

Designing the Student Learning Journey: A Practical Approach to Integrating Generative AI within Higher Education

Michael Grove

Deputy Pro-Vice Chancellor, Education Policy & Academic Standards

Professor of Mathematics and Mathematics Education

The University of Birmingham, UK

m.i.grove@bham.ac.uk, ORCID: 0000-0002-3150-125X

Published Online: 7 July 2025

Abstract

Generative AI technologies are reshaping higher education, transforming how students access knowledge, engage with learning, and complete assignments. While institutional responses have largely focused on academic integrity and assessment security, this paper argues for a proactive, programme-level approach that embeds generative AI thoughtfully and ethically across the student learning journey. Drawing on examples from the mathematical sciences, it presents a practical framework to support curriculum teams in aligning AI use with programme outcomes, disciplinary values, and assessment design. Key recommendations include designing progression from foundational to advanced AI-supported tasks; fostering coherent, programme-wide expectations for ethical and transparent AI use; and developing students' critical AI literacy as a core graduate attribute. The paper also highlights the importance of equitable access to tools, respecting disciplinary contexts, and rethinking assessment formats to promote higher-order thinking. A programme-level checklist is provided to guide planning and implementation. By integrating generative AI with intentionality, institutions can move beyond reactive policies towards learning environments that prepare students for a future in which human and AI capabilities will increasingly work in partnership.

1. Introduction

To date, much of the discourse surrounding generative artificial intelligence (generative AI) in higher education has centred on its implications for assessment: how to detect it, how to mitigate risks, and how to ensure academic integrity. While these are important concerns, a singular focus on assessment risks overlooking the wide, and arguably more transformative, potential of generative AI technologies to support and enhance student learning. These tools, now widely accessible, offer students new ways to explore ideas, test understanding, and personalise their learning experience. Importantly, the ability to use generative AI tools effectively, ethically, and critically will become an increasingly vital graduate attribute.

It would be a mistake to begin by assuming that all students will use generative AI inappropriately or with the intention of gaining unfair advantage. Many are now entering higher education having already experimented with such tools in school or college (Freeman, 2025). They will continue to use these tools to make sense of complex material, generate examples, or check their understanding, especially when they are unsure where else to turn, when support is not available at convenient times, or when it doesn't align with their preferred learning approach. The appeal is clear.

Designing the Student Learning Journey: A Practical Approach to Integrating Generative AI within Higher Education

Generative AI tools offer what many students perceive as effortless content creation, immediate answers to difficult questions, and personalised feedback on demand. They can generate multiple versions of a written task, suggest how to improve grammar and structure, or provide near-instantaneous solutions to mathematical problems. For students facing uncertainty, time pressure, or confidence barriers, generative AI promises speed, clarity, and convenience.

Generative AI tools also offer a form of personalised learning, tailored prompts, interactive dialogue, practice problems, and 24/7 availability, making them feel more accessible than many traditional forms of academic support. Students use them to summarise lecture content, develop research questions, translate texts, or refine presentation materials. For some, they are a creative partner, for others, a non-judgemental tutor. It is this broad appeal, and their growing role in everyday student study patterns, that makes it essential to engage with generative AI thoughtfully and proactively within programme design.

As educators, we therefore have a responsibility not to ignore or restrict these tools entirely, but to help students learn how to use them well. This includes ensuring that all students:

- Understand the significance of generative AI for their studies and their future careers.
- Recognise appropriate and inappropriate uses of generative AI in the context of learning and assessment.
- Appreciate both the strengths and limitations of generative AI tools as part of their educational experience.
- Develop the skills to ethically and critically use generative AI to support learning and appraise their own progress and understanding.

Student support in this area must be scaffolded. All students should be introduced to these tools through a clear and coherent programme of regular guidance and practical activity. But beyond that, they need opportunities to use generative AI within their discipline, with clear expectations and feedback on the appropriateness and success of their use. When used effectively, these tools can also benefit educators by enhancing their teaching practices, whether through generating practice questions, drafting explanations, developing feedback, or supporting differentiated instruction.

Yet to realise these benefits, the use of generative AI must be designed, and designed with intent. This means asking very fundamental questions about teaching and student learning:

- What do we want students to learn, and why?
- How can we design learning experiences that promote deep, connected, and sustained understanding?
- How might generative AI support the development of disciplinary thinking, academic skills, and graduate attributes?

These are not new questions, but the presence of generative AI in the current learning landscape changes how we must approach them. Consideration of its use can no longer sit outside of learning design, it must now be embedded within it, just as we routinely consider how assessment aligns with learning outcomes. This does not mean every module must require or support the use of generative AI tools, nor that their use is always appropriate. In some cases, allowing students to rely on generative AI may risk undermining the very skills and attributes we are seeking to develop, such as

constructing arguments, performing symbolic manipulation, or engaging in sustained problem solving. Decisions by educators about when to use, or not use, generative AI must be intentional and transparent.

While some institutions have looked to mitigate misuse through detection tools, these are often unreliable (Weber-Wulff et al., 2023) and risk fostering a climate of mistrust. A more productive approach lies in intentional curriculum design, clear communication, and proactive support for ethical use of these tools by staff and students. This also underscores the importance of developing staff confidence and capability. Designing for AI-enhanced learning is not solely a technical matter, it also requires academic judgement, disciplinary literacy, and pedagogic intent.

What matters most is coherence. At programme level, students should experience a consistent and well-communicated approach. Where generative AI is encouraged, the rationale should be clear. Where it is limited or restricted, the pedagogical reasons should be explained. Inconsistent messaging, or indeed a lack of communication, only leads to confusion, inequity, or misuse.

This paper explores how programme teams can take a structured and practical approach to integrating generative AI into learning and teaching design. It outlines key design considerations, offers examples of effective practice within the mathematical sciences, and sets out a framework to support students in engaging with these tools confidently and responsibly. The focus is not on replacing teaching or outsourcing thinking, but on how generative AI tools might help us design better learning, preparing students not only for success in higher education, but for the demands of an ever-evolving world of work.

2. Considerations for Design

The integration of generative AI within learning and teaching should not begin with tools or technology, but with the principles of learning design. Specifically, it should be grounded in programme-level learning outcomes (PLOs) and the overarching aims of the curriculum. The following sections outline five key considerations for effective and sustainable integration of generative AI.

Programme-level approaches to integrating generative AI must not only be pedagogically grounded but also aligned with institutional frameworks and guidance. Institutional frameworks provide a shared foundation for practice, ensuring that individual programmes support consistent messaging on academic integrity, ethical use, digital skills, and student support. Aligning with these policies helps ensure students encounter a coherent experience, where the expectations around AI use are both transparent and justifiable across modules and departments. This is particularly important for joint honours or interdisciplinary students, who may otherwise face conflicting guidance across subjects, undermining both equity and clarity.

2.1. Purpose and Progression

Effective design begins by considering the purpose of integrating generative AI and how its use will support student progression across the programme. Programme-level learning outcomes should

guide what is taught at the module level, the skills students are expected to develop, and how those skills are assessed. While generative AI may not feature explicitly within these outcomes, many programmes already include references to digital literacy, independent learning, critical thinking, or effective communication. These provide natural points of alignment.

More broadly, generative AI can support the development of higher-level academic skills such as synthesising ideas, identifying relevant knowledge, evaluating information, and applying concepts to unfamiliar contexts. These align closely with the upper levels of Bloom's revised taxonomy: analysing, evaluating, and creating (Anderson & Krathwohl, 2001). A central goal of higher education is to help students move beyond content reproduction and towards critical engagement and original thought. Used well, generative AI can support this progression by acting as a scaffold for inquiry, reflection, and experimentation.

However, the value of generative AI is not limited to these higher levels. In the earlier stages of a programme, students may use AI tools to support foundational cognitive processes, particularly remembering and understanding. For example, they might generate summaries of lecture content or readings, ask for simplified explanations of unfamiliar concepts, build personalised glossaries or revision cards, or translate technical terms into everyday language to check their understanding. These uses can be especially valuable for students who are new to a subject, returning to study, or lacking confidence in academic language or disciplinary conventions. By supporting the consolidation of foundational knowledge, generative AI can help students begin from a more equitable starting point and build their confidence to engage with more complex ideas. As students progress, they may also begin to use generative AI at the applying level, for instance, by creating practice problems or worked examples, exploring variations on standard methods, or testing their ability to adapt a known process to new conditions. These uses allow for greater personalisation and encourage active engagement with content, particularly when students are encouraged to evaluate the relevance and accuracy of what AI tools generate.

