

GPT-3

Generative Pre-trained Transformer 3 (**GPT-3**; stylized **GPT·3**) is an autoregressive language model that uses deep learning to produce human-like text.

The architecture is a standard Transformer network (with a few engineering tweaks) with the unprecedented size of 2048-token-long context and 175 billion parameters (requiring 800 GB of storage). The training method is "generative pretraining", meaning that it is trained to predict what the next token is. The model demonstrated strong few-shot learning on many text-based tasks.

It is the third-generation language prediction model in the GPT-n series (and the successor to GPT-2) created by OpenAI, a San Francisco-based artificial intelligence research laboratory.^[2] GPT-3's full version has a capacity of 175 billion machine learning parameters. GPT-3, which was introduced in May 2020, and was in beta testing as of July 2020,^[3] is part of a trend in natural language processing (NLP) systems of pre-trained language representations.^[1]

The quality of the text generated by GPT-3 is so high that it can be difficult to determine whether or not it was written by a human, which has both benefits and risks.^[4] Thirty-one OpenAI researchers and engineers presented the original May 28, 2020 paper introducing GPT-3. In their paper, they warned of GPT-3's potential dangers and called for research to mitigate risk.^{[1]:34} David Chalmers, an Australian philosopher, described GPT-3 as "one of the most interesting and important AI systems ever produced."^[5]

Microsoft announced on September 22, 2020, that it had licensed "exclusive" use of GPT-3; others can still use the public API to receive output, but only Microsoft has access to GPT-3's underlying model.^[6]

An April 2022 review in *The New York Times* described GPT-3's capabilities as being able to write original prose with fluency equivalent to that of a human.^[7]

Generative Pre-trained Transformer 3 (GPT-3)

Original author(s)	<u>OpenAI</u> ^[1]
Initial release	June 11, 2020 (beta)
Repository	<u>github.com</u> <u>/openai/gpt-3</u> (<u>https://github.com/openai/gpt-3</u>)
Type	<u>Autoregressive Transformer language model</u>
Website	<u>openai.com</u> <u>/blog/openai-api</u> (<u>https://openai.com/blog/openai-api</u>)

Contents

Background

Training and capabilities

Reception

Applications

Reviews

Criticism

See also

Background

According to *The Economist*, improved algorithms, powerful computers, and an increase in digitized data have fueled a revolution in machine learning, with new techniques in the 2010s resulting in "rapid improvements in tasks" including manipulating language.^[8] Software models are trained to learn by using thousands or millions of examples in a "structure ... loosely based on the neural architecture of the brain".^[8] One architecture used in natural language processing (NLP) is a neural network based on a deep learning model that was first introduced in 2017—the Transformer.^[9] GPT-n models are based on this Transformer-based deep learning neural network architecture. There are a number of NLP systems capable of processing, mining, organizing, connecting and contrasting textual input, as well as correctly answering questions.^[10]

On June 11, 2018, OpenAI researchers and engineers posted their original paper on generative models—language models—artificial intelligence systems—that could be pre-trained with an enormous and diverse corpus of text via datasets, in a process they called generative pre-training (GP).^[11] The authors described how language understanding performances in natural language processing (NLP) were improved in GPT-n through a process of "generative pre-training of a language model on a diverse corpus of unlabeled text, followed by discriminative fine-tuning on each specific task." This eliminated the need for human supervision and for time-intensive hand-labeling.^[11]

In February 2020, Microsoft introduced its Turing Natural Language Generation (T-NLG), which was claimed to be the "largest language model ever published at 17 billion parameters."^[12] It performed better than any other language model at a variety of tasks which included summarizing texts and answering questions.

