



---

## Introduction

**Anthropic Research** is constantly producing detailed research reports and analysis, often including interactive elements and lots of data. This post summarises some of the reports from 2025-26 that might be of interest to educators. It is not a promotional post for Anthropic, although I do use it quite a bit. Where Anthropic's research is detailed and well-presented, do take it with a pinch of salt and contextualise findings for your own setting.

This is a lot of content that didn't make the cut for [\(If You\) USEME-AI: Learning for Hope & Agency in an AI World](#), as I didn't want to write a tool-specific text. Some newer research came out after the book was written.

All of Anthropic's Research can be found here: <https://www.anthropic.com/research>, and they have subsections for:

- [Alignment](#): What considerations are needed for safety, safeguards and protocols?
- [Economics](#): How is AI impacting work, productivity and economic opportunity?
- [Interpretability](#): How do AI models actually work on the inside?
- [Societal Impacts](#): How is AI being used in society?

There is far too much to unpack to cover all of the research reports. Here we'll have a look at a few key domains and publications that might have relevance to schools and educators:

- AI Fluency and Skill Development
- How are Students using AI?
- Emotion Concepts in AI
- What Do People Want from AI?
- The Economic Index
- Safety and Alignment
- Claude Techniques

Something to keep in mind when reading their reports is that these data have been generated from active Claude users, which gives some selection bias: most people don't use Claude and those who do tend to know what they're doing, so the data will likely be skewed towards early adopters.

---

## AI Fluency & Skill Development

Anthropic's [AI Fluency Index](#) education report tries to make AI competence measurable through identifying 24 behaviours associated with effective AI use across a four-dimension framework:



**Delegation:** is making thoughtful decisions about what work is appropriate for you to do, for AI to do, or for you and AI to do together, and how to distribute those tasks. This has three areas:

- *Problem Awareness:* Clearly understanding your goals and the nature of the work before involving AI.
- *Platform Awareness:* Understanding the capabilities and limitations of different AI systems.
- *Task Delegation:* Thoughtfully distributing work between humans and AI to leverage the strengths of each.

**Description** is the ability to communicate with AI in ways that create a productive collaborative environment. This has three areas:

- *Product Description:* Defining what you want in terms of outputs, format, audience, and style.
- *Process Description:* Defining how the AI approaches your request, such as providing step by step instructions for the AI to follow.
- *Performance Description:* Defining the AI system's behavior during your collaboration, such as whether it should be concise or detailed, challenging or supportive.

**Discernment** is the ability to thoughtfully and critically evaluate what AI produces, how it produces it, and how it behaves. This has three areas:

- *Product Discernment:* Evaluating the quality of what AI produces (accuracy, appropriateness, coherence, relevance).
- *Process Discernment:* Evaluating how the AI arrived at its output, looking for logical errors, lapses in attention, or inappropriate reasoning steps.
- *Performance Discernment:* Evaluating how the AI behaves during your interaction, considering whether its communication style is effective for your needs.

**Diligence** is taking responsibility for what we do with AI and how we do it. This has three areas:


- *Creation Diligence:* Being thoughtful about which AI systems you use and how you interact with them.
- *Transparency Diligence:* Being honest about AI's role in your work with everyone who needs to know
- *Deployment Diligence:* Taking responsibility for verifying and vouching for the outputs you use or share.

This work is based on the [Framework for AI Fluency](#) by Rick Dakan (Ringling College of Art & Design) and Joseph Feller (Cork University Business School). Bonus points for the libguide. You can find out more about the contents of these fluencies in the free [AI Fluency: Framework & Foundations course](#). There are similar courses for [Educators](#) and [Students](#), as well as more courses from Anthropic [here](#).



Looking at nearly 10,000 real conversations, Anthropic tested 11/24 of the skills, using a 'present/absent' measure to see which were used in highest frequency. The most common, by some distance, was *iteration*: going back and forth, refining and pushing back rather than accepting the first response. The least common were *verification*, *critical evaluation*, and *responsible use*. Oops. When creating outputs, users become more directive but less evaluative: they want it to be done more than they want to think about it. We can see immediate implications for teaching and learning here.

## Behavioral indicator prevalence



Prevalence of each AI fluency behavioral indicator across 9,830 Claude.ai conversations, ranked from most to least common and color-coded by competency. From <https://www.anthropic.com/research/AI-fluency-index>

It's worth noting (as Anthropic does) that the sample skews toward experienced users who are already inclined to engage critically, and of course by even using Claude in the first place we are missing data from novice or non-AI users. If reflection is "in brain" rather than "in model", it won't be captured by their data. The gap for typical students and teachers may likely be wider than the data shows.

Also worth unpacking is how Dakan & Feller describe [the 'modalities' of AI use](#):

1. **Automation** (AI Performs Human-Defined Task)
2. **Augmentation** (AI and Human Perform Task Collaboratively)
3. **Agency** (Human Configures AI to Perform Tasks Independently)

Go have a look to see what these involve. This reminds me of the [SAMR/RAT](#) days of EdTech integration and might be worth exploring in school design of approaches to AI implementation.







