Feeling lost in an AI barrage? If you like, I can help you think this through: Or towards a way of comprehending genAI

Abstract

Generative AI (genAI) has received significant coverage describing its potential impact on productivity, employment and education. A core weakness of much of this coverage is that it is purely *descriptive* and solutionistic*,* thus assuming a certain value framework that may not be relevant to everyone affected by the uptake of genAI. In this essay, I argue that instead of looking at what genAI *does*, we must instead begin with what genAI *is,* what differentiates it from other implements of social change, and ask ourselves questions about how its outputs impact our realisation of our values. Drawing on a large body of evidence, I suggest how we may be changed by genAI and highlight the danger of allowing corporate entities to be the primary custodians of a tool with the characteristics of genAI.

Introduction and basic definitions

Artificial intelligence (AI), just like organic intelligence, exists in many different forms. We can call the set of computational tools which enable a computer to generate novel outputs on the basis of inputted training data generative AI. When a human being does something similar, we call it creativity. When a person sees patterns and has some ability to predict things based on them, we might say they have insight or perceptiveness. When a computer does the same thing it's likely being done via machine learning regression for prediction, and machine learning classification for pattern recognition. In this essay, I'd like to talk a bit about generative AI (genAI), whose most recognisable form is the large language model (LLM), which is the form of genAI on which chatbots like ChatGPT, Claude and others are based. Tools such as these have the potential for significant social change[1-3].

There are multiple methods in which we may critique social changes, e.g. those due to policy or disruptions like war and technology. The first is from a point of shared aims and values meaning that the only point of debate is *how* a given change affects the realisation of those same shared goals and values. For instance, in the management of a pandemic situation: is mandating vaccination and locking everyone in their homes a

good way of achieving reduction of a disease burden? In this case, committing to answering the question implies that reducing disease burden is good and we can collect data on our interventions and see if they achieved said good. Essentially, the scientific method is being used to test an intervention against an endpoint to inform an expression of a set value in the form of policy.

Despite being a scientist myself, I have to point out the limitations of this approach. First: It only answers the *specific* question you're asking within a given value framework. We may think ourselves able to address questions of values scientifically, but this is misguided. Second: this data-centric approach often leads to ancillary questions which become the object of discussion if one is not careful. This is a problem because the quality of the data, their provenance and the vested interests of those who collected them, all help to test a hypothesis, but never the underlying values of that hypothesis.

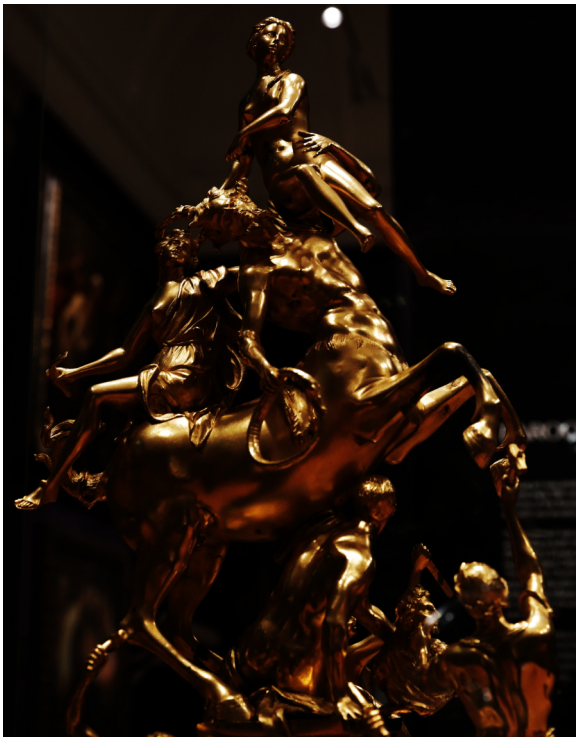
For a more complete, radical, analysis, we must question our values themselves. This is beyond the realm of data and experiment and involves something far more fraught; the arrival at a definition of what is good. With the pandemic example: We might disagree over vaccine mandates and locking people in their homes as effective means of reducing disease burden, but does that mean other coercive actions towards an assumed good of reducing disease burden *are* acceptable? This is a question of values and questions like these can be frightening to many, particularly because there may not be a 'right' or 'wrong' answer. Both the scientific and 'values-first' modes of analysis are important but we often privilege 'scientific' analysis at the expense of thinking about our values. When we do this, we engage in scientism — abrogating moral decision-making to a technical process that was never intended to replace morality.