Training and capabilities

On May 28, 2020, an arXiv preprint by a group of 31 engineers and researchers at OpenAI described the development of GPT-3, a third-generation "state-of-the-art language model".^{[1][4]} The team increased the capacity of GPT-3 by over two orders of magnitude from that of its predecessor, GPT-2,^[13] making GPT-3 the largest non-sparse language model to date. (In a sparse model, many of its parameters are set to a constant value, so even if there are more total parameters, there is less meaningful information.)^{[1]:14[2]} Because GPT-3 is structurally similar to its predecessors,^[1] its greater accuracy is attributed to its increased capacity and greater number of parameters.^[14] GPT-3's capacity is ten times larger than that of Microsoft's Turing NLG, the next largest NLP model.^[4]

Sixty percent of the weighted pre-training dataset for GPT-3 comes from a filtered version of Common Crawl consisting of 410 billion byte-pair-encoded tokens.^{[1]:9} Other sources are 19 billion tokens from WebText2 representing 22% of the weighted total, 12 billion tokens from Books1 representing 8%, 55 billion tokens from Books2 representing 8%, and 3 billion tokens from Wikipedia representing 3%.^{[1]:9} GPT-3 was trained on hundreds of billions of words and is capable of coding in CSS, JSX, Python, among others.^[3] A 2022 review again highlighted that the training continues to include review of Wikipedia.^[7]

GPT-3 Training Data

Dataset	# Tokens	Weight in Training Mix
<u>Common Crawl</u>	410 billion	60%
WebText2	19 billion	22%
Books1	12 billion	8%
Books2	55 billion	8%
Wikipedia	3 billion	3%

Since GPT-3's training data was all-encompassing, it does not require further training for distinct language tasks.^[3] The training data contains occasional toxic language and GPT-3 occasionally generates toxic language as a result of mimicking its training data. A study from the University of Washington found that GPT-3 produced toxic language at a toxicity level comparable to the similar natural language processing models of GPT-2 and CTRL. GPT-3 produced less toxic language compared to its predecessor model, GPT-1, although it produced both more generations and a higher toxicity of toxic language compared to CTRL Wiki, a language model trained entirely on Wikipedia data.^[15]

On June 11, 2020, OpenAI announced that users could request access to its user-friendly GPT-3 API—a "machine learning toolset"—to help OpenAI "explore the strengths and limits" of this new technology.^{[16][17]} The invitation described how this API had a general-purpose "text in, text out" interface that can complete almost "any English language task", instead of the usual single use-case.^[16] According to one user, who had access to a private early release of the OpenAI GPT-3 API, GPT-3 was "eerily good" at writing "amazingly coherent text" with only a few simple prompts.^[18] In an initial experiment 80 US subjects were asked to judge if short ~200 word articles were written by humans or GPT-3. The participants judged incorrectly 48% of the time, doing only slightly better than random guessing.^[1]

Because GPT-3 can "generate news articles which human evaluators have difficulty distinguishing from articles written by humans,"^[4] GPT-3 has the "potential to advance both the beneficial and harmful applications of language models."^{[1]:34} In their May 28, 2020 paper, the researchers described in detail the potential "harmful effects of GPT-3"^[4] which include "misinformation, spam, phishing, abuse of legal and governmental processes, fraudulent academic essay writing and social engineering pretexting".^[1] The authors draw attention to these dangers to call for research on risk mitigation.^{[1][19]:34}

GPT-3 is capable of performing zero-shot, few-shot and one-shot learning.^[1]

In June 2022, Almira Osmanovic Thunström wrote that GPT-3 was the primary author on an article on itself, that they had submitted it for publication,^[20] and that it had been pre-published while waiting for completion of its review.^[21]

Reception

Applications

- GPT-3, specifically the Codex model, is the basis for GitHub Copilot, a code completion and generation software that can be used in various code editors and IDEs.
- GPT-3 is used in certain Microsoft products to translate conventional language into formal computer code.^[22]
- GPT-3 has been used by Andrew Mayne for AI Writer,^[23] which allows people to correspond with historical figures via email.