In this way, generative AI tools can be embedded across all stages of Bloom's taxonomy (Table 1), supporting student learning in different ways depending on their current level of understanding, the learning outcomes being targeted, and the nature of the discipline. What matters is that the use of generative AI is intentional and clearly aligned with the purpose of the learning activity. By final year, the emphasis should shift towards using generative AI critically and selectively, as one tool amongst perhaps many, to support independent research, synthesis of complex ideas, and the development of original outputs. This progression, from supported use for understanding, to critical use for knowledge creation, should be reflected in programme-level planning and curriculum mapping.

2.2. Designing the Learning Environment

The increasing presence of generative AI in students' academic routines brings with it not only new tools, but also new behaviours. As students learn to engage with AI to ask questions, summarise material, or test understanding, there is a real risk that learning becomes more solitary, transactional, or disconnected. Left unchecked, this shift could undermine core features of a strong university experience such as peer collaboration, dialogue, feedback, and community. Designing effectively for generative AI therefore also means designing around it: identifying what matters in a

rich, supportive, and developmental learning environment, and ensuring those features are preserved and prioritised through curriculum structures and learning activities.

Bloom's Level	Examples of Generative AI Use	Purpose in Learning Progression	Example Assignment Brief
Remembering	Generate glossaries, flashcards, or summary notes from lecture content; retrieve definitions or formulas.	Support foundational knowledge and initial engagement with subject content.	<i>Use a generative AI tool to produce definitions for 10 core terms in real analysis (for example, limit point, bounded sequence, convergence, supremum, uniform continuity). Cross-check each with lecture notes and textbooks, and annotate where clarifications or corrections are needed.</i>
Understanding	Ask for simplified explanations of difficult concepts; translate terminology into everyday language; paraphrase key ideas.	Build conceptual understanding and confidence in interpreting core ideas.	<i>Choose two mathematical concepts introduced this week and use generative AI to produce simplified explanations for each. Evaluate the explanations and write a short reflection on how your understanding developed.</i>
Applying	Generate practice questions or step-by-step examples; explore different applications of known techniques.	Enable practice and reinforcement of taught skills through self-directed learning.	<i>Ask a generative AI tool to create three practice problems on integration by substitution. Solve each problem, annotate your working, and evaluate the accuracy of the generated examples.</i>
Analysing	Compare alternative solutions generated by AI; identify flaws or omissions in AI-produced reasoning.	Encourage deeper engagement with content and methods; develop critical thinking.	<i>Use generative AI to solve a first-order differential equation using two different methods. Compare the outputs and analyse which is more complete, rigorous, or appropriate. Identify any mathematical inaccuracies or shortcuts.</i>
Evaluating	Critique AI-generated arguments, explanations, or code; assess reliability and accuracy; choose the most appropriate output.	Promote judgment, reflection, and academic independence in evaluating outputs.	<i>Submit an AI-generated proof of a standard result (for example, the sum of an arithmetic series). Annotate it to highlight correct reasoning, questionable logic, or missing justifications. Suggest improvements and justify your changes.</i>
Creating	Use AI to brainstorm ideas, draft outlines, or develop project scaffolds; refine outputs through iterative prompting.	Support originality, synthesis, and extended inquiry at advanced levels of study.	<i>Use a generative AI tool to help draft a mathematical modelling problem relevant to your discipline. Refine the prompt to include constraints, assumptions, and possible solution strategies. Submit a project outline and a reflective commentary on your use of AI in the design process.</i>

Table 1: Illustrative tasks for AI-integrated mathematics teaching and assessment. Bloom's taxonomy levels are aligned with examples of generative AI use in mathematics, highlighting how AI can support progression from foundational understanding to advanced thinking. Assignment briefs illustrate practical ways to integrate these approaches into teaching and assessment design.

At a programme level, this involves a shift in emphasis. It is not sufficient to focus solely on where and how AI tools are used; it is also necessary to ask: what features of the learning experience do we not want generative AI to replace, replicate, or diminish?

Some essential features of a modern learning environment might include:

- **Peer interaction and collaboration:** Design group tasks, problem-solving activities, and peer review processes that foster co-construction of knowledge.
- **Personalised feedback and dialogue:** Prioritise small group teaching, formative feedback opportunities, and open-ended tutorial discussions.
- **Development of academic identity and voice:** Create space for students to explain decisions, reflect on learning, and take intellectual ownership of their experience.
- **Challenge, uncertainty, and 'messy' thinking:** Encourage open-ended inquiry, problem formulation, and iterative drafts, not just polished and final outputs.

But good learning design also requires coherence and clarity. The integration of generative AI within a programme cannot be left solely to individual modules or their associated leads. While local flexibility is essential, students should experience a consistent and intentional approach across the curriculum. Without this, they may encounter contradictory guidance, unclear expectations, or unintentional inequities in learning opportunities.

Programme-level coherence does not require uniformity. It calls for transparency and shared intent. Programme teams should work together to map where generative AI is used across the curriculum, identify which skills or learning outcomes its use is designed to support, and agree on consistent language and expectations for appropriate use. These shared principles should be clearly communicated to students through handbooks, module guides, and digital platforms, thereby ensuring expectations are understood and reinforced across contexts. For example, one module might explicitly allow students to use generative AI to explore problem structures or generate graphs, while another prohibits its use in take-home assessments to protect independent reasoning. These differences are pedagogically valid, but only if students understand *why* the approaches differ and how they relate to the learning outcomes. Designing for generative AI is not simply a question of access or policy, it is a question of educational design. It is about creating a learning environment in which students use AI to enhance their experience, not escape from it. By embedding these values and structures into programme-level thinking, institutions can ensure that generative AI contributes positively to a vibrant, relational, and coherent student learning experience.

2.3. Ethical Use and Acceptable Behaviour

With the widespread availability of generative AI tools, clear expectations around acceptable and ethical use are essential. These expectations must be communicated at programme level, not left to the discretion of individual module leads. Without a shared understanding of when and how AI use is appropriate, students are likely to encounter inconsistent messages, leading to confusion, anxiety, or unintentional breaches of academic integrity. Programme teams should agree on a common framework for communicating acceptable use, with flexibility for disciplinary nuance. This framework should be introduced to students early, ideally during yearly inductions and core tutorial

sessions, and reinforced through programme handbooks, module virtual learning environment (VLE) pages, and assignment briefs. Tutorials or in-module sessions can also be used to support students in using generative AI to understand marking criteria, interpret assessment briefs, or plan their approach, thereby making the AI use itself a designed part of early-stage preparation.

Each piece of formative or summative assessment should include an explicit statement about the permitted level of generative AI use, using a shared classification system. This enables consistency across modules and clarity for students.

2.3.1. Academic Integrity Frameworks

Some programmes may prefer a three-level ‘traffic light’ model for its clarity and ease of communication, such as that shown within Table 2. While this model is useful for setting broad expectations, it may not offer sufficient detail for complex tasks, particularly where partial use (for example, grammar correction vs. content generation) must be clearly distinguished. In these cases, a five-level model offers more nuance and can help students better understand how to use AI responsibly in both preparation and submission; such an example is shown within Table 3.




Level	Description	Permitted AI Use	Example
 Prohibited	No use of AI tools allowed at any stage.	None	Handwritten closed-book exam.
 Permitted with Limits	AI can support preparation (for example, idea generation, brief analysis) but not any submitted work.	Use to explore topics or understand criteria.	Use AI to help unpack a problem brief but write the solution independently.
 Encouraged	AI is allowed in both preparation and submission, as part of the learning process.	Integrated use expected, with critical engagement.	Draft a reflective blog post with AI assistance, noting prompts and revisions.

Table 2: Three-level model for generative AI use within assessment. A simplified framework for categorising acceptable, cautious, and prohibited uses of generative AI in student assessment, designed to promote clarity and consistency.

While the three-level model may be appropriate in some contexts, the five-level scale provides greater pedagogical transparency and practical flexibility, especially where assessment types vary, or students are required to reflect critically on their use of AI tools. What matters is that students understand *why* the boundaries exist and that expectations are both fair and transparent. For example, if students are permitted to use AI to understand the marking rubric or assessment structure, that should be made explicit. Equally, if students are expected to demonstrate unaided reasoning or construction, that too must be made clear.

