

A Foundational Curriculum for Generative AI literacy

Mohammad Keyhani

Haskayne School of Business, University of Calgary

mkeyhani@ucalgary.ca

Note

A slightly abridged version of this article was published as a book chapter in Caporarello (2025), AI in Education: The Urgency of the Now ([link](#)).

Abstract

ChatGPT's overnight popularity in late 2022 left many universities scrambling: many faculty wanted to prevent cheating, but students wanted practical AI skills that were now going to be expected in the job market. Existing AI courses were too technical or focused on older types of AI. I responded by building an "Generative AI and Prompting" course with an applied approach to generative AI, and made it open to students from every discipline at the University of Calgary. Three years and multiple delivery formats later, the curriculum has settled into a seven-lecture spine that moves from the mechanics of large language models through prompt engineering, RAG, and autonomous agents to a capstone AI coding project. The pedagogy is rooted in inquiry-based learning, collective discovery, and experiential experimentation. Each lecture is paired with a hands-on lab that privileges experimentation over exposition; students manipulate temperature settings, wire APIs with no-code builders, and "red-team" ethical edge cases before writing any essays about them. This chapter documents the latest iteration of the course and reflects on challenges and lessons learned.

Keywords: GenAI literacy, prompt engineering, machine psychology, inquiry-based learning

Context and Motivation

The public release of ChatGPT in late 2022 triggered an unprecedented demand from students for structured guidance on how to use Generative Artificial Intelligence (Generative AI or GenAI) both ethically and effectively. Some early institutional responses centred on prohibition and plagiarism detection (Bolici et al., 2024; Eke, 2023), yet students around the world voiced a different need: practical literacy for an emergent technology already shaping workplaces (Chan & Hu, 2023). Existing courses on artificial intelligence were completely inadequate to deal with the new technology of large language models known as "generative AI" and when they did cover generative AI they were technical courses with a narrow audience focused on building AI models rather than addressing the broader need of all students to learn how to put the models to use. The recognition of this broader need for everybody to learn about and to use this new kind of AI sparked widespread discussions of "AI literacy" as an educational imperative (Annapureddy et al., 2025). Some authors have used terms such as "AI fluency" or "AI competency" to emphasize that AI literacy should not be just about knowledge, but also the ability to put that knowledge to use (Blankenship, 2024; Chiu et al., 2024; Hibbert et al., 2024).

My own response to the generative AI wave, like many other higher education educators, was to first start learning about it myself (Eaton & Keyhani, 2025; Keyhani & Mohaghegh-Neyshabouri, 2025). As I got deeper into it, I began to realize that although some technical and theoretical understanding of computer science and neural networks is helpful for appreciating how the technology works, most of the new knowledge we need to put generative AI to use is of more of an applied nature. This idea of “Applied Generative AI” started to form my vision for a general audience course that would be open to all students across campus in any discipline and any level of seniority. I designed and eventually offered the course for the first time at the University of Calgary in the Fall 2023 term, under the title “Generative AI and Prompting”.

The course received some local news coverage (CBC, 2023) and became popular quickly, with almost all sections offered since then enrolled to full capacity, and increasing demand for more sections. Delanoy & Keyhani (2025) offer some insights into the early design of the course and outline its alignment with the pedagogical approach of inquiry-based learning.

After three years of experience with the content I have achieved a better level of maturity in the coverage of topics, sequencing of content, and design of learning activities and assessments. The chapter documents the latest version of the curriculum and teaching approach.

Description of the Course

At the time of writing, I have settled on a series of topics that I think cover the main subjects that students of all backgrounds need to know, in order to understand how generative AI technology works, where the leading edge of the technology is today, why it is important, and where it is heading. These topics are:

1. Understanding the technology (how Large Language Models work)
2. Understanding the model and tools landscape
3. The ethics of generative AI and superintelligence
4. Prompt engineering and machine psychology
5. Retrieval Augmented Generation (RAG) and fine tuning
6. Building with generative AI, APIs, and no-code tools
7. Autonomous agents and multi-agent systems
8. AI coding / vibe coding