- GPT-3 has been used by Jason Rohrer in a retro-themed chatbot project named "Project December", which is accessible online and allows users to converse with several AIs using GPT-3 technology.^[24]
- GPT-3 was used by The Guardian to write an article about AI being harmless to human beings. It was fed some ideas and produced eight different essays, which were ultimately merged into one article.^[25]
- GPT-3 was used in AI Dungeon, which generates text-based adventure games. Later it was replaced by a competing model after OpenAI changed their policy regarding generated content.^[26]

Reviews

- In a July 2020 review in The New York Times, Farhad Manjoo said that GPT-3's ability to generate computer code, poetry, and prose is not just "amazing", "spooky", and "humbling", but also "more than a little terrifying".^[27]
- Daily Nous presented a series of articles by nine philosophers on GPT-3.^[28] Australian philosopher David Chalmers described GPT-3 as "one of the most interesting and important AI systems ever produced".^[5]
- A review in Wired said that GPT-3 was "provoking chills across Silicon Valley".^[29]
- The National Law Review said that GPT-3 is an "impressive step in the larger process", with OpenAI and others finding "useful applications for all of this power" while continuing to "work toward a more general intelligence".^[30]
- An article in the MIT Technology Review, cowritten by Deep Learning critic Gary Marcus,^[31] stated that GPT-3's "comprehension of the world is often seriously off, which means you can never really trust what it says."^[32] According to the authors, GPT-3 models relationships between words without having an understanding of the meaning behind each word.
- Jerome Pesenti, head of the Facebook AI lab, said GPT-3 is "unsafe," pointing to the sexist, racist and other biased and negative language generated by the system when it was asked to discuss Jews, women, black people, and the Holocaust.^[33]
- Nabla, a French start-up specializing in healthcare technology, tested GPT-3 as a medical chatbot, though OpenAI itself warned against such use. As expected, GPT-3 showed several limitations. For example, while testing GPT-3 responses about mental health issues, the AI advised a simulated patient to commit suicide.^[34]
- Noam Chomsky expressed his skepticism about GPT-3's scientific value: "It's not a language model. It works just as well for impossible languages as for actual languages. It is therefore refuted, if intended as a language model, by normal scientific criteria. [...] Perhaps it's useful for some purpose, but it seems to tell us nothing about language or cognition generally."^[35]
- Luciano Floridi and Massimo Chiriatti highlighted the risk of "cheap production of good, semantic artefacts".^[36]

Criticism

GPT-3's builder, OpenAI, was initially founded as a non-profit in 2015.^[37] In 2019, OpenAI did not publicly release GPT-3's precursor model, breaking from OpenAI's previous open-source practices, citing concerns that the model would perpetuate fake news. OpenAI eventually released a version of GPT-2 that was 8% of the original model's size.^[38] In the same year, OpenAI restructured to be a for-profit company.^[39] In 2020, Microsoft announced the company had exclusive licensing of GPT-3 for Microsoft's

products and services following a multi-billion dollar investment in OpenAI. The agreement permits OpenAI to offer a public-facing API such that users can send text to GPT-3 to receive the model's output, but only Microsoft will have access to GPT-3's source code.^[6]

Large language models, such as GPT-3, have come under criticism from Google's AI ethics researchers for the environmental impact of training and storing the models, detailed in a paper co-authored by Timnit Gebru and Emily M. Bender in 2021.^[40]

The growing use of automated writing technologies based on GPT-3 and other language generators, has raised concerns regarding academic integrity^[41] and raised the stakes of how universities and schools will gauge what constitutes academic misconduct such as plagiarism.^[42]

GPT-3 was criticized for its algorithmic bias; for example, it is more likely to associate Islam with terrorism and Black people with crime.^[43]

In its response to the Request for Comments on Intellectual Property Protection for Artificial Intelligence Innovation from the United States Patent and Trademark Office ("USPTO"), OpenAI acknowledges that "copyright protection arises automatically when an author creates an original work and fixes it in a tangible medium, see 17 U.S.C. § 102, the vast majority of content posted online is protected by U.S. copyright laws.^[44] GPT was built with data from the Common Crawl dataset, a conglomerate of copyrighted articles, internet posts, web pages, and books scraped from 60 million domains over a period of 12 years. TechCrunch reports this training data includes copyrighted material from BBC, The New York Times, Reddit, the full text of online books, and more.^[45] In April 2021, a group of computer scientists used a tool that identifies text generated by GPT in an effort to isolate the reason for strange phrases appearing in scientific papers. Cabanac and colleagues ran a selection of abstracts from the journal *Microprocessors and Microsystems* through this tool and discovered "critical flaws", such as nonsensical text and plagiarized text and images.^[46]

