

The dangers of using AI for coursework: A warning to university students

Prof. Hannah Daly, UCC¹

September 17th 2026

This is a personal essay on my current thinking on the use of GenAI within higher education, targeted at university students. This is not a “policy” for any of my modules, the BSc Sustainability programme or UCC as a whole.

We need to talk about AI

Generative AI (GenAI – “chatbot” tools like ChatGPT, Gemini, Claude) lets you simulate the production of a good assignment, or of the right answer. To a non-expert, the output can look fluent, well-organised and confident, and it can get you a fine grade with little work. The temptation to use GenAI to produce your coursework could be strong, especially if you are under time pressure from other commitments.

But there are many risks of relying on GenAI to support your education. Passing off AI-generated content as your own work is simply cheating, and breaches academic integrity rules. But my essay focuses on a larger concern – that relying on it ubiquitously as a tool can more broadly undermine the skills you need to learn and think, through a process called *cognitive offloading*.

Cognitive offloading

Cognitive offloading is where you reduce the mental effort needed to think, remember, reason, or manage information. We do it constantly, and mostly it's fine — shopping lists, counting on fingers, calculators, satnav, saved phone numbers. But we have to understand how our brains work to

¹ Please cite as Daly, H. (2026). The dangers of using AI for coursework: A warning to university students. Zenodo. <https://doi.org/10.5281/zenodo.22816288>

understand how it impacts how we think. The brain is efficient – it will use the path of least resistance to get to an outcome. If you give a child a calculator, they won't learn their times tables. Cognitive offloading makes mental tasks easier, so we don't have to remember phone numbers any more (they're stored on our phone now), and we don't have to learn the streets in our city (because we use Google Maps). Cognitive offloading can cause a decline of our capabilities, which means we rely on the tool even more. **It creates a weakening of our abilities, and a dependence loop.** Research suggests that people's ability to retain facts declines when they have access to Google search².

As with training for a sport, or learning to play a musical instrument, you have to struggle and make an effort to learn. You don't get fit by using an e-scooter to get from A to B instead of running, even though the outcome on Strava might be the same!

Thinking is the same. You can't offload the hard part and still get the benefit.

GenAI is different from a calculator

A calculator offloads arithmetic, which is a specific, well-defined outcome. Relying on a calculator will reduce your ability to do long multiplication, but you might accept this trade-off.

But GenAI is different and the risks of depending on it, especially during third level education, are much greater. It is a general tool and it's higher order – it can offload a lot of the work of learning – like framing the problem, having to read and understand long texts, deciding what matters, structuring an argument, and judging the answer. It's also frictionless, and it can produce a polished answer, where the style gets mistaken for depth.

Some people are going further and calling this **cognitive surrender**, a term proposed by Shaw and Nave (2026)³. They describe how users of AI adopt its reasoning and answers without playing any part of the problem-solving process themselves: "cognitive surrender represents a deeper abdication of critical evaluation, where the user relinquishes cognitive control and adopts the AI's judgment as their own"

What is the university for?

² <https://pmc.ncbi.nlm.nih.gov/articles/PMC10830778/>

³ https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6097646

The University is not here to teach you how to make a shiny, well-polished report. It's to teach you how to think, how to learn, how to communicate. A University is where truth is gathered, stored, curated and disseminated for society. It's supposed to make graduates who are independent thinkers, responsible citizens, and who are "creators, evaluator and communicators of knowledge" (see [UCC Graduate Attributes](#)).

Assessment, including exams and assignments, is how we test your ability to meet learning objectives for your degree. The problem if you start relying on GenAI to replace the work of learning, or use it as a crutch, is not just that you will be cheating, **but that you'll stop doing the hard work of learning.**

Relying on GenAI for assignments harms study

Here is some evidence:

Bastani et al. (2025)⁴ — "unfettered access to GPT-4 can harm educational outcomes". Field trial, ~1,000 students, maths practice. Unrestricted GPT access raised *practice* grades by about 48% but cut *unaided exam* grades by 17%. A hint-giving "tutor" version raised practice grades by 127% but did not improve exam performance. "Without guardrails, students attempt to use GPT-4 as a "crutch" during practice problem sessions, and subsequently perform worse on their own. "

