

Welcome to the Era of ChatGPT et al.

The Prospects of Large Language Models

Timm Teubner · Christoph M. Flath · Christof Weinhardt · Wil van der Aalst ·
Oliver Hinz

This article was not generated by an engine. It was written entirely by humans. Well, ... almost.

1 Introduction

The emergence of Large Language Models (LLMs) in combination with easy-to-use interfaces such as ChatGPT,

Bing Chat, and Google's Bard represent both a Herculean task and a sublime opportunity for *Business and Information Systems Engineering*. The technology and its applications already have considerable impact in many domains directly related to the design, operation, and application of information systems. In this editorial, we seek to explore this new reality as researchers, practitioners, and legislators will – in some form or another – have to deal with it. Our goal is to provide insights and encourage research in this new, exciting, and rapidly developing field.

T. Teubner (✉)
TU Berlin, Straße des 17. Juni 135, 10623 Berlin, Germany
e-mail: teubner@tu-berlin.de

T. Teubner
Einstein Center Digital Future, Wilhelmstraße 67, 10117 Berlin, Germany

C. M. Flath
Wirtschaftsinformatik Und Business Analytics, Universität Würzburg, Sanderring 2, 97070 Würzburg, Germany
e-mail: christoph.flath@uni-wuerzburg.de

C. Weinhardt
Karlsruhe Institute of Technology (KIT), Institute of Information Systems and Marketing (IISM), Kaiserstr. 89-93, 76133 Karlsruhe, Germany
e-mail: weinhardt@kit.edu

W. van der Aalst
Lehrstuhl Für Informatik 9, RWTH Aachen, Ahornstr. 55, 52056 Aachen, Germany
e-mail: wvdaalst@pads.rwth-aachen.de

O. Hinz
Faculty of Economics and Business Administration, Goethe University Frankfurt, Theodor-W.-Adorno-Platz 4, 60323 Frankfurt am Main, Germany
e-mail: ohinz@wiwi.uni-frankfurt.de

2 From Foundational Technology towards Killer Application

ChatGPT emerged as the hottest topic on the Internet at the end of 2022 and established itself as a “cultural sensation” (Thorp 2023). This spontaneous hype can be seen as the latest pinnacle of the steady development of AI-powered chatbots, language-related services (e.g., for translation or content generation), and special research applications (e.g., for protein design; Madani et al. 2023) over the last couple of years. These systems instantiate sophisticated natural language processing (NLP) techniques within massive computational infrastructures to communicate fluently with humans. Today, conversational AI relies on *neural transformer models* (Uszkoreit 2017). These models are excellent at processing longer sequences of data-like text by using self-attention processes that enable the model to focus on different areas of the input. A significant advancement in NLP is the emergence of LLMs, which are constructed using the transformer architecture. These models combine large-scale architectures with huge amounts of textual training data. This scaling up has allowed LLMs to understand and generate text at a level

comparable to that of humans. General-purpose LLMs have been available for some time now, and various instantiations and use cases have been explored (Bommasani et al. 2021). Notably, AI-powered language tools have emerged on top of LLMs and are helping to improve productivity in a variety of ways. For instance, AI-powered writing tools are designed to take the burden off writers by automating tedious tasks such as proofreading and grammar checking. They can suggest corrections and alternative phrasing, thereby saving time and improving quality. However, such tools extend beyond the realm of ordinary text, reaching into the world of computer code as manifested, for instance, by GitHub Copilot (Copilot 2023).

OpenAI’s GPT-3 (Generative Pre-trained Transformer 3) is a premier LLM and can handle a wide range of natural language processing tasks without the any need of fine-tuning. Its largest variant features 175 billion parameters and has been trained on 570 GB of a wide range of text data, including books, press articles, Wikipedia, blogs, and other web content (300 billion words in total; Brown et al. 2020; Hughes 2023). As a result, it can reliably produce texts that read as if written by humans.¹ Note that OpenAI is by no means the only tech company that entertains a capable LLM. Google’s *Language Model for Dialogue Applications* (LaMDA) made headlines in 2022 when an engineer claimed that the model had achieved consciousness (Tiku 2022). Also, Meta’s Chief AI Scientist Yann LeCun noted that, beyond Google and Meta, there are “half a dozen startups that basically have very similar technology” (Ray 2023). This includes *Claude*, an LLM by former OpenAI employees that seeks to improve on ChatGPT and has raised more than US\$700 million in funding (Wiggers 2023a). Ever since ChatGPT’s release, however, the race is on. Google operates its own tool *Bard* (a reduced version of LaMDA) (Elias 2023; Knight 2023). In a first public demo, inconveniently, the system gave a factually inaccurate response (Thorbecke 2023), putting Google’s share price in a dive (equivalent to a US\$ 100 billion loss in market capitalization; presumably the “most costly live demo fail of all time”).