See also

- Wu Dao
- GPT-2

References

1. Brown, Tom B.; Mann, Benjamin; Ryder, Nick; Subbiah, Melanie; Kaplan, Jared; Dhariwal, Prafulla; Neelakantan, Arvind; Shyam, Pranav; Sastry, Girish; Askell, Amanda; Agarwal, Sandhini; Herbert-Voss, Ariel; Krueger, Gretchen; Henighan, Tom; Child, Rewon; Ramesh, Aditya; Ziegler, Daniel M.; Wu, Jeffrey; Winter, Clemens; Hesse, Christopher; Chen, Mark; Sigler, Eric; Litwin, Mateusz; Gray, Scott; Chess, Benjamin; Clark, Jack; Berner, Christopher; McCandlish, Sam; Radford, Alec; Sutskever, Ilya; Amodei, Dario (July 22, 2020). "Language Models are Few-Shot Learners". [arXiv:2005.14165](https://arxiv.org/abs/2005.14165) (<https://arxiv.org/abs/2005.14165>) [[cs.CL](https://arxiv.org/archive/cs)] (<https://arxiv.org/archive/cs>).^[CL]
2. Shead, Sam (July 23, 2020). "[Why everyone is talking about the A.I. text generator released by an Elon Musk-backed lab](https://www.cnbc.com/2020/07/23/openai-gpt3-explainer.html)" (<https://www.cnbc.com/2020/07/23/openai-gpt3-explainer.html>). *CNBC*. Retrieved July 31, 2020. Four preprints were released between May 28 and July 22, 2020.
3. Bussler, Frederik (July 21, 2020). "[Will GPT-3 Kill Coding?](https://towardsdatascience.com/will-gpt-3-kill-coding-630e4518c04d)" (<https://towardsdatascience.com/will-gpt-3-kill-coding-630e4518c04d>). *Towards Data Science*. Retrieved August 1, 2020.

4. Sagar, Ram (June 3, 2020). "OpenAI Releases GPT-3, The Largest Model So Far" (<https://analyticsindiamag.com/open-ai-gpt-3-language-model/>). *Analytics India Magazine*. Retrieved July 31, 2020.
5. Chalmers, David (July 30, 2020). Weinberg, Justin (ed.). "GPT-3 and General Intelligence" (<https://dailynous.com/2020/07/30/philosophers-gpt-3/#chalmers>). *Daily Nous*. Philosophers On GPT-3 (updated with replies by GPT-3). Retrieved August 4, 2020.
6. Hao, Karen (September 23, 2020). "OpenAI is giving Microsoft exclusive access to its GPT-3 language model" (<https://www.technologyreview.com/2020/09/23/1008729/openai-is-giving-microsoft-exclusive-access-to-its-gpt-3-language-model/>). *MIT Technology Review*. Retrieved September 25, 2020. "The companies say OpenAI will continue to offer its public-facing API, which allows chosen users to send text to GPT-3 or OpenAI's other models and receive its output. Only Microsoft, however, will have access to GPT-3's underlying code, allowing it to embed, repurpose, and modify the model as it pleases."
7. Johnson, Steven; Iziev, Nikita (April 15, 2022). "A.I. Is Mastering Language. Should We Trust What It Says?" (<https://www.nytimes.com/2022/04/15/magazine/ai-language.html>). *The New York Times*.
8. "An understanding of AI's limitations is starting to sink in" (<https://www.economist.com/technology-quarterly/2020/06/11/an-understanding-of-ais-limitations-is-starting-to-sink-in>). *The Economist*. June 11, 2020. ISSN 0013-0613 (<https://www.worldcat.org/issn/0013-0613>). Retrieved July 31, 2020.
9. Polosukhin, Illia; Kaiser, Lukasz; Gomez, Aidan N.; Jones, Llion; Uszkoreit, Jakob; Parmar, Niki; Shazeer, Noam; Vaswani, Ashish (June 12, 2017). "Attention Is All You Need". arXiv:1706.03762 (<https://arxiv.org/abs/1706.03762>) [cs.CL (<https://arxiv.org/archive/cs>)].
10. "Natural Language Processing" (<https://www.thomsonreuters.com/en/artificial-intelligence/natural-language-processing.html>). Retrieved July 31, 2020.
11. Radford, Alec; Narasimhan, Karthik; Salimans, Tim; Sutskever, Ilya (June 11, 2018). "Improving Language Understanding by Generative Pre-Training" (https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf) (PDF). p. 12. Retrieved July 31, 2020.
12. Sterling, Bruce (February 13, 2020). "Web Semantics: Microsoft Project Turing introduces Turing Natural Language Generation (T-NLG)" (<https://www.wired.com/beyond-the-beyond/2020/02/web-semantics-microsoft-project-turing-introduces-turing-natural-language-generation-t-nlg/>). *Wired*. ISSN 1059-1028 (<https://www.worldcat.org/issn/1059-1028>). Retrieved July 31, 2020.
13. "Language Models are Unsupervised Multitask Learners" (https://d4mucfpksywv.cloudfront.net/better-language-models/language_models_are_unsupervised_multitask_learners.pdf) (PDF). *openai.com*. Retrieved December 4, 2019. "GPT-2, is a 1.5B parameter Transformer"
14. Ray, Tiernan (June 1, 2020). "OpenAI's gigantic GPT-3 hints at the limits of language models for AI" (<https://www.zdnet.com/article/openais-gigantic-gpt-3-hints-at-the-limits-of-language-models-for-ai/>). *ZDNet*. Retrieved July 31, 2020.
15. Gehman, Samuel; Gururangan, Suchin; Sap, Maarten; Choi, Yejin; Smith, Noah A. (November 16–20, 2020), *REALTOXICITYPROMPTS: Evaluating Neural Toxic Degeneration in Language Models* (<https://arxiv.org/abs/2009.11462>), Association for Computational Linguistics, pp. 3356–3369, arXiv:2009.11462 (<https://arxiv.org/abs/2009.11462>), retrieved June 2, 2021
16. "OpenAI API" (<https://openai.com/blog/openai-api/>). *OpenAI*. June 11, 2020.