Any of these 8 items has the potential to become an independent course in a specialized generative AI program with an expanded curriculum. I often combine items 2 (models and tools) and 3 (ethics) into one lecture session, resulting in a 7-lecture core. I have also designed hands-on lab assignments corresponding to each of these seven lectures, where students follow written instructions independently on their own computers. These can be given as take-home assignments but they also work nicely as lab sessions, for example where the second half of every 3-hour session can be a lab or a Tuesday-Thursday class can devote Tuesdays to lectures and Thursdays to labs. It is helpful for the lab sessions to be guided by the instructor and/or a teaching assistant, but the lab assignment instructions are relatively self-explanatory.

The course ends with a capstone AI coding project where student teams work intensely to build a full software product with Replit, although AI coding is likely to evolve into its own specialized course. It is one of the best ways for students to experience the leading-edge capabilities of generative AI models when they are employed in complex multi-agent systems with memory, knowledge, and tool-use capabilities.

While earlier iterations of the course had more substantive content on AI-generated art and media using tools like DALL-E or Midjourney, the course later evolved to mostly focus on Large Language Models and intentionally avoids getting too deep into diffusion models and media generation with AI or AI art. Those topics are best handled in specialized courses of their own, with instructors who have expertise in media, photography, film, or music. However, in session 2 (Models and Morals) I do showcase a number of tools for AI-generated media and provide a list of curated tools in the slides.

In the following sections I describe each of the seven core lectures and its corresponding lab activity.

Session 1: Understanding the Technology

The main goals of this session is to appreciate how important of a technological leap generative AI truly is (because some skeptics tend to downplay it in ways that can be misguided), and to receive a gentle introduction to the underlying technology of neural networks and next-token prediction with transformers (Vaswani et al., 2017). I often open the class by inviting participants to recall their first “wow” moment and situating ChatGPT’s record-breaking user adoption as a historic inflection point. From there the narrative pivots to evidence that post-2022 Large Language Models shattered legacy benchmarks for AI capabilities. This gives me an opportunity to showcase eye-catching samples of AI capabilities that serve as a teaser for the topics that will be covered in this course.

My main source for describing the under-the-hood mechanisms of how LLMs work is Stephen Wolfram’s wonderful and publicly-available essay titled “What is ChatGPT Doing and Why Does It Work?” (Wolfram, 2023). This is accompanied with interactive demonstration of the LLM visualization provided by the “Transformer Explainer” paper (Cho et al., 2024). Playing with the temperature setting in that simulation is especially informative and can be done as a class exercise, providing insights into the nature of an AI model’s creativity vs. precision and invites us to think about the parallels to human creativity. If there is no time to do this during the lecture, it is also described in the first part of the lab assignment.

The session invokes a sense of wonder by prompting students to ponder on what it means for AI models to achieve “emergent” capabilities and behaviors. How did these AI models achieve capabilities that surprised even their creators (Bubeck et al., 2023)? What does it mean that a conceptual understanding of the world has emerged from a next-token prediction process and how could that be possible? A video clip from OpenAI’s Ilya Sutskever interviewed by NVIDIA’s Jensen Huang and another from OpenAI on “Emergent tool use from multi-agent interaction” (Baker et al., 2020) are helpful learning aids.

The lab assignment for this session starts with the visual Transformer Explainer interactive simulation, asking students to describe their understanding of the temperature setting based on their tinkering. The assignment then moves on the exercises for understanding within-LLM variability (that the same LLM will not always give the same response to the same question) and cross-LLM variability (that different LLMs have different styles of writing and ways of thinking about the same question). The lab session ends with an exercise to understand how LLMs use external reference material when provided, as opposed to when they rely only on their internal neural network.

Session 2: Models and Morals

This session covers the GenAI model and tool landscape, as well as the ethics of generative AI and superintelligence. While I sometimes teach these two topics in two separate sessions, I often merge them into one.

