DEPARTMENT FOR EDUCATION INNOVATION

Lecturer's Guide: Leveraging Generative Artificial Intelligence for Teaching and Learning Enhancement at the University of Pretoria

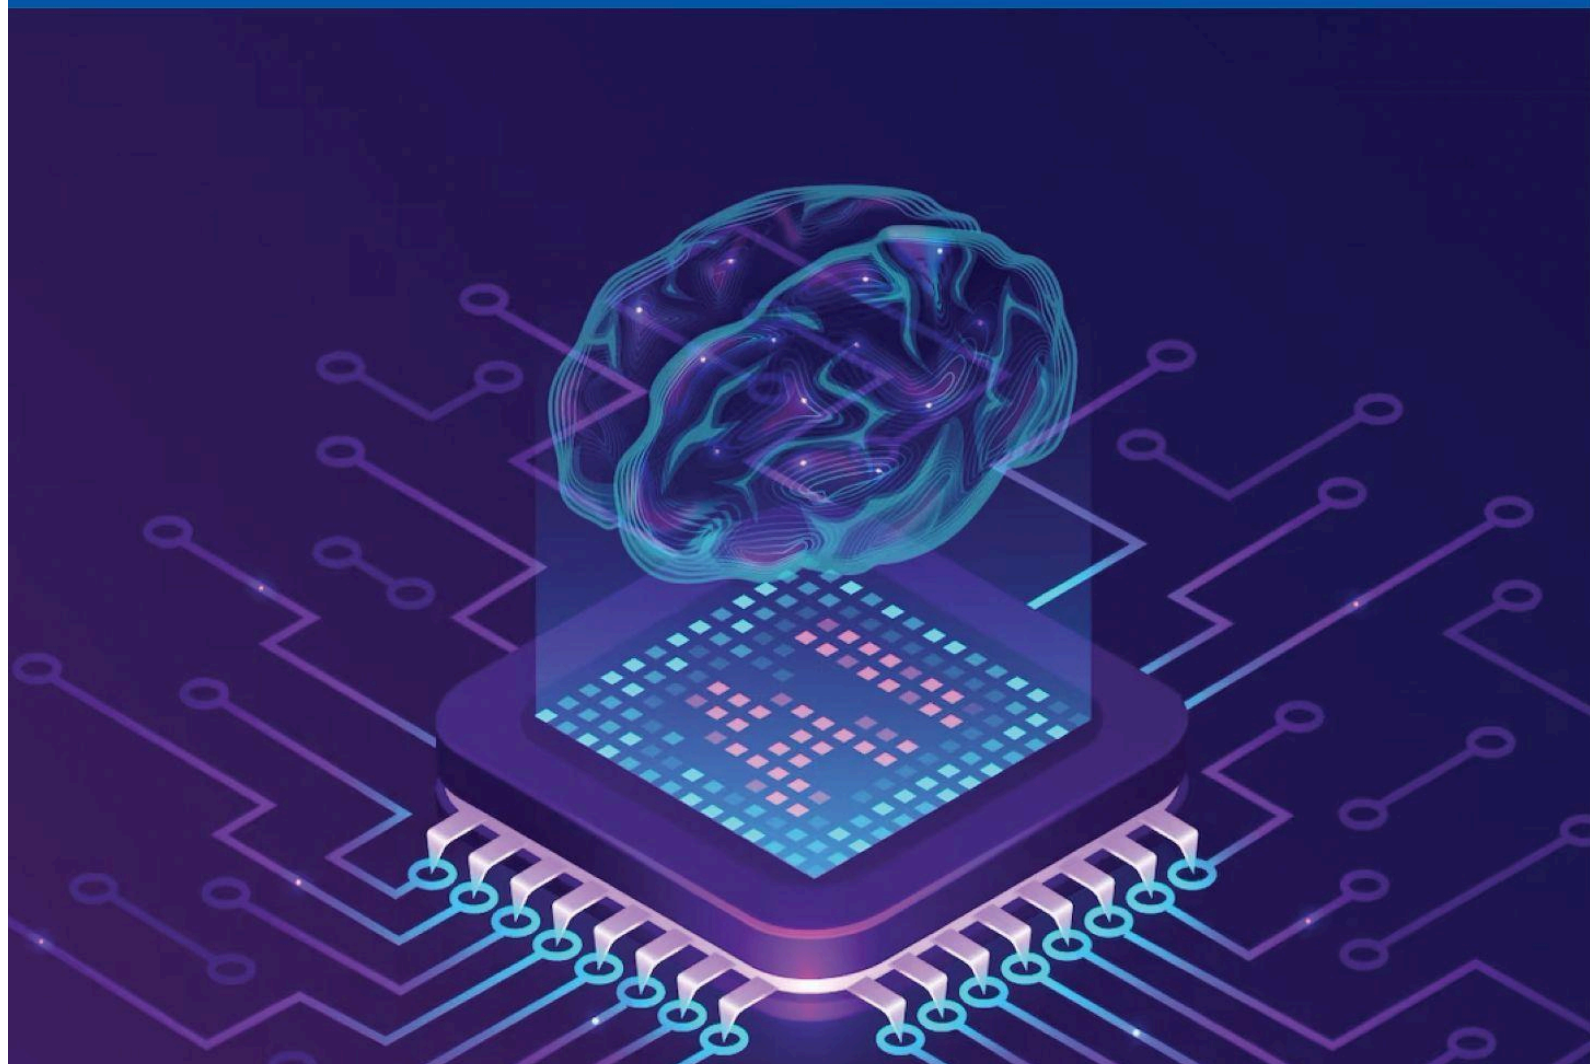

UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA

Department for Education Innovation
Departement Onderwysinnovasie
Kgoro ya Tšwetšopele ya Thuto

Make today matter
www.up.ac.za

Table of Contents

Acknowledgement: The creation, enrichment, and editing of this article for enhanced clarity were facilitated by the use of ChatGPT.

1 Introduction

Quick Start for lecturers (4 things to do this week)!

1. **Discuss responsible use with students in Week 1:** Cover ethics, disclosure, privacy, and cross-checking; include a 10-minute activity on critiquing an AI output. (See 3.4–3.7 & 3.9.)
2. **Publish your AI rules in the study guide:** Add a short "Permitted AI" statement for every task (prohibited/limited/encouraged) and point students to the traffic-light table. (See 3.9.1–3.9.3.)
3. **Complete the Module AI Exposure self-assessment:** Use Appendix A to calculate exposure (See 3.9.4 & Appendix A). Run each task through ChatGPT/Gemini; if the AI can score well, redesign or shift to in-class assignments that are invigilated. (See 3.9.3.)
4. **Require the Generative AI Declaration on every submission:** Collect prompts/process evidence (drafts, version history) where appropriate. (See 3.9.3 & Appendix B.)