Importantly, if the use of generative AI is prohibited, it is not sufficient to simply state this in guidance and place the responsibility for its non-use entirely on students. The assessment must be designed in such a way that AI use is not practically possible. This may involve using proctored examinations, timed in-person assessments, or Vivas that require explanation and justification. Prohibition, like permission, is a *design decision*, and it carries a responsibility to ensure that conditions support the intended learning and outcomes. Educators should also model transparency

by being open about their own use of AI in teaching and feedback and by encouraging dialogue around its use. Just as we scaffold the development of academic writing, we must now scaffold students' capacity to engage with generative AI responsibly.

Level	Description	Permitted AI Use	Example of Generative AI Use
1. Prohibited	Generative AI must not be used at any stage, including preparation.	None	None allowed: In-person exam with no internet access.
2. Preparation Only	Students may use AI in preparing for an assessment but not in any submitted work.	Use to clarify the brief, understand criteria, or explore task structure.	Use AI to summarise the brief and generate initial ideas, but write the essay independently.
3. Basic Skills Support	AI may assist with surface-level editing or rephrasing, but not substantive content.	Grammar checking, spelling, formatting.	Use Grammarly or rewording tools to improve clarity.
4. Research and Exploration	AI may support summarising content, identifying key ideas, or proposing structures, but students must produce original work.	Summarising sources, suggesting proof structures.	Use AI to generate a solution outline for a mathematics problem, then develop your own formal proof.
5. Collaborative Partner	Generative AI is used as a co-creator or content partner throughout the task. Students are expected to engage critically, reflect on their use, and evaluate the quality and appropriateness of AI contributions.	Co-produced content, iterative drafting, experimentation and critique with full documentation.	Use generative AI to support the development of a mathematical model. Submit the AI-supported work along with a reflective commentary evaluating its accuracy, limitations, and your decisions during the process.

Table 3: Five-level model for generative AI use within assessment. An extended version of the traffic light model that offers greater nuance in defining levels of permitted generative AI use, from preparatory support to full collaboration.

In the context of a tiered permissions model, institutions should clearly distinguish between undeclared but permitted use of AI and inappropriate use of AI in restricted assessments. Where students fail to reference AI use in a permitted task (for example a 'green' assessment), this would be better addressed through grading criteria and feedback, rather than formal academic integrity processes. Expectations around AI use, such as the inclusion of prompts, outputs, or a reflective commentary, should be clearly stated in the assignment brief and rubric, allowing educators to respond transparently and proportionately.

While tiered frameworks can provide valuable clarity and promote consistency, they are not a comprehensive solution. In practice, students may interpret expectations differently or misjudge the boundaries between levels, and simplified models such as traffic-light systems risk oversimplifying the nuanced realities of responsible AI use. These frameworks should therefore be regarded as tools to support dialogue and reflection, rather than as definitive mechanisms for ensuring compliance.

Programmes should also consider student development and confidence when applying these expectations, recognising that Year 1 students may require greater scaffolding and more formative feedback, while final-year students are expected to demonstrate mature, accurate, and transparent use.

Formal academic integrity investigations should be reserved for cases where students attempt to present AI-generated work as their own in contexts where its use is explicitly prohibited, and where there is evidence of deliberate intent to deceive. This pragmatic distinction helps ensure that students are held accountable for integrity breaches, without penalising poor documentation or unintentional misjudgements in otherwise open tasks.

2.3.2. Citing and Acknowledging Generative AI Use

Where students are permitted to use generative AI in assessments, they should be required to clearly acknowledge this use. Transparent attribution supports academic integrity, helps tutors understand how students have engaged with AI, and promotes reflective practice. Citation practices should be agreed and communicated at programme level so that students receive a consistent message across modules. Options include:

- A brief declaration at the end of a submission (for example, “*I used ChatGPT to generate a draft of the introduction, which I then revised*”).
- A dedicated section or appendix for more substantial use (for example, reflecting on how AI supported model development or argument structure).
- Including prompts or AI outputs (for example, screenshots) where appropriate.

Some institutions recommend referencing AI tools according to standard academic styles (for example, APA or Harvard), but consistency is more important than formality. What matters is that students understand when and how to declare use, and that staff are equipped to fairly assess the appropriateness of that use (if any).

2.3.3. Assured and Exploratory Credits: A Programme-Level Approach

As generative AI becomes increasingly embedded within the learning environment, ensuring that assessments can genuinely demonstrate students’ unaided capabilities is both challenging and necessary. While prohibiting AI use outright may be impractical or undesirable across an entire programme, there remains value in maintaining a secure foundation of independent academic performance. To address this, programmes may adopt a balanced approach, distinguishing between *Assured Credits* and *Exploratory Credits*.

Assured Credits refer to a defined portion of a programme, recommended as a minimum of one-third of total credits, where assessments are deliberately designed to ensure students demonstrate knowledge and skills without the use of generative AI tools. For a 120-credit per year programme, this would equate to 40 credits annually. These assessments provide assurance that part of the degree has been completed independently, supporting progression and award decisions. This approach shifts the responsibility for securing unaided assessment from individual students to programme design. By embedding *Assured Credits* within the programme structure, teams can balance innovation and integrity, enabling the remaining credits, so-called *Exploratory Credits*, to

support more open, creative, and AI-enabled assessment approaches. This distinction creates space for students to develop critical, ethical, and effective use of generative AI, without compromising the validity of academic achievement.

Assured Credits do not prescribe specific assessment types but define a minimum secured volume of assessment. Departments can determine the most appropriate methods, which might include proctored in-person examinations, supervised practicals, interactive oral assessments, or other formats suited to the discipline. The emphasis is on ensuring that these credits are assessed under conditions that prevent unauthorised AI use.

By incorporating both Assured and Exploratory Credits at programme level, institutions can foster a transparent and coherent approach to generative AI. Students benefit from clarity, equity, and a structured environment in which they can both demonstrate independent academic performance and develop the confidence to use AI tools responsibly and reflectively as part of their broader learning journey.

2.4. Respecting the Discipline: AI Use in Context

The appropriate use of generative AI is not universal, it must be shaped by the norms, values, and methods of each discipline. What is considered useful or ethical in one subject may be inappropriate or even counterproductive in another. For this reason, decisions about how and when generative AI is integrated into learning and assessment must be grounded in disciplinary thinking and pedagogical purpose.

In the mathematical sciences, for example, the focus is often on formal reasoning, symbolic manipulation, conceptual clarity, and rigour. While generative AI tools can produce worked examples, draft solutions, or simplified explanations, their outputs frequently lack the logical transparency, precision, or notation required in formal mathematics. They may present incorrect arguments confidently, omit crucial reasoning steps, or misrepresent the structure of a proof. As such, they cannot replicate the process of mathematical thinking, abstraction, or justification that underpins deep understanding. Because of these limitations, generative AI is often more appropriate in exploratory, formative, or diagnostic tasks, rather than in summative assessments where students must demonstrate reasoning for themselves. For instance, students might use AI to compare solution strategies and identify inconsistencies, test informal AI explanations against lecture-based formal proofs, explore prompts to refine definitions, or critique mathematical writing generated by AI for errors or ambiguity.

By contrast, in disciplines such as design, education, or the humanities, generative AI may more naturally support brainstorming, planning, or synthesis, provided students critically evaluate outputs and integrate them into their own intellectual framework. Here, fluency with AI tools might form part of disciplinary development in a more explicit, and perhaps convenient, way.

To support students effectively, programme teams should consider:

- The types of thinking, knowledge construction, and communication that the discipline values.

- How generative AI can support, extend, or potentially undermine those capabilities.
- How tasks can be designed to reveal the limitations of AI as well as its potential, encouraging critique, comparison, and reflection.

The aim is not simply to decide whether generative AI is ‘allowed’, but to help students understand what kind of learning tool it is, and when, how, and why it may or may not be appropriate. This is a key part of developing disciplinary judgement and academic identity. In mathematics, this means helping students understand that while generative AI cannot replace active engagement with proofs, problem solving, or symbolic reasoning, it may play a useful role in surfacing misconceptions, provoking dialogue, and sharpening their thinking through critique.

