HOW TO USE AI

INTRODUCTION

Artificial Intelligence is advancing very rapidly right now, and nobody is sure what effect it will have on scientific research.

Some people believe that the current forms of AI (large learning models/generative AI) are very useful at doing tasks that they have been trained on, but will never be able to do truly original ground-breaking research.

Others disagree and think that all scientists will be replaced by millions of Einstein-level AI genius researchers in data centres.

My personal opinion (this is Paul Francis writing) is that AI can currently do about 30% of what researchers do. It's great at literature searches, helping you learn unfamiliar topics, doing algebra and writing computer code. It can certainly help us do research a little faster. In the next few years that will increase. Now let's imagine that (say) ten years from now, AI can do 90% of what researchers currently do.

If we have the same number of researchers, that means that science can go ten times faster! Each human researcher will have a team of AI assistants working with them. Humans and AIs will each specialise in what they do best, and working together, they can get 100 years of progress in only a decade.

THE SITUATION TODAY

Whatever you think about the future of AI, it is certainly true right now that:

- AI can do pretty much any homework question we can set, better than most students.
- But AI cannot currently do cutting-edge research.

If we let students use AI to do all their work for them, then we could end up with a generation of science graduates who know no science - and science research will grind to a halt. So we need to make sure that our students really can do science themselves.

In courses, we use in-person assessment, such as workshops, labs and exams, to make sure that you can actually do the science yourself.

You can use AI to help you with learning the content of the course. In fact, we encourage you to do so. But you must use it to help you understand the material, not just to do the work for you. Otherwise you will fail the exams and hence fail the course.

GOALS FOR AI USE IN THIS COURSE

We have two goals for AI use in this course.

1. AI as a learning assistant

There is now considerable evidence that if you use AI as a person learning assistant, it will help you learn physics faster and better. A recent study of first year Physics students at Harvard showed that students who used AI correctly - as a learning buddy, did about 15% better than those taught in normal ways.

We hope to help you learn how to use AI to improve your learning.

2. Keeping Track of AI

Nobody knows where AI is going - but everyone knows it will be very important. So we want you all to be using AI a lot, and keeping a close watch on how it is changing. So you will be ready to use it in the most effective way possible.

WHICH AI TO USE?

There are many AI models out there, with new and better ones released every few days. You can use whatever model or models you prefer. The following notes are correct as of June 2025 but will go out of date very soon!

- **Microsoft Copilot:** https://m365.cloud.microsoft/chat/ This is the AI model that the ANU likes students to use because it is secure and doesn't use your data to train models. Basically, this is ChatGPT 4. You get access to the paid version of this for free if you log in using your ANU credentials.

- **OpenAI ChatGPT:** https://chatgpt.com/ : The most famous and widely used set of AI models. They currently have a wide range of models with varying capabilities. Some require paid subscriptions.

- **Anthropic Claude:** https://claude.ai/ : Another very widely used model. Once again there are different versions, some requiring paid subscriptions. Excellent general-purpose model and widely thought to be the best for writing computer code.

- **Google Gemini:** https://gemini.google.com/ another excellent general-purpose model, once again with different subscription levels. Offers family subscriptions.

- **Google NotebookLM:** https://notebooklm.google.com/ You to give it notes and readings (e.g. multiple pdf papers), and it can turn them into a podcast for you, summarise the material, ask you questions about it and have a discussion about it.

- **DeepSeek** (https://www.deepseek.com/) and **LeChat** (https://chat.mistral.ai/chat): recently released models that people say are seriously impressive and free to use.

- **xAI Grok** (x.com), built into x (former Twitter). The just-released Grok 3 is seriously powerful.

- **Perplexity** (https://www.perplexity.ai/): Good for searching the web and giving sources for what it finds.
- **Cursor** (https://www.cursor.com/): AI code editor for programming.
- And many others...

Right now the models can be divided into three classes.

- **Standard Large Language Models.** Trained on vast amounts of information. Good for text. Not always so good for complex physics problems.
- **"Deep Thinking" models.** These will go through several iterations of thinking before giving you an answer. So they check their own conclusions, double-guess themselves, argue with themselves etc, before giving you the final answer. Much better for complex multi-step physics problems, but take longer to run.
- **"Deep Research" models.** These will go away and using a combination of web search and their own reasoning, come up with a comprehensive literature analysis of your requested topic.

The main AI labs say that they plan to merge all these different types of model together - so you will see a single model that will do deep thinking or deep research when appropriate. But for the moment they are separate models.