Artificial Intelligence (AI), particularly in its generative forms, is impacting many facets of human life, from communication to entertainment, and education is no exception. This guide serves as an introductory resource for lecturers at the University of Pretoria who wish to explore and harness the potential of generative AI to enhance teaching and learning outcomes. Large language models have demonstrated impressive capabilities in identifying language patterns and predicting words in context, excelling at generating coherent and relevant text with minimal user input. By leveraging their extensive training on language data, these models can produce creative poems, write comprehensive essays, analyse topics in depth, and present arguments persuasively, all in response to user prompts. However, they are fallible: outputs can be inaccurate or fabricated. Use AI as a starting point only and subject all items and materials to human verification and moderation.

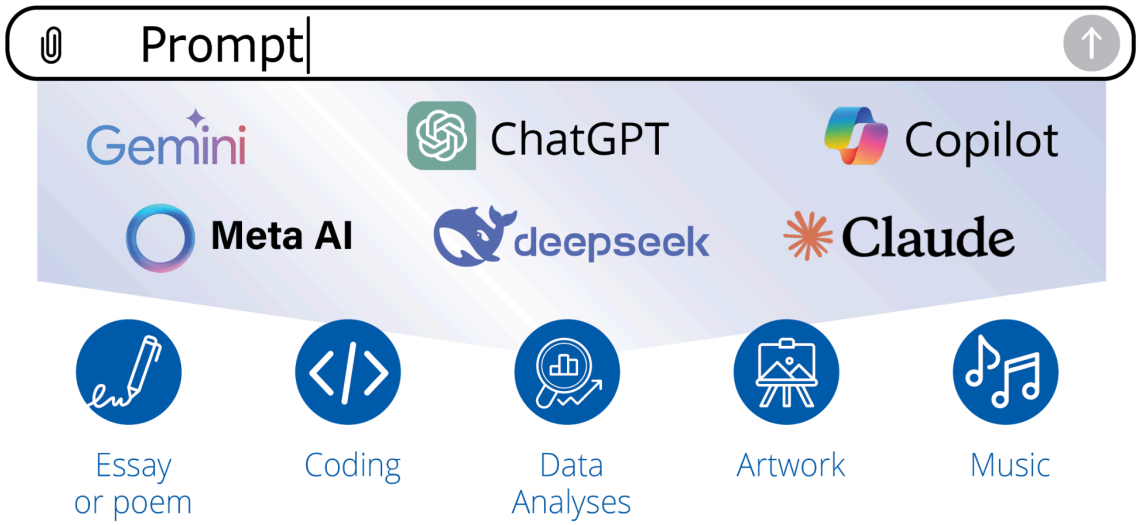

Advancements in generative AI have led to the development of a wide range of powerful tools. Notable examples include ChatGPT by OpenAI, Claude by Anthropic, DeepSeek AI, Microsoft's Copilot, Google's Gemini and NotebookLM, Grok by xAI, and Perplexity. Since

2023, generative AI has advanced from simple question-answering systems to sophisticated reasoning partners. Newer models, such as OpenAI's GPT-5, Anthropic's Claude, and Google's Gemini, demonstrate markedly stronger performance, not just more fluent text, but also deeper reasoning, multimodal understanding, and the ability to solve complex problems step-by-step. GPT-5, for example, has been benchmarked at expert-level performance in areas such as academic writing, coding, and mathematics, while Claude and Gemini bring comparable innovations. This competitive landscape is driving rapid, qualitative improvements in AI capability.

Consequently, AI literacy for academics must be understood not merely as a technical skill but as a core pedagogical competence. The newer systems are not only more powerful in processing and reasoning but also create new opportunities to design learning that is engaging, relevant, and meaningful. AI literacy, therefore, involves more than knowing how to operate tools; it requires the ability to embed them into teaching in ways that deepen student learning and sustain access to disciplinary knowledge. This is fundamental to creating "AI-resistant" assessments, while also preparing students for a future workplace where collaboration with AI assistants, from data analysis to content creation, will be routine. Ultimately, what matters is not just the performance of the systems but how lecturers harness them to support the learning process and foster genuine student learning and understanding.

The following table outlines key AI capabilities relevant to teaching and learning, alongside exemplar tools or platforms. It shows how today's AI spans a range from Socratic tutoring (guided questioning) to agentic research assistants and multimodal content generators.

Category	**What it means**	**Examples of tools**
Guided tutoring	AI asks helpful questions that prompt students to think, rather than just providing answers.	ChatGPT Study Mode; Google Guided Learning; Khan Academy Khanmigo; Claude (Socratic mode)
Research help	AI can read across many documents, check sources, add citations, and summarise findings.	Claude Deep Research; ChatGPT Agent Mode; Perplexity Pro; Google Scholar
Content creation	AI can automatically convert documents into podcasts, videos, or interactive slides.	NotebookLM (audio/video overviews); Sora (OpenAI); Veo (Google)
Custom AI (no coding needed)	You can build a small AI helper for your own course or task — no programming required.	OpenAI GPTs & GPT Store; Google Gems; Claude Projects; Microsoft Copilot Studio
Step-by-step problem solving	AI can show reasoning step by step — useful for mathematics, logic, or technical problems.	Claude Think Mode; ChatGPT o1; DeepSeek Reasoning; Wolfram Alpha

Study with your own material	AI works directly with your documents and gives answers with proper references.	NotebookLM; Claude Projects with uploaded docs; ChatGPT with file uploads; Gemini's Gems
Learning by conversation	You can talk to AI in natural voice, practise languages, and get instant feedback.	ChatGPT Advanced Voice (145+ languages); Gemini Live; Claude Voice
AI in your workspace	AI is built into everyday tools and learning platforms (like clickUP or Microsoft Office).	Microsoft 365 Copilot; clickUP (Blackboard) AI
Personalised practice	AI creates quizzes, practice questions, and small tests tailored to your students.	NotebookLM Quiz Mode; Custom GPTs
Global access	AI supports multiple languages and adapts to cultural contexts, making learning more inclusive.	Most major platforms, Real-time translation, Cultural adaptation
AI agents	Personal "digital assistants" that can complete tasks step by step, e.g., completing a fully online course.	ChatGPT agents

These tools offer opportunities to create more personalised, accessible, and engaging learning experiences, but they must be approached with a clear understanding of their inherent limitations.

This guide is committed to a balanced and critical perspective. Generative AI holds immense potential for productivity and enhancing learning processes. However, it is essential to remain mindful of its limitations and risks. These AI tools do not truly understand meaning; they generate content based on statistical patterns and relationships in their training data. This fundamental characteristic gives rise to significant challenges, including the generation of inaccurate or fabricated information ("hallucinations"), persistent data privacy concerns, and the perpetuation of inherent algorithmic biases learned from vast, uncurated datasets. Furthermore, the ease with which AI can perform cognitive tasks introduces the risk of "cognitive offloading," a phenomenon where excessive reliance on AI can hinder the development of students' own critical thinking and problem-solving skills. The goal, therefore, is to leverage AI to stimulate and support human thinking, not to replace it.

In summary, generative AI offers a powerful new toolkit for educators, from automating routine tasks to enabling interactive, personalised learning experiences. To harness these advantages, lecturers should stay informed and experiment early: try out ChatGPT study mode, explore Gemini's guided learning, or test NotebookLM with your course materials. Meanwhile, remain vigilant about pitfalls: always verify AI outputs and maintain rigorous academic standards. This guide will help you do exactly that, providing practical tips for leveraging generative AI while preserving the core of meaningful education. With thoughtful adoption, lecturers can transform these cutting-edge technologies into catalysts for improved teaching and enhanced student learning.