2.4.1. Examples from the Mathematical Sciences

The following examples seek to illustrate how generative AI can be integrated into learning of the mathematical sciences within higher education in ways that reflect and complement disciplinary thinking:

- **Error identification task:** Students are given an AI-generated solution to a calculus problem (for example, finding a local maximum using the second derivative). They must identify conceptual errors or procedural shortcuts, and rewrite the solution in full, justifying each step.
Focus: understanding critical features of differentiation and reasoning structure.
- **Compare and contrast proofs:** Students prompt a generative AI tool to produce a proof of the Cauchy Integral Theorem, then compare this with the formal proof provided in lectures or a recommended text. They annotate both versions to highlight missing assumptions, issues in logical progression, or differences in formalism and explanatory clarity.
Focus: developing proof fluency, critical evaluation, and understanding of complex analysis.
- **AI as an exploration tool:** Students use generative AI to explore different formulations of a mathematical concept (for example sequences vs. series), then critique which explanation is most helpful, precise, or misleading. This might be used as part of a tutorial or paired activity to generate discussion.
Focus: conceptual understanding and meta-cognition.
- **Model generation with reflection:** In a mathematical modelling task, students may use AI to generate initial ideas or explore variable relationships. They must document their use of AI, justify modelling decisions, and reflect on how the AI outputs influenced their thinking.
Focus: transparency, applied problem solving, reflective practice.

These activities position generative AI not as a shortcut to achieving ‘correct’ answers but as a dialogue partner, a source of challenge, or a thinking scaffold. They also help develop critical awareness of where AI fails to meet disciplinary expectations, turning limitations into new learning opportunities.

2.5. Digital and AI Literacy

Generative AI skills are now a fundamental part of broader digital and academic literacies. Students need support not only in accessing tools, but in understanding how to use them thoughtfully, critically, and appropriately. This includes recognising where AI tools are embedded in common

platforms and how they may shape learning behaviours, skill development, and academic outputs. A structured approach to developing AI literacy might include:

- **Introductory sessions** on how generative AI works, and its strengths and limitations.
- **Practical workshops** on prompting, critiquing outputs, and recognising misuse or over-reliance.
- **Embedded learning tasks** that require students to reflect on their own use of generative AI.
- **Comparative exercises** analysing human- vs AI-generated work to explore quality, rigour, and disciplinary fit.

Importantly, these opportunities should be embedded throughout a programme, not confined to induction events or optional study skills modules. Digital literacy is developmental, and students need repeated, supported experiences over time to become confident and critical users.

2.5.1. Awareness of Embedded AI Tools

An increasingly complex challenge is that generative AI is becoming invisible, integrated into tools students already use daily, often without them, or educators, realising AI is involved. For example:

- Grammarly now offers rephrasing, content suggestions, and tone control. All functions powered by generative AI.
- Microsoft Copilot in Word, PowerPoint, and Excel provides AI-generated summaries, auto-generated text, and data insights.
- Overleaf, a common platform in mathematics and technical disciplines, has introduced AI-based LaTeX content generation and document suggestions.
- Google Workspace includes generative features in Docs and Slides.

This blurring of boundaries raises critical questions: Are students aware when they are using generative AI? Do they know when and how to declare it? Are staff able to distinguish between tools that are permitted and those that aren't?

As generative AI becomes increasingly embedded, and often hidden, within common software platforms, supporting students to use these tools responsibly is more important than ever. Programme teams have a responsibility to help students recognise when generative AI is being used, even if it is not explicitly labelled, and to understand how such use aligns with institutional policies and expectations for assessment. Crucially, students also need space to reflect on how these tools influence their own thinking, writing, or problem-solving processes. This means supporting students in making informed decisions about the tools they already use, especially where the lines between traditional functionality and AI-generated content is becoming increasingly blurred. Programme teams should therefore provide clear guidance on:

- Whether such tools are permitted for preparation or submission.
- How to differentiate between surface-level features (such as formatting or syntax) and AI-generated content.
- Where assessment briefs and module handbooks should be updated to reflect these evolving capabilities.

Supporting AI and digital literacy is no longer an optional enhancement, it is part of the shared responsibility of curriculum design. The goal is not only to ensure technical competence, but to help students develop ethical, reflective, and academically grounded approaches to AI as part of their learning journey.

2.5.2. Access and Equity

As generative AI tools become more embedded in higher education, equity of access and confidence must be a core design consideration. Not all students begin from the same place. Some may be unable to purchase premium tools, while others may lack regular access to suitable devices or browsers. For many, the challenge lies not in access alone, but in navigating unfamiliar platforms, interpreting complex outputs, or using the tools effectively, particularly if they are working in a second language or are less confident with academic conventions. These disparities risk compounding existing inequalities unless addressed through inclusive programme and institutional design. Crucially, equity in AI use is not just about access, it is also about confidence, transparency, and support.

To promote equity and inclusive participation, programme teams and institutions should:

- **Prioritise free and accessible tools:** Where possible, design learning tasks that can be completed using open-access platforms such as ChatGPT (free tier), Microsoft Copilot (available through the Edge browser), or Google Gemini. Tasks should not reward students for having access to more advanced or premium tools but instead focus on how well they engage with the learning process itself.
- **Promote institutional licences:** Ensure students are aware of centrally supported tools, such as Grammarly (for writing support), GitHub Copilot (for code generation), or subject-specific tools like Wolfram Alpha or Wolfram Chat. A mathematical sciences programme might offer workshops showing how Wolfram tools can be used to explore symbolic algebra or graph functions safely and effectively.
- **Support onboarding:** Provide students with guidance documents, annotated screenshots, or short demonstration videos that walk through how AI tools can be used within academic tasks. For instance, a screencast might show how to use Copilot in Overleaf to generate LaTeX-based mathematical expressions, highlighting what's appropriate for preparation and what must be original.
- **Offer alternatives:** Design tasks that allow students to meet learning outcomes with or without AI. For example, if one option involves prompting an AI tool to generate model solutions, an alternative might allow students to use worked examples from lecture notes or textbooks, combined with their own commentary or analysis.
- **Avoid hidden advantage:** When setting assignments, consider whether access to premium tools (for example, GPT-4, paid statistical plugins, or advanced coding assistants) might confer unfair advantage. This might mean standardising the tool expected for a task, or clearly stating that outputs must be human-authored, even if AI is used during preparation.

These principles apply not only in formal assessments, but in everyday learning. Embedding low-stakes, supported opportunities for experimentation, such as peer-led discussions, tutorial activities, or scaffolded practice, can help normalise AI use and build collective confidence. Wherever possible,

AI-related activities should be designed so that any permitted tool can be used effectively. This not only promotes inclusion but ensures that equity concerns don't unintentionally reinforce existing gaps in access, confidence, or engagement.

Illustrative Example 1: Building a Glossary Using AI in Year 1 Mathematics

As a foundation-level task, students use a generative AI tool to produce draft definitions for ten key mathematical terms relevant to their course, such as *surjection*, *convergence*, and *basis*. Students are encouraged to treat the AI output as a starting point for critical engagement.

They then:

1. Refine each definition using lecture notes, textbooks, or other trusted resources.
2. Provide examples to illustrate each term in context.
3. Submit both the revised glossary and a brief reflection describing how the AI-generated definitions were improved or clarified.

This task supports the development of mathematical language, conceptual understanding, and early AI literacy in a low-stakes, formative setting.

3. Designing Learning Experiences

Integrating generative AI into learning design is not simply about introducing new tools, it is about reshaping how students engage with knowledge, build understanding, and develop academic confidence. Just as assessment design must be reconsidered in the context of AI, so too must the learning experiences that underpin it. When used well, generative AI can support students in asking better questions, exploring alternative approaches, receiving personalised feedback, and practising at their own pace. However, these benefits are only realised when learning activities are deliberately designed to take advantage of what AI can offer. As with any educational tool, its value depends entirely on how it is used.

Effective design of AI-supported learning should be intentional, transparent, and clearly aligned with learning outcomes. Students need to understand why generative AI is being used in a particular task, and how it supports the development of specific knowledge or skills. Activities should prompt critical engagement with AI outputs, encouraging students to question, interpret, and evaluate what they are given, rather than passively accept it. Crucially, the use of AI must not replace the essential human elements of meaningful learning: dialogue, feedback, collaboration, and productive struggle must remain central to the learning experience.

At programme level, coordination is key. Educators should not assume that AI use in learning will emerge organically or be consistent across modules. Programme teams should work collaboratively to map where AI-supported learning is already happening, identify gaps or opportunities for experimentation, and ensure that students encounter a variety of AI interactions throughout the curriculum, from exploratory practice to critical analysis to co-creation.