Most of the analysis of genAI has focused on technical pros and cons such as: the impact on task productivity[4] (lines of code written per employee per unit time, or tasks completed per unit time), cost (savings on employee wages, or the cost:benefit ratio of running datacentres); the environmental impact of datacentres[5], and the controversial financing methods used by many of the largest genAI companies[6].

There are assumptions implicit in the above: higher productivity and lower costs are good and that environmental damage and circular financing are bad. These assumptions do not address how deeper societal and human values, to the extent we can measure or assess them, will be affected by the rampant adoption of genAI. To do that, we need to reframe our analysis.

As such, this article does not seek to meta-analyse effects on productivity, environment, or any economic shock caused by genAI uptake. Reading the news, or engaging with ‘content’ on the subject will give you better measures on those things better than I ever could. My aim is to look at some of the use cases of genAI, beyond a simple evaluation of costs and benefits, and ask what the implications for these use cases are for us as Human, economic and principled beings.

Tools and the uniqueness of genAI



Like the computers on which they run, genAI models are tools[7]. A tool can be understood as anything that enables an idea in the mind to be realised in the real world. Even something as simple as the human hand can be understood as a tool[8]. A tool has no intrinsic moral attributes. For instance a hand has great creative potential in the form of the ‘hand that feeds’ or ‘the painter’s hand,’ but may also deliver destruction, e.g. ‘died at the hands of.’ If we put a knife into a hand, then the consciousness in control of that hand may either prepare food or kill someone more efficiently. The choice is influenced by a value system.

Just as we make and use tools, use of them also remakes us[9]. A good example is the smartphone, which has reshaped western life in its image. Its influence on our attention, consciousness and even our mating patterns are obvious for all to see. There are many other examples from many other spheres of life such as motorised transport redefining our understanding of distance, or written language transforming knowledge from something recalled from memory to something which can be retrieved in a book. As the tool changes us, we find it more difficult to conceive of life before the tool, even if that life may have been more aligned with our values. If the invention of tools is to be accompanied by the retention of desirable values, we should consider the direction and magnitude of the changes new tools impose upon us. This concept is called counterfoil research[9].

What differentiates genAI from the tools I have described earlier is that, unlike the hand, the mnemonic, the car or the smartphone, genAI can imitate the outputs of human cognition[6] in a more general sense than anything which came before. Thus using genAI, it is now possible to generate outputs which were previously the sole domain of human thought, e.g. a piece of writing, art, music or code. In certain cases, agentic AI automates action and intention, such as coming up with ways to structure a codebase on the basis of suitability and then doing that same structuring, meaning that we now have tools to which we may outsource intention and thought. Thus, intellectual labour or knowledge work, something understood to be profoundly human, may be performed by these tools. The synthesis of information from multiple sources into an argument, rapid retrieval of facts unencumbered by an inconsistent memory, and the prospect of never having to struggle with difficult subject matter are now possibilities for anyone with access to a web browser. So, how do we understand this tool within the paradigm of counterfoil research?

What happens when mental processing doesn't matter?



In an essay called 'The Gentle Singularity' by Sam Altman, the CEO of OpenAI, we are promised a world of plenty because of the ability of genAI to mimic intelligence. He posits that a lack of intelligence was a 'fundamental limiter on human progress' and with the advent of genAI 'whole classes of job are about to be wiped out.' According to Altman, this may not be a bad thing though, since 'we'll build ever more wonderful things for each other' and that 'brand new occupations will emerge'[1,10] to replace those jobs. Stories about a well-prompted LLM 'help[ing] to expand the frontier of theoretical physics[11]' and how 'AI can help us become less lonely[12]' have also been published.

In the interests of avoiding argumentation based on technicalities, as opposed to values, let's start by assuming these statements are true. Once the theoretical physicist no longer theorises but prompts her way to an equation, her practice is transformed from one of mathematical cognition to one of skilful prompting. Thus her intellectual labour acquires a different quality, one where her mathematical intelligence is subordinated to her intent and ability to ask the right questions. Over time, even that may be subordinated to something else.