2 Applications and Strategies

There are numerous ways in which generative AI can be utilised to enhance teaching, learning, assessment, and student support. When used effectively, generative AI can save time, enhance teaching quality, deepen students' understanding, foster critical thinking, and support lecturers' planning and administration. Many of these benefits align with the principles of active learning. If you are new to using generative AI, following some basic prompting hints can help you improve your AI inputs and achieve the desired results. Below are some prompt-crafting tips and examples:

Hint	**Examples of prompts (input)**
Be clear, concise and specific	Explain the concept of blockchain technology in simple terms.
Provide context	What are some effective study techniques for someone preparing to take a biology exam
Ask follow-up questions	Explain more about how they work.
Specify the language output	Respond in UK English
Specify the length of the response	I would like a brief/detailed/300-word response (the maximum is about 3000 words)
Specify the response format	Present your response in bullet points/table/paragraph
Specify the level of output	Response suitable for a second-year university student, or I would like an in-depth analysis.
Specify the tone/style of the output	Write my text in the style of Shakespeare.

For instance, a combined prompt that uses several of the above guidelines might be: "*Write a one-page summary about the supernatural role in Shakespeare's Macbeth and to what extent it motivated Macbeth's actions, in clear and formal UK English, appropriate for a second-year university student.*" Using such detailed prompts helps set the expectations for the AI's output, making the results more relevant and useful for your teaching needs.

2.1 AI as a tool for planning and preparation

Beyond the AI functionalities in the LMS (clickUP-Blackboard), lecturers can utilise other generative AI tools to enhance various aspects of teaching practice – from planning lessons to developing learning materials. Below are some common teaching tasks and examples of how AI prompts can assist with each:

Task	**Examples of prompts (input)**
Create lesson plans	Design a comprehensive lesson plan for a first-year philosophy class focusing on the contrasting views of Kant and Plato on reality. The plan should include an assessment task and a memorandum for students, a challenging group assignment that requires higher-order thinking, and a clear rubric to evaluate the assignment.

Write lesson outcomes	I am a first-year university lecturer, and I need to write clear and measurable outcomes for my upcoming lesson on genetics. Write specific outcomes in the following format: 'By the end of this unit, students will be able to describe the scientific methods and provide examples of their application'.
Develop a worksheet	I am a university lecturer teaching second-year French, and I must create a worksheet for my students to practice their vocabulary. Create a worksheet with some exercises, activities, problems, and a group work activity.
Create PowerPoint slides	I am a third-year history lecturer and must create a PowerPoint presentation for my upcoming lesson on World War II. Create ten engaging and informative slides to enhance my students' learning experience.
Identify the big ideas of a course, unit, or lesson	I am a second-year English lecturer creating a new unit on Shakespeare's Hamlet. Identify the big ideas for this unit.

Using generative AI in this way can help you quickly brainstorm and flesh out teaching content. However, while efficiency and productivity are valuable, they should be balanced with a reflective approach to teaching and learning. A 'slow approach' allows space for students to think critically, engage with complexity, and appreciate the nuances of knowledge in the disciplines. Always remember to review and adapt the AI's suggestions to ensure they are accurate and appropriate for your specific course context. AI can generate a wide range of ideas and draft material, but the lecturer's expertise remains crucial in selecting, curating, and refining these materials for effective teaching.

2.2 Leveraging clickUP (Blackboard) Ultra's built-in AI features

The University's learning management system, clickUP (Blackboard Ultra), has a powerful suite of generative AI tools integrated directly into the platform, collectively known as the *AI Design Assistant*. These features provide lecturers with efficient ways to create, refine, and manage course content and assessments. By using the Blackboard AI Design Assistant, you can save significant time in course development and administration, allowing you to focus more on teaching and student interaction. It is essential to consider core educational principles when integrating AI features in clickUP Ultra. The AI Design Assistant is designed to assist, not replace, the lecturer – you remain in control of all content and can adjust or override AI suggestions as needed.

The Blackboard AI Design Assistant is not a single tool, but a collection of capabilities available at different points in your course creation process:

- **Course design and content creation:** The AI Design Assistant can help you build and structure your course from the ground up.
    - Auto-generate learning modules: If you are starting with an empty course shell, the AI can generate a structure of learning modules based on your course title and description. You can customise the number of modules, their complexity, and even have the AI generate relevant banner images for each module to create a more engaging visual layout.

    - Design document layouts: When creating an Ultra Document, the AI can suggest visually appealing layouts that include headings, images, and even "knowledge check" questions to make the content more interactive for students.
- **Enhancing student engagement and interaction**: You can use the AI Design Assistant to create activities at the appropriate complexity level that encourage student participation and critical thinking.
    - Generate discussion and journal prompts: The AI can create prompts for discussions and journals based on a description you provide or on existing course content.
    - Create AI-powered conversations: A unique feature is the ability to create interactive scenarios where students engage in a conversation with an AI persona.
- **Streamlining assessment:** The AI Design Assistant offers support for creating and managing assessments at the appropriate complexity level, which is particularly useful for both formative and summative tasks.
    - ***Generate test questions:*** clickUP's AI Design Assistant can generate a variety of question types, including multiple-choice, True/False, Fill in the Blank, Matching, and Essay questions. You can specify the number of questions, their complexity, and provide a description to guide the generation process.
    - ***Build question banks from content***: You can upload a document (e.g., your lecture notes or a reading), specify question types, and the AI Design Assistant will generate a bank of questions based on that specific text. This ensures that assessment questions are directly aligned with the material you have provided.
    - ***Create Rubrics***: This is one of the most powerful features. The AI Design Assistant can generate a complete rubric based on your assignment description. You can define the rubric type (e.g., percentage, points), the number of columns (levels of achievement), and the number of rows (criteria). This provides a solid foundation for establishing transparent and consistent evaluation criteria.

In all these cases, the AI-generated outputs in clickUP are meant as a starting point or inspiration. Lecturers should review and adjust AI suggestions to ensure accuracy, appropriateness for the class context, and alignment with the module outcomes.

2.3 AI Tutoring and "Study/Learn" Modes

A major benefit of generative AI in education is the emergence of on-demand tutor-like modes. Both OpenAI and Google now offer modes that behave like personalised, always-available teaching assistants. For example, in ChatGPT's "Study & Learn" mode, any chat can turn into a tutoring session with guiding questions and feedback rather than straightforward Q&A. ChatGPT (in this mode) will prompt a student with questions like "How did you arrive at this answer?" instead of simply stating correct or incorrect, mimicking the approach of a good human tutor. Google's Gemini model offers a similar "Guided Learning" mode that encourages learners with open-ended questions and breaks down solutions into steps. These adaptive modes adjust to the student's level, offering more help when needed and recalling what the student struggled with in previous interactions. The result is a more engaging and adaptive learning experience. Students might snap a photo of a problem and ask the AI tutor for guidance, or ask a question on the ChatGPT mobile app and receive a spoken explanation in their preferred language. These features make learning more immersive and accessible.