DEEP THINKING MODELS:

These are the models which think through a problem in many steps. Can dramatically improve the performance for complex multi-step physics/maths/computing problems.

- Deepseek is the best known and is free to use. It shows its reasoning as it works through a problem.
- ChatGPT. The free versions are not currently deep thinking, but if you have a paid subscription you can use o3-mini, which is an excellent deep thinking model.
- Grok. It has a "Think" button which puts it in deep thinking mode. There are probably limits on how often you can use this without paying for a premium subscription.
- Gemini. Free account has an option to use "Flash Thinking Experimental", which seems to be a basic deep thinking model. With a paid subscription you can use the more deep thinking model.

DEEP RESEARCH MODELS

These models go away for some minutes and do a detailed literature review on some topic for you, combining web search with the model's intelligence. Sometimes they come up with a research plan first, ask you to approve or refine it, and then

execute the plan. They can summarise hundreds of papers and web pages to give you a coherent analysis.

- ChatGPT Deep Research. Probably the best, but can only be used a limited number of times unless you are paying $300 per month.
- Gemini Deep Research. Only available for paid subscriptions (but these are only $30 per month). Tends to give somewhat shallower analyses with lots of references.
- Perplexity Deep Research. New. Free to use (with limits) or subscribe for more use. Seems to be pretty powerful - better than Gemini for most purposes, but maybe not quite as good as ChatGPT?
- Grok Deepsearch. New. Free to use (with limits) or subscribe for more use. Seems to be very powerful: better-written responses than Perplexity but perhaps less deep in their analysis?

Personally, I currently pay for subscriptions to Gemini and ChatGPT.

Any of these models will do 90% of what you need to help you with this course, so use whichever you like. It is definitely worth playing around with a few different models to see which you prefer.

THINGS TO WATCH OUT FOR IN DIFFERENT AIS:

- Can the model understand images? With many models, you can input a graph, an image or even a photo of handwritten equations, and it will understand them. But some models can't do this.
- Speach Mode. Many models allow you to enter your prompts by talking. A few allow you do have a conversation where you talk to the AI and it talks back. Google Gemini and ChatGPT are good at this. Useful if you don't like typing, or prefer to learn with your ears more than your eyes.
- Import PDFs: some models allow you to import large PDF files and then ask questions about the contents or summarise them. The size limits you can import vary widely.

COMMON MISTAKES

These are the most common mistakes I see students make when using AI.

- **Don't use AI enough.** Most people don't use AI enough. The current AI systems, though flawed, are still amazing and can be incredibly helpful. But many people don't realise this, or perhaps they played with AI a few months ago and were unimpressed - not realising how rapidly the AI systems have improved.
- **Treat AI like a search engine** - give it a single request, like "explain buoyancy to me", or cut-and-paste a question into the chat box. The current AI systems work best with a back-and-forth conversation. Ask it something, look at the response,

then refine your question and ask for more details, tweak your request, or ask follow-up questions. Only then will you get good responses.

- **Prompts too short.** The more detail you give the AI, the better its response will be. See examples in the prompting section - but good prompts are often many paragraphs long. If you don't like typing so much, you can store your favourite prompts in a text file and copy them into the AI chat window. Or use voice mode and dictate.

- **Mix content in the same chat.** The AI systems remember what is in each chat. So if you change topics, start a new chat. If you are using AI to plan your dinner menu, and then decide to use it for your chemistry homework, start a new chat as otherwise the AI will get very confused!

- **Just give AI the problem you want it to solve.** Normally it works better if you talk the AI through the problem step-by-step. Ask it to outline its approach in words. Clarify anything you don't understand. Then get it to go through the solution, one step at a time.

- **What to do when the AI persists in something stupid.** I've seen cases where an AI has got completely the wrong end of the stick, but persists in telling you something completely wrong, despite your best efforts to correct it. In this case, best to start a new chat and try again.

PROMPTING ADVICE

A commonly used framework for prompting is RTRI: Role, Task, Requirements and Instructions.

- **Role:** “You are an expert astrophysics professor and prize-winning educator”

- **Task:** “Explain the observations that led to the discovery of the Hubble Law”

- **Requirements:** “Explain everything at a level appropriate for a first-year university science student, who is not very confident about maths. You should pose questions to test both my conceptual and mathematical understanding”

- **Instructions:** “Break up your explanation into steps, and after each step, pose me a question to test my understanding, and wait for my response”

Goal-Focussed Prompting

Some of the most recent AI models (ChatGPT o1,o3,o4, Deepseek, Gemini Flash Thinking) have the AI go through multiple steps of reasoning. Some people think you should prompt them differently.