The incentives for working in this way (increased productivity relative to her academic rivals, recognition and so on) mean that in a work environment where productivity is usually the primary measure of success, engaging genAI heavily will become more common. One possible implication is that over time individual deficiencies in thinking in

certain ‘field appropriate’ ways (e.g. mathematically for science and engineering or visually for interior design and visual arts) may cease being impediments to entering a field, provided a person who wants to be in said field can ask questions appropriate to the field *and* the right computational crutches are available.

Superficially, none of this is problematic. We value abundance and genAI may enable that abundance. However, we are connected to our outputs, be they for work or for play. If we change the nature of our intellectual work so much that appropriate prompting is the primary mental skill we need, then despite our AI-enabled abilities to do things we couldn’t before (e.g. write elegant prose, edit images to a high standard or develop an app), we may find our connection with our outputs will likely become shallower. In fact some academic research has demonstrated that study participants who rely heavily on genAI to write essays do not feel as strong a sense of ownership compared to those who do not[13].

Given the capitalist drive towards increasing productivity, be that in lines of code shipped, support tickets closed, articles published, widgets made and so on, workers and industries must adopt ever more efficient mental (mnemonics, stock phrases, snippets of code) or physical (power) tools. If they don’t, they’re outcompeted by those who do. When these tools are combined with the worker’s skill in wielding them, we’d expect more output per worker per unit time. Perhaps then, in the drive for productivity, we will become even less engaged than we already are at work in the genAI age. One may make the counterargument that we can use genAI to only automate the bits of work we dislike leaving aspects which engage modes of thinking we find more agreeable untouched, but this is naïve[14]. The aspects of a job which are bottlenecks to productivity will be the first to be automated, not necessarily the least stimulating.

At this point it’s important to bring up learning. We are learning creatures. We have to learn to have any function at all in the world (such as how to walk, eat, or move). Furthermore, we learn as an end in itself, to make external knowledge a part of us. Yet, we are also incredibly lazy. It’s been established we humans are ‘cognitive misers,’ keen to avoid mental expenditure wherever we can. If the outcomes of such mental expenditure can reasonably be mimicked by genAI, accelerated cognitive miserliness, and reduced intellectual capacity, are likely to ensue. Unfortunately, cognitive effort, struggle and failure enhance learning and knowledge retention[15]. Given so much learning and education is directed towards work (e.g. preparation for it or training while

in it), it seems credible that there will be changes in what we *need* to learn to survive. However, if we end up reducing the mental struggle in education, we'll probably end up learning less overall since *knowing* will become less relevant[16].

Returning to values, let us now ask ourselves what the implications of interacting with genAI as it stands are. We must ask ourselves many things. What is more important? genAI enabled productivity or a feeling of ownership over our outputs? Do we run the risk of being productive but perpetually unfulfilled, or will how we understand ownership itself change?

If thinking in a certain way is unnecessary to perform a certain activity, is the activity then opened up to those who can ask the most appropriate questions? What does this mean for the esteem in which these activities are held, how they are rewarded, and how changes in that esteem affect social structure?

If genAI remains embedded in our work patterns as I imagine it will, what styles of thinking will become endangered or extinct through their irrelevance to modern life? What then do we have to learn and what do we choose to learn to live well[17]? Will the ability to think deeply be relegated from a survival advantage to a leisure activity[18], in the same way that hunting for food was?

Remoteness



Skill, thought, education and work aren't just ends in themselves. They are closely tied to how we relate to one another. When we work, we often work *for* others and in acting we *think* of others. There is always an "other" in our inner world. Education and skill acquisition often involve an other too, be it a teacher or a student. Interactions and chains of interactions between people have a profound effect on our experience of life. Life affects relations and relations affect life, and a lack of close relations is a source of suffering[19]. It's therefore important to think about technological progress, particularly in how genAI could affect our relations.

