



Policy for Use of AI in Teaching and Learning

University Policy Workgroup for AI in Teaching and Learning

LOY Hui Chieh (Chair; Philosophy), LIOU 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), Stella TAN (Student Conduct)

7 Aug 2024

© 2024 National University of Singapore

Policy for 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), A/P Natalie Pang (NUS Libraries), A/P Stella Tan (Student Conduct)

Interim version 5 Feb 2023

Approved 5 Feb 2023

Updated and expanded 7 Aug 2024

Approved 7 Aug 2024

1. Introduction

On 30 November 2022, the US based research laboratory OpenAI launched "ChatGPT", a generative artificial intelligence (Gen AI) model capable of processing and producing text in natural language, as a free service for the public. Observers soon noticed the challenge that such tools posed to assessment in higher education, as they are fully capable of producing output difficult to distinguish from student-produced work. Modalities affected include essays, solutions to mathematical problems, computer code, and even multiple-choice questions (AI tools can also generate explanations for why each option is correct or incorrect). The impression is that current Generative AI tools can produce work that is passable, especially in introductory classes, even if it is not uniformly of a quality approximating what our better students can produce or what is required of them in more advanced classes.

Given the potential misuse of such tools for academic dishonesty, the University Policy Workgroup for AI in Teaching and Learning was formed in January 2023, 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. The Workgroup's remit centers on compliance and discipline, while other university units focus on positive uses of AI tools both to support the delivery of education and to enhance our students' learning. Later in 2023, the newly formed [NUS AI Community-of-Practice](#) brings together work being done by different units in the university, and the Workgroup now comes under its purview.

This document updates the Workgroup's first Interim Policy for the Use of AI in Teaching and Learning published in February 2023. The Policy is articulated on the assumption that, as a university, we will embrace the ethical and appropriate use of AI tools among students and staff. Beyond this Introduction, the Policy consists of these sections:

Section 2. **Guidelines for Instructors regarding use of AI in NUS courses;**

Section 3. A **General Message to Students** about broad principles, which we recommend that educators make use of when explaining to students the rationale for our guidelines; and

Section 4. **Guidelines governing students' use of AI tools** in typical scenarios.

We recommend that this document be disseminated to all NUS teaching staff to guide their own course design and assessment planning. Instructors can also broadcast suitable portions from it to their own students, as called for by their situation. To support this, Sections 3 and 4 are framed as we, the educators, addressing our students.

2. Guidelines for Instructors and the Use of AI

NUS instructors are encouraged to explore how AI can be used to enhance the efficiency and quality of their education efforts. Sections 2.1-2.4 are guidelines governing instructors' use of AI when delivering courses. Instructors are also encouraged to incorporate the use of AI in course activities when doing so would make students' learning more effective, and not to issue blanket and unenforceable prohibitions against the use of AI by students. Sections 2.5-2.8 relate to students' use of AI in NUS courses.

2.1 Be Transparent About the Use of AI

Instructors should be transparent about where and how they deploy AI in NUS courses. This is especially important where AI is deployed to generate course content (including assessment questions), to function as virtual tutors to answer student queries, or to help with assessment feedback and grading. A succinct declaration can be made via the course page on the Learning Management System.

2.2 Instructors have Final Responsibility

Instructors have the final responsibility for the quality of their courses. They are liable for any errors whether AI is utilized or not to create content for the course or to deliver it. This includes the course content (including assessment questions), responses to queries relating to the content, and the accuracy of the assessment scores and any feedback given to students, especially for graded work.

2.3 Use Approved Tools

Wherever NUS data is involved, use NUS approved AI tools (see the list [here](#)).

2.4 Deploying AI Requires Explicit Approval

The use of AI tools to provide instruction to learners in the form of responses, feedback and/or marks, whether as virtual tutors or as markers, requires prior approval by Head of Department or relevant Deanery, under the oversight of Chair of the AI-COP. This is meant as due diligence to balance innovation with ethical considerations. Institutional approval also shields the individual instructors from complaints. Approval must be sought through submission of an [AI Risk Assessment](#). Approval can be granted only when supported by sound pedagogy, and evidence of the tool's accuracy and trustworthiness.

In general, except for AI marking of objective or close-ended assessments, a human instructor or reviewer should always be in the loop, e.g., to review, revise, and finalize the output of the AI. Non-compliance may be deemed academic misconduct.

In general, where AI tools are used to provide instruction, it is preferable to use them to generate first-pass evaluations that instructors can then review and make any necessary adjustments.

Note that even where there is confidence that an AI tool can reliably generate accurate scores or comments without human intervention, spot checks should still be conducted on the output. In all cases, the instructor remains ultimately responsible for the use of the tool and its impact on learners (see 2.2 above).

