivalent conditions to support this comparability. For example, if Ancient History and Modern History at Lake Burley Griffin College were in the *same* ACT Certification System markbook, using marks generated from the same rubric would be problematic if the essay from Ancient History was in-class and from Modern History was a take home. If, however, the conditions were the same, this comparability would do much to mitigate any issues from a marking perspective.

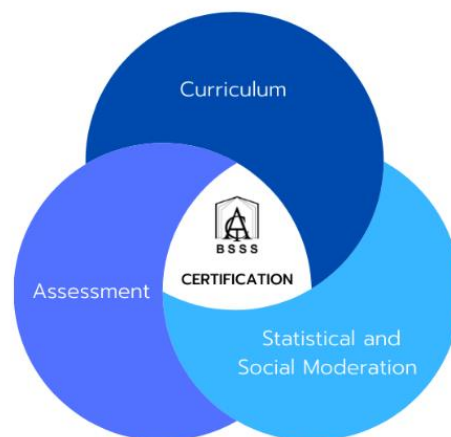
The comparison between schools is part of what is called “scaling”. Before ATARs are calculated, a number of complex mathematical processes take place that account for any differences between schools and courses, bringing all ATAR-seeking student results onto the same scale. See the [first student newsletter](#) for more information about scaling and the ATAR.

What else is the BSSS doing in this area?

The OBSSS has been closely involved with the development of the National AI Ethics Framework for Education, representing the ACT on the National Taskforce. We took feedback from students as part of this process; encouragingly, students were very interested in any AI use being substantive (not “sparkle” added to lessons) and with educational benefits supported by quantitative data once the use of AI becomes widespread enough that this data can be gathered. Students were also highly interested in how technology and AI can be more broadly incorporated into curriculum across all subjects.

We are pleased to announce that the Framework, which includes a range of aspirational principles for the use of AI in all Australian schools, was approved at the Education Ministers Meeting in October. Additionally, the Education Ministers have allocated funding for Education Services Australia to create a national set of product expectations.

We anticipate that work will continue between jurisdictions to share how AI is being used, and that once the Framework is formally launched by the federal Education Minister, jurisdictions and federal Education will be able to work on how the Framework, and AI itself, will be operationalised in schools.



Bing's Crochet Mushroom – an AI object lesson

We asked Bing AI to create a pattern for a crochet mushroom as part of seeking an example of how it references, and quickly realised that this was an interesting practical lesson in how some of the concerns and considerations about AI can play out in the most mundane of objects.

In the past few years, crochet mushrooms have been popping up all over the internet as part of a group of aesthetics; for example, 'cottagecore' or 'goblincore'. This makes it unsurprising that a Bing crochet mushroom is a reasonable pattern because there are multiple examples of crochet mushroom patterns that likely form part of its training data. This is, obviously, controversial; as with AI generated art, the use of someone's creative work in this way is, if not illegal, at least taboo amongst makers.



Algorithms are a set of instructions for computers. When we talk about “algorithmic bias”, this is usually a reflection of how the unseen and often unarticulated biases of the programmer affect this set of instructions. With machine learning, these biases may also come from the training data.

Bing's crochet mushroom clearly shows the presence of algorithmic bias:

1. The pattern is written in US crochet terms. This points to the majority of its training data using US crochet terms for amigurumi³.
2. The pattern for a “crochet mushroom” gives an amanita muscaria (fly agaric) mushroom. This distinctive mushroom is the archetypal “fairy tale” mushroom and likely familiar to you from children's books and cartoons.
3. The pattern for a crochet mushroom uses the same stitch counts for both pieces; this would seem to indicate that the training data biases to a 6-stitch base. This leads the stem to be somewhat out of proportion to the cap.
4. The pattern doesn't include any extra information – e.g., no information about how to crochet, or that in this style of crochet, the work is not turned. There is a reasonable amount of subject knowledge required to interpret the pattern and get the item you want from it. This reflects the training data coming from completed patterns using this jargon (e.g., terminology like “magic loop”).

How does this relate to what students and teachers do in the classroom?