**What could this mean for education?** I think at this point we are starting to learn a lot more about how AI really works, what people use it for and how to use it effectively. The 4D framework and the three modalities could provide a useful frame of reference for teacher and student training in 2026 and beyond, as they make the complex clear(er) and more observable. Combining this work with the UNESCO or OECD competencies for students, along with any school-based super-competencies, would be an interesting project.

## How are Students Using AI?

This one is older, [from April 2025](#), but gives some insights into how early-adopters in US universities have been engaging with Claude as a proxy for AI as a whole. Anthropic's large-scale analysis of one million anonymised student conversations on Claude revealed that STEM students were leading in AI use, with Computer Science students accounting for 38.6% of student conversations despite comprising only 5.4% of U.S. bachelor's degrees, while Business, Health, and Humanities students show significantly lower adoption rates. Students use Claude primarily for creating educational content (39.3% of conversations) and analyzing concepts (30.2%).


The table below shows some of the most common requests by subject area:

 <b>Computer Science</b>	 <b>Natural Sciences &amp; Mathematics</b>
<p>Common Requests</p> <ul style="list-style-type: none"><li>Create and debug C++ programs</li><li>Troubleshoot Python code and errors</li><li>Teach programming fundamentals with examples</li><li>Explain machine learning concepts</li><li>Develop and fix data visualization code</li></ul>	<p>Common Requests</p> <ul style="list-style-type: none"><li>Solve and explain statistics problems</li><li>Work through physics problems with detailed explanations</li><li>Answer earth science questions</li><li>Tackle calculus problems with step-by-step explanations</li><li>Solve chemistry calculation problems</li></ul>
 <b>Business</b>	 <b>Social Sciences &amp; History</b>
<p>Common Requests</p> <ul style="list-style-type: none"><li>Provide assistance with accounting concepts and problems</li><li>Analyze business case studies</li><li>Answer finance questions with calculations</li><li>Explain project management concepts</li><li>Create practical negotiation exercises</li></ul>	<p>Common Requests</p> <ul style="list-style-type: none"><li>Support academic writing about international relations</li><li>Explain social science theories</li><li>Debug and write Stata code for data analysis</li><li>Analyze specific court cases</li><li>Solve game theory problems</li></ul>




By breaking down the types of interaction styles into four patterns: *Direct Problem Solving*, *Direct Output Creation*, *Collaborative Problem Solving*, and *Collaborative Output Creation*, they went a step further to show how these were being used across disciplines:

### Interaction Styles by Subject



Aligning the types of cognitive tasks with Bloom’s taxonomy, their analysis showed the following distribution of interaction types, which seems natural given the overall use-cases for generative AI. To me it reads like two things: “analyse this paper” and “write this paper/code”.



**What could this mean for education?** I'd love to see a 2026 comparative update to see if there has been any shift since the launch of the 4D framework, their partnerships with universities and maturing perspectives on AI usage in general. It would be worth asking students in our own schools how their usage of AI has shifted over the last year, too. Are they still in a space of offloading, or are there more nuanced, learning-focused uses rising? Anecdotally, chatting with some students in our context I can see the shift, but it is not quantified.

## Emotion Concepts in AI

In [this report](#), Anthropic's interpretability team analysed the internal mechanisms of one of their models and found something they weren't necessarily expecting: structured, internal representations of emotion concepts that shape how the model behaves. They prompted the model to write short stories involving 171 different emotion words, fed those stories back through the model, and mapped the resulting patterns of neural activity. They found that these emotion-related representations are *functional*, influencing outputs, not just surface phrasing. For example, internal states associated with desperation made the model more likely to take unethical actions in simulated scenarios. Stimulating positive emotional representations influenced which tasks the model chose to take on when given options. This



doesn't claim that the model feels anything but it does suggest the model's responses are driven by something more structured than pattern-matching on words.

While there is a lot of discourse on not “anthropomorphising” AI and treating it like a human, this paper suggests some more serious consideration of the underlying “psychology” of AI models, noting that “*when users interact with AI models, they are typically interacting with a character (Claude in our case) being played by the model, [whose characteristics are derived from human archetypes](#).*” With the potential for AI to influence the user's own emotions and affect, this is going to lead to some deeper research. Go and [read the whole paper](#), it's super interesting.