2.5 Assessments Forbidding the Use of AI Should be Backed by the Appropriate Assessment Setting

Course designers will need to decide whether students should be allowed to use AI tools in an assessment, depending on whether the use of that tool defeats or advances the learning goal of the assessment.

If the decision is that students should be forbidden from using AI tools for an assessment (for pedagogical reasons), then crucial aspects of that assessment should be conducted in-person and instructor-supervised, to ensure that students do not access those tools.

Conversely, the default assumption for any unsupervised (e.g., “take home”) assessment task is that the use of AI tools is **permitted** so long as the use is duly acknowledged. (Whether the use of such tools is **advisable** is a different issue; see 2.7 below.)

For unsupervised assessments, instructors should aim to set tasks that require higher-order skills going beyond the capabilities of generative AI using simple prompts. Submissions for unsupervised assessments that duly acknowledge the use of AI should be graded on the quality of the work alone and not penalized just for having used AI.

2.6 Academic Dishonesty and AI Detector Verdicts

A student **found** to have submitted work generated by AI but **failed to acknowledge their use of AI** can still be sanctioned for plagiarism, assuming the case can be made. Such a student has committed academic dishonesty in misrepresenting the nature and source of their work.

The verdicts of current AI tools purported to determine whether an analyzed input has been generated by AI are **not admissible as conclusive evidence** in a disciplinary process to charge a student with academic dishonesty or as justification to penalize student work.

(See also Appendix 3: Uniform Treatment of Plagiarism and Misuse of AI in Written Assessment Work, in [NUS Plagiarism Policy](#))

2.7 Communicate with Your Students About the Advisable and Inadvisable Uses of AI for Learning Purposes

Instructors should communicate with students about the advantages and pitfalls of using AI tools in the context of the courses they are teaching. One way to do so is to publish a document that goes through various typical ways in which students might want to use AI tools in the course, laying out the pros and cons for each, with actual test examples. Typical scenarios include: “Using AI to better understand a reading or concept”, “Using AI to generate different assessment answers”, “Using AI to brainstorm ideas”, “Using AI to improve one’s writing”, “Using AI to generate code”, and so on. Some samples can be found on the AI COP site [here](#).

2.8 AI Tools can contribute to Digital Divide

While some AI tools are essentially free, better classes of AI tools can be utilized by students with more resources at their disposal. Instructors should thus be aware that the unfettered use of AI

tools by students without guidance may inadvertently result in a differentiation between students caused by what AI tools they are able to access, rather than student's mastery of the subject.

2.9 Final Remarks

AI has the potential to help with student learning just as it will also 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 make use of these tools in an ethical way that does not compromise their own longer-term learning or the integrity of our academic processes. In addition, the capabilities of AI tools will necessarily change over time as they improve. Instructors may want to spend effort assessing the quality of assistance rendered by specific AI tools and develop local policies about which tools are advisable for which tasks or assessments.

3. A General Message to Our Students

Artificial intelligence tools such as ChatGPT are now essentially costless and do not 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 you grasp the actual subject matter for yourself. 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](#):

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.

Given the above, representing an AI's output as your own work, without any acknowledgement that that you have used such a tool, is plagiarism. (See "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.

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

4.1 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 without attribution.
- Generating an output, paraphrasing it, and then presenting the output as your own work or idea without attribution.
- 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 another AI tool, and then presenting the final output as your own work or idea) without attribution.

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 is still wrong of you to represent yourself as having produced something when you did not produce it.

Note that students do not have to paraphrase AI output where its use is permitted. In such cases, so long as they acknowledge the use of AI, they have not committed plagiarism, and they will not be penalized for declaring its use. The assignment will be marked based solely on the quality of the submission. However, it is the students' responsibility to check that the AI output properly engages with the assignment prompt and has the appropriate tone and style.

4.2 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.

4.3 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. In fact, if you are ever in doubt, it is always a good idea to declare your use of a tool. Using the output 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 is 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 can also preserve a full transcript of the relevant interactions with the AI as an appendix for submission with your assignment. Your instructor may also require that if AI tools were not used in a specific part of your assignment, you should declare that explicitly. In all cases, seek advice from your course instructor.

4.4 You Are Responsible for Your Work

Remember the limitations of current generative 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 may not produce opinions or judgment calls that are always 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.

4.5 Start a Conversation with your Instructors about the Use of AI

These guidelines are framed with typical scenarios in mind, and there're bound to be uncertainties as the field of Artificial Intelligence continues to evolve over the years to come. 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 our university's 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. There are also often better resources you can access. For instance, if you need help with proofreading and editing, you can turn to the [NUS Libraries Writers' Centre](#); you will 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. The instructor facing side of these documents tells them to do the same—start a conversation with you!