- **Goal:** I want to understand how to use spectroscopy in a wide variety of astronomical contents.

- **Return Format:** Take me step-by-step through a one-hour lesson, stopping after every stage to ask me questions.

- **Warnings:** Be careful to explain the astrophysics correctly, and in a way that will help me learn how to solve exam problems.

- **Context Dump:** I'm a first-year university student studying an introductory astrophysics course. I am having trouble understanding spectroscopy – both what it is, but more importantly how to use it. I'm OK if given specific problems to solve, but have trouble with open-ended problems.

EXAMPLES:

1. Learning the Content

You can use AI to teach you the content, rather than watching the videos/reading the notes.

Here is an example prompt that works well for me:

"You are a university physics professor and a prize-winning educator. Please explain to me the physics of buoyancy, at a level suitable for a first-year undergraduate physics student. Break your explanation up into short steps, and test my understanding after each step. You should test both my conceptual understanding and my ability to do calculations, where appropriate."

The AI will explain the content to me bit by bit, and ask questions at every stage. I can try to answer the questions - if I get them wrong, the AI will explain my mistake. If anything in the explanation doesn't make sense, I can ask the AI to clarify or tell it that I disagree.

At the end of my chat I ask:

"Looking back over the conversation so far, how would you rate my learning approach?"

If you want make sure that the AI instruction covers all the same points as my lectures, you can type something like

"I would like you to teach me about the contents of this attached document."

and then upload a pdf of the notes, or copy-and paste the notes into the chat window. Or copy the video transcript into the chat window. You can get it to go through any file or set of notes step by step, explaining, elaborating and probing your understanding with questions.

2. Homework Problems

Imagine you are doing some homework and you have trouble with a question. You can type something like:

You are an expert physics professor and prize-winning educator. I'm a first-year physics student and I'm having trouble understanding how to do the following homework question. Please carefully explain, step-by-step, how I can solve this problem. Stop after each step and ask me questions to prove whether I'm understanding so far. Spell out all the mathematical steps in detail. (paste the problem into the chat window)"

It will start giving you an explanation. Stop and argue with it if you disagree with any part of the explanation, or if anything is not clear.

At the end, type something like

"OK, I think I understand now. Please ask me a similar question that tests the same knowledge, so I can see if I really understand this material".

or

"What do you think was the biggest difficulty I had with this problem? Do you have suggestions for how I could improve my ability to do similar problems in the future?"

You can also ask the AI to give you a hint, rather than just give you the answer to a question.

3. Exploring New Topics

If you are investigating something new, it is often worth asking the AI a very general prompt to begin with and then focussing in steadily on interesting topics.

"I am interested in figuring out theoretical limits on the capabilities of submarine drones. I would like to focus on what limits are placed by the laws of physics, and how close current drones are to these limits - i.e. how much room there is theoretically for improvement. Please outline the possible directions this research might go in."

This is particularly useful when writing computer code. Rather than just ask for your code straight away, start off by asking for an overall outline, and then get the AI to flesh in the details:

"I would like to write computer code to simulate the motion of a projectile exposed to realistic drag and wind. The simulation should run in a web browser, and allow users to choose the projectile properties and wind speed. What would be a good way for me to proceed?"

4. Getting the AI to probe deeper.

Often the AI will give you a very shallow and generic consensus response to a prompt, but you want it to think critically and probe below the surface. This can be quite difficult to achieve. You can just ask the AI to think critically, but there are other techniques you can use, like this:

"Using a probabilistic approach, weigh up the arguments for and against nuclear power being more expensive than firmed renewables for Australia, ranking the importance of each argument, and give an estimate of the probable price difference between nuclear and firmed renewables."

5. Writing Styles

You can ask the AI to write something in a very specific style. This is a good way to learn good styles yourself.

"What is the best practice style for writing ministerial briefs?"

"You are an acclaimed public servant, with great expertise at writing deep and insightful ministerial briefs. You have been asked to brief the energy minister about whether it makes economic sense for Australia to develop nuclear power, as compared to firmed renewables. The minister has heard contradictory statements about the likely price of nuclear power compared to renewables, and wants to know the truth. Write a brief for the minister on this topic."

APPENDIX: ENTERING MATHS

The AIs understand maths in many formats. You can use something like

y = 4 x^2 - sin and it will be understood.

You can use LaTeX format - which is also the format that the AIs typically output.

So something like 4 \pi r^2 \sqrt(x^2-2) will be understood.

For the AIs that understand images (e.g. ChatGPT) you can handwrite equations and upload a photo of them.