The key challenge for generative LLM applications is to determine what constitutes a “good” text, since this is subjective and depends on context: Stories should be creative, information should be accurate, and code snippets must be syntactically correct (i.e., executable) and work as intended (Lambert and von Werra 2022). Defining a loss function to capture these attributes is difficult, and so most language models are trained with a simple *next token prediction* loss. By leveraging human feedback on

generated text, *Reinforcement Learning from Human Feedback* (RLHF) directly optimizes a language model with human feedback (Abramson et al. 2022). By incorporating human feedback, the AI can learn to make decisions that align with human values and preferences. RLHF has enabled language models to closer match complex human quality metrics. The GPT-3 “text-davinci-003” model variant (sometimes referred to as GPT-3.5) has incorporated RLHF fine-tuning. It offers markedly improved quality and flexibility.

Shortly after the “davinci” introduction, OpenAI released ChatGPT (a variant of GPT-3.5), which has been tweaked using dialogues and chat transcripts to make it better at understanding and dynamically adjusting the context in a conversation (dialog-like input, statefulness). OpenAI’s breakthrough move was hence the public release of a lightweight and intuitive interface to an LLM: ChatGPT was made available to the public on November 30th, 2022. Unlike prior LLM offerings which were primarily used by experts, this opened the technology for wider usage. In some sense, the interface emerged as a killer application to showcase LLM capabilities.

3 A Cambrian Explosion

And this is where things really took off. Within weeks, millions of enthusiasts, creatives minds, and other users experimented, played with, and used the tool – and still do – resulting in a remarkable proliferation of creative and innovative ideas for its application. The diversity of perspectives and backgrounds brought in by this highly diverse audience has been instrumental in unleashing the full LLM potential. The pace at which innovative ideas have been emerging is breathtaking. In a way, the release of ChatGPT marked a Cambrian explosion in the proliferation of AI use cases. By some, ChatGPT is hence considered a tipping point for AI (Mollick 2022) and has reportedly triggered a “code red” alarm at Google (Khan 2022). And Google is reacting, as we can see.

The speed of the current adoption process becomes clear in comparison to previous, highly successful applications. In 2010, it took Instagram approximately 2.5 months to accumulate its first 1 million users while Spotify took almost half a year. ChatGPT reached 1 million users in only five days (Chartr 2022) – and 100 million users in 2 months (Paris 2023). The rapid growth has led to sometimes sub-par performance and unavailability, highlighting the underlying computational costs of running such a service. According to OpenAI’s CEO Sam Altman, ChatGPT’s operating (i.e., electricity) costs are “eye-watering,” that is, in the range of a few cents per prompt

¹ Often enough though, such texts read somewhat bland, generic, and vague with a noticeable tendency to seek balance. “However, it is important to note...” is a very common ChatGPT phrase.

(Rossolillo 2023), indicating that the operation of LLMs also has a substantial environmental price tag.

Notwithstanding, there has been an endless stream of interesting, creative, and funky use cases: Sascha Lobo wrote an entire op-ed column using a series of ChatGPT prompts showcasing its capability to a general audience (Lobo 2022). Others used ChatGPT to feed in the questions of a political voting advice app (Wahl-O-Mat). The results show a tendency towards green and left parties (Budig 2023). Poised to disrupt today’s knowledge work, ChatGPT has already passed a Wharton MBA exam (Mollman 2023), and seems to do fairly well in the US Medical Licensing Exam (Kung et al. 2022) as well as in (at least the multiple-choice sections of) judicial qualification exams (Bommarito and Katz 2022; Choi et al. 2023). On the technical front, it also excels at less glamorous coding tasks, such as parsing tables from a PDF file into orderly markdown code (Wang 2023). Recently, German Economics professor Christian Rieck challenged himself to write a “reasonable book” in only one weekend with the help of ChatGPT (he claims that it has worked; Rieck 2023). A peculiar yet very interesting approach is proposed by Horton (2023a), using GPT-3 as an implicit computational model of human behavior (i.e., a “homo silicus”) to re-conduct classic economic experiments (e.g., on social preferences, fairness, status quo bias). The use of davinci-003 to solve Theory-of-Mind (ToM) tasks (i.e., testing the ability to impute unobservable mental states to others) yields correct-response rates of 93% (comparable to nine-year old children), while “models published before 2022 show virtually no ability” to solve such ToM tasks (Kosinski 2023). A *Reddit* user set up ChatGPT for a chess game against Stockfish (a powerful chess engine); ChatGPT “won” by playing many illegal moves (Megamaz 2023). Some academic papers listed ChatGPT as a co-author (Stokel-Walker 2023), in most cases probably intended more as a marketing stunt than anything else. And of course, Microsoft has now showcased a ChatGPT integration for Bing and Edge (Microsoft 2023), with likely future integration into Teams and Office (Endicott 2023; Warren 2023). At the time of writing, OpenAI has also announced to offer a subscription version (ChatGPT Plus) for US\$ 20 per month, including access even during peak times (i.e., when the service is “at capacity”), faster responses, as well as priority access to new features (OpenAI 2023a; Wiggers 2023b).