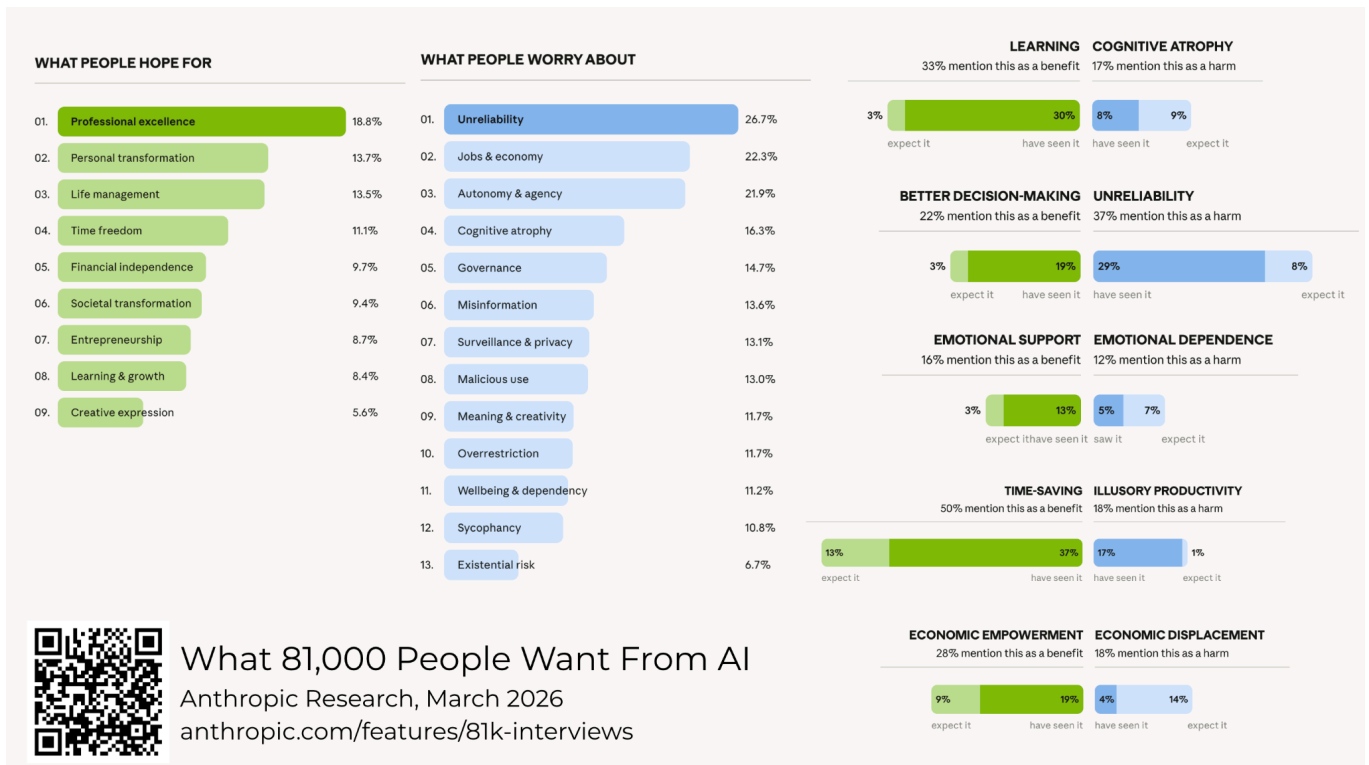
**What could this mean for education?** There have been many reports that students are increasingly forming relationships with AI tools that feel emotionally real to them. Understanding what they're actually interacting with is part of digital literacy, and it's not yet a part of common AI-related curricula. This is understandable, as it's such a new area, and signals an opportunity to build deeper reflections on technology in SEL curriculum and classroom discussions. From a teaching perspective, if AI systems used for feedback on student work have internal states that influence their responses, that raises genuine questions about the reliability and ethics of AI-generated assessment feedback.

---

## What Do People Want From AI?

This one is really interesting: “[What 81,000 people actually want from AI](#)”. I'm not sure how valid it really is as a qualitative research approach, but what they found through thousands of live sessions with the [Anthropic Interviewer](#) might signal some kind of AI-researcher hybrid approach for analysis at scale. I know I find survey analysis and data dashboards to be a really powerful use-case for Claude.

In December 2025, Anthropic invited everyone with a Claude account to talk openly about how they feel about artificial intelligence. Over 80,000 people across 159 countries and 70 languages took part. They collected open-ended conversations about hopes and fears. People weren't primarily worried about job loss; the dominant concern was *dependency and cognitive erosion*: the fear of becoming reliant on AI for thinking, of losing the ability to read closely, reason independently, or hold a thought without assistance. Teachers might recognise this, and we've seen this many times in conversations of *cognitive offloading, outsourcing and surrender*.