1. US-centrism in training data may not be desirable in an Australian classroom or context. In this instance, it's technical jargon that has a different meaning (SC in US terms is a different stitch to SC in UK terms). In the case of using this for classes, this US-centrism can show up as:
 - a. Values and attitudes
 - b. Comparisons, examples, and historical events (e.g., the mushroom we got is not native to Australia)
 - c. Language (e.g., math vs. maths; or words that are acceptable in US English which may not be in Australian English such as 'gotten')

³ “Amigurumi” is a loanword from Japanese used within the online crochet community to mean a stuffed toy. They are typically crochet but can also be knitted.

2. In this instance, the presentation of the stereotyped fairy tale mushroom is fairly harmless; however, the same processes that gave us this cheerful redcap also serve to continue to marginalise cultures, ethnicities, gender expressions and other human differences as the algorithm works towards a “more probable” answer. The fly agaric mushroom points to the problem of AI language models presenting cases that are drawn from dominant cultural perspectives. There are so, so many more species of mushroom out there! We might ask ourselves why we didn’t get a button mushroom, an amethyst deceiver mushroom, a latticed stinkhorn, or a bunch of enoki. The reason is because when we asked Bing for a mushroom, it decided the mushroom we were most likely thinking of was red-and-white.

Students may need to be led to understand that they need to prompt an AI for edge cases, the margins, or perhaps even... use a source that isn’t an AI. If teachers use AI to write lesson plans or create material, it is good practice to consider whose beliefs, experiences, ways of being and ways of thinking are covered in the material, so that gaps can more easily be filled in as needed.

3. The stitch counts show us that Bing doesn’t recognise numbers in the way we might expect it to. It’s not a calculator! While it has some programming in the background that deals with the numbers more intuitively as far as its human users are concerned (and this will continue to be improved upon), it still offers the most probable solution. This being: in its training data, there is information that says six stitches into a magic loop is the way to start a piece of amigurumi. This highlights the fact that the AI isn’t the best tool for every situation; even though it’s a reasonable one in this situation, one might be better off paying for a published pattern or working out their own for a nicer mushroom.
4. This element shows two diverging implications for students, showing how different classes and different expertise levels can produce dramatically different results:
 - a. The lack of extraneous information shows that expert knowledge is still needed. A novice could generate the pattern but would be unable to work out easily if the pattern seemed reasonable, may struggle with the instructions, and would need to access supplementary information to assist them in their completion of the pattern. It points to the fact that a student who is not well-versed in a subject area won’t be able to rely on AI alone and may have to do a significant amount of other work to use the AI output successfully.
 - b. The reasonableness of the pattern, however, means that for an expert, it is a useful steppingstone. An experienced crocheter can easily use the pattern as a base for modifications/finessing without having to do the mathematics for themselves – for example, adding a few more rows in the middle of the stem to make it less like a ball, making the stem slimmer, or adding a stitch pattern to add interest to the cap are all simple modifications that could be done on the fly (agaric). It points to the fact that a student may use AI for the basics of a task, and then use their expertise to finesse and improve upon an existing product, giving them a significant advantage if they are being compared to novices.

Yes, the pattern works. We made it! (See next page)

Some the notes from our tester for any other mushroom makers:

I did not use cotton yarn; I used red and white acrylic 8-ply. This pattern is simple and will work for a small amount of any smooth yarn of your choice with an appropriate hook size (listed on the yarn ball band; go for the smallest listed or just below it for a nice firm fabric).

All crochet terms in the pattern are [US terms](#). The pattern is worked in rounds without turning, so you may wish to use a stitch marker to keep track of the start of your row.

The stem of the mushroom as written is a ball. You may like to add some more rounds around rounds 4-7 to lengthen the stem, and add some more rounds between the decrease rounds 8 and 9 to make the taper of the stem slimmer as it meets the cap. The mushroom as written will not stand up on its own. Put some weights (e.g., marbles, toy weighting pellets, a few small clean rocks) in the bottom before you stuff it if you want it to stand up. Do not use beans as weights, because if they get wet, your mushroom might grow some friends.

Bing's Crochet Mushroom

Materials: Cotton yarn in red and white 3.5 mm crochet hook Toy stuffing Yarn needle Scissors	Abbreviations: ch = chain sc = single crochet inc = increase (2 sc in one stitch) dec = decrease (sc 2 stitches together) sl st = slip stitch st(s) = stitch(es)
---	---



Pattern

Stem:

With white yarn, make a magic ring.

