

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), A/P Martin Buist (Biomedical Engineering), A/P Chen Zhixiong (Physiology), A/P Kan Min-Yen (Computer Science), Prof Ng Hwee Tou (Computer Science), A/P Natalie Pang (NUS Libraries), Prof Stella Tan (Student Conduct), Prof Christian Witting (Law)

First version 5 Feb 2023; Second version 7 Aug 2024; Current version 15 Sept 2026

1. Introduction

1.1 Background

Artificial intelligence (AI), including generative AI, is increasingly shaping how knowledge is accessed, produced, communicated, and evaluated. Its rapid development has created both opportunities and challenges for higher education.

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 pose to assessment in higher education, as they are fully capable of producing output difficult to distinguish from student-produced work. There is also the fear that students will be harmed by their use of such tools in ways that bypass the intellectual work through which learning occurs.

Given the potential misuse of such tools for academic dishonesty and the more general risk they pose to learning, 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. The Workgroup's purpose is 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 centres 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. In 2023, the newly-formed NUS AI Community-of-Practice (hereafter, AI-COP) brought together work being done by different units in the University, and the Workgroup now comes under its purview.

This document presents the latest revision to the Workgroup's Policy for the Use of AI in Teaching and Learning. The Interim Policy was first published in February 2023 and subsequently revised in August 2024. 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.

1.2 Scope of the Policy

This policy covers the use of AI Tools in all NUS teaching and learning contexts, and applies to all instructors, teaching assistants, program staff, and AI-enabled systems used in support of teaching and learning.

“AI Tools” refers to software applications or systems that use artificial intelligence techniques to perform tasks that typically require human intelligence, including but not limited to analysis, decision-making, pattern recognition, and automation. Within this larger umbrella, **“Generative AI (GenAI) Tools”** refers to AI-powered applications or systems that create, generate, or produce new content—including text, images, audio, video, or code—in response to user inputs or prompts.

1.3 Parts of the Policy

Beyond this Introduction, the Policy consists of these sections:

2. Guidelines for Instructors on the Use of AI;

3. 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

4. Guidelines on the Use of AI Tools for Academic Work 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 appropriate. To support this approach, Sections 3 and 4 are written from the perspective of educators addressing their students.

2. Guidelines for Instructors on 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 may be placed on the course web page, syllabus, assignment brief, or other appropriate communication channel. Such declarations should indicate the nature of the AI use and, where relevant, any implications for students.

2.2 Instructors have Final Responsibility

Instructors have the final responsibility for the quality of their courses. They are responsible for any errors, whether or not AI is used 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. The use of AI does not diminish or transfer the instructor’s responsibility for these matters.

2.3 Use Approved Tools

Instructors should use NUS-approved AI tools (see the list [here](#)) whenever NUS data is involved. “**NUS data**” refers to information or data that are created, processed, stored, or transmitted by NUS staff, students, or agents, including all data which the University has a legal or contractual obligation to manage, protect, or control.

2.4 Deploying AI to Provide Instruction, Feedback, or Marks 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 the Head of Department or relevant Deanery, under the oversight of the 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 obtained through a submission to the [AI System for Assurance and Evaluation](#). Approval will 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 closed-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. Failure to comply with this section may constitute a breach of University policy and may be addressed through the relevant staff governance processes. 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 to. 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 their use defeats or advances the assessment’s learning goals.

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 under instructor supervision, 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 such as reasoning, judgement, critique and application, rather than tasks that reward simple prompting. 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.

Note that an extended assessment can contain components conducted in a supervised setting where AI is not allowed and also unsupervised components where the use of AI is permitted, and the student’s work is graded holistically.

2.6 Academic Dishonesty and AI Detector Verdicts

A student **found** to have submitted work generated by AI who has **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 purporting to determine whether an analysed input has been generated by AI are **not admissible as 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. Samples can be found on the AI-COP site [here](#).

2.8 AI Tools can contribute to the Digital Divide

NUS provides students with access to an institutionally supported baseline of AI tools for learning, with specialised or advanced tools provided according to curricular need. Instructors should nevertheless be mindful that differences in access to other AI tools may create or amplify inequities among students. Where the use of AI tools is relevant to a course or assessment, instructors should ensure that differences in access do not advantage students primarily because of the tools they can afford or obtain, rather than their 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 key 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 to regularly re-assess the quality of assistance rendered by specific AI tools and develop local policies about which tools are advisable for each task or assessment.