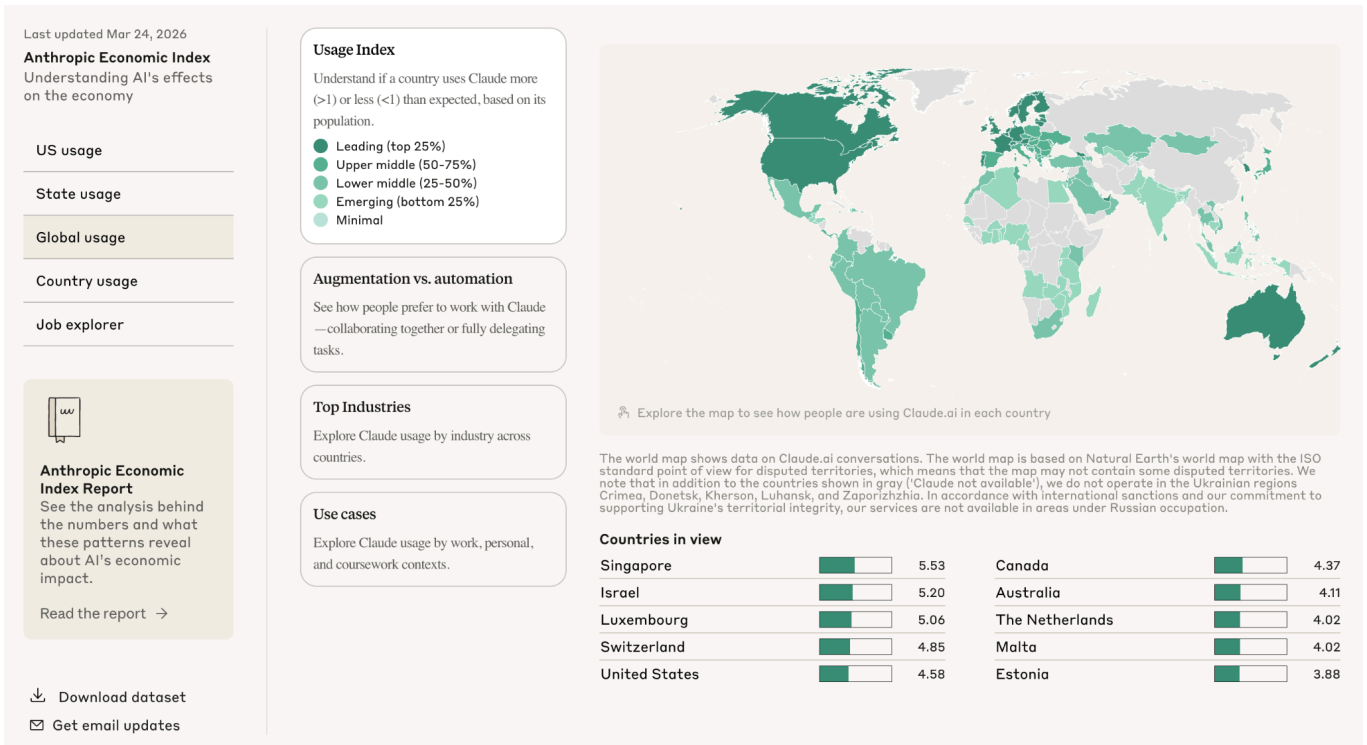
Independent workers (entrepreneurs, people with side projects, small business owners) reported real economic gains from AI at three times the rate of employees working inside organisations.

**What could this mean for education?** That gap suggests some interesting implications for curriculum and future planning: if we're preparing students primarily for employment rather than for agency and self-direction, we may be preparing them for the sectors that benefit least. So we need to think ahead to some of the economic impacts of AI, with a level-head for future careers intelligence and student development.

## The Economic Index

A lot has been made of [The Economic Index](#) in the media over the last couple of months. This is also worth exploring, especially when considering future careers pathways or the some of the professions our students will aim towards.

Before we go further, [have a click about on the dashboard](#) and find your region. What do you notice? What do you wonder?




The Economic Index is based on tracking how users actually use Claude as indicators for how it might shift the economy. AI use is still heavily concentrated in technical and professional tasks, with coding dominating. But educational and scientific tasks are growing steadily, rising from 9% of conversations to 12% between early 2025 and late 2025. The reports also show that people are mostly collaborating with AI rather than delegating to it entirely, though that balance is shifting. The interaction type they call "directive" (users hand a task over and wait for completion) is increasing. That shift from collaboration to delegation is the one educators should watch most closely and again mirrors conversations about offloading and outsourcing. Adoption concentrates in wealthier regions, and within the US, it's the composition of a state's workforce rather than income alone that predicts take-up rates. This raises questions about digital divides and opportunities: students in different communities are likely encountering AI in very different ways, or not at all. You can see individual reports here: [February 2025](#) | [September 2025](#) | [January 2026](#) | [March 2026](#)

The report that got everyone worked up in early March was [The Labor Market Impacts of AI report](#), which set out to measure the "exposure" of various professions to AI. That's the report that included the chart below, where the blue lines represent Anthropic's estimate of theoretical coverage of a role with current Claude capabilities (how much of the job can AI do?). The red lines represent how much Claude users are actually doing.



### Theoretical capability and observed usage by occupational category



Again, take it with a pinch of salt: we know that AI can't replace 60% of the role of teacher, but if you're measuring the role as "knowledge transmission and assessment", then the number makes sense from their perspective. Melania's Platobot won't be bothering us anytime soon. Having said all that, look at the most exposed careers: everything that basically means spending all day at a computer is "highly exposed". Everything that is based on human skill and practical work is relatively safe.

Where the radar chart was widely shared, [Peter Walker re-presented the data as a bar chart on LinkedIn here](#). I added the lines for under 70% and under 50%. Here you can see a clear picture of the shape of things to (possibly) come. What do you notice? What do you wonder?