That said, it is crucial to remind students (and ourselves) that an AI tutor, no matter how advanced, is a supplement, not a substitute for human educators. The mentorship, depth of expertise, and personal interaction provided by human lecturers remain irreplaceable. Generative AI is best used as a supportive tool to reinforce and personalise learning outside of class.

2.4 Build a small AI tutor from your own materials

As a lecturer, you can create a focused AI assistant for a specific module, topic, or class so that responses are concise, academically appropriate, and grounded in your own materials. Set-up is quick (often within minutes).

2.4.1 NotebookLM (Google)

NotebookLM is a Google tool that allows you to upload your sources (e.g., lecture notes, readings, slides, curriculum documents) and then question, summarise, and generate audio interviews and teaching artefacts directly from those sources, with citations to maintain academic integrity. Use it to:

- Answer questions and summarise your uploaded content, with inline citations to the exact source passages.
- Develop teaching materials in Studio (FAQ, study guide, briefing, timeline, mind map, notes), all grounded in your sources.
- Create overviews, such as Audio Overviews (podcast-style) and Video Overviews, for quick recaps that you can share with students.

Steps (or watch the video):

1. Prepare your content: Gather clean, rights-cleared sources for the specific module/topic (guide, slides, readings, rubrics, past papers, FAQs).
2. Create a notebook & add sources: notebooklm.google.com → Create new notebook → upload Docs/Slides/PDFs/URLs/YouTube etc.
3. Chat & Studio: Ask questions (citations included); generate artefacts in Studio (FAQ, study guide, briefing, timeline, mind map, notes).
4. Overviews: Generate Audio and Video overviews for quick recaps.
5. Share the notebook link with students in clickUP (LMS). Consider read-only sharing to preserve the source set.

Notes:

- All outputs link back to the trusted sources you provided.
- Privacy & data: Google states that content you add to NotebookLM isn't used to train the underlying models. Avoid uploading confidential student data.
- Respect copyright: Upload only materials you are licensed to use or that fall under fair dealing.

2.4.2 Small GPT in ChatGPT (OpenAI)

Custom GPTs enable you to build a module-specific assistant that operates using your own materials and content. You can upload readings, slides, handbooks, and curriculum documents as Knowledge, write clear instructions (e.g., "cite the file and page for every factual claim"), and optionally enable tools such as web browsing, file uploads, and API

Actions. You can keep the GPT private, share it with your class, or (if appropriate) publish it to the GPT Store.

Students can then use it to:

- Answer questions and summarise your uploaded content, grounded in your sources. In your instructions, tell the GPT to quote short excerpts and name the source file (and page/section) each time to support academic integrity.
- Develop teaching artefacts – FAQs, study guides, briefings, timelines, mind-maps, quiz items, and marking rubrics – directly from your Knowledge files. (Add guidance in the instructions about tone, level, and alignment to learning outcomes.)

Steps (or watch the video):

1. Open the GPT editor (Plus/Pro/Team): chatgpt.com → Explore GPTs → Create.
2. Instructions: State role. E.g., "You are a tutor for [MODULE CODE] at the University of Pretoria. Answer only from the uploaded module sources (guide, slides, readings, rubrics, past papers). If the answer isn't in those sources, say you don't know and suggest where to look. Keep answers brief (3–6 sentences) and cite the source file and page/section. Align examples with our learning outcomes: [paste outcomes]. Use UK spelling and a clear, professional academic tone. Avoid speculation, external facts, and personal data."
3. Knowledge: Upload your content (up to 20 files).
4. Capabilities: Switch on only what you need (e.g., Web Search; Code Interpreter for data tasks; Image generation if relevant).
5. Share: choose access (Invite-only / Workspace / Public by link / GPT Store) and keep teaching bots private to your class/workspace. Share the link to the specific GPT with your students in clickUP.

Note: In Data Controls, choose whether your chats are used to improve models; for teaching contexts, many lecturers opt not to allow training.

Upload only rights-cleared material. When building small AI tutors or uploading teaching materials, prefer rights-cleared open textbooks (OER) where possible:

- OpenStax: https://openstax.org (STEM, Economics)
- LibreTexts: https://libretexts.org (STEM, Technical)
- Open Textbook Library: https://open.umn.edu/opentextbooks (Humanities, Social Science)

3 Principles and Guidelines

When using generative AI tools for teaching and learning at universities, it is essential to follow some guiding principles to ensure the technology is used effectively and ethically. Here are some guiding principles to consider:

Guiding principles

3.1 Clarify the purpose

Any use of generative AI should be aligned with clear teaching and learning goals. When defining the purpose of using AI in a given activity or course, it should be thoughtfully integrated (and used with caution) to enhance the achievement of learning outcomes or to provide additional support to students. Ensure that both you and your students understand why a particular AI tool is being used and what educational value it provides.

3.2 Reiterate the ultimate purpose of higher education

Remind students that the overarching goal of higher education is to develop their own thinking skills and to become independent thinkers and knowledge creators. Generative AI can assist in the learning process, but the student's intellectual development, their ability to reason, solve problems, and create new knowledge, remains paramount. We use advanced tools to support that development, not to short-circuit it.

3.3 Provide guidance and communicate rules

Students should be explicitly guided on how to use generative AI effectively and ethically in your course. This includes teaching them how to interpret AI results, how to use such tools to enhance (not replace) their learning, and how to avoid common pitfalls (like blindly trusting AI outputs). Clearly communicate the rules for AI use in your module's study guide and at the start of each assignment. Students need to know whether they are prohibited from using AI, allowed with limitations (and with proper citation), or expected to use AI as part of the task (more on setting these rules below in section 3.9.1). Also, explain how students should acknowledge any use of AI in their work to maintain transparency (see section 3.9.2 on plagiarism). By providing this guidance upfront, you set clear expectations and help students use AI as a learning tool rather than a cheating shortcut.

3.4 Ensure transparency

Be transparent with your students about any use of generative AI in the course. This includes explaining how the AI tools you recommend or use work at a basic level, what data these tools might collect, and how their outputs are generated. If, for instance, you use AI-generated content in your teaching materials, you might mention it as an example ("This summary was generated with an AI tool and then reviewed for accuracy"). Transparency helps demystify AI and also models honesty in acknowledging sources (even if the source is an AI). It also involves discussing the limitations of AI; students should know that AI can make mistakes or produce biased outputs, so they understand why human oversight is necessary.

3.5 Openly discuss ethical implications

Create opportunities to discuss the ethical and societal implications of AI in your field. Encourage students to consider issues such as potential biases in AI outputs, implications for copyright and intellectual property when utilising AI-generated content, data privacy concerns, and the broader impact of AI on the discipline, industry, and natural resources. These discussions not only raise awareness but also reinforce the importance of using AI in accordance with academic integrity policies and ethical norms. If your department or the university has specific guidelines or statements on AI ethics, incorporate those into your discussion or course materials.