3.1. Roles Generative AI Can Play in Enhancing Student Learning

One helpful way to think about how generative AI can be integrated into learning is to consider the role it can potentially play within a given task. Sharples (2023), and expanded within Sabzalieva and Valentini (2023), outlines a set of pedagogical roles that AI might adopt, depending on how

educators design learning activities. These roles are not mutually exclusive, students might move between several within a single task, but asking “What role might (or do) I want AI to play in this activity?” is a useful initial design prompt for educators.

AI Role	What it Does	Mathematics- Specific Example
Possibility Engine	Suggests alternative ideas or expressions	A student exploring methods to solve a system of equations prompts AI to suggest alternative approaches (for example, substitution, matrix methods, graphical). In statistics, AI proposes different visualisations (boxplot, histogram, violin plot) for summarising a dataset.
Socratic Opponent	Challenges thinking with counterpoints or questions	A student preparing a proof involving irrational numbers tests it against AI by asking for potential flaws or counterexamples. In a statistics context, students use AI to generate critiques of a sampling method or challenge assumptions in an experimental design.
Collaboration Coach	Supports group problem solving and information gathering	A student group working on a final-year project asks AI to suggest types of regression models suitable for predicting housing prices, then compares AI's suggestions with academic sources.
Guide on the Side	Provides scaffolding or suggestions	While tackling an unfamiliar integration problem, students ask AI for a hint on which substitution might be useful, without being shown the full solution.
Personal Tutor	Offers feedback on progress or draft responses	A student submits a worked solution to a proof involving induction and asks AI to spot gaps in the logical progression. In statistics, AI reviews a student's draft report and suggests improvements to the interpretation of p-values.
Co-designer	Assists in developing plans or tasks	A student working on a mini project in applied mathematics uses AI to structure a comparison between exponential and logistic growth models, identifying variables and expected behaviour. In statistics, students use AI to design a small survey and plan how they will clean and visualise the data.
Exploratorium	Prompts experimentation and discovery	Students vary parameters in a function and use AI to help visualise the resulting graphs, investigating how changes affect continuity or convergence. In statistics, students use AI to simulate repeated sampling from different distributions and observe variability.
Study Buddy	Supports revision and retrieval	AI quizzes a student on key theorems in real analysis, generating practice problems with varying difficulty. In statistics, it generates true/false questions on hypothesis testing assumptions.
Motivator	Offers challenges and gamified learning prompts	AI provides a set of logic puzzles that become progressively more complex and aligned with the week's learning. In statistics, students ask AI to create a challenge sequence for interpreting confidence intervals under time pressure.
Dynamic Assessor	Helps track learning and identify gaps	A student asks AI to generate a concept map of their current knowledge in vector calculus. In statistics, a student reviews a sequence of AI-generated summary questions to identify weak areas in their understanding of correlation and causation.

Table 4: Roles of Generative AI in Mathematics and Statistics Learning. *Adapted and extended from the work of Sharples (2023) and as presented in Sabzalieva and Valentini (2023), the roles outlined here demonstrate how generative AI can support mathematics and statistics learning by acting as exploratory partner, feedback provider, or dynamic assessor.*

3.2. Designing AI-Supported Learning Activities

The roles outlined in Section 3.1 can be translated into patterns of learning activity, that is structured ways in which students interact with content, ideas, and one another, with AI acting as a supporting presence. The following examples serve as adaptable templates for designing effective AI-enhanced learning tasks. Programme teams might use these to map existing AI-supported learning opportunities, identify new areas for exploration or reflection, and ensure alignment between learning activities and assessment practices.

3.2.1. Prompt-led Learning Tasks

(Related roles: Possibility engine, Personal tutor, Study buddy)

These tasks involve students prompting AI tools directly and interpreting the outputs. They support the development of procedural fluency and conceptual understanding, encouraging students to evaluate the clarity, accuracy, and educational value of generated content. Example activities might include:

- **Calculus:** Students use AI to generate five integrals involving substitution, solve each one, and explain which were well-constructed and which were misleading.
- **Set Theory:** Students prompt AI for an explanation of the difference between injective and surjective functions, then rewrite it to support peer understanding.
- **Real Analysis:** Students request definitions and examples of uniform convergence from AI, annotate the outputs, and evaluate their mathematical accuracy.

Illustrative Example 2: Prompt-Led Practice in Calculus

Students are introduced to substitution as a method for evaluating integrals and then tasked with using generative AI to create a set of integration problems designed to reinforce this skill. They prompt the AI to generate five definite or indefinite integrals that require substitution and solve each by hand.

Following this, students complete a short evaluative commentary where they:

1. Assess whether each integral is solvable using standard substitution techniques and identify any that were poorly constructed, ambiguous, or beyond the intended scope.
2. Reflect on whether the AI-generated questions matched the complexity and structure of problems typically encountered in class or assessments.
3. Highlight any potential misconceptions a learner might develop if they relied solely on the AI-generated questions, such as missing key constraints or misapplying substitution.

This task is used in a formative setting to support both procedural fluency and the ability to critically assess learning resources, helping students take ownership of their practice and develop a more evaluative relationship with the tools they use.

3.2.2. Reverse Engineering and ‘Flipped’ Learning

(Related roles: Socratic opponent, Dynamic assessor, Guide on the side)

These tasks involve students analysing, critiquing, or correcting AI-generated responses. They develop deeper reasoning, logical clarity, and confidence in academic argumentation. Example activities might include:

- **Mathematical Logic:** Students prompt AI to prove De Morgan’s Laws, critique the logic, and write a corrected version with explanations.

- **Probability Theory:** Students prompt AI to calculate and explain the expected value of a discrete random variable. They then critique the explanation, correct any misconceptions, and rewrite the reasoning to meet formal statistical standard.
- **Differential Equations:** Students prompt AI to solve a first-order differential equation using an integrating factor. They are then asked to verify the steps, identify any shortcuts or incorrect assumptions, and rewrite the solution to include all working and justifications expected at university level.

Illustrative Example 3: Clarifying Statistical Misconceptions

Students are asked to use generative AI to produce explanations of fundamental statistical concepts, such as *p-values*, *confidence intervals*, or *correlation vs. causation*. They are then required to critique and refine the output to develop a deeper understanding of both the concept and its appropriate communication. Students complete the following steps:

1. Identify inaccuracies, oversimplifications, or common misconceptions present in the AI-generated explanation (for example, implying that a *p-value* indicates the probability the null hypothesis is true).
2. Rewrite the explanation using statistically rigorous language suitable for a first-year audience, drawing on taught materials and trusted resources.
3. Reflect on why the original explanation was misleading, and explain how their revision provides a clearer, more accurate interpretation.

The final output includes both the original AI response and the revised version, annotated with brief justifications. This activity supports the development of statistical reasoning, precision in communication, and the critical appraisal of AI-generated content, particularly in contexts where inaccurate explanations may be persuasive but flawed.

3.2.3. Personalised Exploration and Practice

(Related roles: Exploratorium, Study buddy, Personal tutor)

These tasks help students tailor AI use to their individual learning needs, generating practice material, summarising concepts, or quizzing themselves on areas of weakness. They support metacognitive development and build learner independence. Example activities might include:

- **Differential Equations:** Students prompt AI to explain the difference between homogeneous and non-homogeneous ODEs, then generate and solve related problems.
- **Mathematical Logic:** Students use AI to test different truth table structures for logical equivalence and evaluate which best support their learning.
- **Descriptive Statistics:** Students ask AI to generate small datasets that illustrate specific statistical concepts, such as skewness, outliers, or variance. They then calculate summary statistics by hand, interpret the results, and reflect on how well the AI-generated data meets the original brief.

Illustrative Example 4: Personalised Practice with Generative AI

Students identify specific areas of weakness or uncertainty in their mathematical understanding, such as techniques in integration, properties of sequences, or interpreting statistical output. Using a generative AI tool, they prompt the system to produce relevant practice questions tailored to those areas. They then select at least two examples to solve independently, followed by a structured evaluation in which they:

1. Assess whether the generated material was accurate, relevant, and appropriately challenging for their level.
2. Reflect on how helpful the AI was in reinforcing concepts or clarifying misunderstandings.
3. Identify any misconceptions or over-simplifications introduced by the AI, and consider how they might affect learning if left unchallenged.

This task encourages metacognition, supports independent study, and helps students take greater ownership of their learning. It also introduces opportunities for dialogue in tutorials - comparing questions, evaluating usefulness, or exploring how prompts shape outcomes.