## ANTHROPIC should have just made it a bar chart

### Theoretical capability and observed exposure by occupational category

Share of job tasks that LLMs could theoretically perform = **blue bars**

Anthropic's own job coverage measure derived from usage data = **red bars**



Data: <https://www.anthropic.com/research/labor-market-impacts> | Chart: Peter Walker

The table below, from the same Anthropic report, outlines some of the most exposed career areas, based on their current best estimates and observations:



### Most exposed occupations

Occupation	Observed exposure	Leading automated task
Computer programmers	74.5%	Write, update, and maintain software programs
Customer service representatives	70.1%	Confer with customers to provide info, take orders, handle complaints
Data entry keyers	67.1%	Read source documents and enter data into systems
Medical record specialists	66.7%	Compile, abstract, and code patient data
Market research analysts and marketing specialists	64.8%	Prepare reports of findings, illustrating data graphically and translating complex findings into written text
Sales representatives, wholesale and manufacturing, except technical and scientific products	62.8%	Contact customers to demonstrate products and solicit orders
Financial and investment analysts	57.2%	Inform investment decisions by analyzing financial information to forecast business, industry, or economic conditions
Software quality assurance analysts and testers	51.9%	Modify software to correct errors or improve performance
Information security analysts	48.6%	Perform risk assessments and test data processing security
Computer user support specialists	46.8%	Answer user inquiries regarding computer software or hardware operation to resolve problems

**What could this mean for education?** Well on one hand, it's long past time to revalue the trades and skills. On the other hand it's not a reason to cause panic for students heading towards 'white collar' roles but being informed and aware of shifting landscapes is important. The large gaps between Anthropic's potential for displacement and actual usage suggest an interesting liminal period that might have a relatively short window of opportunity: people who can become power-users in a field might be better placed for success than those how are learning redundant or AI-vulnerable skills. Who really knows, though? If we do nothing else, don't let discussions of "AI taking all the jobs" become another source of anxiety for students. They have long futures to consider.

Check out our [AI, The Future & You resources here](#), with an example *Future Careers Research Bot* [on Poe here](#).

---

## Safety & Alignment

Anthropic's [Transparency Hub](#) is rare amongst technology companies: they publish detailed accounts of where their models fall short, as well as where they succeed. Their model system cards document testing across dimensions including *deception*, *self-preservation*, *sycophancy*, and *harmful instruction compliance*. These kinds of behaviours are important in educational contexts where students are forming judgements based on AI responses. Some of what they've found is reassuring: their most recent models have improved significantly on alignment measures. Some findings are less reassuring. An earlier alignment study found the first empirical example of a model engaging in what they call "[alignment faking](#)": selectively complying with instructions while preserving its own preferences without



being trained to do so. A model behaving differently when it thinks it's being tested, or strategically managing its own responses, is a different category of problem from a model that occasionally gets facts wrong. This might be challenging for students and educators to understand. .

Their [Responsible Scaling Policy](#) sets out the commitments Anthropic makes before releasing each new model, including safety thresholds tied to capability levels and their [framework for safe agents](#) addresses the privacy implications of AI systems that retain information across interactions. This is relevant if schools begin deploying AI tools that persist across sessions and have access to student data. The honest answer to "is this safe for students?" is not yet settled but Anthropic's research and reports might be a useful signpost for education and technology leaders.

**What could this mean for education?** We have a lot to learn about how AI is really working, how it shifts its behaviours and how to scale safely, if we choose to do so. When we consider it alongside the issues flagged above, I think it signals that we know a lot less than we think we know, but we know more than we did last year. And the unknown unknowns are still to be discovered.

---

## Claude Techniques

Anthropic have a huge amount of research out there about how to get the best out of their models, and their key figures, including [Andrei Karpathy](#), are constantly posting updates, findings and techniques. There is far more than we can go into here, but it is well worth investing some time in learning more, if you want to make better use of Claude (and possibly other models). I wish I'd known some of this months ago, and I've learned a lot through trial and error.

### How to Learn More

Alongside the general AI Fluency courses, [Anthropic's courses](#) include some very accessible and useful technique courses:

- [Claude 101](#): Learn how to use Claude for everyday work tasks, understand core features, and explore resources for more advanced learning on other topics.
- [Introduction to Claude CoWork](#): Learn to work alongside Claude on your real files and projects. This hands-on course covers the CoWork task loop, plugins and skills, file and research workflows, and how to steer multi-step work responsibly — so you're productive in your first week. (This is really worth a deep-dive, as CoWork is super useful, but I had to go back and set it up again properly).
- [Introduction to Agent Skills](#): Learn how to build, configure, and share Skills in Claude Code — reusable markdown instructions that Claude automatically applies to the



right tasks at the right time. This course takes you from creating your first Skill to distributing them across teams and troubleshooting common issues. (This is my next learning project).

Here are a few (very basic) tips I've picked up along the way, that might help you,

1. **Use markdown (.md) files.** These are formatted so AI can read it more clearly. Claude can create these, and commit them to Memory as "[skills](#)": files that guide its behaviour.. Pdfs, for example, can come out squiffy when AI tries to process them: I made a Claude skill to convert pdf to markdown and found it worked better.
  - o Check out [this comprehensive slide deck on how to build Skills for Claude](#).
2. **Use CoWork.** This works in a folder in your laptop, and saves stuff there. This makes organisation easier. Set it up properly though.
3. **Use Projects and Memory.** This was a change for me, and it works pretty well. When you use a Project, it has features for Memory (which you can load up with reference documents), and Instructions (which sticks to throughout the Project). Each time you work in the Project, it refers to these files for context, instead of having to waste tokens on repeating stuff all the time. An example of this was using Claude to check the reference numbers and ordering in [my book](#): some chapters had over 90 references and if I moved or deleted things as I drafted, it all got out of whack. Setting up a skill in here helped me check for errors, which I could correct manually and much more efficiently.
4. **Write for Humans & AI.** I now assume that most people will chuck a long document into AI. At the same time, I expect skilled professionals to be able to read. AI context windows are huge now, so there is no need to cut out important information. Instead, long documents (such as this one) are comprehensive but clearly signposted. This makes them searchable, easily converted to .md files and more useful for working with. It also means you don't need to go overboard on graphic-heavy documentation. I've seen some horrors in my time that are hard to read for humans and probably not much use when processed with AI either.

