



Interim Policy for the Use of AI in Teaching and Learning

University Policy Workgroup for AI in Teaching and Learning

LOY Hui Chieh (Chair; Philosophy), LIOW Chee Hsiang (deputy Chair; Public Health), NG Hwee Tou (Computer Science), KAN Min-Yen (Computer Science), Martin BUIST (Biomedical Engineering), CHEN Zhixiong (Physiology), Natalie PANG (NUS Libraries)

5 Feb 2023

© 2023 National University of Singapore

Interim Policy for the Use of AI in Teaching and Learning

University Policy Workgroup for AI in Teaching and Learning

A/P Loy Hui Chieh (Philosophy) (Chair), A/P Liow Chee Hsiang (Public Health) (Deputy Chair), Prof Ng Hwee Tou (Computer Science), A/P Kan Min Yen (Computer Science), A/P Martin Buist (Biomedical Engineering), A/P Chen Zhixiong (Physiology), Dr Natalie Pang (NUS Libraries)

Final draft 5 Feb 2023

Approved 5 Feb 2023

1. Briefing for Teaching Staff

1.1 Background

On 30 November 2022, the US-based artificial intelligence research laboratory OpenAI made a beta version of their chatbot “ChatGPT” freely available to the public. The AI chatbot is named after the language model it is based on—“Generative Pre-trained Transformer” (GPT; currently version 3.5), an artificial intelligence (AI) model capable of processing and producing text in natural language. When fed a text prompt, ChatGPT produces text that responds to the prompt. Even though ChatGPT is trained on data up to 2021, users can also feed it a substantial amount of additional information, both in text and in table form.

While the technology that ChatGPT is based on is not completely new, and there are competitors from other companies, its launch gave the public wide and almost costless access to a recent generation AI. Observers soon noticed the implications tools such as ChatGPT have on assessment in higher education. This is because ChatGPT is fully capable of producing output that is not easily distinguished from student-produced work, including essays, solutions to mathematical problems, computer code, and even answers to multiple-choice questions complete with explanations for why each option is correct or incorrect. With careful prompt crafting, it can even generate decent reflective or introspective essays.

The general impression is that ChatGPT can produce work that is at least passable, especially in introductory classes, even if it is not of a quality approximating what our better students can produce or what is required of them in more advanced classes. (In the US, the output of ChatGPT has been shown to pass the Bar exam, the MBA exam, and even a medical licensing exam, although the quality of the answers is judged “mediocre”.)

Some resources are appended to this briefing as an appendix for educators wishing to know more about ChatGPT and its capabilities. However, it is probably easier for colleagues to sign up for a free account and try out the chatbot for themselves at <https://chat.openai.com/>. Note that users must agree to terms and conditions—which NUS will not be able to endorse or indemnify—when signing up.

1.2 NUS Responses

Given ChatGPT’s potential as an instrument for academic dishonesty if not used properly, the University Policy Workgroup for AI in Teaching and Learning was formed with core members from the NUS Board of Undergraduate Studies, to recommend policies and actions to the University in response

to evolving threats and opportunities presented by AI to teaching and learning, as embodied in the first instance by ChatGPT.

The Workgroup's first recommendation consists in the **Interim Policy for the Use of AI in Teaching and Learning**. The Interim Policy is articulated on the assumption that, as a university, we will not ban the use of AI tools among students and staff: doing so is neither feasible nor desirable. On the other hand, the NUS community will benefit from clear and consistent guidance regarding the proper uses of such a technology for teaching and learning.

With the above in mind, the rest of the Policy consists of these two sections, both framed as we, the educators, addressing our students:

- **A General Message to Our Students** about broad principles, which we recommend that educators make use of when explaining to students the rationale for our guidelines; and
- **Guidelines on the Use of AI Tools for Academic Work** in typical scenarios, which do not rule out more (or less) restrictive policies that may be imposed by instructors and units given specific pedagogical needs.

We recommend that this document be disseminated to all NUS teaching staff to guide their own course design, assessment planning, and engagement with students, and the two sections addressed to students be broadcast to the NUS student community. The information from the two sections should also be hosted online. We also recommend that NUS invest in infographics and other media to present the interim policy to students and staff in an attractive and concise manner.