Due to the march of technology, physical distance is no longer an impediment to communication. Yet, we appear to be more lonely than our technologies would suggest we should be. Over time, the fraction of time a person spends alone has increased over the past century and this is in no small part due to the impact of that same technology, e.g. the smartphone which has been linked to young people being 'less likely than in previous decades to get their driver's license, or to go on a date, or to have more than one close friend, or even to hang out with their friends at all.[20]' More concerning still: 'the share of boys and girls who say they meet up with friends almost daily outside school hours has declined by nearly 50 percent since the early 1990s, with the sharpest downturn occurring in the 2010s.[20]' The situation is somewhat more nuanced since physically and emotionally close connections (e.g. spouses) are generally more frequently contacted using these new technologies than they would be before their adoption. What is disintegrating is interaction in 'third places,' where one might come across those they may not necessarily agree with, or gather with friends[20]. Technology seems to dissolve some borders, but we are rather good at putting other ones up. And it's not as if we don't understand what technology is doing to us. 73% of respondents to a recent survey on loneliness cited technology as one of its biggest causes[21].

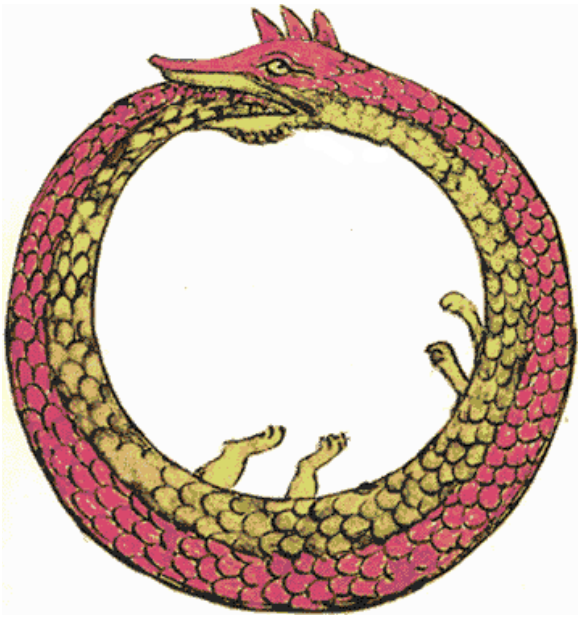
Where does genAI come into all of this? It's in the motivations for isolating ourselves. One can see why the virtual world is seductive. There's no need to travel to loved ones, no need to reconcile disagreements, the screen is available any time, all the time and the effort required to form new, if superficial, bonds is minimised. Living virtually is easy, but it's unfulfilling. LLMs seem to provide knowledgeable, seemingly intelligent, tactful companions, who won't judge you (they cannot), who are perpetually available, and won't ever challenge you if you don't want to be challenged. We're also able to develop a warped form of intimacy with machines (e.g. humanoid robots) in a similar way we do with other humans, revealing more of ourselves across repeated interactions[22]. Maybe it is for those reasons LLMs have been touted as a way to alleviate loneliness[23], but their effects are time limited[12]. Those reasons may also be why those who live with (social) anxiety, depression or loneliness are at higher risk of developing a dependence on chatting with an LLM[24, 25]. LLM dependence is not well characterised but does seem to involve an increase in a person's isolation[26] as they withdraw from the 'real' world.

My suspicion is that a greater volume of communication, whether with humans or LLMs alone, won't help with loneliness and isolation. They aren't solely problems of

communication; they are problems of connection. As the data show, we are relating less and less. As genAI encroaches on relational work such as consoling, teaching, or counselling, I imagine this tendency towards lowered relation will increase, compounding the loneliness problem.

If our world is one of mutual interactions between one another, then we must seek to preserve not only the joy in communal experience, but also its agony, disagreement and resolution. I believe any long-term outsourcing to genAI of our need for connection is foolish at best and dangerous at worst. We have to be prepared to provide human connection just as we wish to receive it. Our experience of life improves when we give as well as take[27] and giving is most impactful when perceptible by a sentient, receiving, human party.

A bigger danger



Over the past sections we've covered the nature of genAI as a tool, and asked ourselves what the consequences for thinking and relations between ourselves are. However, there is a more imminent concern: The corporate power behind the genAI available to most people. The vast majority of LLM prompts are made to models belonging to just a few companies[28]. These companies rely almost exclusively on one company — NVIDIA — to provide the components on which the models run[29]. These companies operate on a profit motive. The profit motive does not serve human welfare.