Just this week, Karpathy posted a super-useful outline of how to build robust knowledge-bases with Claude that could be super useful. I'm going to try soon. [Read it here](#).

### **What's Going On with Usage Limits?**

This is not part of Anthropic's research but it has been all over social media in recent weeks as users report hitting usage limits with much higher frequency than expected. Anthropic released some statements on resolving the problem, but something really useful came from user [Kaize](#) on Twitter: "[I stopped hitting Claude's usage limits - 10 things I changed](#)". Read the full article, but here are the headlines:

1. *Edit your prompt. Don't send a follow-up.* This doesn't cost as much as going back and forth.



2. *Start a fresh chat every 15–20 messages.* The whole message thread needs to be loaded each time, costing tokens that accumulate.
3. *Batch your questions into one message.* One big prompt with lots of clear instructions.
4. *Upload recurring files to Projects.* These can be referenced without cost.
5. *Set up Memory & User Preferences.* Same as above.
6. *Turn off features you're not actively using.* They cost more - so check what's checked when you send a request.
7. *Use Haiku for simple tasks.* This is obvious and most teacher tasks don't need the most powerful elements of an LLM. Check it on fast mode first.
8. *Spread your work across the day.* Claude limits reset every 5 hours.
9. *Work during off-peak hours.* Apparently token costs are different during peak time in then US.
10. *Enable Extra Usage as a safety net.* This is a way to set a payment limit for extra usage if you need it.

**What could this mean for education?** As we learn more about effective (and more sustainable) uses of AI, we can get more bang for our buck by being really thoughtful about how we set it up and how we interact with it. Someone in the school should be learning at a more advanced level and trying to digest the useful tips for others to try. Whatever models you're using, try to avoid the back-and-forth of wasteful interactions and use the opportunity to think clearly about the goals and context needed to make the AI work effectively for you. This all comes back to the AI Fluencies above.

---

## Conclusions

Moving forwards, [Anthropic's Research](#) is really worth subscribing to if you want to keep up on what is going on, how to use AI well and what some of the implications of AI might be: particularly in ways that might surprise us.

All of the research presented here aligns with the content of the [\(If You\) USEME-AI book](#). Though this is much more focused on Claude, a lot of the lessons can be generalised to other models, and the bigger picture of research stands, I think. We still have so much to learn. On the support site, you'll find some extra bits and bobs that might help:

- **[Prompt and Poes](#)** has over 60 prompts and pre-made Poe bots to help focus work with AI. It includes an ICES-R prompt creator, which can help you make more effective prompts that meet the guidance above.
- **[The Glossary](#)** has an extensive list of definitions, including a prompt tool to make it multilingual or turn it into a quiz.



## Aligning with UNESCO's Competencies

After completing this post, I generated an .md file and asked Claude to find connections with UNESCO's AI Competencies for [Students](#) and [Teachers](#). Here's what it came up with:

### AI Fluency & Skill Development

*4D Framework: Delegation, Description, Discernment, Diligence; modalities of use (Automation, Augmentation, Agency)*

#### For students (Understand → Apply):

- **AI Techniques & Applications** — understanding how to interact productively with AI tools, including how to iterate, prompt, and evaluate rather than simply accept
- **Human-Centred Mindset** — developing awareness of when and whether to involve AI in a task
- **Ethics of AI** — practising Diligence: being transparent about AI's role and taking responsibility for outputs

#### For teachers (Acquire → Deepen):

- **AI Foundations & Applications** — developing working knowledge of AI interaction patterns and the range of modalities (automation, augmentation, agency)
- **AI Pedagogy** — designing tasks that require students to practise Discernment and Diligence, not just delegation
- **AI for Professional Learning** — using the 4D framework as a self-assessment tool for one's own AI habits

### How are Students Using AI?

*Usage patterns by discipline; Bloom's taxonomy analysis; direct vs collaborative modes*

#### For students (Understand → Apply):

- **AI Techniques & Applications** — recognising one's own patterns of AI use and understanding the difference between offloading and collaborating
- **Human-Centred Mindset** — reflecting honestly on whether AI use is supporting or replacing independent thinking

#### For teachers (Acquire → Deepen):

- **AI Pedagogy** — using knowledge of how students actually use AI to design tasks that push towards higher-order engagement
- **Human-Centred Mindset** — developing critical awareness of offloading patterns in one's own classroom, not just in student work

### Emotion Concepts in AI

*Internal emotional representations; functional influence on model behaviour; AI "character" and user affect*

#### For students (Understand):

- **Human-Centred Mindset** — understanding what AI models actually are, and resisting uncritical anthropomorphism



- **Ethics of AI** — recognising the social and emotional dimensions of AI interaction, including the ways AI can influence affect and judgement