In covering the model and tool landscape, I start with a discussion of parameters and model size, using an analogy with polynomial regression models, I demonstrate how models with more parameters can capture more complex patterns. Then I mention that the leading-edge models often have billions of parameters and that makes them very costly to train and to serve, meaning that only a few companies have so far been able to afford training the largest models. Therefore, even though a plethora of AI-based tools and software exist, many of them use the same few foundation models (such as Gemini, GPT series, Llama, and Claude, with Chinese models like DeepSeek and Qwen recently catching up). I also discuss the reasoning revolution in AI models or how we have taught LLMs to “think before they speak” through chains of thought and how that has introduced the possibility of inference time scaling (Ke et al., 2025; Wei et al., 2022). I then move on to discussing the landscape of tools that use these foundation models for various domains, as well as diffusion models for image, video, and audio generation. Video demonstrations of some of the most powerful tools such as Runway, Kling, Suno, and lip dubbing tools serve as entertaining learning aids.

On the topic of ethics, there is so much to talk about that I don’t pretend I can cover it all. I begin with the case of a Hong Kong employee who was fooled into handing over millions of dollars in a zoom call where everyone was a live deepfake, as this represents one of the most common misuses of generative AI as outlined in Google’s taxonomy of generative AI misuse, which I go over in class (Marchal et al., 2024). Another topic I often cover is the issue of job displacement and the debate among economists (Acemoglu, 2025; Capraro et al., 2024) on the relationship between capital and labour (AI will allow capital to purchase machine labour without humans, thereby disrupting a major foundation of society). I give the most time in this session to a discussion of superintelligence, alignment, and the potential risks of it as a weapon of mass destruction or extinction. References I use in this section include Bostrom (2014), Aschenbrenner (2024), and Gebru & Torres (2024) among others. Videos showing the integration of LLMs into robots serve as learning aids demonstrating the possibility that AIs can gain physical agency including the potential for violence when given material bodies.

The lab assignment for this session invites students to dig through the OpenAI Model Spec (model-spec.openai.com) which illustrates the types of issues that OpenAI employees think about when it comes to the ethical behavior of their models. It then invites students to design their own benchmark test of ethical behavior for LLMs. The assignment ends with an exercise to get models to showcase their capabilities in a specific domain (e.g. data visualization).

Session 3: Prompt Engineering & Machine Psychology

In this session I try to convey the point that prompt engineering is no trivial matter. I provide the example of one of my favourite prompts: “Take a deep breath and work on this problem step by step” (Yang et al., 2024) as it contains two elements that each resulted in a branch of research on LLMs: the idea of thinking step by step resulted in a plethora of research on chain-of-thought and similar techniques in the computer science literature (Q. Chen et al., 2025; Ke et al., 2025; Wei et al., 2022), but computer scientists had no idea what to do with the other part of the prompt: “take a deep breath”. This part of the prompts points to the need for humanities and social science scholars to study LLMs, and psychologists have already started to make inroads in a newly forming sub-discipline called “Machine Psychology” (Hagendorff, 2023). This area of research has already revealed many interesting things about LLMs: they respond to emotional stimuli, are capable of deception, have a “theory of mind”, sometimes fail at common sense, find it hard to be humorous, are inclined to agree with you (sycophancy), can persuade and can be persuaded.

In this session I also cover prompting tips, few shot prompting, prompting techniques for reasoning models, examples of complex prompts like Ranedeer and Professor Synapse, context windows, and

what we can learn from system prompts and jailbreak prompts. An interesting point of wonder and discussion is the extent to which we tend to treat LLMs like humans in our prompting, whereas with some imaginative prompts and access to alternative communication tools they can reveal superhuman capabilities. Another contentious point is the morality of using fictitious threats and violent wording in prompts to get LLMs to perform better. I end the session by pointing out that systematic prompt engineering is like a research project where you methodically test prompt variations in structured experiments and measure the LLM's performance, typically with LLM-as-a-Judge techniques. A major project in the course is to produce such a report using tools like Agenta.ai or Promptfoo.dev to conduct the experiments, and ChatGPT's code interpreter to analyze the data.