3.2.4. Simulated Dialogue and Peer Learning

(Related roles: Socratic opponent, Collaboration coach, Motivator)

These activities simulate academic dialogue or collaborative reasoning, enabling students to practice articulating ideas, debating alternatives, and responding to critique. They prepare students for tutorials, group work, or oral assessments. Example activities might include:

- **Number Theory:** Students ask AI to 'disagree' with their conjecture about a divisibility rule and use the resulting dialogue to identify limitations or counterexamples. They then revise their conjecture and test it with peers.
- **Mathematical Modelling:** Groups use AI to generate and evaluate assumptions for a population model, deciding collaboratively which to adopt and why.
- **Abstract Algebra:** Students role-play a dialogue with AI around whether a non-zero ring with no multiplicative identity can still be a ring. The conversation becomes the basis for class discussion and clarification of formal axioms.

Illustrative Example 5: Simulated Dialogue in Abstract Algebra

Students use AI to investigate whether a particular mathematical structure satisfies the group axioms. They engage the AI in a structured dialogue about a proposed set and operation, for example, the set of 2×2 invertible matrices under matrix multiplication. Students are asked to:

1. Prompt AI to assess the group properties (closure, associativity, identity, and inverse) for the chosen structure.
2. Identify and annotate points of agreement or disagreement with formal mathematical reasoning.
3. Rewrite the conversation as a structured proof, highlighting where AI responses were helpful, incomplete, or misleading.

This task develops deeper understanding of abstract structures, encourages critical evaluation of mathematical reasoning, and promotes confidence in formal proof-writing through dialogic exploration.

3.3. Framing AI as a Learning Partner

Designing learning experiences that make effective use of generative AI is not about automating teaching, it is about expanding the ways students can practise, test, and deepen their understanding. When used well, these tools offer new spaces for exploration, experimentation, and feedback. The goal is not to replace the learning process but to enhance it, making space for students to engage more meaningfully with ideas and develop confidence through practice.

Critically, this requires more than ad hoc integration at the module level. It calls for intentional programme-level design, where the use of AI is scaffolded, aligned with learning outcomes, and

supported through dialogue, reflection, and inclusive practices. When designed in this way, generative AI becomes not a threat to teaching, but a partner in learning, helping students build knowledge, question assumptions, and take ownership of their academic journey.

4. Designing Assessment

Effective assessment is a cornerstone of good programme design. In an era of generative AI, the need for diverse, inclusive, and well-aligned assessment strategies has never been more urgent. While much of the national and institutional focus has centred on the risks AI poses to traditional forms of assessment, there is also significant opportunity to rethink the purpose and value of assessment within a programme.

Not every assessment needs to be written. Not every task needs to be individual or unseen. A well-designed programme will expose students to a range of assessment types, oral, visual, practical, reflective, collaborative, aligned to the knowledge and skills the programme aims to develop. This variety encourages different modes of thinking and expression, offers more inclusive pathways for students to demonstrate their learning, and reduces over-reliance on any single format that may be particularly vulnerable to automation.

The detailed pedagogical challenges and design implications of assessment in the context of generative AI are explored in greater detail within Grove (2024). This section draws on and complements that work, focusing on practical ways to support programme-level coherence and the integration of AI into assessment practices in mathematics and beyond.

4.1. Assessment Design and Bloom's Taxonomy

Bloom's taxonomy offers a useful framework for understanding the types of thinking that assessment tasks aim to elicit. While tasks at the lower levels, such as remembering or basic comprehension, can now be completed more easily using generative AI tools, this only strengthens the case for designing assessments that promote higher-order thinking. Tasks that require students to apply their knowledge in new contexts, solve problems, evaluate alternatives, make reasoned decisions, or produce original work are increasingly important in an AI-enabled learning environment. At the same time, focusing on these higher-level skills does not mean abandoning the lower levels; rather, it affords opportunities to consolidate foundational knowledge through meaningful application, helping students reinforce core concepts in more authentic and challenging contexts.

Table 5 shows a set of mathematics-focused assessment examples aligned with different levels of Bloom's taxonomy. Each example illustrates how AI might be used either as a tool within the task or as a feature for students to critique and build upon.

These examples demonstrate that the use of AI in assessment is not inherently problematic, but the purpose of the assessment must be clearly communicated. Students should know whether AI use is permitted, what form that use can take, and how it should be acknowledged (see Section 2). Just as importantly, they must understand what learning outcomes the task is designed to assess. This includes knowing whether the focus is on accuracy, reasoning, conceptual understanding,

communication, or reflection. Without that clarity, students may unintentionally misuse AI or fail to demonstrate the very skills the assessment is intended to develop.

Cognitive Level	Assessment Task (Mathematics Example)	Use of Generative AI	Learning Purpose
Remembering	Define and explain key terms from Real Analysis using your own examples.	AI can provide initial definitions; students must personalise, extend with examples, and check accuracy.	Support terminology recall, concept clarification, and confidence-building particularly in early-stage learning.
Understanding	Explain the difference between pointwise and uniform convergence, with annotated diagrams.	Students use AI to draft an explanation, then refine it using lecture notes and annotate errors or omissions.	Promote conceptual understanding, diagrammatic reasoning, and the ability to identify nuance in formal explanations.
Applying	Solve a differential equation and apply it to a physical model (for example, cooling of an object).	AI may assist in exploring solution strategies during preparation; final submission must include full working and interpretation.	Encourage procedural fluency, application of methods to real-world contexts, and awareness of modelling assumptions.
Analysing	Compare two AI-generated solutions to a matrix problem. Identify strengths, flaws, and missing steps.	AI-generated solutions are integrated into the task; students focus on critique, logical coherence, and comparative reasoning.	Develop critical thinking, error detection, and understanding of valid mathematical argument structure.
Evaluating	Select a method to approximate an integral numerically. Justify your choice and discuss its limitations.	AI can suggest possible methods; students evaluate these, select the most appropriate, and explain the rationale in a structured report.	Foster evaluative judgement and decision-making between alternative mathematical techniques.
Creating	Design a mathematical model to represent population growth, stating assumptions and constraints.	Students brainstorm with AI to generate possible model forms, then document decisions, reflect on assumptions, and justify their final approach.	Support creative modelling, mathematical justification, and reflection on the use and limitations of AI in exploratory tasks.

Table 5: Generative AI Use Across Levels of Bloom’s Taxonomy in Mathematics Assessment

This overview illustrates how assessment tasks in mathematics can be designed to align with Bloom’s taxonomy while integrating generative AI in purposeful and pedagogically appropriate ways

4.2. Flipped Assessment: An Example Approach

While assessment design has been discussed in more detail in previous work (Grove, 2024), it is important to reinforce here that integrating generative AI into learning requires a corresponding shift in how we design assessment. Assessments should not be disconnected from the tools and strategies students are using throughout their studies. If AI is part of their learning process, whether to generate examples, explain concepts, or simulate problem-solving, then assessment must evolve to account for that. The aim is not simply to permit or prohibit AI use, but to create assessment tasks that require students to think critically, engage deeply with content, and demonstrate intellectual ownership of their work.

One possible approach is a *‘flipped assessment’* model. In this design, students begin with AI-generated content but are assessed on their ability to interrogate, adapt, and improve that content. Rather than focusing solely on producing work from scratch, students are asked to demonstrate higher-order understanding through critique, transformation, and reflective commentary. This

model mirrors professional and academic practice. Rarely are problems in research or the workplace solved in isolation or from a completely original starting point. Mathematicians and scientists often refine flawed solutions, test assumptions, adapt known structures, or improve clarity and precision. These are valuable academic and graduate skills, and ones difficult to outsource to AI.

Illustrative Example 6: Validating GAI Solutions in Linear Algebra

Students are provided with a generative AI-produced solution to an inverse matrix problem. The AI's output may include notational errors, omitted justifications, or incorrect interpretations.

Students are tasked to:

1. Identify any inaccuracies or inconsistencies in the solution.
2. Annotate the solution using correct mathematical notation and terminology.
3. Compare the AI's method to their own and provide a written justification of any differences in approach or interpretation.
4. Identify the key underpinning mathematical ideas that are fundamental to successful understanding (for example, if the inverse matrix is multiplied by the original matrix, the identity matrix should be obtained).

This task helps students practise rigorous notation, clarify common misconceptions, and strengthen their understanding of eigenvalue problems through critique and comparison.

Flipped assessment is not about reducing expectations and standards; it is about shifting them. It positions students as critical users of AI, not passive consumers, and it rewards intellectual control, insight, and precision. At a programme level, this approach can help scaffold students' engagement with generative AI ahead of summative tasks, place greater emphasis on reasoning, analysis, and revision within assessment criteria, and prepare students for the evaluative demands of research and professional practice.