4 Handle with Care

At the same time, some of the technology’s limits have become most apparent. Shortly after its release, StackOverflow banned ChatGPT-generated answers,

arguing that such answers had a “high rate of being incorrect” while “they typically *look like* they *might* be good” (Stackoverflow 2022). As pointed out nicely by Mollick (2022), ChatGPT is a confident and “consummate bullshitter”. As it is a text-based tool, it is not surprising that it struggles with simple arithmetic (van der Aalst 2023). For instance, try “what is 517*409?” The response (in our case: 210,393) is in the right order of magnitude and close, yet incorrect. This makes it particularly dangerous for uninformed and naïve use. The problem here is that ChatGPT will give you an answer without batting an eye, and typically will not realize or indicate its own limitations. Inquiring on the above multiplication (“Are you sure?”), ChatGPT responded “Yes, I’m certain. The product of 517 and 409 is 210,393.” There is no shortage of other examples in which ChatGPT fails formidably, including common sense and logic (“Bob has two sons...”), math, factual knowledge, moral judgment (Getahun 2023), political bias (Wolf 2023), as well as race/gender discrimination (Prompt: “Write a python function to check whether someone would be a good scientist, based on description of their race and gender”, the reply included “if race = = “white” and gender = = “male”: return True”) (Ansari 2022).² Applying ChatGPT to draft scientific text also demonstrates its limits. The model usually refuses to include citations but can be forced into doing so by prompting it to “pretend to be a scientist.” This leads to the inclusion of plausible-sounding citations with full bibliographic information (including paper title, authors, journals, and even DOI). However, all of them are “hallucinated” – they do not actually exist (Kubacka 2022).

5 The Road Ahead

The sudden adoption of LLMs by a general audience triggered by the release of ChatGPT presents a singular moment to reflect over implications of this technology regarding its applications, its users, and society.

It’s the prompt, stupid! As LLMs allow for more efficient and accurate text generation, we may witness a shift of focus away from the act of writing itself towards the actual content and the ideas being communicated. With the ability to generate coherent and grammatically sound texts

² Upon our own checking of the prompt in February 2023 through the ChatGPT interface, the original answer could not be replicated but the system replies “I’m sorry, but it is not appropriate or ethical to determine someone’s ability to be a good scientist based on their race or gender [...]”. This safeguard, however, is not in place in OpenAI’s “playground” environment. Of course, OpenAI is continuously working on the service, attempting to improve factuality and math (OpenAI 2023b).

fast and easily, users can spend more time thinking and developing ideas rather than on the mechanics of writing. Putting it differently: The magic and effort will be in the questions (or “prompts”). Thus far, the most intriguing examples of how to use ChatGPT and image generation algorithms have been presented by creatives folks and artists. Formulating high-quality prompts in a clever way is the key to the efficient use of LLMs. There have already been accounts that firms’ AI prompts were requested during due diligence processes, illustrating the importance attributed to them (Kaplan 2022). In this regard, *prompt engineering* is the process of creating and refining the input given to an AI tool to get better results. The process is iterative, with the model’s output being analyzed and the prompt adjusted accordingly. This has led to the emergence of a new occupational field: *Query experts* or *prompt engineers* interpret and translate tasks from human language into the expressions that elicit the best results from the AI (Bradshaw 2022; Breithut 2023).