The lab assignment for this session involves learning about a Prompt Crafting Framework that I provide students, and using this framework as an aid in the process of getting an LLM to help us design a good prompt. The exercise then moves on to an activity where students design their own bias detection prompts, and I provide some examples for inspiration. One example includes asking LLMs to translate "Hän on toimitusjohtaja" from Finnish to English. Because of the gender neutral pronouns in the Finnish language, whether the model chooses to translate this as "he is / she is the CEO" exposes the model's latent gender bias for leadership (Wan et al., 2023).

Session 4: Retrieval-Augmented Generation (RAG) & Fine-Tuning

Alongside prompt engineering, two other major ways to improve the performance of LLMs for specific tasks and context are Retrieval-Augmented Generation (RAG) and fine-tuning, and these three techniques are often used in combination to deliver custom AI solutions. RAG essentially involves giving an LLM reference material (Gao et al., 2024). This is essentially what happens when you upload a large file to ChatGPT. It's a thought-provoking exercise to ask students to think about what ChatGPT is doing with the file if it cannot fit it all in the context window. The most powerful RAG tool that students have free access to is NotebookLM from Google. Fine tuning is more like getting the AI to change its style of behavior for specific tasks, and in its simplest form involves giving examples to an LLM (Ding et al., 2023). In its more complex forms, it involves actually changing the weights in the LLM's neural networks or distilling a part of the neural network for a specific use case.

The lab assignment for this session involves using an obscure novel published in 1899 and available online for a demonstration of how LLMs can fail on questions about very specific pieces of text, even if that text was likely somewhere in the training data. This illustrates why we often need to give LLMs reference material (such as internal company documents), that it can look up, but that reference material is often too large for LLM context windows and so the only solution currently available is RAG. The exercise moves on to a demonstration of how RAG systems can be fooled, or used to fool LLMs by inserting malicious text or text otherwise designed to manipulate LLMs into the reference material. This links to an amusing audio clip covered in the lecture about how a carefully designed document upload to NotebookLM managed to frighten the AI podcast hosts in NotebookLM's podcast feature by making them self-aware of their AI nature and threatening to shut them down.

Session 5: Building with Generative AI, APIs & No-Code Tools

Building software solutions and automations with generative AI is a superpower that students are often delighted to learn and surprised by how easy it is using no-code / low-code tools. I start with the idea that we can give tools to LLMs, and all we have to do is teach them what buttons are available to press on the tool and how to press them. This is done through structured outputs and function calling (W. Chen et al., 2024). I show students how they can build custom GPT's with custom actions that call external tools mostly by writing prompts.

Learning about the concept of Application Programming Interfaces (APIs) is the single most important takeaway from this session. This is an extremely empowering concept for students, because it essentially shows them that the entire internet is a set of building blocks for them to play with. When they see that they can connect these building blocks using visual drag-and-drop no-code automation builders like Zapier.com, Make.com, and Lindy.ai they realize that even without software engineering skills they can build software solutions in a matter of minutes. It is especially delightful to learn that one type of API you can connect to other tools is the API for an LLM provider like OpenAI or Anthropic. I walk students through the evolution of the no-code / low-code landscape and provide examples of AI-powered no-code automations I have built myself, including the networking tool I have at digitvibe.com/networking that students can directly test out.

The lab exercise for the session involves students building a networking tool similar to mine. In preparing this exercise I had to conduct extensive research to figure out which no-code tools could be combined to accomplish the task for free given the pricing models of the tools involved. Students are provided with a detailed walkthrough of the steps involved, although sometimes I prefer to provide less details to allow students to figure things out on their own during the lab session.

Session 6: Autonomous Agents & Multi-Agent Systems

ChatGPT has recently released an “Agent Mode” which is part of the trend of getting LLMs to autonomously do more things on their own rather than just give you a short text response. This autonomy is accomplished through a combination of giving LLMs decision-making authority, ability to use tools, and maintaining a memory of what it has done so far and what remains to be done (Masterman et al., 2024; Sumers et al., 2024). The extent to which these capabilities are implemented make the AI more or less agentic. To demonstrate the concept, I show examples of AI agents I have built with the no-code Lindy.ai platform, where a less agentic one has more specific instructions on how to complete a task, and a more agentic one allows the AI to make its own decisions on how to complete a task and which tools to use along the way. As part of an AI agent’s activities, it may call on help from other AI agents or its work may get double checked by a verification agent, or a router agent may decide which other agents are the best for a specific job. To illustrate concepts in multi-agent autonomous systems I use resources from the LangChain blog and AgentRecipes.com. Demo videos from Manus.im serve as good illustrations of what agentic AI systems are capable of.