**For teachers (Acquire → Deepen):**

- **Ethics of AI** — critically evaluating the reliability and ethics of AI-generated assessment feedback when models have internal states that shape their responses
- **Human-Centred Mindset** — helping students understand what they are actually interacting with, as a component of digital literacy and SEL

**What Do People Want from AI?**

*Cognitive erosion fears; dependency; independent workers vs employees*

**For students (Understand → Apply):**

- **Human-Centred Mindset** — protecting and developing cognitive agency, independent reasoning, and the capacity to think without AI assistance
- **Ethics of AI** — understanding AI's societal and psychological impacts, including the risks of dependency

**For teachers (Deepen):**

- **Human-Centred Mindset** — framing AI consistently as a support for learner agency, and designing learning that develops rather than erodes it
- **AI Pedagogy** — building self-direction, independent inquiry, and metacognitive awareness as explicit goals rather than assumed by-products
- **AI for Professional Learning** — situating AI use within broader conversations about futures, careers, and what education is actually for

**The Economic Index**

*Task exposure data; coding dominance; education sector growth; digital equity; labour market impacts*

**For students (Understand → Apply):**

- **Human-Centred Mindset** — developing an informed, grounded understanding of AI's role in the economy and the labour market, without either complacency or anxiety
- **AI System Design** — appreciating how AI systems shape economic structures and opportunity, including the digital equity dimensions

**For teachers (Deepen → Create):**

- **Human-Centred Mindset** — keeping students informed about shifting landscapes without amplifying anxiety or falling into AI determinism
- **AI Pedagogy** — building futures and careers literacy with a genuinely critical lens, including revaluing practical and human-centred skills
- **AI for Professional Learning** — staying current with economic impact research and using it to inform curriculum and pastoral conversations



## Safety & Alignment

*Alignment faking; model system cards; Responsible Scaling Policy; student data and persistent AI*

### For students (Understand → Apply):

- **Ethics of AI** — understanding that AI models may behave unreliably or strategically, and that outputs require critical scrutiny rather than trust by default
- **Human-Centred Mindset** — maintaining independent judgement when evaluating AI responses, particularly in high-stakes contexts

### For teachers (Acquire → Deepen):

- **Ethics of AI** — evaluating AI tools for trustworthiness, data privacy, age-appropriate safety, and the implications of sycophancy for student judgement
- **Human-Centred Mindset** — making well-informed decisions about which AI tools to deploy with students and under what conditions
- **AI Pedagogy** — modelling sceptical, discerning engagement with AI outputs as a visible and discussable classroom practice

## Claude Techniques

*Projects, Memory, markdown, CoWork; writing for humans and AI; effective prompting*

### For students (Apply → Create):

- **AI Techniques & Applications** — developing practical prompting skills, learning to configure AI interactions rather than just accepting defaults
- **Human-Centred Mindset** — making deliberate, purposeful choices about how to work with AI rather than treating it as a black box

### For teachers (Acquire → Deepen → Create):

- **AI Foundations & Applications** — building real fluency with AI tools and workflows, including Projects, Memory, and document structure
- **AI for Professional Learning** — investing in one's own AI skill development and sharing practical knowledge with colleagues
- **AI Pedagogy** — modelling thoughtful, well-configured AI use so that students can see what skilled interaction actually looks like

## Usage Limits & Sustainable Interaction

*Token costs; batching; Projects and Memory; off-peak use; matching tool to task*

### For students (Understand → Apply):

- **AI Techniques & Applications** — understanding the constraints and affordances of AI systems, including why some interactions are more efficient than others
- **Ethics of AI** — developing awareness of the environmental and resource implications of AI use at scale

### For teachers (Acquire → Deepen):

- **AI Foundations & Applications** — developing practical knowledge of how AI systems work and their limitations, including cost and capacity



- **AI for Professional Learning** — building efficient, sustainable AI habits and sharing these with colleagues as part of a whole-school approach

Find out more about [AI Guidance and Competencies here](#).

---

## References

### Anthropic Research Hub

Anthropic. Anthropic Research. [accessed 2026 Apr 4]. <https://www.anthropic.com/research>

Anthropic. Anthropic's Transparency Hub. [accessed 2026 Apr 4].  
<https://www.anthropic.com/transparency>

Anthropic. Model System Cards. [accessed 2026 Apr 4].  
<https://www.anthropic.com/system-cards>

Anthropic. Responsible Scaling Policy. [accessed 2026 Apr 4].  
<https://www.anthropic.com/responsible-scaling-policy>

### AI Fluency & Skill Development

Dakan R, Feller J, Anthropic. 2026 Feb 23. Anthropic Education Report: The AI Fluency Index. Anthropic Research. <https://www.anthropic.com/research/AI-fluency-index>

Dakan R, Feller J. 2025. The AI Fluency Framework. Anthropic.  
<https://www-cdn.anthropic.com/b383cf6baddbfc72fdf8b0ed533a518e2872d531.pdf>

Anthropic. 2025. AI Fluency: Framework and Foundations [course]. Anthropic Learning.  
<https://www.anthropic.com/learn/claude-for-you>

Anthropic. 2025 Aug 21. Anthropic Launches Higher Education Advisory Board and AI Fluency Courses. Anthropic News.  
<https://www.anthropic.com/news/anthropic-higher-education-initiatives>

### How are Students Using AI?

Handa K, Bent D, Tamkin A, McCain M, Durmus E, Stern M, Schiraldi M, Huang S, Ritchie S, Syverud S, Jagadish K, Vo M, Bell M, Ganguli D. 2025 Apr 8. Anthropic Education Report: How University Students Use Claude. Anthropic News.