Another part of the Workgroup's task will be to guide NUS Centre for Instructional Technologies in their development of responses, including the implementation of AI detection tools, and these efforts are already underway. However, even if we can develop a good detector for ChatGPT, other AI tools will emerge, and we cannot be developing detectors all the time. The evidentiary value of detector results in a disciplinary hearing is also weak. This is because the typical detector expresses its results as statistical likelihoods, are not fully accurate (they can generate both false negative and false positive results), and often have a bias against non-native English speakers. Consequently, they are wide open to challenges, and we cannot use them as the sole basis to impose penalties on students. At best, they provide a preliminary indication of students who might require additional scrutiny, but even so, colleagues need to be aware that there may be students using AI tools who were not detected at all (false negatives). In short, our response as educators cannot rely on detectors. Rather, we should be aiming to set questions and grading schemes that reward higher order skills outside of the current reach of AI, or require students to sit for in-person exams if we are testing for basic skills.

1.3 Pedagogical Considerations

In general, the Workgroup is of the opinion that instructors and course designers should decide whether students are allowed to use ChatGPT (or any other tool) in an assessment or task depending on whether the use of that tool defeats or advances the learning goal of the assessment or task. However, the feasibility of enforcement must be kept in mind. That is, if we need to be sure that students are not using such tools in an assessment or task, the prudent thing is to arrange for a supervised setting, for instance, a small group in-person meeting where it is possible to visually monitor whether students are using connected devices, or an on-site proctored exam, rather than assign a take-home, unsupervised task.

Conversely, if a supervised arrangement is not feasible, it probably means that the assessment or task is no longer appropriate or suitable due to the availability of tools such as ChatGPT. In that case, colleagues should be targeting higher-order assessment outcomes, such that students cannot simply use an AI tool to generate a passable answer without knowing what he or she is looking for in the first place and without extensive prompt engineering. Alternatively, we can design assignments which directly integrate the use of AI. A sample of such assignments are attached as an appendix. CDTL has also provided training to colleagues about constructing good assessments that reward higher-order skills, and colleagues are encouraged to upskill themselves through these resources.

In general, course designers should endeavor to craft assignments with care, taking into account the specific circumstances of individual courses, curriculum, assessment objectives, class sizes, and distribution of learners' needs and abilities. They should also be ready to explain the rationale behind them to students in a thoughtful way. Where appropriate, they should explain more explicitly the higher-order outcomes they are targeting and how AI tools may be insufficient for the demands of the assessments.

In closing, we encourage colleagues not to shy away from conversations with students about the pitfalls of using AI tools, but also how they can be used to enhance student learning, and what students should and should not do—the Interim Policy is meant to give you a starting point. At the end of the day, colleagues can temper both our excitement and anxiety. Most of our students are not as dishonest (or indolent) as we are sometimes led to fear, even though we need plans to guard against the errant minority degrading the integrity of our assessment systems. AI also has the potential to help students formulate ideas, develop new questions, improve writing, and beyond that, push us as educators to clarify what we are trying to do in our own teaching. The balance is to make sure students learn to use these tools in an ethical way that does not compromise their own longer-term learning or the integrity of our academic processes.

2. A General Message to Our Students

Artificial Intelligence tools such as ChatGPT are now essentially costless and no longer require special expertise to use. As part of the repertoire of tools available in the modern world, we understand how you—our students—might be drawn to using them. It is also likely employers and stakeholders will soon assume that graduates have a basic competency for using such tools, just like how they now assume that new hires are able to use calculators, spreadsheets, and word processing software.

That said, there are some things we need you to take to heart.

First, while we are fine with you experimenting with new tools, we are also wary you will end up taking shortcuts that will disadvantage you in the longer run. Consider, for instance, the different levels of capability involved in the following three scenarios:

- (1) Using a tool to generate an output wholesale from inputs.
- (2) Using a tool to generate intermediate outputs that are then developed into a final output through further human intervention without the use of the tool.
- (3) Evaluating the output of a tool to confirm its accuracy, relevance, objectivity, and completeness.

You will need a higher level of capability in yourself to be able to do (2) and (3), as compared with doing (1), and conversely, functions that only require (1) are at a higher risk of being completely replaced by AI. The implication is that you are only cheating your future selves if you go straight to

using such tools before learning the actual subject matter. If you, as a learner, take shortcuts today, you risk becoming first in line to be replaced by bots!