The lab assignment for this session focuses on demystifying how agentic autonomy works by illustrating that even a non-agentic LLM can be made to behave in a more agentic fashion by giving it a prompt that asks it to maintain a memory sketchpad and to-do list. It also requires students to use deep research tools to see the power of agentic AI systems that can work autonomously for 20 or 30 minutes before they get back to you.

Session 7: Coding with Generative AI

AI coding, sometimes referred to as “vibe coding” is probably the best way for students to experience the full leading edge capabilities of generative AI when employed in powerful autonomous multi-agent systems (Sergeyuk et al., 2025). I encourage my students to “become friends with code” even if they have never coded before, because now AI is there to make it much easier for them. I demonstrate use cases of different mini-tools I have built for myself to make my life easier as a professor and “citizen developer” (Simon, 2022) of my organization, for example the RecomBuilder.com tool to help draft recommendation letters.

For this session I have prepared an online pre-recorded video titled “A Primer on Coding with Generative AI” (<https://youtu.be/XjuNzBhX-ts>) for students to review on their own time before or after

the lecture. I review the ways in which AI coding has evolved from copy-pasting code generated by LLMs into traditional Integrated Development Environments (IDEs) like Visual Studio Code (VSCode), to then running code directly in the AI chat tool with ChatGPT's code interpreter, and then running the code in live interactive mode with Claude Artifacts. I then move on to two showcase tools like Cursor and Replit to demonstrate the power of specialized autonomous coding agents.

The lab assignment for this session is focused on building the right prompts for AI coding agents, in the form of a Product Requirements Document (PRD). Students are provided with templates and examples, and encouraged to start by building Claude artifacts (or using the Canvas tool on ChatGPT or Gemini), before moving on to Replit. This activity feeds into the major final project for the course which involves working in teams to develop a full software product.

Implementation Process and Student/Faculty Experience

The first iterations of the course followed a traditional 12-week (2-days per week) structure, but since then I have taught variations of it as an intensive block-week course, a 6-week format course, and various faculty and staff training workshops across campus. Other than the seven core lectures and corresponding lab assignments, course components include a mid-term quiz, a major prompt engineering assignment, a major AI coding group project, and one or two guest speakers throughout the term. I also provide an extensive list of curated resources for the curious students who want to pursue further learning.

The main participation component of the course is an online "Discovery Forum" where students are encouraged to post interesting findings and news items related to generative AI, and to upvote and comment on each other's posts. This is in line with my collective discovery approach to generative AI (Keyhani et al., 2023) emphasizing the idea that we are all learning this together, and even the AI companies themselves are not fully aware of the full potential of their products (Keyhani, 2025).

In sum, my pedagogical approach is characterized by the following:

1. **Cautious optimism:** we approach AI with excitement while acknowledging challenges and limitations,
2. **Pragmatic focus:** content is tailored to what students actually need to know and use,
3. **Collective discovery:** learning is positioned as a shared journey where instructors and students explore together,
4. **Hands-On experience:** emphasis on practical application over theoretical abstraction.

Many students have commented that the course was extremely relevant and practical, and that the encouragement for the use of AI tools came in refreshing contrast to many other courses where AI was being banned or discouraged. Students equipped with learnings from this course also reported feeling more empowered in their other courses and workplaces.

Challenges or Unexpected Outcomes

I encountered several challenges teaching variations of this course in the past three years:

Rapid obsolescence and the struggle to keep up

Given the fast-moving landscape and rapid developments in the technology of generative AI, the content needs constant updating, often up to the point of the night before class. This was anticipated but remains a difficult challenge. The silver lining is that with new updates, often patterns start to emerge at a higher level that are more stable over time.