Stadler et al. (2024)⁵ "students using LLMs [for a short research task] experienced significantly lower cognitive load. However, despite this reduction, these students demonstrated lower-quality reasoning and argumentation in their final recommendations compared to those who used traditional search engines" ... "This outcome supports the findings of prior studies suggesting that a higher degree of interaction with diverse and sometimes challenging information—as often encountered in traditional searches—may promote better comprehension and processing of learning material... This interaction likely encourages students to engage more deeply with the content, thus enhancing their learning experience and leading to more sophisticated justifications in their conclusions."

Melumad & Yun (2025)⁶. Learning from AI summaries produces shallower knowledge than learning from search, because the model performs the synthesis you would otherwise have done. "those who learn from LLM

⁴ <https://www.pnas.org/doi/10.1073/pnas.2422633122>

⁵ <https://www.sciencedirect.com/science/article/pii/S0747563224002541>

⁶ <https://academic.oup.com/pnasnexus/article/4/10/pgaf316/8303888>

syntheses (vs. traditional web links) feel less invested in forming their advice, and, more importantly, create advice that is sparser, less original, and ultimately less likely to be adopted by recipients”

Kosmyna et al. (2025)⁷ used EEGs to scan the cognitive load of students writing essays with only their brains, with a search engine, or with an LLM. “Brain-only participants exhibited the strongest, most distributed networks; .. LLM users displayed the weakest connectivity. Cognitive activity scaled down in relation to external tool use.” “Over four months, LLM users consistently underperformed at neural, linguistic, and behavioral levels. These results raise concerns about the long-term educational implications of LLM reliance and underscore the need for deeper inquiry into AI's role in learning.”

There is also emerging evidence that professionals can become deskilled if they rely on AI as part of their job⁸.

So the costs are:

- Shallower work, if you're letting GenAI do the synthesis. This is the work of thinking.
- Reduced critical thinking
- Poor memory and ownership of your own work
- And outsourcing thinking robs you of the deep satisfaction of creating something yourself, like an idea or argument.
- Costs to your employability – employers will not want to hire people who just surrender their critical thinking skills to a tool. Your independent thinking, judgement, and oversight of anything produced by AI are going to be highly valuable in the future we're facing.

What you need to practise and not outsource

- **READING.** You can feed the PDF in and get a clean summary — I do this myself when I need to summarise and get specific information from a long report in my field. But you need to train your ability for reading actual books and long articles. Lose the attention span and you lose the ability to hold a complex argument in your head. Reading the writing of another human being is an essential part of learning, don't outsource this.
- **WRITING.** Writing is thinking. You can't delegate the brainstorming, self-reflection, reasoning, etc.

⁷ <https://www.media.mit.edu/publications/your-brain-on-chatgpt/>

⁸ <https://pubmed.ncbi.nlm.nih.gov/40816301/>

- **JUDGEMENT.** Distinguishing what matters from what doesn't – this is a critical skill to develop over time.

All of these skills require effort, in the same way that learning to play a musical instrument requires that you struggle with various skills, and mastering a sport requires physical exertion.

If you rely on GenAI to take shortcuts in these skills, your ability to think, to learn, and to judge will be compromised, and what's more, you'll end up dependent on these tools.

It is especially important that you do not build a dependency in the early years of your degree. Much of the assessment in this programme is continuous assessment, i.e., take-home assignments, where the temptation to use GenAI to produce your work will be present. But if you rely on GenAI to do the work of thinking for you in these early years, you won't be equipped to do the deeper learning and work of later years. These include your Capstone Project, Work Placement (no employer wants someone who just feeds their tasks into a chatbot), and modules with large final exams. Assessment is likely to become more "AI-proof", with oral exams, traditional closed-book exams, and the Q&A part of a presentation being weighted highly.