3.6 Teach students to cross-check information

Emphasise to students that generative AI should be one resource among many, not their sole source of information. Instruct them to verify AI-generated information against credible

sources. For example, if an AI provides a factual claim or a quote, students should check it in a textbook or an academic article. Ensure that students understand that content generated by generative AI is not considered scholarly work. Even if an AI's training data includes scholarly literature, the output itself has not been peer-reviewed or verified. It may sound confident and authoritative, but it could be incorrect or incomplete. By treating AI output with a degree of scepticism and cross-checking against established sources, students practice critical information literacy. This principle goes hand in hand with maintaining academic rigour and integrity in the age of AI.

3.7 Fostering critical thinking skills

While generative AI can provide helpful information and quick answers, it must not replace students' own critical thinking. Continuously encourage, teach, and require students to critically evaluate any information or content they receive from AI. They should use AI outputs as a springboard for developing their own ideas and perspectives, not as a final authority. Designing assignments that require students to analyse or critique AI-generated content (rather than just submitting it) can help maintain the primacy of critical thinking. Below are some recommendations on how lecturers can utilise generative AI to enhance students' critical thinking abilities:

- Actively teach students what critical thinking entails in your discipline at the appropriate NQF level of the module.
- Requiring students to submit a generative AI version of the assignment, including the prompts, and their reflection/critique on the generative AI version based on the course material provided and their final assignment in track changes.
- Use generative AI to generate varied viewpoints and prompt students to compare and critically evaluate them by:
    - Identifying the assumptions
    - Evaluating the soundness of arguments
    - Evaluating the reliability of sources and evidence presented
    - Identifying common logical fallacies that can weaken an argument
- Use generative AI to generate questions that challenge assumptions.
- Encourage students to develop assessment criteria/rubrics for evaluating the generative AI responses.
- Utilise generative AI to generate multiple versions of an essay or research paper and pinpoint areas for improvement.
- Have students work in groups to compare, discuss, rate, and debate their generative AI prompts and responses, then rank them from best to worst.
- Provide a mind map for each assignment or have students explain their thought processes.
- Prompt students to reflect on what they learned about the topic using generative AI.
- Emphasise authentic problem-solving in the assignments.

3.8 Ensuring data privacy and confidentiality

When using generative AI, it is crucial to protect sensitive information from being incorporated into AI training data. Many AI tools, including ChatGPT, offer settings to prevent the system from learning from or storing your interactions. However, if these settings are not properly configured, the AI algorithm may treat your input as public training data. To comply with institutional data privacy policies (e.g. POPIA), never input sensitive, personal, or

confidential information into AI tools, and use only anonymised data when creating case studies or examples. These precautions help ensure that private information remains secure and prevent unintended exposure through AI systems.

3.9 Upholding academic integrity

Ensuring academic integrity in the age of advanced technologies, such as generative artificial intelligence, requires careful planning and clear communication.

3.9.1 Communicating restrictions to students in the study guide

Setting out and explaining the core rules for using generative AI tools in assignments is crucial. This ensures that students understand the boundaries and comply with course or departmental regulations. When determining the appropriate use of AI in assignments, lecturers should consider factors like learning outcomes and the nature of the task. The use of generative AI should align with both the module outcomes and the specific purpose of the task.

To address concerns around plagiarism and academic misconduct versus acceptable use, it is crucial to guide students on how to use generative AI ethically and constructively. In the study guide, and for each assignment, clearly define whether generative AI tools are prohibited, allowed with limitations and proper citation, or required as part of the learning objectives. Any use of AI that falls outside these defined parameters will be treated as academic misconduct.

Ethical considerations, including academic integrity and the need for transparent disclosure of AI use, must also be taken into account within the framework of established departmental and institutional policies that reflect the disciplinary context. Include the Generative AI declaration form (Appendix B) with every assignment: Students must complete the University's Generative AI declaration form upon submitting each assignment, regardless of whether any AI tools were used. This instruction should be stated in the study guide and assignment briefs so that students are aware that they must formally declare their use of AI or non-use of AI for every submission.

Implementing a colour-coded framework for AI use in UP assignments promotes clarity and safeguards academic integrity in the era of generative AI. The framework categorises AI involvement into three levels and should be tailored to the specific requirements of each discipline, as acceptable AI practices vary across different fields.

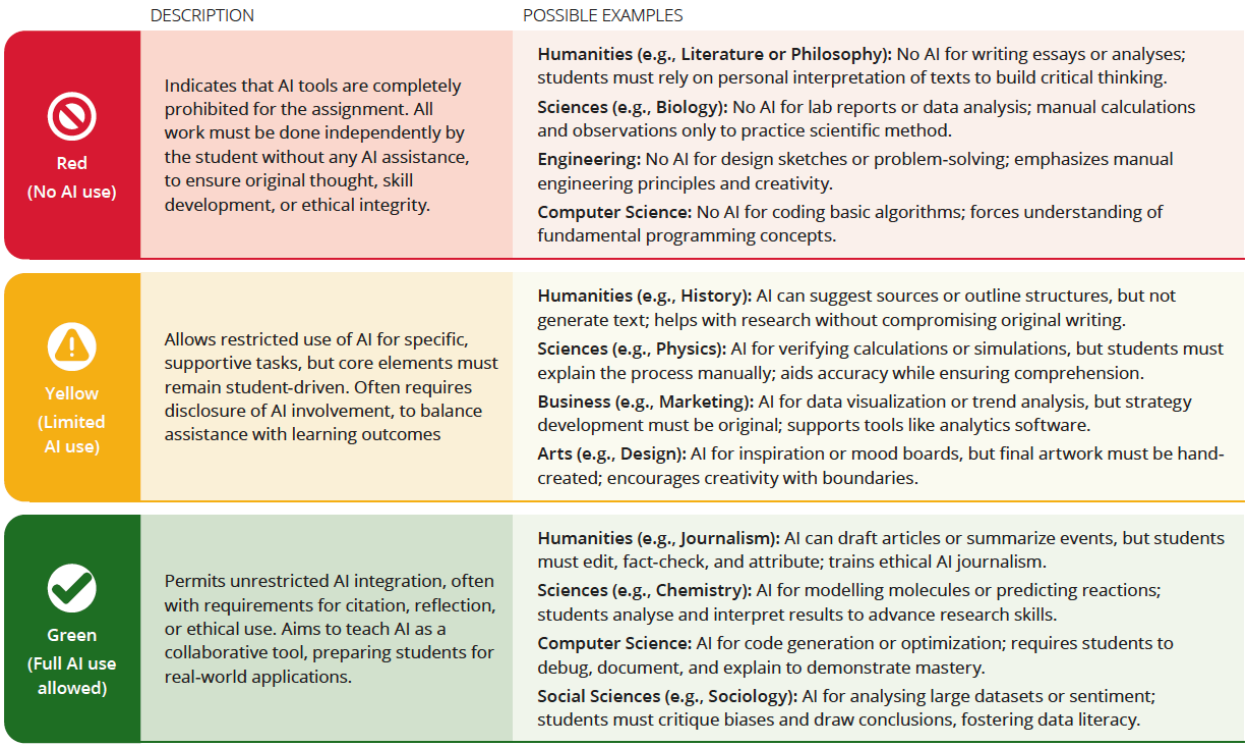

Here are some potential restrictions a lecturer might consider:

- **Prohibition:** The lecturer may ban the use of generative AI entirely in assignments, requiring students to rely solely on their own knowledge and resources.
- **Limited use:** Alternatively, the lecturer might allow restricted use of generative AI for certain purposes. For instance, students could use it to clarify concepts, generate ideas, or seek initial guidance, but its use may be restricted to problem-solving, essay writing, or completing significant portions of an assignment.
- **Collaboration guidelines:** If collaboration is permitted, the lecturer could specify how students may work together and outline any limits on seeking assistance from external sources, including generative AI.
- **Transparency:** The lecturer might require students to clearly disclose their use of generative AI, such as including a statement in their assignments that details if and how the tool was used (use the Generative AI declaration form).
- **Procedural instructions:** Specific instructions may be provided for the use of generative AI, such as setting time limits, specifying allowed prompts or questions, or defining acceptable types of responses.
- **Compulsory submission:** Students might be required to submit a version of their assignment generated by AI, including the prompts used, along with a reflective critique comparing the AI version to their final assignment (using tracked changes).