3. A General Message to Our Students

Artificial intelligence tools such as ChatGPT are now widely accessible 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. In many workplaces and professional settings, graduates are increasingly expected to know how to use such tools appropriately, just

as they are expected to know how 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 using the 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 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 because you will not have developed any independent skill set!

Since we target learning outcomes that hone higher-level capabilities, and we want our students to demonstrate originality, voice, intellectual engagement with content, rigour, and contribution to discourse, we will sometimes need you to demonstrate these capabilities independently of AI through appropriately supervised or otherwise verifiable 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 mathematical and computational processes. And even though we accept your use of search engines, spelling and grammar checkers, 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' outputs.

Second, the ordinary rules of academic integrity 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 you have used such a tool, is plagiarism. (See "[Guidelines on the Use of AI Tools for Academic Work](#)".)

Third, do be mindful that the tools typically require you to sign up for an account and accept various terms and conditions. We encourage you to be clear about what you are accepting if you sign up for an account. The legal and intellectual property position surrounding generative AI continues to evolve, which may affect how particular tools or outputs may be used. And finally, don't forget that like all online tools, platforms can become inaccessible without warning. It is thus unwise to assume that you will always have access.

In no circumstances should AI tool accounts be shared with others. The use of a shared account may raise concerns about collusion and, where misconduct is established, may result in disciplinary action. All in all, we encourage you to use AI thoughtfully, responsibly, and in ways that support rather than diminish your learning.

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. Anyone found to have infringed these provisions will be dealt with accordingly. Keep in mind that even though AI tools are not authors (in the conventional sense) and thus cannot be harmed by someone stealing an idea from them, it is still wrong to represent yourself as having produced something when you did not produce it. More importantly, teachers can only assess properly when they know the extent of the student contribution.

Where the use of AI is permitted, students are not required to paraphrase AI-generated material and certainly shouldn't be doing so to conceal its source. What matters is that the use of AI is properly acknowledged, and that the submitted work is fit for purpose, accurate, and responsive to the assignment. In such cases, so long as students 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.

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 written work.

The above is not exhaustive. 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 knowledge or a certain capability independent of 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 may impose varying levels of restriction for the use of AI tools depending on the learning outcomes targeted.

Whenever you have any doubts about whether an AI tool may be used for a specific assignment, or about how it may be used, clarify them directly with your course instructor.

4.3 Acknowledging your use of AI

If you completed any work with the aid of an AI tool, assuming that such use is permitted under the course or assessment rules, 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 and of any other external assistance. Using the output of an AI tool without proper acknowledgement is equivalent to copying or paraphrasing a paragraph from a source without citation and attracts the same sanctions.

A simple declaration may take the following form:

I used [AI tool name, e.g., GPT-4.1] to [describe specific uses: e.g., generate ideas, format paragraphs, improve expression, analyse effectiveness, create images and illustrations, produce drafts, refine, and/or finalise my assignment]. I am responsible for the content, accuracy, and quality of the submitted work.

Individual course instructors may also provide their own forms to facilitate declaration. Or they may require that if AI tools were not used in a specific part of your assignment, you should declare that explicitly.

In all cases of doubt, seek advice from your course instructor.

4.4 You Are Responsible for Your Work

Always remember that even in cases where the use of AI is permitted, the student submitting a piece of work is the one accountable for its quality and integrity, and for its engagement with the assignment prompt and having the appropriate tone and style.

A student who used AI when creating a piece of work and properly acknowledged that use is not liable for having plagiarised by using AI (see 4.3 above). But they may still be penalized for submitting inaccurate or shoddy work. Egregious cases involving non-existent data may also be treated as data falsification.

Remember the limitations of current generative AI tools:

- Output quality is dependent on the quality of the users' prompts.
- Output may be out of date.
- Output may be inaccurate (e.g., information produced might not be true, the 'citations' generated may be made-up and incorporate non-existent sources, or the 'data' provided may be the product of hallucination).
- Output may present dominant values and opinions as truth, not because other views are incorrect, but because dominant claims are more commonly found in the literature.
- Output may be offensive or discriminatory, as AI tools may produce opinions or judgment calls that are not 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 might be usable, while the details provided are more dubious.

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

These guidelines are framed with typical scenarios in mind, and there are bound to be uncertainties as AI continues to evolve and its uses multiply. 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 in the course of a learning process which requires individual effort in order to succeed. Furthermore, if you don't already have the subject matter knowledge and understanding, you might not even be able to tell if the output is accurate or relevant. Often, there are better resources you can access. For instance, if you need help with proofreading and editing, you can schedule a writing consultation [here](#); 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. Instructors are advised to do the same—that is, to start a conversation with you!