Illustrative Example 7: Flipping the Proof with Fermat

Students are asked to use AI to generate a worked solution to a question involving Fermat's Little Theorem (which in the notation of modular arithmetic is written as $a^p \equiv a \pmod{p}$) and which asks them to show that if a is not divisible by p , then Fermat's Little Theorem is equivalent to $a^{p-1} \equiv 1 \pmod{p}$. The AI output often omits that a must be coprime to the prime p . The student's task is to:

1. Review the AI-generated proof, appraising its accuracy, conciseness, and level of detail, and identifying any missing assumptions - especially the requirement that a and p be coprime.
2. Rewrite the solution correctly, including all necessary conditions and ensuring that the logic is clear and formally valid.
3. Reflect on why such errors are common in AI-generated mathematics and how they relate to formal mathematical reasoning.

This task assesses the student's understanding of proof structure, logical conditions, and the importance of precision - skills central to mathematical practice. Rather than producing a proof from scratch, which can readily be found online anyway, the student is challenged to diagnose, critique, and improve, demonstrating deeper engagement with the core mathematical concepts, ideas, and subtleties.

4.3. Aligning Assessment Criteria

As generative AI becomes increasingly embedded in higher education, our assessment criteria must evolve to reflect the changing demands of the discipline and the wider contexts in which disciplinary thinking is applied. At a programme level, this means re-evaluating the weight given to procedural tasks, such as routine calculations, symbolic manipulation, or reproducing standard techniques, and placing greater emphasis on deeper understanding, strategic thinking, and the ability to apply mathematical ideas flexibly and in increasingly novel contexts.

While fluency with core methods remains essential, assessment should highlight students' capacity to explain their reasoning, justify the choice of particular techniques, and apply concepts to unfamiliar or complex situations. Clarity of mathematical communication, through structured argument, precise notation, and effective use of visual representations, should be recognised explicitly in assessment criteria, reflecting the importance of conveying ideas clearly in both academic and professional settings. The use of real or simulated data offers opportunities for students to demonstrate interpretation and analysis, moving beyond abstract manipulation to more applied, context-rich problems. Criteria should also acknowledge the value of identifying and addressing errors, contradictions, or limitations in reasoning, whether through individual work or with support from tools such as generative AI. These tasks develop students' critical thinking and can help surface deeper mathematical understanding.

Incorporating reflective components, such as commentary on the decision-making process, the interpretation of AI-generated outputs, or the evaluation of multiple solution strategies, can further enhance students' metacognitive awareness. Project-based assessments and open-ended problems that emphasise problem formulation, modelling, and exploration provide authentic opportunities to assess how students think mathematically, not just what they can compute or remember. At a programme level, a consistent and transparent approach to updating assessment criteria can help students build confidence in what is valued across their learning journey and prepare them for evolving mathematical practice beyond university.

4.4. Coherence and Transparency at Programme Level

At a programme level, assessment design should be guided by shared principles that ensure coherence, progression, and transparency. Students should not encounter contradictory guidance about the use of generative AI in different modules; rather, programmes should present a consistent and clearly communicated stance. Alongside diverse assessment types, programmes may also include a defined proportion of Assured Credits (see Section 2.3.3) to ensure a baseline of unaided student achievement. These are complemented by Exploratory Credits, where students can engage more openly and reflectively with generative AI, supported by clear expectations and scaffolded practice. Assessment formats should evolve over the course of the programme, offering increasing complexity, independence, and opportunities for critical reflection. Alongside academic knowledge and disciplinary skills, students should also be supported in developing metacognitive awareness and ethical judgement, particularly in relation to how they engage with generative AI technologies.

Assessment remains one of the most powerful tools we have as educators to shape student learning. In an AI-enabled era, it is no longer enough to protect assessment from generative AI, we must

instead design assessment in ways that respond to its presence and potential. When thoughtfully aligned to programme aims, assessment can encourage deeper engagement, foster independence, and prepare students for the intellectual and ethical demands of academic and professional life.

5. Supporting Students

Designing for generative AI at programme level involves more than embedding tools into teaching and assessment. It requires a commitment to supporting students as they develop the capacity to use these technologies independently, critically, and responsibly, both within formal learning tasks and beyond. This support must evolve over time. Generative AI is not static, and the ways students engage with it will change as tools develop, expectations shift, and confidence grows.

Programmes therefore need to provide sustained support that recognises generative AI as both a learning tool and a literacy. This includes helping students develop the judgement to evaluate when and how to use AI effectively, and when not to use it at all.

5.1. Independent and Responsible Use

Students across disciplines, including mathematics, are already using generative AI to support their learning. These tools offer instant explanations, walkthroughs of problems, auto-generated revision resources, and help with structuring answers. Used with discernment, they can encourage self-directed study, boost confidence, and provide flexible support. But the risks are real. Over-reliance on AI tools can inhibit the development of reasoning skills, mask misconceptions, or lead to the uncritical acceptance of flawed or superficial responses. To address this, programme teams should actively support students in using generative AI in independent study, not through prohibition, but through design.

Activities such as reflective logs, structured prompts, or tutorial discussions can help students consider when AI tools add value, and when they obscure understanding. Tutorials can incorporate short tasks that ask students to compare AI-generated summaries to lecture notes, explore whether an AI explanation would be suitable for a peer, or critique step-by-step solutions for gaps in reasoning. These small interventions help students treat AI as a companion to learning, not a shortcut.

Illustrative Example 8: Problem Deconstruction from AI Solutions

Students are given a differential equation generated and solved by an AI tool, and are challenged to work backwards to reconstruct the context in which the equation might have arisen.

Their task is to:

1. Infer the original problem scenario that could lead to the given equation (for example, population growth, thermal change).
2. Identify any missing boundary conditions, assumptions, or modelling steps.
3. Explore alternative ways the problem could be modelled and explain the reasoning behind those choices.

This activity supports higher-order modelling skills, critical engagement with AI outputs, and the development of a more intuitive understanding of differential equations in real-world contexts.

Programmes might also include short reflections in portfolios or module assessments, asking students to describe how they used AI, what limitations they found, and what strategies they would recommend to others. These kinds of reflective habits encourage ethical and thoughtful engagement with AI and offer staff valuable insight into student learning practices.

As part of developing students' understanding of generative AI, it is important to introduce the environmental and ethical dimensions associated with its use. This includes helping students critically reflect on the sustainability implications of large-scale AI models, such as their energy demands, data usage, and broader societal impacts. Students should also be encouraged to consider how to engage with generative AI in a responsible and sustainable manner, for example, using tools purposefully rather than excessively, and evaluating when their use adds value to learning or problem-solving.

Programmes should also acknowledge that some students will choose not to use AI. This may be due to uncertainty, ethical concerns, or a preference for traditional approaches. These students should be supported with reassurance rather than pressure. For some, the sheer number of tools can be overwhelming; others may feel the need to be constantly 'optimising' their study. Promoting thoughtful, balanced engagement with technology is an important part of supporting student wellbeing.

5.2. Developing AI Literacy and Agility

Generative AI is not a fixed technology. New tools, interfaces, and expectations are constantly emerging. Students need more than a one-off induction; they need a developmental approach to AI literacy, and one that treats it as part of their academic and professional skillset. Programmes can support this through a scaffolded model that revisits key principles throughout the curriculum:

- **Orientation:** Early in the programme, students should learn what generative AI is, how it works, and where it supports or conflicts with disciplinary values. This includes understanding ethical use, citation practices, and institutional expectations (for example, through the traffic-light frameworks outlined in Section 2).
- **Skill-building:** As students progress, they can be introduced to more advanced practices: writing purposeful prompts, comparing tool outputs, or identifying the strengths and weaknesses of AI-generated responses.
- **Critical engagement:** Advanced activities should involve not just using AI, but challenging it. For example, students might test the robustness of an AI-generated proof, annotate a flawed response, or reflect on how AI shaped their thinking.

Illustrative Example 9: Dialogue with a Mathematician

To explore mathematical concepts in a creative and reflective way, students are asked to simulate a scripted dialogue with an historical mathematician (for example, Ramanujan, Noether, Euler, or perhaps an exchange of letters between de Fermat, Wiles, and Taylor), using a generative AI tool to generate a first draft of the exchange.