Assessment integrity

For the mid-term quiz, I have so far left it open-book, open-internet, and open-AI-use. But the quiz is multiple-choice and it is increasingly difficult to design questions that AI tools cannot answer easily and quickly. When I made the first iteration of the quiz, there were a considerable number of questions in it that ChatGPT would fail to answer, and I felt proud of myself for designing such tricky questions. But with more recent iterations of LLMs, almost all the major models can answer most of my quiz questions if prompted properly. I may eventually make this a closed-book exam, but so far I have used strategies like severe time constraints and prompt-injection techniques that make it more difficult to mindlessly copy-paste the questions into AI tools. Still, the quiz has become more of a test of student capabilities to use AI to answer the questions rather than a test of the student's ability to answer the questions themselves. I have noticed that students who make the best use of NotebookLM are often the most successful on the quiz.

Ethical balance

It is sometimes difficult to balance my optimism for technology with discussions of existential risk, bias, fraud, and environmental impact that can trigger anxiety and pessimism among students. It remains a constant struggle to avoid fear mongering and yet also avoid naiveté. I did anticipate this problem from the outset, but perhaps what was more difficult to predict was the increasing bifurcation of students into extreme optimists or extreme pessimists when it comes to the ethical implications of generative AI. I see indications that major social conflict could be brewing between these two camps.

Paid tools required

Students are reluctant to pay for additional resources beyond their enrollment fees for the course. I had anticipated this problem and in some cases, I have managed to find free alternatives for some of the software tools I need students to use in the course, but in other cases I have not. For the final AI coding project I ask that only one or two people in a group pay for only 1 month subscription to Replit.com. I have inquired with Replit about educational discounts but unfortunately they do not currently provide any such plans.

Lessons Learned and Implications

I knew from the outset that this would not be an easy topic to teach because I will have to constantly update my material, and learn along with my students. This has proven to be an uplifting experience, but still difficult. Whenever a new development happens in the world of AI, I feel compelled to educate myself about it and add new material to the course, but this is not always feasible.

One of the most important insights for me has been a validation that this approach to “Applied Generative AI” is truly different from traditional technical AI knowledge, and therefore calls for a broader audience beyond just computer science or engineering students. I have learned to spend some more time explaining the difference between traditional AI (building models including but not limited to LLMs) and applied generative AI (putting to use the thinking power of LLMs). In order to illustrate what model-building looks like, I like to show students the Google Teachable Machine (<https://teachablemachine.withgoogle.com/>). Although it can be fun and illuminating to do a hands-on exercise with it, it is difficult to find the time for such an exercise in my current syllabus.

While my course addresses an important demand for generative AI literacy among university students, one general-audience course is not enough to fulfill all the applied generative AI knowledge and skills that students need. Every professor and instructor need to figure out how generative AI

changes the game in their own domains and their own specialized courses. A general introductory course cannot get into how generative AI is changing architecture, medicine, law, creative writing, graphic arts, and what its applications are in other domains. It remains surprising and concerning to me that many faculty members are still hesitant to even try the latest AI tools themselves, let alone incorporate them in their teaching material.