17. Coldewey, Devin (June 11, 2020). "OpenAI makes an all-purpose API for its text-based AI capabilities" (<https://techcrunch.com/2020/06/11/openai-makes-an-all-purpose-api-for-its-text-based-ai-capabilities/>). *TechCrunch*. Archived (<https://web.archive.org/web/20211027000059/https://techcrunch.com/2020/06/11/openai-makes-an-all-purpose-api-for-its-text-based-ai-capabilities/>) from the original on October 27, 2021. Retrieved July 31, 2020. "If you've ever wanted to try out OpenAI's vaunted machine learning toolset, it just got a lot easier. The company has released an API that lets developers call its AI tools in on "virtually any English language task.""
18. Arram (July 9, 2020). "GPT-3: An AI that's eerily good at writing almost anything" (<https://arr.am/2020/07/09/gpt-3-an-ai-thats-eerily-good-at-writing-almost-anything/>). *Arram Sabeti*. Retrieved July 31, 2020.
19. Brown, Tom B.; et al. (2020). "Language Models are Few-Shot Learners". [arXiv:2005.14165](https://arxiv.org/abs/2005.14165) (<https://arxiv.org/abs/2005.14165>) [cs.CL (<https://arxiv.org/archive/cs.CL>)].
20. Thunström, Almira Osmanovic (June 30, 2022). "We Asked GPT-3 to Write an Academic Paper about Itself—Then We Tried to Get It Published" (<https://www.scientificamerican.com/article/we-asked-gpt-3-to-write-an-academic-paper-about-itself-then-we-tried-to-get-it-published/>). *Scientific American*. Retrieved June 30, 2022.
21. Transformer, Gpt Generative Pretrained; Thunström, Almira Osmanovic; Steingrimsson, Steinn (June 21, 2022). "Can GPT-3 write an academic paper on itself, with minimal human input?" (<https://hal.archives-ouvertes.fr/hal-03701250>). *Archive ouverte HAL* (in French). Retrieved June 30, 2022.
22. "Microsoft announced its first customer product features powered by GPT-3 and @Azure" (<https://blogs.microsoft.com/ai/from-conversation-to-code-microsoft-introduces-its-first-product-features-powered-by-gpt-3/>). *The AI Blog*. May 25, 2021.
23. "AI|Writer" (<https://www.aiwriter.app/>). *www.aiwriter.app*.
24. Fagone, Jason (July 23, 2021). "The Jessica Simulation: Love and loss in the age of A.I." (<https://www.sfchronicle.com/projects/2021/jessica-simulation-artificial-intelligence/>) *San Francisco Chronicle*. Retrieved July 29, 2021.
25. GPT-3 (September 8, 2020). "A robot wrote this entire article. Are you scared yet, human? | GPT-3" (<https://www.theguardian.com/commentisfree/2020/sep/08/robot-wrote-this-article-gpt-3>). *The Guardian*. ISSN 0261-3077 (<https://www.worldcat.org/issn/0261-3077>). Retrieved September 15, 2020.
26. "Update: Language Models and Dragon" (<https://latitude.io/blog/update-language-models>). *Latitude blog*. December 8, 2021.
27. Manjoo, Farhad (July 29, 2020). "How Do You Know a Human Wrote This?" (<https://www.nytimes.com/2020/07/29/opinion/gpt-3-ai-automation.html?>). *The New York Times*. ISSN 0362-4331 (<https://www.worldcat.org/issn/0362-4331>). Retrieved August 4, 2020.
28. Weinberg, Justin, ed. (July 30, 2020). "Philosophers On GPT-3 (updated with replies by GPT-3)" (<http://dailynous.com/2020/07/30/philosophers-gpt-3/>). *Daily Nous*. Retrieved July 31, 2020.
29. Simonite, Tom (July 22, 2020). "Did a Person Write This Headline, or a Machine?" (<https://www.wired.com/story/ai-text-generator-gpt-3-learning-language-fitfully/>). *Wired*. ISSN 1059-1028 (<https://www.worldcat.org/issn/1059-1028>). Retrieved July 31, 2020.
30. Claypoole, Theodore (July 30, 2020). "New AI Tool GPT-3 Ascends to New Peaks, But Proves How Far We Still Need to Travel" (<https://www.natlawreview.com/article/new-ai-tool-gpt-3-ascends-to-new-peaks-proves-how-far-we-still-need-to-travel>). *The National Law Review*. Retrieved August 4, 2020.
31. Marcus, Gary (December 1, 2018). "The deepest problem with deep learning" (<https://medium.com/@GaryMarcus/the-deepest-problem-with-deep-learning-91c5991f5695>). *Medium*. Retrieved September 29, 2020.