They are then asked to:

1. Edit and expand the dialogue to ensure historical accuracy and mathematical clarity.
2. Include at least one moment where a misconception is corrected or challenged.

3. Write a short reflective commentary on what they learned, what surprised them, and how the exercise deepened their understanding of the concept or figure.

This task helps students humanise mathematical thinking, clarify misconceptions, and practise articulating mathematical ideas in dialogue form.

These opportunities can be integrated into discipline-based teaching or offered through co-curricular activities such as peer-led workshops or embedded tasks in tutorials and labs. Key is that they are iterative, revisited at multiple stages, and aligned to increasing expectations for independence, accuracy, and ethical awareness.

As discussed earlier in Section 2.5.2, equity and access must be central considerations when designing learning tasks that incorporate generative AI. These principles apply not only in structured teaching activities but also in the independent and informal ways students engage with AI tools. Programmes should continue to embed opportunities for low-stakes experimentation, peer learning, and reflection, ensuring that all students feel confident and supported, regardless of prior experience or access.

5.3. Supporting Students Who Choose Not to Use Generative AI

While much of this framework focuses on enabling students to use generative AI effectively and responsibly, it is equally important to support those who choose *not* to engage with such tools. Some students may have ethical or environmental concerns, others may be cautious about academic integrity risks, and some may simply prefer to develop skills through more traditional approaches. These choices should be respected and supported as part of an inclusive learning environment.

One of the most practical ways to ensure inclusivity is through flexible assessment design. Tasks should allow students to meet the intended learning outcomes without requiring generative AI. For example, where one option invites students to prompt an AI tool to produce content for critique, an alternative should be available that draws on lecture materials, textbooks, or staff-provided examples. Time and workload assumptions also require careful consideration. Some students may choose not to use AI tools to accelerate or automate stages of their work, such as research, summarising, or drafting. Assessment briefs should avoid implicitly privileging AI-enabled efficiency. Timelines and expectations should be realistic for students completing all tasks manually.

Where AI-generated materials are used in teaching or assessment, it can be helpful to provide optional resources so that students are not required to use a tool themselves; pre-prepared examples, for instance, can ensure learning parity while preserving autonomy. This principle also extends to feedback and peer interactions, where it is important not to assume all students have used AI tools. In both formative and summative contexts, staff should recognise and respect different approaches, including deliberate non-use, and avoid presenting AI use as inherently more advanced or effective. While showcasing examples of productive AI use can be helpful, this should be balanced with recognition of strong work produced without it. More broadly, acknowledging that non-use is a valid position can support confidence and wellbeing. Students should feel able to articulate their approach, whether in reflective writing, one-to-one settings, or portfolio commentary, and thoughtful non-use should be positioned as an intellectually engaged choice rather than a deficit.

Finally, teaching about the broader context of generative AI, including its ethical, environmental, and social implications, creates space for students to explore their concerns. Supporting critical engagement in this way enables all students to make informed, reflective decisions about whether and how they wish to engage with AI tools as part of their learning.

6. Programme-Level Checklist: AI-Integrated Programme Design

This checklist is intended to support programme teams as they review, refresh, or redesign their curriculum in light of generative AI technologies. It poses a series of questions that encourage reflection upon the principles, strategies, and examples presented within this paper.

1. Purpose and Progression

- Have we reviewed programme learning outcomes for opportunities to embed AI literacy or align with digital graduate attributes?
- Are students supported to progress from basic to critical use of AI across the years of study?
- Have we aligned AI-enhanced tasks and assessments to different levels of Bloom's taxonomy?

2. Designing the Learning Environment

- Have we protected time and space for discussion, feedback, and collaboration?
- Are we actively designing learning opportunities that AI cannot replace, for example peer work, dialogic teaching, reflective engagement?
- Do students understand the value of these human elements in a tech-enabled learning environment?

3. Ethical Use and Acceptable Practice

- Is there a clear, shared framework (for example, a 'traffic-light' model) applied to every summative assessment?
- Are students taught when AI may be used in preparation - even if not in submission?
- Where AI is prohibited, have we designed conditions that make this meaningful (for example, Vivas, in-class work)?
- Have we established clear guidance on citing AI use, including tools like Grammarly or Overleaf?
- Are students introduced to the ethical, social, and environmental implications of generative AI, including how to use it in a responsible and sustainable way?

4. Respecting the Discipline: AI Use in Context

- Have we explored how generative AI aligns, or conflicts, with disciplinary ways of thinking or our values as a subject area?
- Are students given opportunities to compare AI-generated outputs to academic or professional standards in our field?
- Have we provided examples where AI use is pedagogically valuable, and where it's pedagogically limiting?

5. Digital and AI Literacy

- Are there scaffolded opportunities across the programme to develop AI-related skills?
- Do students know which tools are free, supported, and appropriate for use in our context?

- Have we acknowledged and addressed differences in access and prior experience?

6. Assessment Design

- Have we reviewed assessments for over-reliance on formats vulnerable to AI automation?
- Are we experimenting with new formats (for example, flipped assessment, critique, meta-analysis) to assess deeper learning?
- Is assessment varied, inclusive, and aligned with programme-level principles around AI use?
- Does the programme include Assured Credits or equivalent secured assessments to ensure a baseline of unaided student achievement, alongside opportunities for open and exploratory AI use?
- Have marking schemes been reviewed to reflect permitted AI use, including expectations for documentation, critical engagement, and citation where required?

7. Supporting Students

- Are students supported in their independent use of generative AI, including how to use it responsibly outside of taught sessions?
- Do we treat AI literacy as a skill to be revisited and extended across the programme?
- Are we actively working to reduce inequities in tool access, confidence, and support?
- Do we support students who choose not to use AI through clear design, balanced expectations, and reassurance that non-use is a valid academic choice?

7. Conclusion

The integration of generative AI into higher education is not a one-time activity, but an ongoing pedagogical and strategic consideration that must evolve alongside technological, institutional, and disciplinary developments. This paper has presented a programme-level approach to embedding generative AI across learning, teaching, assessment, and student support, with a particular emphasis on the mathematical sciences.

We have argued that generative AI cannot be introduced in a piecemeal or opportunistic manner. Its use must be aligned with programme outcomes, assessment strategy, disciplinary identity, and the wider student learning journey. Designing for AI requires intentionality, not only in identifying where its use is permitted, but in articulating how it supports learning and what forms of engagement are educationally meaningful.

When used thoughtfully, generative AI can enhance student confidence, foster independence, and support deeper forms of reflection and enquiry. But this potential is only realised when its use is transparent, principled, and scaffolded. Programme teams must protect what matters most in human learning, dialogue, collaboration, criticality, and intellectual struggle, while helping students develop the literacy, ethics, and confidence to use AI responsibly across varied contexts.

A programme designed for learning with generative AI is not one that integrates it everywhere. It is one that uses it purposefully, to extend thinking, to enrich engagement, and to prepare students for a world in which human and machine intelligence operate in partnership. As tools evolve, so too must our approaches to curriculum, assessment, and support. There is no single blueprint. But there is a shared responsibility, to ensure our programmes are coherent, inclusive, and future-facing.

The checklist within section 6 offers a practical tool for programme teams. It is designed to prompt discussion, guide planning, and support continuous reflection as institutions navigate the opportunities and challenges of learning with generative AI.

References

Anderson, L. W. & Krathwohl, D. R. (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives: complete edition*. Addison Wesley Longman, Inc.

Freeman, J. (2025). *Student Generative AI Survey 2025*. Technical report, Higher Education Policy Institute (HEPI), <https://www.hepi.ac.uk/2025/02/26/student-generative-ai-survey-2025>.

Grove, M.J. (2024). Generative AI technologies and their role within assessment design. *Education in Practice*. Vol.5, Issue 1, pp 15 – 35. Available at <https://www.birmingham.ac.uk/documents/hefi/gai/eip-gai-and-assessment-design-publisher-version-august-2024.pdf>.

Sabzalieva, E. & Valentini, A. (2023). *ChatGPT and artificial intelligence in higher education: Quick start guide*. UNESCO. Available at <https://unesdoc.unesco.org/ark:/48223/pf0000385146>

Sharples, M. (2023). Towards social generative AI for education: theory, practices and ethics. *Learning: Research and Practice*, 9(2), 159-167.

Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., ... & Waddington, L. (2023). Testing of detection tools for AI-generated text. *International Journal for Educational Integrity*, 19(1), 26. <https://doi.org/10.1007/s40979-023-00146-z>.