Since we target learning outcomes that hone higher-level capabilities, and we want our students to demonstrate originality, voice, intellectual engagement with content, rigor, and contribution to discourse, we will often need you to produce unaided work as part of your assessment. This is no different from how, even though we expect you to use calculators and spreadsheets, we still want you to learn the underlying math and computational processes. And even though we are fine with you using search engines, spelling and grammar checkers, or for that matter, lexicons and thesauri, we still want you to be able to independently organize your thoughts and produce your own compositions—so that you can assess the tools’ deliverances.

Second, note that normal academic rules continue to apply. As provided in the Code of Student Conduct (<https://studentconduct.nus.edu.sg/wp-content/uploads/NUS-Code-of-Student-Conduct.pdf>):

The University takes a strict view of cheating in any form, deceptive fabrication, plagiarism and violation of intellectual property and copyright laws. Any student who is found to have engaged in such misconduct will be subject to disciplinary action by the University.

In this regard, representing an AI’s output as your own work is plagiarism. (See the below “3. Guidelines on the Use of AI Tools for Academic Work”.)

Third, do be aware that the tools typically require you to sign up for an account and accept various terms and conditions. We encourage you to be very clear about what you are accepting if you sign up for an account. You should also be aware that there are lawsuits underway challenging the intellectual property provenance of many of these tools. If they are ruled illegal (while this is unlikely, it is not impossible), there may be implications for your use as well. And finally, don’t forget that like all online tools, platforms can sometimes go down without warning, or are under such high load that their responses become too slow for use. It is thus unwise to assume that you will always have access.

All in all, we hope that you will be thoughtful about how you can advance your learning and make the most of your university education.

3. Guidelines on the Use of AI Tools for Academic Work

3.1 Don’t Use AI to Plagiarize

The following are always improper uses of AI tools:

- Generating an output and presenting it as your own work or idea.
- Generating an output, paraphrasing it, and then presenting the output as your own work or idea.
- Processing an original source not created by yourself to plagiarize it (e.g., using an AI paraphrasing tool to disguise someone else’s original work, or even the output of an AI tool, and then presenting the final output as your own work or idea).

All of the above violate NUS policies on academic honesty and anyone found to have done any of them will be dealt with accordingly. Keep in mind that even though AI tools are not authors and thus cannot be harmed by someone stealing an idea from them, it’s still wrong of you to represent yourself as having produced something when you didn’t produce it.

3.2 Check with Your Instructors on Proper Uses of AI Tools

Whether or not using an AI tool in a particular way is allowable depends on the learning purposes of the course and the targeted outcomes of the assignment. Some possible legitimate uses include:

- Gathering information and looking up explanations for basic concepts.
- Generating output for critique and analysis, for self-learning, or to compare against one’s work for self-evaluation and improvement.
- Help with proofreading and editing writing work.

The above is not meant to be comprehensive. An assignment designed to integrate the use of an AI tool, for instance, may require you to use that tool more extensively. Conversely, if there is a need to test whether you possess a certain knowledge or capability without access to AI tools or other resources, your instructors will continue to arrange for appropriate assessment settings (e.g., an on-site proctored exam or oral interview). In general, course instructors will need to impose varying restriction levels for the use of AI tools depending on the learning outcomes targeted. Whenever you have any doubts about whether an AI tool could be used for a specific assignment, or how it could be used, clarify them directly with your course instructors.

3.3 Acknowledging Your Use of AI

If you completed any work with the aid of an AI tool, assuming a setting in which the instructor gave permission for such tools to be used, you should always acknowledge the use. Using the outputs of an AI tool without proper acknowledgement is equivalent to lifting or paraphrasing a paragraph from a source without citation and attracts the same sanctions. You can give this acknowledgement through a note or “methods section” at the end of the assignment explaining, e.g., which AI tools were used, in which parts of the process they were used, what were the prompts used to generate results, and what you did with the outputs to add value.

One way this can be done is in a tabular form as shown below:

AI Tool used	Prompt and output	How the output was used in the assignment

Alternatively, if an AI tool was used to generate a more extensive set of intermediate outputs that were then developed into a final product, you should preserve a full transcript of the relevant interactions with the AI as an appendix for submission with your assignment. Your instructor may also require you to explicitly declare if AI tools were not used in a part of your assignment, or for the whole assignment. In all cases, seek advice from your course instructor.