32. Marcus, Gary; Davis, Ernest (August 22, 2020). "GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about" (<https://www.technologyreview.com/2020/08/22/1007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion>). *MIT Technology Review*. Retrieved August 23, 2020.
33. Metz, Cade (November 24, 2020). "Meet GPT-3. It Has Learned to Code (and Blog and Argue)" (<https://www.nytimes.com/2020/11/24/science/artificial-intelligence-ai-gpt3.html>). *The New York Times*. ISSN 0362-4331 (<https://www.worldcat.org/issn/0362-4331>). Retrieved November 24, 2020.
34. "Medical chatbot using OpenAI's GPT-3 told a fake patient to kill themselves" (<https://artificial-intelligence-news.com/2020/10/28/medical-chatbot-openai-gpt3-patient-kill-themselves/>). *AI News*. October 28, 2020. Retrieved January 8, 2021.
35. *Chomsky on Terence McKenna, Sam Harris, GPT3, Cryptocurrencies, Kierkegaard, Neuralink, & Hofstadter* (<https://www.youtube.com/watch?v=c6MU5zQwtT4>). March 24, 2021. Event occurs at 1:11:44.
36. Floridi, Luciano; Chiriatti, Massimo (November 1, 2020). "GPT-3: Its Nature, Scope, Limits, and Consequences" (<https://link.springer.com/article/10.1007/s11023-020-09548-1>). *Minds and Machines*. **30** (4): 681–694. doi:10.1007/s11023-020-09548-1 (<https://doi.org/10.1007%2Fs11023-020-09548-1>). S2CID 228954221 (<https://api.semanticscholar.org/CorpusID:228954221>).
37. Olanoff, Drew (December 11, 2015). "Artificial Intelligence Nonprofit OpenAI Launches With Backing From Elon Musk And Sam Altman" (<https://techcrunch.com/2015/12/11/non-profit-openai-launches-with-backing-from-elon-musk-and-sam-altman/>). Tech Crunch. Retrieved May 31, 2021.
38. Hao, Karen (August 29, 2019). "OpenAI has released the largest version yet of its fake-news-spewing AI" (<https://www.technologyreview.com/2019/08/29/133218/openai-released-its-fake-news-ai-gpt-2/>). MIT Technology Review. Retrieved May 31, 2021.
39. Coldewey, Devin (March 11, 2019). "OpenAI shifts from nonprofit to 'capped-profit' to attract capital" (<https://techcrunch.com/2019/03/11/openai-shifts-from-nonprofit-to-capped-profit-to-attract-capital/>). Tech Crunch. Retrieved May 31, 2021.
40. Bender, Emily M.; Gebru, Timnit; McMillan-Major, Angelina; Shmitchell, Shmargaret (March 3, 2021). *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*. FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. pp. 610–623. doi:10.1145/3442188.3445922 (<https://doi.org/10.1145%2F3442188.3445922>).
41. Mindzak, Michael; Eaton, Sarah Elaine. "Artificial intelligence is getting better at writing, and universities should worry about plagiarism" (<https://theconversation.com/artificial-intelligence-is-getting-better-at-writing-and-universities-should-worry-about-plagiarism-160481>). *The Conversation*. Retrieved November 6, 2021.
42. Rogerson, Ann M.; McCarthy, Grace (December 2017). "Using Internet based paraphrasing tools: Original work, patchwriting or facilitated plagiarism?" (<https://edintegrity.biomedcentral.com/articles/10.1007/s40979-016-0013-y>). *International Journal for Educational Integrity*. **13** (1): 1–15. doi:10.1007/s40979-016-0013-y (<https://doi.org/10.1007%2Fs40979-016-0013-y>). ISSN 1833-2595 (<https://www.worldcat.org/issn/1833-2595>). S2CID 9473217 (<https://api.semanticscholar.org/CorpusID:9473217>).
43. O'Sullivan, Liz; Dickerson, John P. (August 7, 2020). "Here are a few ways GPT-3 can go wrong" (<https://techcrunch.com/2020/08/07/here-are-a-few-ways-gpt-3-can-go-wrong/>). *TechCrunch*. Archived (<https://web.archive.org/web/20211126192240/https://techcrunch.com/2020/08/07/here-are-a-few-ways-gpt-3-can-go-wrong/>) from the original on November 26, 2021. Retrieved November 28, 2021.