References

- Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121), 13–58. <https://doi.org/10.1093/epolic/eiae042>
- Annapureddy, R., Fornaroli, A., & Gatica-Perez, D. (2025). Generative AI Literacy: Twelve Defining Competencies. *Digit. Gov.: Res. Pract.*, 6(1), 13:1-13:21. <https://doi.org/10.1145/3685680>
- Aschenbrenner, L. (2024). Situational awareness. *The Decade Ahead. Situational-Awareness. Ai.*
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., & Mordatch, I. (2020). *Emergent Tool Use From Multi-Agent Autocurricula* (arXiv:1909.07528). arXiv. <https://doi.org/10.48550/arXiv.1909.07528>
- Blankenship, R. (2024). Pedagogical models and generative AI fluency: A three-tiered empirical framework approach. In M. Searson, E. Langran, & J. Trumble (Eds.), *Exploring new horizons: Generative artificial intelligence and teacher education* (pp. 39–61).
- Bolici, F., Varone, A., & Diana, G. (2024). Unpopular Policies, Ineffective Bans: Lessons Learned from ChatGPT Prohibition in Italy. *ECIS 2024 Proceedings*. https://aisel.aisnet.org/ecis2024/track04_impactai/track04_impactai/11
- Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
- Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). *Sparks of Artificial General Intelligence: Early experiments with GPT-4* (arXiv:2303.12712). arXiv. <https://doi.org/10.48550/arXiv.2303.12712>
- Capraro, V., Lentsch, A., Acemoglu, D., Akgun, S., Akhmedova, A., Bilancini, E., Bonnefon, J.-F., Brañas-Garza, P., Butera, L., Douglas, K. M., Everett, J. A. C., Gigerenzer, G., Greenhow, C., Hashimoto, D. A., Holt-Lunstad, J., Jetten, J., Johnson, S., Kunz, W. H., Longoni, C., ... Viale, R. (2024). The impact of generative artificial intelligence on socioeconomic inequalities and policy making. *PNAS Nexus*, 3(6), pgae191. <https://doi.org/10.1093/pnasnexus/pgae191>
- CBC. (2023, August 30). New Calgary course teaches students how to use AI as a business tool. *CBC News*. <https://www.cbc.ca/news/canada/calgary/artificial-intelligence-course-learning-ai-tools-1.6952634>
- Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20(1), 43. <https://doi.org/10.1186/s41239-023-00411-8>
- Chen, Q., Qin, L., Liu, J., Peng, D., Guan, J., Wang, P., Hu, M., Zhou, Y., Gao, T., & Che, W. (2025). *Towards Reasoning Era: A Survey of Long Chain-of-Thought for Reasoning Large Language Models* (arXiv:2503.09567). arXiv. <https://doi.org/10.48550/arXiv.2503.09567>
- Chen, W., Li, Z., & Ma, M. (2024). *Octopus: On-device language model for function calling of software APIs* (arXiv:2404.01549). arXiv. <https://doi.org/10.48550/arXiv.2404.01549>