Round 1: 6 sc in the ring. (6 sts)

Round 2: inc in each st around. (12 sts)

Round 3: (sc in next st, inc in next st) repeat 6 times. (18 sts)

Round 4-7: sc in each st around. (18 sts)

Round 8: (sc in next st, dec in next st) repeat 6 times. (12 sts)

Round 9: dec in each st around. (6 sts)

Fasten off, leaving a long tail for sewing. Stuff the stem firmly.

Cap:

With red yarn, make a magic ring.

Round 1: 6 sc in the ring. (6 sts)

Round 2: inc in each st around. (12 sts)

Round 3: (sc in next st, inc in next st) repeat 6 times. (18 sts)

Round 4: (sc in next 2 sts, inc in next st) repeat 6 times. (24 sts)

Round 5: (sc in next 3 sts, inc in next st) repeat 6 times. (30 sts)

Round 6: (sc in next 4 sts, inc in next st) repeat 6 times. (36 sts)

Round 7: sc in each st around. (36 sts)

Round 8: (sc in next 5 sts, inc in next st) repeat 6 times. (42 sts)

Round 9: sc in each st around. (42 sts)

Fasten off, leaving a long tail for sewing.

Assembly:

Sew the cap to the top of the stem, making sure to align the centre of the cap with the centre of the stem.

With white yarn, embroider some spots on the cap as desired.

Weave in all the ends and trim the excess yarn.

Your crochet mushroom is done! Enjoy! 🍄

Refresher: How do Large Language models like ChatGPT, BingAI, Perplexity and Bard work?

Generative AI works by making use of very large bodies of text and very large amounts of computing power. Training and running an AI can cost millions of dollars. [One estimate places the running costs of OpenAI](#) at up to US\$700,000.00 a day.

Most Large Language Models are “trained” on data sourced from places like Reddit, Wikipedia, and Facebook. The program is given a large body of text, and some instructions. These instructions help it to analyse this text by breaking it up into abstract parts called “tokens”.

These tokens aren’t just letters and numbers – they’re a virtual representation of all the things that make up types of text, like the qualities that would make a text “Shakespearean” or the qualities that make a text a haiku, sonnet, formal letter, or eulogy. The model does not understand the training data – it uses statistical processes and probability to learn the patterns that are in the training data.

The program is then trained to learn when it should use these “tokens”. This is done using a combination of software programs and human intervention; it works on probability, and again, not the probability that a specific *letter* or *number* will appear, but that a piece of information that the Large Language Model has designated as a token will appear. For example: Depending on the context, the token “slan-“ might be part of “slant” or “slang” – as the model runs, it learns when it’s more probable that the slan- will finish with a -g (making “slang”), and when it will finish with the token -derous (making “slanderous”); or even when it will be sandwiched between parts to make i-slan-d.

This applies on the sentence level, too; for example, the large language model learns that a sentence like “the quick brown fox jumps over the lazy...” is likely to finish with “dog”, but it doesn’t know what a dog or a fox is. It just knows that when the combination of tokens that look like “fox jumps over” are given to it to respond to, then the tokens that look like “dog” are most likely to follow.

As the program works, it uses the rules that were set by the program designers to move from gibberish to words, until it gets really, really good at predicting what the most likely “next word” is. The large language models we deal with aren’t just a pure access to the model. There are sets of information and instructions that take our prompts and ‘translate’ them for the model, and then take the model’s response and ‘translate’ them for us.

That’s why sometimes you get a warning saying “I’m not going to answer that question” – because in the instructions, it’s been told not to show things that are potentially dangerous. This is where we can get into issues of what is/isn’t allowed by a particular creator. These types of programs are usually very good at specific tasks – you have likely seen examples out there in the popular press that do things like suggest movies to you on your favourite streaming service or find patterns that point to medical issues in hospitals.

Unfortunately, the nature of probability means that the models are also prone to making things up – sometimes called “hallucinating” or “confabulating”. Just because it’s probable doesn’t mean it’s right! For example, in Orwell’s novel *1984*, a character is told that “ $2+2=5$ ” is correct. It’s likely that *1984* and criticism or discussion of the novel appear in the training data for many large language models. That means that sometimes, the program has seen “ $2+2=5$ ” as a legitimate expression that may, on the balance of probability, come up every so often.