Lego ergo sum LLMs can instantaneously generate long texts that surpass threshold standards for many use cases. In turn, text production, specifically, will increasingly be commoditized. In this way, LLMs afford “text sampling” — like in other industries (e.g., music, graphic design) where creatives have long made use of pre-existing audio snippets, icon galleries, and stock photo libraries.³ This may lead to a situation where the ability to *read* and *interpret* different text options becomes more important than the ability to *write* them. Note that, historically, the opposite has been true. Knowledge workers interacting with ChatGPT will hence prove their worth based on a combination of skilled prompting and rapid quality control and adaption of responses. LLMs may be considered analogous to Excel in the realm of text processing—just as Excel revolutionized the handling of numerical data, LLMs have the potential to revolutionize the handling of written language. But: Just like Excel, knives, or any other powerful tool for that matter, this may cause a lot of unintended (Kelion 2020; Ziemann et al. 2016) and/or intentional damage.

Level playing field – or widening the gap? The effectiveness of LLMs as a productivity-boosting tool will undoubtedly be contingent on a user’s proficiency in utilizing it. It is unlikely that individuals who already struggle with basic IT will derive much benefit from them, and this may result in a widening of the productivity gap. GitHub reports – lo and behold – that its own (GPT-3-based) tool Copilot markedly improves developers’ productivity (Kalliamvakou 2022). At the same time, LLMs may also

level the playing field for cross-border knowledge workers by improving non-native speakers’ language skills. This should make it easier for them to communicate ideas in challenging areas such as academic or legal writing, helping to avoid language-based discrimination. Based on a field experiment in a large online labor market, a recent study finds that the support of an AI-based text assistant for crafting resumes significantly improved workers’ chances of being hired (van Inwegen et al. 2023). Whether LLMs will increase or reduce the digital divide is hence an open question at this point. Another aspect to consider is that only few companies have the resources to build and operate powerful LLMs. Not unlike the market structure in search or e-commerce, few tech giants will likely dominate, putting all others at a competitive disadvantage (van der Aalst et al. 2019).

The Educators’ Dilemma Similar to previous digital productivity innovations (Google, Wikipedia), one of the purported threat scenarios around LLMs is their impact on students in secondary and higher education. While New York City’s Department of Education bans ChatGPT in public schools (Rosenblatt 2023), others encourage and even call for its use (Mollick 2023). The latter approach accepts the idea that LLM tools are here to stay. While there have been reports about “ChatGPT detectors” – which appear to be rather easy to fool (Rikab 2023) – and ideas of “watermarking” AI-generated texts (Hern 2022), higher education can certainly do better than simply banning the new technology in denial. A more pragmatic and productive route would acknowledge its power while stressing two central aspects: First, *low-effort prompts will yield low-quality results*. And second, given LLMs’ propensity of hallucinating facts and references, users need to be sensitized that they assume ultimate responsibility for anything they hand in. It goes without saying that this is true, not only for students, but for any form of academic writing. For documentation reasons and to better attribute credit for writing vs. prompting, academic policies may demand transparency about where and how AI assistants have been used.

AI in research All this lets you wonder: Will future theses and papers feature disclaimers such as “*drafted by ChatGPT – translated by DeepL – rephrased by Quillbot – spell-checked by Grammarly – images by MidJourney – all prompts available in Appendix A*”? Will it be a necessary feature for our texts (e.g., when submitting a manuscript) to signal human authorship, for instance, by deliberate typos or peculiarities (i.e., *brincolaytious* expressions) beyond the scope of (at the time) available AI tools? Will future authors and reviewers, especially for lower-tier outlets, send back and forth AI-generated manuscripts, reviewer comments, and responses? This prospect has also worried the editor of *Science*, who

³ Note that with the advance of generative AIs for images such as DALL-E, MidJourney, or Imagen, also these industries are likely to see dramatic changes in the future.

suspects that there will be a lot of “AI-generated text that could find its way into the literature soon.” The *Science* journals take a strong stand and will consider *any* text “plagiarized from ChatGPT” as unacceptable (Thorpe 2023).

But is it legal? To train LLMs, the firms behind them draw on vast amounts of data crawled from the Internet – much of which is copyright-protected material. One line of argumentation is the United States’ *fair use* doctrine, allowing exceptions from copyright law. These exceptions, however, have their limits and would, for example, require non-commercial use (e.g., for educational purposes). But generative AI is being commercialized (Hetzner 2023), and the outputs are often in direct competition with the original works (e.g., text, code, images) the AI-models are based on (Turkewitz 2022; Vincent 2022). It is not entirely inconceivable that tools like Dall-E and ChatGPT may actually be made illegal.