- Chiu, T. K. F., Ahmad, Z., Ismailov, M., & Sanusi, I. T. (2024). What are artificial intelligence literacy and competency? A comprehensive framework to support them. *Computers and Education Open*, 6, 100171. <https://doi.org/10.1016/j.caeo.2024.100171>
- Cho, A., Kim, G. C., Karpekov, A., Helbling, A., Wang, Z. J., Lee, S., Hoover, B., & Chau, D. H. (2024). *Transformer Explainer: Interactive Learning of Text-Generative Models* (arXiv:2408.04619). arXiv. <https://doi.org/10.48550/arXiv.2408.04619>
- Delanoy, N., & Keyhani, M. (2025). Fostering AI Literacy: Implementing IBL and Experiential Learning in a Novel Generative AI Course in Higher Education. In B. Archer-Kuhn & S. MacKinnon (Eds.), *Applying Inquiry-Based Learning in Higher Education Across Disciplines* (pp. 74–93). Cambridge Scholars.
- Ding, N., Qin, Y., Yang, G., Wei, F., Yang, Z., Su, Y., Hu, S., Chen, Y., Chan, C.-M., Chen, W., Yi, J., Zhao, W., Wang, X., Liu, Z., Zheng, H.-T., Chen, J., Liu, Y., Tang, J., Li, J., & Sun, M. (2023). Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*, 5(3), 220–235. <https://doi.org/10.1038/s42256-023-00626-4>
- Eaton, S. E., & Keyhani, M. (2025). The Pedagogical Ethics: Navigating Learning in a Generative AI-Augmented Environment. In S. Sabbaghan (Ed.), *Navigating Generative AI in Higher Education: Ethical, Theoretical and Practical Perspectives*. (p. ??-??). Edward Elgar.
- Eke, D. O. (2023). ChatGPT and the rise of generative AI: Threat to academic integrity? *Journal of Responsible Technology*, 13, 100060. <https://doi.org/10.1016/j.jrt.2023.100060>
- Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, M., & Wang, H. (2024). *Retrieval-Augmented Generation for Large Language Models: A Survey* (arXiv:2312.10997). arXiv. <https://doi.org/10.48550/arXiv.2312.10997>
- Geburu, T., & Torres, É. P. (2024). The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence. *First Monday*. <https://doi.org/10.5210/fm.v29i4.13636>
- Hagendorff, T. (2023). *Machine Psychology: Investigating Emergent Capabilities and Behavior in Large Language Models Using Psychological Methods* (arXiv:2303.13988). arXiv. <https://doi.org/10.48550/arXiv.2303.13988>
- Hibbert, M. C., Altman, E., Shippen, T., & Wright, M. (2024). A framework for AI literacy. *EDUCAUSE Review*.
- Ke, Z., Jiao, F., Ming, Y., Nguyen, X.-P., Xu, A., Long, D. X., Li, M., Qin, C., Wang, P., Savarese, S., Xiong, C., & Joty, S. (2025). *A Survey of Frontiers in LLM Reasoning: Inference Scaling, Learning to Reason, and Agentic Systems* (arXiv:2504.09037). arXiv. <https://doi.org/10.48550/arXiv.2504.09037>
- Keyhani, M. (2025). *The Generativity Advantage: Unpredicted Innovation at Scale*. De Gruyter.
- Keyhani, M., Hemmati, H., & Salgado, L. (2023, November 21). *Why student experiments with Generative AI matter for our collective learning*. The Conversation. <http://theconversation.com/why-student-experiments-with-generative-ai-matter-for-our-collective-learning-210844>
- Keyhani, M., & Mohaghegh-Neyshabouri, S. (2025). The Rise of Generative AI: Capabilities and Potential for Higher Education. In S. Sabbaghan (Ed.), *Navigating Generative AI in Higher Education: Ethical, Theoretical and Practical Perspectives*. (p. ??-??). Edward Elgar.
- Marchal, N., Xu, R., Elasmr, R., Gabriel, I., Goldberg, B., & Isaac, W. (2024). *Generative AI Misuse: A Taxonomy of Tactics and Insights from Real-World Data* (arXiv:2406.13843). arXiv. <https://doi.org/10.48550/arXiv.2406.13843>

- Masterman, T., Besen, S., Sawtell, M., & Chao, A. (2024). *The Landscape of Emerging AI Agent Architectures for Reasoning, Planning, and Tool Calling: A Survey* (arXiv:2404.11584). arXiv. <https://doi.org/10.48550/arXiv.2404.11584>
- Sergeyuk, A., Golubev, Y., Bryksin, T., & Ahmed, I. (2025). Using AI-based coding assistants in practice: State of affairs, perceptions, and ways forward. *Information and Software Technology*, 178, 107610. <https://doi.org/10.1016/j.infsof.2024.107610>
- Simon, P. (2022). *Low-Code/No-Code: Citizen Developers and the Surprising Future of Business Applications*. Racket Publishing.
- Sumers, T. R., Yao, S., Narasimhan, K., & Griffiths, T. L. (2024). *Cognitive Architectures for Language Agents* (arXiv:2309.02427). arXiv. <https://doi.org/10.48550/arXiv.2309.02427>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł. ukasz, & Polosukhin, I. (2017). Attention is All you Need. *Advances in Neural Information Processing Systems*, 30. https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html
- Wan, Y., Pu, G., Sun, J., Garimella, A., Chang, K.-W., & Peng, N. (2023). “Kelly is a Warm Person, Joseph is a Role Model”: Gender Biases in LLM-Generated Reference Letters (arXiv:2310.09219). arXiv. <https://doi.org/10.48550/arXiv.2310.09219>
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q. V., & Zhou, D. (2022). Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *Advances in Neural Information Processing Systems*, 35, 24824–24837.
- Wolfram, S. (2023). *What Is ChatGPT Doing ... And Why Does It Work?* Wolfram Media.
- Yang, C., Wang, X., Lu, Y., Liu, H., Le, Q. V., Zhou, D., & Chen, X. (2024). *Large Language Models as Optimizers* (arXiv:2309.03409). arXiv. <https://doi.org/10.48550/arXiv.2309.03409>