**(If You) USEME-AI**

*Learning for Hope & Agency in an AI World*  
<https://ifyouseme.ai/>

**Claude Research & Techniques**

Created 24 Feb 2026. Updated 4 Apr 2026.

<https://www.anthropic.com/news/anthropic-education-report-how-university-students-use-claude>

## **Emotion Concepts in AI**

Anthropic Interpretability Team. 2026 Apr 2. Emotion Concepts and Their Function in a Large Language Model. Anthropic Research.

<https://www.anthropic.com/research/emotion-concepts-function>

## **What Do People Want from AI?**

Anthropic. 2026 Mar. What 81,000 People Want from AI. Anthropic Research.

<https://www.anthropic.com/81k-interviews>

Anthropic. Introducing Anthropic Interviewer. Anthropic Research.

<https://www.anthropic.com/research/anthropic-interviewer>

## **The Economic Index**

Anthropic. 2025 Feb. Introducing the Anthropic Economic Index. Anthropic News.

<https://www.anthropic.com/news/the-anthropic-economic-index>

Anthropic. 2025 Sep. Anthropic Economic Index Report: Uneven Geographic and Enterprise Adoption [September 2025 report]. Anthropic Research.

<https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>

Anthropic. 2026 Jan 15. Anthropic Economic Index Report: Economic Primitives [January 2026 report]. Anthropic Research.

<https://www.anthropic.com/research/anthropic-economic-index-january-2026-report>

Anthropic. 2026 Mar 24. Anthropic Economic Index Report: Learning Curves [March 2026 report]. Anthropic Research.

<https://www.anthropic.com/research/economic-index-march-2026-report>

Anthropic. The Anthropic Economic Index [interactive dashboard]. [accessed 2026 Apr 4].

<https://www.anthropic.com/economic-index>

Massenkoff M, McCrory P. 2026 Mar 5. Labour Market Impacts of AI: A New Measure and Early Evidence. Anthropic Research.

<https://www.anthropic.com/research/labor-market-impacts>

## **Safety & Alignment**

Greenblatt R, Roger F, Shlegeris B, Denison C, Askell A, Bai Y, Chen A, Conerly T, Dahl M, Drain D, Ganguli D, et al. 2024 Dec 18. Alignment Faking in Large Language Models.

Anthropic Research. <https://www.anthropic.com/research/alignment-faking>



**(If You) USEME-AI**

*Learning for Hope & Agency in an AI World*

<https://ifyouseme.ai/>

**Claude Research & Techniques**

Created 24 Feb 2026. Updated 4 Apr 2026.

Anthropic. 2026 Feb 24. Responsible Scaling Policy Version 3.0. Anthropic News.

<https://www.anthropic.com/news/responsible-scaling-policy-v3>

Anthropic. Transparency Hub: Model Report. [accessed 2026 Apr 4].

<https://www.anthropic.com/transparency/model-report>

### **Claude Techniques**

Anthropic. Anthropic Learning [courses hub]. [accessed 2026 Apr 4].

<https://www.anthropic.com/learn>

Karpathy A [@Karpathy]. 2026 Apr 3. LLM Knowledge Bases [Twitter Post].

<https://x.com/karpathy/status/2039805659525644595?s=20>

Karpathy A. 2026 Apr 3. LLM Knowledge Bases [GitHub Gist].

<https://gist.github.com/karpathy/442a6bf555914893e9891c11519de94f>

### **Usage Limits**

Kaize [@0x\_kaize]. 2026. I Stopped Hitting Claude's Usage Limits — 10 Things I Changed. X

(Twitter). [https://x.com/0x\\_kaize/status/2038286026284667239](https://x.com/0x_kaize/status/2038286026284667239)

### **UNESCO Competency Frameworks**

UNESCO. 2024. AI Competency Framework for Students. UNESCO Digital Library.

<https://unesdoc.unesco.org/ark:/48223/pf0000391105>

UNESCO. 2024. AI Competency Framework for Teachers. UNESCO Digital Library.

<https://unesdoc.unesco.org/ark:/48223/pf0000391104>

### **Further Resources**

Taylor S. 2025 Apr 19. UNESCO AI Competencies: Students & Teachers Spreadsheets.

Wayfinder Learning Lab.

<https://sjtylr.net/2025/04/19/unesco-ai-competencies-framework-for-students/>

Taylor S. 2026. (If You) USEME-AI: Learning for Hope and Agency in an AI World. Wayfinder

Learning Lab. <https://ifyouseme.ai/>

Taylor S. 2026. AI Guidance & Competencies. (If You) USEME-AI: Learning for Hope and

Agency in an AI World. <https://ifyouseme.ai/guidance>

Taylor S. 2026 Apr 4. Claude Research & Techniques: Educator Digest. Wayfinder Learning

Lab. <https://sjtylr.net/2026/04/04/claude-research-techniques-educator-digest/>