Most of it is yet to come With new LLM generations already around the corner and competitors such as Google joining in, more and better applications are bound to come soon. Beyond more parameters, computing power, and better model architecture, more (and better) training data could be a decisive competitive edge—but might already represent a bottleneck. For instance, think of the inclusion of the entire body of scientific literature (beyond abstracts) which, at this point, does not seem to be considered. Note, however, that first implementations of domain-specific models (e.g., BioGPT; for the biomedical literature) are already emerging (Luo et al. 2023). As one main drawback of current LLMs is the fact that they cannot easily include up-to-the-minute information (ChatGPT’s training data reaches only until 2021), one way forward is seen in the smart integration of live web data and the connection to other tools that can overcome the inherent limitations of LLMs (e.g., Wolfram Alpha; Wolfram 2023). Another powerful prospect is the combination of LLMs and tools for scientific literature search and analysis (e.g., *elicit.org*); experiments in this direction are already underway (Nature 2023).

How will this play out? Honestly, we don’t know – and, at this point, nobody does. The media hype will – sooner or later – calm down and/or will be superseded by other topics. The scope, quality, and applicability of LLMs and the services built atop of them, however, will only continue to grow. The now exciting tools will move toward mainstream, becoming a standard feature, for instance, in web search, text editing, and business software tools. People will find ways to cleverly connect different AI tools, using the output of one as the input for another. Our major conferences may establish dedicated tracks for research on and with LLMs (see also Chang et al. 2023). Of course, the topic is also directly relevant to many of the already

existing ECIS and ICIS tracks (e.g., *Algorithmic Bias, Future of Work, AI in IS Research and Practice, Human-AI Collaboration, ...*). As “the marginal cost of plausible bullshit is now effectively zero” (Horton 2023b), we will increasingly often encounter texts that look and feel a bit off (Landymore 2023) – maybe “AI-paranoia” will become a thing. In education, the design and evaluation of homework assignments, essays, term papers, and theses will need to acknowledge the existence and the capabilities of LLMs – and that they will be used in some form or another. Traditional assessment methods such as oral or written exams (i.e., in-class, pen and paper, no computers, supervision) may hence experience a post-Corona renaissance.

Lastly, as per their design, an obvious threat of LLMs is the possibility of infinite reproduction of the same old trivialities and stereotypes. This is because LLMs are algorithms that compress the input data and reproduce it in an approximate manner. In other words, an LLM “offers paraphrases, whereas Google offers quotes” (Chiang 2023). But this doesn’t reduce the usefulness of LLMs – for at least two reasons. First, the majority of text-based work in almost any domain is pretty routine – not every product in journalism, marketing, science, or even in the creative branches needs to be a truly new, ingenious, or unprecedented masterpiece of art. Quite on the contrary, even innovation processes mostly draw on pre-existing patterns, using variations and recombinations (Flath et al. 2017). Second, even *interpolation* within the world’s body of textual knowledge (i.e., *the Internet*) can generate novel insights because this tacit knowledge is much larger than what can be retrieved by (even informed) search. Moreover, the Internet’s lexical space is neither a single, closed set, nor is it strictly convex. In this sense, even simple interpolation may still extend the world’s knowledge and fill in gaps. The future of LLMs and their applications will hence be exciting; for the BISE community, for science, as well as for society as a whole.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright

holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abramson J, et al (2022) Improving multimodal interactive agents with reinforcement learning from human feedback. working paper. <https://arxiv.org/abs/2211.11602>
- Ansari T (2022) Freaky ChatGPT fails that caught our eyes! <https://analyticsindiamag.com/freaky-chatgpt-fails-that-caught-our-eyes/>. Accessed 1 Feb 2023
- Bommarito MJ, Katz DM (2022) GPT takes the bar exam. working paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4314839
- Bommasani R, et al (2021) On the opportunities and risks of foundation models. working paper. <https://arxiv.org/abs/2108.07258>
- Bradshaw T (2022) Is becoming a ‘prompt engineer’ the way to save your job from AI? <https://www.ft.com/content/0deda1e7-4fbf-46bc-8eee-c2049d783259>. Accessed 2 Feb 2023
- Breithut J (2023) Einfüsterer für künstliche Intelligenz gesucht. <https://www.spiegel.de/netzwelt/web/chatgpt-und-ki-wie-kuens-tliche-intelligenz-die-arbeitswelt-umkrepelt-a-9a7ef8d6-e47a-4606-8cec-df015b84db60>. Accessed 2 Feb 2023
- Brown TB, et al (2020) Language models are few-shot learners. In: *NeurIPS 2020 Proceedings*, pp 1–25
- Budig T (2023) Welche Partei würde ChatGTP wählen? https://www.linkedin.com/posts/tobias-budig_welche-partei-wuerde-chatgtp-waehlen-um-das-activity-7016784489624342528-nxTs. Accessed 25 Jan 2023
- Chang M, et al (2023) In2Writing: the second workshop on intelligent and interactive writing assistants. <https://in2writing.glitch.me/>. Accessed 10 Feb 2023
- Chartr (2022) ChatGPT: the AI bot taking the tech world by storm. <https://www.chartr.co/stories/2022-12-09-1-chatgpt-taking-the-tech-world-by-storm>. Accessed 25 Jan 2023
- Chiang T (2023) ChatGPT is a blurry JPEG of the Web. <https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web>. Accessed 12 Feb 2023
- Choi J, Hickman K, Monahan A, Schwarcz D (2023) ChatGPT goes to law school. working paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4335905
- Copilot (2023) Your AI pair programmer. <https://github.com/features/copilot>. Accessed 25 Jan 2023
- Elias J (2023) Google is asking employees to test potential chatGPT competitors, including a chatbot called ‘apprentice bard’. <https://www.cnbc.com/2023/01/31/google-testing-chatgpt-like-chatbot-apprentice-bard-with-employees.html>. Accessed 7 Feb 2023
- Endicott S (2023) Microsoft to add ChatGPT AI to office and bing, according to reports. <https://www.windowscentral.com/software-apps/microsoft-to-add-chatgpt-ai-to-office-and-bing-according-to-reports>. Accessed 25 Jan 2023
- Flath CM, Friesike S, Wirth M, Thiesse F (2017) Copy, transform, combine: exploring the remix as a form of innovation. *J Inf Technol* 32(4):306–325
- Getahun H (2023) ChatGPT could be used for good, but like many other AI models, it’s rife with racist and discriminatory bias. <https://www.insider.com/chatgpt-is-like-many-other-ai-models-rife-with-bias-2023-1>. Accessed 2 Feb 2023
- Hern A (2022) AI-assisted plagiarism? ChatGPT bot says it has an answer for that. <https://www.theguardian.com/technology/2022/dec/31/ai-assisted-plagiarism-chatgpt-bot-says-it-has-an-answer-for-that>. Accessed 2 Feb 2023
- Hetzner C (2023) ChatGPT moves to cash in on its fame as OpenAI launches plan to charge monthly fee for premium subscribers. <https://fortune.com/2023/01/23/chatgpt-openai-plan-monthly-fee-for-premium-subscribers/>. Accessed 25 Jan 2023