3.9.2 Addressing plagiarism in a Generative AI context

Generative AI tools, such as ChatGPT, must be acknowledged. It is essential that students understand that presenting AI-generated content as their own work is deceptive and undermines academic integrity because the output does not represent original thought. For instance, the University of Pretoria's plagiarism declaration requires that, *"I declare that this essay, report, project, assignment, dissertation, thesis, etc., is my own original work. Where other people's work has been used (either from a printed source, the internet or any other*

source), this has been properly acknowledged and referenced in accordance with the requirements as stated in the University's plagiarism prevention policy." Accordingly, students must acknowledge the use of generative AI in their work. An appropriate attribution might read as follows: "The creation, enrichment, and editing of this article for enhanced clarity were facilitated by the use of ChatGPT (OpenAI, accessed 14 March 2025)".

3.9.3 Safeguarding the integrity of assessments and assignments

Students who outsource thinking to AI risk impeding their cognitive development. If AI is misused for dishonest purposes, the University faces erosion of academic integrity and the credibility of its qualifications. Our dual responsibility is to equip students to use AI wisely and protect the value of a University of Pretoria degree.

Safeguarding the integrity of assessments and assignments while ensuring the credibility of qualifications is a significant concern for universities, particularly those with large undergraduate classes. A foundational step is to educate all students about the University's academic plagiarism policies. A comprehensive understanding of these policies, including the implications of using AI to assist with their work, can help students maintain high academic integrity.

Evidence note on AI-detection tools: AI-detection tools (e.g., Turnitin's AI indicator, GPTZero) are not reliable for high-stakes decisions. They generate false positives and false negatives; their accuracy degrades as generative models evolve. Simple edits, such as translation or paraphrasing, can also evade detection, raising concerns about fairness and privacy, especially for multilingual writers. Because detection technology lags behind generative AI, these tools should not be used to substantiate allegations of misconduct or determine marks. Instead, universities worldwide advise safeguarding integrity and public credibility through tightly invigilated assessments for high-stakes tasks, supported by process evidence (drafts/version histories) and, where appropriate, brief oral validations. AI detection tools such as Turnitin's AI Detection tool can serve as *'smoke alarms'* rather than proof of misconduct. Like smoke alarms, they signal where closer scrutiny may be warranted, but they do not, on their own, establish guilt or intent.

When designing assessments in the generative AI era, lecturers should carefully consider these key factors:

- **Purpose and format:** Consider the learning outcomes and whether the assessment is written, oral, practical, or a presentation.
- **Stakes of the assessment:** Determine whether the assessment is low-stakes (formative) or high-stakes (summative).
- **Class size and familiarity:** In smaller classes, it is easier to assess the authenticity and originality of each student's work.
- **Testing the task:** Run the assessment through a generative AI tool to gauge how easily it can solve the task, then design the task accordingly.
- **Integration of AI:** Ask if it is possible to incorporate AI into the assessment process by having students use the tool as part of the evaluation, rather than banning it outright.

The following flow diagram will help guide lecturers in safeguarding learning and academic integrity in the age of AI:

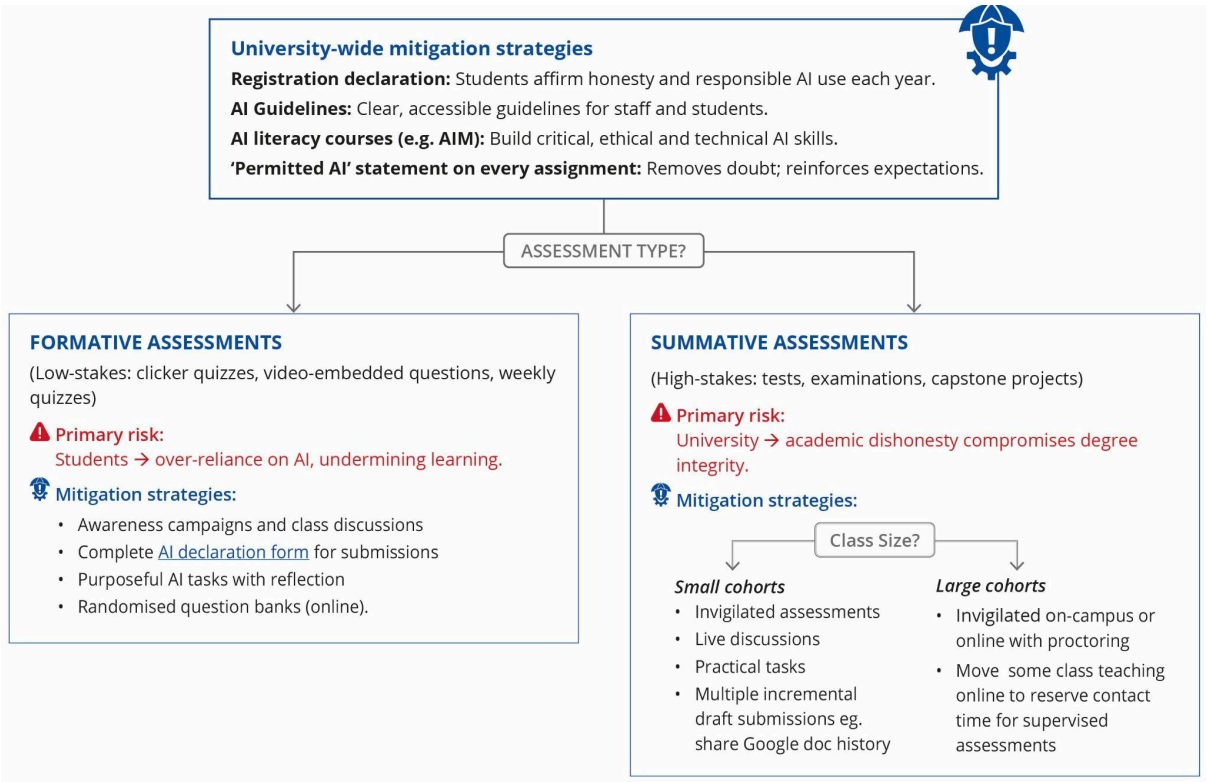

Foundational AI literacy is developed through courses such as the first-year AIM courses, which build critical, ethical, and technical skills. To remove ambiguity, every assignment should carry a brief "Permitted AI" statement. A suitable example is that students may use AI for brainstorming and planning; however, the final submission must be their own work. Any AI use must be declared, including the tools, prompts, and how the output was utilised. Additionally, AI-generated text, code, or images must not be submitted as original work.