3.4 You are Responsible for Your Work

Remember the limitations of current generation AI tools:

- Output’s quality is dependent on the quality of the users’ prompts.
- Output may be out of date, as they are dependent on the available training data.
- Output may not be accurate (e.g., they don’t always present information that is true, the ‘citations’ they may generate may be made-up and point to non-existent sources).

- Output may present dominant values and opinions as truth not because other views are incorrect, but simply because dominant claims are more common in the training data.
- Output may be offensive or discriminatory, as AI tools cannot make opinions or judgment calls aligned with legal and social norms.

You should thus always assume that the AI's output is incorrect until you have separately checked it against reliable sources (citing those sources properly) or have gone through the workings yourself. You also cannot assume that the AI's output is relevant and sufficiently contextualized for your purposes. In some cases, the rhetorical structure of the AI's output is usable, but the details of the content are not. Always remember that you, rather than the AI tool, are responsible for the quality and integrity of the work you submit. AI tools are tools, and as such, cannot take responsibility for any information or text that they produce.

3.5 Start a Conversation with Your Instructors about the Use of AI

These guidelines are framed with typical scenarios in mind, and there's bound to be uncertainties during this transition period. So, whenever you are in doubt, clarify directly with your course instructors. If you are going for an overseas exchange, find out what the host university's policy on the use of AI is and clarify any doubts with instructors there. Don't assume that NUS policy is universally applicable.

Remember that just because there are legitimate uses for AI tools in your academic work, it does not mean you should resort to them at every turn, especially if you are still learning the subject matter. By jumping straight to using the tools, you may end up missing an opportunity to learn the subject matter for yourself. Furthermore, if you don't already have the subject matter knowledge yourself, you might not even be able to tell if the output is accurate or relevant. In many cases, there are also better resources you can access. For instance, if you need help with proofreading and editing, you can turn to the NUS Writers' Centre (<https://nus.mywconline.com>); you will likely learn more that way too!

More generally, do not be shy to approach your instructors to start conversations about how the learning outcomes they are targeting go beyond what the AI tools can deliver, and how you can use AI tools in ways that will enhance your own learning in your courses.

Acknowledgements: The University Policy Workgroup for AI in Teaching and Learning extends its thanks to individuals who contributed suggestions: Prof Simon Chesterman, Prof Bernard Tan, Prof Wong Lim Soon, A/P Melvin Yap, A/P Stuart Derbyshire, Dr Zeehan Jaafar, Dr Jonathan Frome, Dr Yew Kong Leong, Mr Jonathan Sim, and other staff and student members of the NUS Community.

Appendix 1: Resources and Further Reading

- Sabrina Ortiz, “What is ChatGPT and why does it matter? Here’s everything you need to know,” ZDNet (Jan 27, 2023)
<https://www.zdnet.com/article/what-is-chatgpt-and-why-does-it-matter-heres-everything-you-need-to-know/>
- Marco Ramponi, “How ChatGPT actually works,” AssemblyAI Blog (Dec 23, 2022)
<https://www.assemblyai.com/blog/how-chatgpt-actually-works/>
- Stephen Marche, “The College Essay Is Dead: Nobody is prepared for how AI will transform academia,” The Atlantic (Dec 22, 2022)
<https://www.theatlantic.com/technology/archive/2022/12/chatgpt-ai-writing-college-student-essays/672371/>
- Osmond Chia, “Teachers v ChatGPT: Schools face new challenge in fight against plagiarism,” Straits Times (Jan 6, 2023)
<https://www.straitstimes.com/tech/teachers-v-chatgpt-schools-face-new-challenge-in-fight-against-plagiarism>
- Kevin Roose, “Don’t ban ChatGPT in schools, but teach with it,” Straits Times (Jan 16, 2023)
<https://www.straitstimes.com/tech/tech-news/don-t-ban-chatgpt-in-schools-but-teach-with-it>
- Irene Tham, “5 plagiarism detection tools to tell if content is written by a bot like ChatGPT,” Straits Times (Jan 26, 2023)
<https://www.straitstimes.com/tech/tech-news/5-plagiarism-detection-tools-to-tell-if-content-is-written-by-a-bot-like-chatgpt>
- Sam Illingworth, “Students could use ChatGPT to cheat, but it's a chance to rethink assessment altogether,” Channel News Asia (Jan 23, 2023)
<https://www.channelnewsasia.com/commentary/chatgpt-ai-students-cheat-homework-rethink-assessment-3220386>
- “Heart of the Matter - S3E23: Is ChatGPT the start of a big revolution in education?” Channel News Asia (Jan 20, 2023)
<https://www.channelnewsasia.com/listen/heart-matter/chatgpt-start-big-revolution-education-3219431>
- Blayne Haggart, “ChatGPT is a dagger aimed directly at academia and the news media”, Straits Times (Feb 1, 2023)
<https://www.straitstimes.com/opinion/chatgpt-is-a-dagger-aimed-directly-at-academia-and-the-news-media>