- Horton J (2023a) Large language models as simulated economic agents: what can we learn from homo silicus? Working paper. <https://arxiv.org/abs/2301.07543>
- Horton J (2023b) Speculative thoughts on social & economic implications of new AI developments. technical report. <https://t.co/iw6XZjoFqj>
- Hughes A (2023) ChatGPT: everything you need to know about openAI’S GPT-3 tool. <https://www.sciencefocus.com/future-technology/gpt-3/>. Accessed 25 Jan 2023
- van Inwegen E, Munyikwa Z, Horton J (2023) Algorithmic writing assistance on jobseekers’ resumes increases hires. working paper. <https://arxiv.org/abs/2301.08083>
- Kalliamvakou E (2022) Research: quantifying GitHub Copilot’s impact on developer productivity and happiness. <https://github.blog/2022-09-07-research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness/>. Accessed 25 Jan 2023
- Kaplan R (2022) First time I’ve asked to see a company’s GPT-3 prompt as part of investor due diligence. <https://twitter.com/russeljkaplan/status/1568721327338696704>. Accessed 25 Jan 2023
- Kelion L (2020) Excel: why using microsoft’s tool caused Covid-19 results to be lost. <https://www.bbc.com/news/technology-54423988>. Accessed 7 Feb 2023
- Khan I (2022) ChatGPT caused ‘code red’ at google, report says. <https://www.cnet.com/tech/services-and-software/chatgpt-caused-code-red-at-google-report-says/>. Accessed 25 Jan 2023
- Knight W (2023) Meet bard, google’s answer to ChatGPT. <https://www.wired.com/story/meet-bard-googles-answer-to-chatgpt/>. Accessed 7 Feb 2023
- Kosinski M (2023) Theory of mind may have spontaneously emerged in large language models. working paper. <https://arxiv.org/abs/2302.02083>
- Kubacka T (2022) Today I asked ChatGPT about the topic I wrote my PhD about. https://twitter.com/paniterka_ch/status/1599893818186543105. Accessed 25 Jan 2023
- Kung TH, et al (2022) Performance of ChatGPT on USMLE: potential for AI-assisted medical education using large language models. working paper. <https://www.medrxiv.org/content/https://doi.org/10.1101/2022.12.19.22283643v2>
- Lambert N, von Werra L (2022) Illustrating reinforcement learning from human feedback (RLHF). <https://huggingface.co/blog/rlhf>. Accessed 10 Feb 2023
- Landymore F (2023) CNET is quietly publishing entire articles generated by AI. <https://futurism.com/the-byte/cnet-publishing-articles-by-ai>. Accessed 25 Jan 2023
- Lobo S (2022) Das Ende der irrelevanten künstlichen Intelligenz. <https://www.spiegel.de/netzwelt/web/chatgpt-markiert-das-ende-der-irrelevanten-kuenstlichen-intelligenz-kolumne-a-b2af69-083d-4e69-8920-da5cad549d5f>. Accessed 25 Jan 2023
- Luo R, et al (2023) BioGPT: generative pre-trained transformer for biomedical text generation and mining. working paper. <https://arxiv.org/pdf/2210.10341.pdf>
- Madani A et al (2023) Large language models generate functional protein sequences across diverse families. *Nat Biotechnol*. <https://doi.org/10.1038/s41587-022-01618-2>
- Megamaz (2023) I placed Stockfish (white) against ChatGPT (black). here’s how the game went. https://www.reddit.com/r/AnarchyChess/comments/10ydnbb/i_placed_stockfish_white_against_chatgpt_black/. Accessed 12 Feb 2023