Also, your lecturers will probably have a good idea if you're handing up AI generated assignments. It's much easier for experts in the field to detect this. It's hard for us to formally prove and take cases against students handing up AI-generated work, so you'll probably get a fine grade, but the gap between your continuous assessment and your exam performance or in-person engagement will make clear what is going on. Good relationships with academics (and your peers) are essential for your career.

That said, GenAI is genuinely useful for some things.

I use GenAI myself – it is a very powerful tool, and it's becoming ubiquitous. But I went through all my education and most of my professional life without it, so I couldn't use it to take learning shortcuts. Here are some areas where I suggest it can be useful, based on my own experience.

- Nice-looking PowerPoints. Claude can do a decent job of turning a set of notes into a PowerPoint presentation. I have to make a lot of these in my job, and it's a drudgery. Having a tool do this is great. But if I need to make an important presentation, I do most of the work.
- Improving basic understanding of concepts, things I would otherwise Google "Explain what a TWh is and how it differs from a PJ."

- Producing reading lists. "Give me a reading list so I can understand X."
- Finding original sources — then checking and actually reading them. AI is better than Google for this, and it can do a decent first literature search.
- GenAI can be a useful critical reviewer of your own work. "Find the faults in my reasoning." "What are some unintended consequences of my proposal?"
- Coaching, not editing. Instead of "improve the language of this essay," say "coach me on how to improve the language of this essay — what principles should I follow?"
- Self-testing and quizzing.
- Read first, then unpack with AI the parts you didn't understand.

With any of these applications, you should not mainly rely on GenAI. For example, if you use it to quiz you on material, you may lose valuable cognitive work.

UCC has [a very useful learning hub for GenAI](#), where it includes examples of using the tool as a study aid. It also includes information on UCC's Academic Integrity policy, which students should be familiar with.

For most assignments, the following uses of GenAI are examples of cognitive offloading that are probably harmful to your education:

- Writing a first draft of anything
- Doing your brainstorming
- Synthesizing information
- Rely on it for factual claims – it hallucinates. Always verify the source.
- Building a dependence on it
- Use it directly to meet any of the learning outcomes of your coursework (so for example, if a LO is to build a slide deck, you should do it yourself!).

In summary, outsource defined tasks, use it as a learning aid rather than a shortcut, and if you find it's helping save time in doing some drudgery for you, then re-invest that time in some cognitively demanding task, like reading a book.

It's very important to note that this list of "dos and don'ts" is not based on a concrete pedagogical theory of GenAI's role in teaching and learning or a deep evidence based, it's mainly from my experience and intuition. We are at the

foothills of learning about this technology; our understanding of how its use supports or hinders learning (and other goals) will evolve. Disciplines will approach this differently; my list above may not be applicable to yours. There are documented cases showing evidence of positive educational effects, such as increasing accessibility.

Risks of dependence beyond your ability to learn

What is GenAI? This requires a much longer article, and I'm not an expert in the technology. In layman's terms (as far as I understand it), it can be understood as theory-free prediction. These models have been trained on an enormous amount of text from the internet and books. This is why they're called Large Language Models (LLMs). It has "learned" (through the training) the statistical patterns of text and what it produces is a kind of prediction. It doesn't think, it doesn't "reason" in the human sense – it very convincingly **simulates** the process of thinking.

It is "theory free" in that it's not trained on a specific subject or set of data, like earlier machine learning tools, which have been around for decades. LLMs are trained on **everything**. That has made it useful as a general-purpose tool, but also has made it enormously costly and inefficient.

It is enormously data- and energy-hungry. Some of the largest power plants in US history are currently being built to feed it⁹. Ireland is a particular hub for data centres (which are largely for digital services, rather than AI), but is a harbinger of what could happen if it continues to scale. Data centres in Ireland now place such strain on the electricity grid that they are building their own power plants, which could pose a major threat to our climate obligations¹⁰. Harms to local communities – from noise, water extraction, and air pollution – have become a prominent political issue in the US. The chips powering AI need to be replaced every few years, and some predict a wave of e-waste¹¹.