Appendix 2: Possible Assessments Involving ChatGPT

By Mr Jonathan Sim (Department of Philosophy)

Assessments for Students to See the Limits of AI and Go Beyond It

1. Students are to generate an essay written by an AI. They will then copy the entire essay into Microsoft Word, enable track changes, and make edits. They are also required to add comments about specific sentences in the AI-generated essay to evaluate its strengths/weaknesses. As an exercise, this will allow students to better understand what makes an essay a good essay, and how to avoid features of poor essay writing. Furthermore, students will see the limits of AI, and we will be able to push them to think and write better.
2. Students are to get the AI to generate a summary of a reading. They will then comment on areas where they feel the summary lacks precision, or how it missed out on certain nuances in the reading. As an exercise, this will enable students to engage in deep reading and to spot subtle nuances in the assigned reading.
3. Students are to prompt the AI to generate an essay about a particular subject (ideally something controversial). They then research and find out the possible sources that led to the AI's output. Students will then assess how biased the output is (e.g., is it leaning more towards a liberal view of the topic, or might it have a certain anti-liberal view on the topic?). To do this, they will have to research positions contrary to the AI's output. As an exercise, it will compel students to broaden their horizon to include views filtered out by biases in the AI's training data, and they can assess how objective or biased the AI might be on a particular subject.

Assessments for Students to Learn Together with the AI

1. Students are to prompt the AI to teach a particular concept. Given that ChatGPT will not be able to provide too much depth, students will then need to enhance the explanation of that concept and tie in relevant examples/applications/case studies to make a point about a specific issue. As an exercise, this will enable students to see ChatGPT as a learning tool for providing them with a base to work from, while recognizing limitations in its ability to recall/explain. Students will need to add value to the AI-generated material so that they and their peers can learn more insightfully from the exercise or sharing (if this is a class presentation).
2. Students are given a few pre-written essays containing a mix of truths and factual errors, as well as a mix of sound and flawed arguments (instructors can generate a few directly using ChatGPT, or edit the output to produce them). They are to pick one of these essays and feed it into the AI and prompt it to generate evaluations and feedback for improvement. Students will then comment on every feedback point (whether they agree/disagree, or if they think the AI's feedback missed the point made in the pre-written essay—a process similar to how researchers reply to peer reviewers for journal publications). They will need to include additional feedback on any remaining factual errors and flawed reasoning that the AI missed. As an exercise, students will have to conduct their own fact-checking research, while also evaluating the arguments of the pre-written essay, and thus learn that they cannot rely solely on the AI to get the job done. On top of that, by evaluating the AI's feedback, students will learn how to actively engage with feedback to critically assess how relevant a piece of

feedback is given the objective of the pre-written essay. And this can prepare them to be better researchers and students.

3. Students are to engage in a conversation with the AI to teach it a niche topic, with the goal of producing a short write-up (it can be an essay, poem, play, or even a song). Through the conversation, students are to figure out what ChatGPT does not know (or what it misunderstands) and attempt to teach it facts or share with the AI additional insights that they might have on the matter. A significant portion of the grading should be based on the quality of information and the clarity of communication when the student teaches the AI (they will have to submit the entire conversation transcript). As an exercise, this will be an opportunity for students to engage in peer teaching, but with an AI as a peer learner in place of a human. As they learn to teach the AI, they will have to engage in research and reflection to feed it with relevant facts and insights. Since students can redo the entire conversation from the start, they will get instant feedback about the clarity of their writing and teaching, and thus learn and improve through these multiple attempts.