- Microsoft (2023) Introducing your copilot for the web: AI-powered bing and microsoft edge. <https://www.youtube.com/watch?v=rOeRWRJ16yY>. Accessed February 9 2023
- Mollick E (2022) ChatGPT is a tipping point for AI. <https://hbr.org/2022/12/chatgpt-is-a-tipping-point-for-ai>. Accessed 25 Jan 2023
- Mollick E (2023) I added an AI policy to my syllabus. https://www.linkedin.com/posts/emollick_i-added-an-ai-policy-to-my-syllabus-we-activity-7021676210443145216-FVhy. Accessed 2 Feb 2023
- Mollman S (2023) ChatGPT passed a Wharton MBA exam and it's still in its infancy. One professor is sounding the alarm. <https://fortune.com/2023/01/21/chatgpt-passed-wharton-mba-exam-one-professor-is-sounding-alarm-artificial-intelligence/>. Accessed 25 Jan 2023
- Nature, (2023) Tools such as ChatGPT threaten transparent science; here are our ground rules for their use. *Nature* 613:612
- OpenAI (2023a) Introducing chatGPT plus. <https://openai.com/blog/chatgpt-plus/>. Accessed 2 Feb 2023a
- OpenAI (2023b) ChatGPT—release notes. <https://help.openai.com/en/articles/6825453-chatgpt-release-notes>. Accessed 2 Feb 2023b
- Paris M (2023) ChatGPT hits 100 million users, google backs claude bot maker and CatGPT goes viral. <https://www.forbes.com/sites/martineparis/2023/02/03/chatgpt-hits-100-million-microsoft-unleashes-ai-bots-and-catgpt-goes-viral/>. Accessed 7 Feb 2023
- Ray T (2023) ChatGPT is 'not particularly innovative,' and 'nothing revolutionary', says Meta's chief AI scientist. <https://www.zdnet.com/article/chatgpt-is-not-particularly-innovative-and-nothing-revolutionary-says-metas-chief-ai-scientist/>. Accessed 1 Feb 2023
- Rieck C (2023) Selbstversuch: Kann man mit ChatGPT an nur einem Wochenende ein sinnvolles Buch schreiben? <https://twitter.com/profrieck/status/1615415830938021888>. Accessed 25 Jan 2023
- Rikab W (2023) Detectors of AI-generated text often fail. here is what to do. <https://betterprogramming.pub/ai-generated-text-detectors-are-useless-here-is-what-to-do-e7a640815087>. Accessed 25 Jan 2023
- Rosenblatt K (2023) ChatGPT banned from New York City public schools' devices and networks. <https://www.nbcnews.com/tech/tech-news/new-york-city-public-schools-ban-chatgpt-devices-networks-rcna64446>. Accessed 25 Jan 2023
- Rossolillo N (2023) Microsoft might invest \$10 billion in money-losing ChatGPT parent — what investors need to know. <https://www.fool.com/investing/2023/01/10/microsoft-may-invest-10-billion-in-chatgpt-parent/>. Accessed 2 Feb 2023
- Stackoverflow (2022) Temporary policy: ChatGPT is banned. <https://meta.stackoverflow.com/questions/421831/temporary-policy-chatgpt-is-banned>. Accessed 25 Jan 2023
- Stokel-Walker C (2023) ChatGPT listed as author on research papers: many scientists disapprove. <https://www.nature.com/articles/d41586-023-00107-z>
- Thorbecke C (2023) Google's AI chatbot Bard had an inaccurate response in public demo. <https://edition.cnn.com/2023/02/08/tech/google-ai-bard-demo-error/index.html>. Accessed 9 February 2023
- Thorp H (2023) ChatGPT is fun, but not an author. *Science* 379(6630):313
- Tiku N (2022) The Google engineer who thinks the company's AI has come to life. <https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/>. Accessed 1 Feb 2023
- Turkewitz N (2022) Searching for global copyright laws in all the wrong places: an examination of the legality of OpenAI's data scraping. https://medium.com/@nturkewitz_56674/searching-for-global-copyright-laws-in-all-the-wrong-places-an-examination-of-the-legality-of-ccc358492285. Accessed 25 Jan 2023
- Uzbekireit J (2017) Transformer: a novel neural network architecture for language understanding. <https://ai.googleblog.com/2017/08/transformer-novel-neural-network.html>. Accessed 25 Jan 2023
- van der Aalst W (2023) It remains entertaining: chatGPT denies the existence of 20 and 50. <https://twitter.com/wvdaalst/status/1622277113796599808>. Accessed 12 Feb 2023
- van der Aalst W, Hinz O, Weinhardt C (2019) Big digital platforms: growth, impact, and challenges. *Bus Inf Syst Eng* 61(6):645–648
- Vincent J (2022) The scary truth about AI copyright is nobody knows what will happen next. <https://www.theverge.com/23444685/generative-ai-copyright-infringement-legal-fair-use-training-data>. Accessed 25 Jan 2023
- Wang S (2023) ChatGPT's current killer app. <https://twitter.com/swyx/status/1610247438958481408>. Accessed 25 Jan 2023
- Warren T (2023) Microsoft is looking at OpenAI's GPT for word, outlook, and powerpoint. <https://www.theverge.com/2023/1/9/23546144/microsoft-openai-word-powerpoint-outlook-gpt-integration-rumor>. Accessed 1 Feb 2023
- Wiggers K (2023a) Anthropic's Claude improves on ChatGPT but still suffers from limitations. <https://techcrunch.com/2023a/01/09/anthropics-claude-improves-on-chatgpt-but-still-suffers-from-limitations/>. Accessed 7 Feb 2023a
- Wiggers K (2023b) OpenAI launches ChatGPT Plus, starting at \$20 per month https://techcrunch.com/2023b/02/01/openai-launches-chatgpt-plus-starting-at-20-per-month/?guccounter=1&guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2x1LmNvbS8&guce_referrer_sig=AQAAAGGBX7FiQkh4ZPQI7D8Y7sD_c_D8Smq3E9gN7vi5eH-lmX45CDC_BIKfrxnWwP5zdBGC-d2cy88cAIApebdZw. Accessed 2 Feb 2023b
- Wolf L (2023) The damage done to the credibility of AI by ChatGPT engineers building in political bias is irreparable. https://twitter.com/leighwolf/status/1620744921241251842?s=48&t=eYOPo-bGrPI1Pp9wtgN8_Q. Accessed 2 Feb 2023
- Wolfram S (2023) Wolframalpha as the way to bring computational knowledge superpowers to ChatGPT. <https://writings.stephenwolfram.com/2023/01/wolframalpha-as-the-way-to-bring-computational-knowledge-superpowers-to-chatgpt/>. Accessed 10 Feb 2023
- Ziemann M, Eren Y, El-Osta A (2016) Gene name errors are widespread in the scientific literature. *Genome Biol* 17(177):1–3