If we choose to outsource skills, modes of thinking, work and relating to these

corporate entities, we also submit to the centralisation they embody. In getting the machines to think for us, we flirt with the prospect of the companies who own those machines owning us. If we persist with an unconsidered adoption of genAI, we may end up unable to reverse any stranglehold put upon us in the pursuit of a higher profit margin. This is roughly how 'enshittification' works. Introduce a product, make people reliant on the product, and then exploit that reliance for profit[30]. I submit that thinking, learning and relations are too important to outsource to these entities.

It may feel impossible to swim against this tide. But, there are some ways to resist. If you absolutely must use genAI, consider how you can do so in a decentralised fashion. Models which run locally exist, as do ones with open weights. Be mindful that those who use open software are free to do things with that software compared with those who don't[31]. Also keep in mind what is at stake when you use genAI, some of which is described here; some of it is better described elsewhere[32]. Try to be a better person to talk to than an LLM for anyone in need. Aspire to better outputs than any genAI model can generate. See the man behind the curtain; much of the press on genAI centres around imagined magical properties. **genAI is not magical.** It works on defined principles and understanding those principles will empower you to know the scope of the models and their limitations[33, 34]. Understand that while you may be able to 'think' or 'write' faster using an LLM, or draw things you'd never be able to using a diffusion or similar model, understand also that their very design gives outputs based on training inputs. What is used for training is decided by those who own and operate those same models, and there may be 'strategic' omissions from training data on their part. You may end up with polished, professional outputs but they will lack the imperfections that only *you* can give them. The imperfections that make you, *you.*

References

[1] Job apocalypse? Not yet. AI is creating brand new occupations. https://archive.ph/dfbuD, December 2025.

[2] Why AI won't wipe out white-collar jobs. https://archive.ph/8o49V, January 2026.

[3] Cristina B. Gibson, Kevin S. Groves, Jaclyn Margolis, Carol Lusk, and Kevin Sakamoto. The Crisis of Dignity at Work: Dehumanization during Digital Transformation, Societal Upheaval, and Workforce Uncertainty. Academy of Management Discoveries, 11(4):479–484, December 2025. doi: 10.5465/amd.2025.0359.

[4] Real-world gen AI use cases from the world's leading organizations. https://cloud.google.com/transform/101-real-world-generative-ai-use-cases-from-industry-leaders, .

[5] "Empire of AI": Karen Hao on How AI Is Threatening Democracy & Creating a New Colonial World, January 2026. https://www.democracynow.org/2026/1/1/empire_of_ai_karen_hao_on

[6] Platforming Journalists and Their Sources. "The Gen AI industry is a Ponzi scheme" Patrick Galey, Investigative Journalist for Global Witness., June 2026. https://www.youtube.com/watch?v=kmcnO8s68ec

[7] Fish Belk. On the Nature of Tools, August 2024. https://fishbelk.substack.com/p/on-the-nature-of-tools

[8] The Tool of Tools and the Form of Forms – 3 Wisdoms | Scott Randall Paine, .https://3wisdoms.com/2019/10/20/the-tool-of-tools-and-the-form-of-forms/

[9] Ivan Illich. Tools for Conviviality. Calder and Boyars, London, 1973. ISBN 978-0-7145-0973-0.

[10] The Gentle Singularity. https://blog.samaltman.com/the-gentle-singularity, .

[11] AI is helping expand the frontier of theoretical physics. https://archive.ph/dp6CK, March 2026.

[12] Feeling Lonely? An Attentive Listener Is an AI Prompt Away | Working Knowledge. https://www.library.hbs.edu/working-knowledge/feeling-lonely-an-attentive-listener-is-an-ai-prompt-away, December 2025.

[13] Nataliya Kosmyna, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task, December 2025.

[14] In praise of grunt work. https://archive.ph/R2VT3, March 2026.

[15] Elizabeth Ligon Bjork and Robert A. Bjork. Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In Psychology and the Real World: Essays Illustrating Fundamental Contributions to Society, pages 56–64. Worth Publishers, New York, NY, US, 2011. ISBN 978-1-4292-3043-8.