44. *Comment Regarding Request for Comments on Intellectual Property Protection for Artificial Intelligence Innovation* (https://www.uspto.gov/sites/default/files/documents/OpenAI_RFC-84-FR-58141.pdf) (PDF). USPTO.
45. *Here are a few ways GPT-3 can go wrong* (<https://techcrunch.com/2020/08/07/here-are-a-few-ways-gpt-3-can-go-wrong/>). *TechCrunch*.
46. Else, Holly (August 19, 2021). "'Tortured phrases' give away fabricated research papers" (<https://doi.org/10.1038%2Fd41586-021-02134-0>). *Nature*. **596** (7872): 328–329. Bibcode:2021Natur.596..328E (<https://ui.adsabs.harvard.edu/abs/2021Natur.596..328E>). doi:10.1038/d41586-021-02134-0 (<https://doi.org/10.1038%2Fd41586-021-02134-0>). PMID 34354273 (<https://pubmed.ncbi.nlm.nih.gov/34354273>).

Retrieved from "<https://en.wikipedia.org/w/index.php?title=GPT-3&oldid=1101347177>"

This page was last edited on 30 July 2022, at 14:55 (UTC).

Text is available under the Creative Commons Attribution-ShareAlike License 3.0; additional terms may apply. By using this site, you agree to the Terms of Use and Privacy Policy. Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.