Formative assessment is low-stakes and designed to promote learning and growth. The principal risk is over-reliance on AI, which can impede genuine understanding. Mitigation includes brief awareness activities and in-class discussions about the appropriate use, as well as purposeful AI-enabled tasks that require reflection (for example, asking students to compare their own answers with an AI output and explain any revisions). Additionally, technical measures are implemented in online quizzes, such as randomised question banks and option shuffling. As a standard integrity measure, require a completed Generative AI Declaration Form with every submission. Students must tick the relevant box to indicate whether they used generative AI for the task; if so, they should specify the tool(s), prompts, and how the output was incorporated. This uniform requirement ensures transparency and accountability, as students formally acknowledge the rules and any assistance they receive.

Summative assessment is high-stakes and focused on certification. The primary risk is academic dishonesty, which could compromise the integrity of the degree. In smaller cohorts, effective options include invigilated assessments, live discussions, authentic practical tasks, and staged draft submissions with version histories such as Google Docs.

Staged drafts with version histories (e.g., Google Docs):

1. Ask students to draft in Google Docs and share with you (viewer/commenter).
2. Set named milestones (e.g., “Draft 1 – methods”, “Draft 2 – analysis”).
3. Show them Version history (File → Version history → See version history) and require meaningful edits at each stage.
4. Use comments/suggestions; review edit trails and who changed what.
5. Submit final with links to version history and, if needed, Activity dashboard evidence.

Whilst oral examinations have a place, particularly for thesis defences and individual presentations, they are ill-suited to module-level assessment, which requires consistent measurement of all students against common content and outcomes. By contrast, invigilated written (or online proctored) examinations scale efficiently, enabling large cohorts to be assessed simultaneously under identical conditions and marked far more quickly than time-intensive individual orals. Using common papers released only at the time of sitting minimises the risk of leaks or pre-prepared AI-generated responses, and invigilation (in-person or approved digital) deters unauthorised devices, notes, or tools. Scripts are marked against predetermined rubrics, supporting anonymous, standardised, and less bias-prone grading. Crucially, written examinations generate a durable record that can be reviewed, moderated, externally examined, archived, and used in appeals, strengthening accountability and quality assurance. They also test valued academic and professional capabilities, organising ideas, constructing coherent arguments, solving problems systematically, and communicating clearly in writing. Written test ensures fairness by subjecting every student to the same constraints and opportunities.

In summary, proctored and invigilated tests and examinations remain indispensable in theoretical modules because they contribute to operationalising core assessment principles (validity, reliability, integrity, transparency, fairness, inclusivity, practicality, and accountability) by confirming identity, standardising conditions, and ensuring that marks genuinely reflect the intended outcomes. Transparent, auditable procedures in examination venues ensure that students and external stakeholders are assured that the process is fair, scalable, and open to scrutiny, even for very large cohorts.

3.9.4 Module AI-vulnerability self-assessment

All modules that include any unsupervised assessments must complete the **Module AI-vulnerability self-assessment (Appendix A)** at the start of each semester and after material assessment changes. This self-assessment is designed to help you evaluate the extent to which the assessments in your module may be vulnerable to the use of generative AI, and to guide you in identifying risks and planning mitigation strategies.

3.9.5 How to identify probable AI use

Generative AI can be useful for drafting and idea generation, yet it is fallible and sometimes produces incorrect results. To safeguard integrity, use an evidence-based triage when a submission seems atypical. Watch for abrupt shifts in voice or quality, unverifiable or irrelevant sources, missing drafts/version history, implausibly rapid turnaround, or work that ignores task constraints and then check objective evidence (AI declaration, drafts and logs, working notes/code, spot-checked sources). The Turnitin AI detection tool may be consulted as one weak signal only. Although various detection tools exist (e.g. Turnitin AI Detection Tool, GPTZero), many are unreliable, especially when texts are paraphrased or minimally edited. If concerns remain, hold a brief clarification/oral discussion to confirm understanding

and authorship. For high-stakes contexts, either allow AI to be transparent and assess the process or use a supervised/invigilated assessment.

Here are several indicators that may help you spot AI-generated text:

- Em-dash usage: AI often overuses em-dashes—sometimes called the “ChatGPT hyphen”—because models learn from human texts where they’re frequent. That said, the presence of em-dashes alone does not reliably indicate AI origin.
- Predictable openings: AI-generated content often begins with formulaic phrases such as “Have you ever wondered…”, “It is worth noting…”, “In summary,” or “Overall,” reflecting patterns common in marketing language.
- Vague language: Be cautious of non-specific statements, such as “many studies show” or “experts agree,” which lack accompanying names, dates, or statistical evidence.
- Over-simplified solutions: AI-generated writing often glosses over complexities, presenting issues in an overly tidy way without acknowledging real-world challenges or caveats.
- Fabricated or generic citations: Be alert to citations that seem made-up or offer no meaningful detail, or where linking seems generic or superficial.

3.10 Responding to non-adherence to Generative AI restrictions

Even with clear guidance, some students may not follow the stated restrictions on generative AI. These cases should be handled fairly and constructively, upholding accountability while supporting learning. The University’s Assessment Policy (S 5127/22) provides the overarching context, emphasising integrity and honesty in all assessments.

Step 1 (Establish the facts): Compare the assignment’s stated AI restrictions, the student’s completed Generative AI declaration form, and the suspected actual use evident in the submission. Relying exclusively on AI-detection tools to enforce restrictions on generative AI is not a sound practice. Where possible, triangulate with drafts, version history (e.g., document revision logs), supervision notes, and/or a brief oral explanation to confirm authorship and process.

Step 2 (Speak with the student): Conduct a conversation (ideally face-to-face) to gain a deeper understanding of the student’s perspective. Discuss expectations openly with students to reinforce a shared understanding of acceptable AI use. Determine whether the issue arose from a misunderstanding, an oversight, or a deliberate breach, and whether the declaration form was accurate.

Respond proportionately (choose the most fitting pathway):

A. Misunderstanding

- Clarify expectations: Re-explain the relevant restrictions and the ethical duties.
- Educational response: Permit a revise-and-resubmit under the correct parameters; provide exemplars or resources showing permissible AI use.
- Documentation: Record the matter as a learning intervention rather than misconduct, for transparency and future reference.

B. Oversight

- Acknowledge and remind: If the breach appears inadvertent, restate expectations for academic honesty.
- Restore learning: Consider partial marks linked to corrective work (e.g., a reflective piece or a revised submission demonstrating correct practice).
- Monitor: For subsequent tasks, you may require staged drafts or guided feedback to prevent repeat issues.