[16] Mariana Lenharo. Is AI ruining our skills? Early results are in — and they're not good. Nature, 655(8121):15–16, July 2026. ISSN 0028-0836, 1476-4687. doi: 10.1038/d41586-026-01947-1.

[17] Daisy Christodoulou. What is the point of learning to write in a world with AI? https://substack.nomoremarking.com/p/what-is-the-point-of-learning-to, August 2025.

[18] Timothy Daly. A 'low-tech' Academic Virtue Ethics in the Age of Generative AI. Journal of Academic Ethics, 24(1): 13, October 2025. ISSN 1572-8544. doi: 10.1007/s10805-025-09683-3.

[19] CDC. Health Effects of Social Isolation and Loneliness. https://www.cdc.gov/social-connectedness/risk-factors/index.html, February 2025.
[20] The Anti-Social Century - The Atlantic. https://archive.li/wxY8p, January 2025.
[21] Milena Batanova, Richard Weissbourd, and Joseph McIntyre. Loneliness in America: Just the Tip of the Iceberg? https://relationalwealthadvisors.com/wp-content/uploads/2025/03/Loneliness-In-America_Batanova-Weissbourd-McIntyre-Harvard-Brief-Report-2024.pdf
[22] Guy Laban, Arvid Kappas, Val Morrison, and Emily S. Cross. Building Long-Term Human–Robot Relationships:Examining Disclosure, Perception and Well-Being Across Time. International Journal of Social Robotics, 16(5):1–27, May 2024. ISSN 1875-4805. doi: 10.1007/s12369-023-01076-z.
[23] Julian De Freitas, Ahmet Kaan Uğuralp, Zeliha Uğuralp, and Stefano Puntoni. AI Companions Reduce Loneliness. SSRN Electronic Journal, 2024. ISSN 1556-5068. doi: 10.2139/ssrn.4893097.
[24] Shunsen Huang, Xiaoxiong Lai, Li Ke, Yajun Li, Huanlei Wang, Xinmei Zhao, Xinran Dai, and Yun Wang. AI Technology panic—is AI Dependence Bad for Mental Health? A Cross-Lagged Panel Model and the Mediating Roles of Motivations for AI Use Among Adolescents. Psychology Research and Behavior Management, 17:1087–1102, March 2024. ISSN 1179-1578. doi: 10.2147/PRBM.S440889.
[25] Minds in Crisis: How the AI Revolution is Impacting Mental Health. Journal of Mental Health & Clinical Psychology, 9(3), May 2025.
[26] Cathy Mengying Fang, Auren R. Liu, Valdemar Danry, Eunhae Lee, Samantha W. T. Chan, Pat Pataranutaporn, Pattie Maes, Jason Phang, Michael Lampe, Lama Ahmad, and Sandhini Agarwal. How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use: A Longitudinal Randomized Controlled Study, October 2025. https://arxiv.org/abs/2503.17473
[27] Elizabeth W. Dunn, Lara B. Aknin, and Michael I. Norton. Spending Money on Others Promotes Happiness. Science, 319(5870):1687–1688, March 2008. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1150952.
[28] Ivan Mehta. ChatGPT’s market share slips below 50% for first time, June 2026. https://techcrunch.com/2026/06/16/chatgpts-market-share-slips-below-50-for-first-time/
[29] Silicon Analysts. NVIDIA AI GPU Market Share 2026: ~80% of AI Accelerators. https://siliconanalysts.com/analysis/nvidia-ai-accelerator-market-share-2024-2026, February 2026.
[30] Democracy Now! "Enshittification": Cory Doctorow on Why Big Tech Sucks, Keeps Getting Worse & What to Do About It, October 2025. https://www.democracynow.org/2025/10/10/cory_doctorow
[31] What is Free Software? - GNU Project - Free Software Foundation. https://www.gnu.org/philosophy/free-sw.en.html, .
[32] Offloading Ourselves, https://www.thenewatlantis.com/publications/offloading-ourselves-staying-human-in-the-age-of-ai.
[33] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, pages 610–623, Virtual Event Canada, March 2021. ACM. ISBN 978-1-4503-8309-7. doi: 10.1145/3442188.3445922.
[34] Deep Dive into LLMs like ChatGPT - YouTube. https://www.youtube.com/watch?v=7xTGNNLPyMI, .