There is also a huge ethical problem in that the technology was trained on vast amounts of stolen data, like pirated books and articles, without paying

⁹ <https://globalenergymonitor.org/research/betting-big-data-centers-us-now-leads-world-new-gas-power-development>
<https://www.nytimes.com/2026/08/08/climate/amazon-data-center-texas-pollution.html>

¹⁰ Bonserio, T., & Daly, H. (2026). Data centre expansion in Ireland: the conflict between grid connection policy and national energy and climate targets. Working Paper. Zenodo.
<https://doi.org/10.5281/zenodo.20585450>

¹¹ https://wiki.ban.org/images/1/16/AI_Waste_Wave_Part_1.pdf

the original creators¹². It could end up undermining many creative industries to such an extent that new “content” stops being created. Many people object to the technology, and opt out of using it, for this fact alone.

There are also risks to individuals from relying on it, beyond impairing their ability to learn. There are very harmful features. For one, chatbots tend to be sycophantic. They generally agree with you, a seductive feature. There are alarming cases of people suffering from “AI Psychosis”¹³ and entirely losing their grip on reality.

There are also major security and economic risks from the industry, which are too big to get into here. Speculation that AI could literally kill off the human race is probably (in my inexpert view) over-hyped. Climate change poses a much larger, concrete threat to the planet, and I’m constantly disappointed that AI has distracted from this. That being said, the implications of this technology could be profoundly harmful to global stability, security and development.

It’s also worth knowing that it could get worse and more expensive – it’s currently heavily subsidised¹⁴, so we’re probably in the “golden early days” like social media was before it became enshittified¹⁵. It’s also now relatively unbiased, but that could easily change (again, see how certain social media platforms reacted to the political views of owners¹⁶).

Another ethical concern is the fact that LLMs have largely been trained on Western media, and therefore are unlikely to represent knowledge outside of these norms. I am also conscious that unequal access to premium/paid GenAI may place those on lower incomes at a disadvantage.

All that is to say, it’s entirely justifiable to opt out of using this technology entirely, on any one of these many grounds. There are many reasons why this technology should not be forced on any student.

As a student, protect yourself:

¹² Anthropic settled a €1.5 billion lawsuit with authors in 2025

¹³ https://en.wikipedia.org/wiki/AI-induced_psychosis

¹⁴ <https://ssir.org/articles/entry/low-cost-ai-illusion-nonprofits>

¹⁵ <https://en.wikipedia.org/wiki/Enshittification>

¹⁶ I recommend the book “Careless People: A Cautionary Tale of Power, Greed, and Lost Idealism: by Sarah Wynn-Williams as a shocking insider account of how Facebook cynically degraded its platform in pursuit of growth and profit, with profound global consequences. I understand the book “Empire of AI” by Karen Hao points to similar themes in the GenAI industry – it’s on my reading list.

That being said, I have somewhat of a pragmatic perspective. From using these tools, I can see their enormous power, and I believe that the technology will probably only grow in use. So our job should be to make sure we harness the benefits and try to make sure that the harms are minimised (such as making sure the data centres powering GenAI do not become a major new source of fossil fuel demand, which unfortunately seems to be happening).

As for avoiding educational harms, here are my final thoughts:

Getting a university education depends on doing things that are cognitively difficult, and GenAI presents an unprecedented opportunity to produce the typical outputs of learning (assignments, reports), while shortcutting many of the essential cognitive processes through which learning happens (reading, thinking, editing, writing).

Don't use it to cheat. Have integrity. Don't hand up work that's not yours. Your lecturer will probably prefer to see your messy, complex paper that shows that you struggled with a topic rather than a polished, generic report¹⁷. If you use it, declare its use according to module guidelines.

If you use it, do so to reduce mundane tasks that you are not required to master, use it to challenge you and push you further, and spend your freed-up time on something else cognitively demanding. Offload what you have already mastered, or do not need to master.

Read – don't lose the ability to read long-form, complex articles and books.

In summary, don't let it do your thinking for you. **Making your own meaning from information is an essential part of being an independent thinker.** Use it to amplify your learning, not to take shortcuts.

¹⁷ That being said, I'm also wary of an over-correction – educators can't assume that something that seems AI-produced is not mainly the product of a person's hard work, with the grammar polished.