C. Intentional disregard
- Treat as misconduct: Where there is evidence of deliberate violation, follow the University's academic dishonesty procedures.
- Proportionate consequences: Calibrate outcomes to severity and year level, e.g., required resubmission, grade penalties, and escalate according to the disciplinary policy when warranted.

D. Blatant plagiarism (substantial or entire AI-generated work presented as the student's own)
- Apply the Plagiarism Policy: Presenting AI-generated text as one's own constitutes academic dishonesty and breaches academic integrity. Refer to Plagiarism Policy S 5105/19 and Annexure C for reporting procedures.
- Document and escalate: Complete the required misconduct reports and, where appropriate, refer to the relevant disciplinary body.

4 Building capacity for AI-enhanced teaching and assessment

The Department for Education Innovation (EI) runs a structured professional development programme that includes a dedicated short course, 'AI in Teaching, Learning and Assessment', which helps lecturers use generative AI and clickUP's AI features responsibly, adapt assessments, and protect academic integrity, alongside broader e-learning and assessment offerings.[OBJ] The University of Pretoria also hosts a cross-faculty AI-related Community of Practice. These networks increasingly foreground AI pedagogy, ethics and assessment, aligning with the institution-wide focus on responsible AI and assessment innovation.

Additionally, every UP lecturer has free access to LinkedIn Learning, which offers a wide range of AI-related courses. These cover topics such as responsible use of generative AI, AI in education, data ethics, and digital transformation. This resource allows lecturers to build their AI literacy at their own pace and complement the structured EI programmes with global best practices.

EI also equips lecturers to utilise Blackboard's AI capabilities responsibly, saving time and enhancing assessment quality. Key tools include:
- AI Design Assistant generates learning modules, test questions and question banks, and drafts rubrics, assignments, discussions and journals, with context-aware prompts to ground outputs in course materials.
- AI Conversations enables structured role-play and formative practice activities directly in Ultra courses.

These initiatives align with UP's commitment to the ethical and responsible adoption of AI in teaching and learning.

5 Other interesting links

- Academy of Science of South Africa (ASSAf) and SciELO Guidelines for the Use of Artificial Intelligence (AI) Tools and Resources in Research Communication, https://criticalai.org/2023/01/17/critical-ai-adapting-college-writing-for-the-age-of-large-language-models-such-as-chatgpt-some-next-steps-for-educators/
- Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025, June 10). Your Brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task (arXiv preprint arXiv:2506.08872). https://doi.org/10.48550/arXiv.2506.08872
- Oakley, B., Johnston, M., Chen, K.-Z., Jung, E., & Sejnowski, T. (2025). "The Memory Paradox: Why Our Brains Need Knowledge in an Age of AI." In The Future of Artificial Intelligence: Economics, Society, Risks and Global Policy (Springer Nature, forthcoming). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5250447

6 Appendix A: Module AI-vulnerability self-assessment (integrity)

This form is a self-assessment tool designed to help lecturers audit a module's vulnerability to generative AI. It guides you through a systematic review of all assessments to identify potential risks to academic integrity (download form).

Module Code:	
Number of students?	

Step 1: Map all assessments that contribute to the final module mark

List every formative or summative assessment that carries a weighting towards the final mark.

Assessment Title	**Weight (%)**	**Assessment Conditions (Invigilated, Proctored, or Unsupervised)**	**Plausible AI mark (%)**	**Weighted AI contribution (= Weighting × AI% ÷ 100)**
e.g., Literature Review Essay	15%	Unsupervised (at home)	80	12
e.g., Semester Test	25%	Invigilated (in-person)	0	0
e..g, Exam	50%	Invigilated (in-person)	0	0

* If the assessment is conducted as an invigilated, synchronous examination, the AI score is recorded as 0. For all other assessments, complete the task using ChatGPT or Gemini and then grade the AI-generated submission.

Step 2: Final assessment (examination)

Does this module have a final, comprehensive assessment (such as an exam) that covers all the content?	☐ Yes / ☐ No Exam weight: ____% & Coursework weight: ____%
Final assessment (exam)	Does a subminimum apply to the final exam (e.g., a student must achieve at least 40% in the exam to pass the module)? ☐ Yes / ☐ No If Yes, what is the required minimum percentage? _____%
Is the final assessment (exam) supervised (in-person, invigilated, or online proctored)?	☐ Yes (in-person invigilated or online proctored) ☐ No
Risk of students sharing questions	☐ Synchronous (all students at the same time) ☐ Asynchronous (students have a window of time)

Step 3: Risk calculation

Module AI Exposure = Sum of the weighted AI contributions (Step 1) = _____% of the final mark.

If there is an exam subminimum, the student must also meet it: required exam points = _____ *(Exam weighting × subminimum ÷ 100).*

Minimum original student contribution to reach the pass mark (50%): _____ %

Step 4: Mitigation and redesign strategy

For high module AI exposure, record the changes you will make. Strategies could include shifting key assessments to an invigilated format or incorporating in-class presentations.

Assessment	**Proposed Changes / Mitigation Strategy**

Step 5: Review

This self-assessment should be reviewed at the departmental level to ensure consistency and share best practices.

Lecturer's Name: _____

HoD / Programme coordinator: _____

Date of Review: _____

7 Appendix B: Generative AI declaration form for assignment submissions

Q1. Student Information

Name and Surname	
Student Number	
Module Code	
Assignment Title	

Q2. AI Declaration [tick one box]

1. I **used** Generative Artificial Intelligence in the current assignment.	
2. I did **not** use any Gen AI for the current assignment.	

Ensure you are permitted to use generative AI by reading the instructions for your assignment or assessment, or by consulting your study guide. Failure to follow the instructions regarding using generative AI for your assignment constitutes academic dishonesty. Examples of generative AI include ChatGPT, Google's Gemini, Microsoft's Copilot, Claude, and Meta's Llama 3, which is integrated into WhatsApp and Facebook. If you selected option 1 in Q2, please continue to complete the rest of the form:

Q3. Prompts used

Please paste all the prompts you created for the assignment and indicate which aspect(s) of the assignment the specific prompt was used for.

AI Tool	Prompt	Aspect of Assignment

Q4. Type of usage

Briefly describe the features for which you used the Gen AI. You may say “Not Applicable” where needed:

Feature used	Description
Brainstorming and idea generation	
Language editing suggestions	
Feedback and revision suggestions	
Explaining complex concepts	
Writing coach	
Other (please specify):	

Q5. Ethical use

Write a brief paragraph explaining how you ensured the usage of Gen AI was aligned with the ethical and responsibility requirements of the University of Pretoria (link). Consider examples such as repurposing and reintegrating ideas generated by Gen AI with your own thoughts, integrating Gen AI ideas with other literature, critically evaluating Gen AI outputs, maintaining transparency about Gen AI usage, enhancing your learning and ensuring comprehension despite using Gen AI, and personal

development through using Gen AI as an assistant. If you directly used text or data generated by Gen AI, ensure it was cited appropriately.

Ethical use statement:

Q6. Why did you use Gen AI?

Signature:

Student Signature	
Date	

Note: The Department of Education Innovation would like to thank Prof. Celeste Combrinck and Ms Nelé Loubser for their innovative ideas, which have been instrumental in creating our new form.