Generative Artificial Intelligence: A Social Construction of Technology Perspective

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Hanzhi Zhou

Spring 2022

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Sean M. Ferguson, Department of Engineering and Society

Generative Artificial Intelligence: A Social Construction of Technology Perspective

Introduction

Generative artificial intelligence (GAI) refers to the set of algorithms that can take in a prompt (text, audio, images, random noise, or nothing) and output something that looks realistic and reasonable to humans. It was not until recent years that this concept becomes reality, thanks to the combination of advancements in machine learning research, blooming public datasets, and computing power. In 2016, generative adversarial networks (GAN) were introduced with the stunning ability to generate realistic images from random noise (Goodfellow et al., 2014). The capabilities of GAN were further explored in a number of recent works, introducing even more interesting applications such as image to image translation and style transfer. The former seamlessly merges one part of an image with another, while the latter involves adapting the style of one kind to another image. In 2017, the Transformer architecture for text comprehension gave birth to a series of large-scale models, called pre-trained language models, that demonstrated impressive abilities in text understanding when trained with a huge amount of data such as Wikipedia or even the entire internet. In this family of models, GPT-3 (generative pre-trained transformer 3) is particularly well-known due to its ability to generate fluent text based on a short prompt (Brown et al., 2020).

Besides the fact that AI can generate a wide variety of data, these generated data can look incredibly realistic to humans too. A study conducted by Caporusso et al. (2020) asked 551 participants to classify 7000 real and machine-generated images. They concluded that the accuracy of humans classifying fake images is close to random guesses, thus demonstrating the power of GAI. The same reality holds for text generation. Kirubarajan and Dugan (2020) summarized their finding of AI's text generation ability: "by presenting sentences one at a time to human annotators, we show that text generation machines can make it an average of 2.14 sentences before a human reader notices the text has switched to a machine".

Although GAI is a relatively new technology, its intriguing properties draw widespread interest, and therefore user-friendly products quickly emerge in the market. Thus, GAI is not just

accessible to researchers, but also to a much wider audience such as hobbyists and other ordinary users. The widespread adoption and application of GAI received mixed reviews from society. While the technology may be immensely useful and impactful to various stakeholders, malicious uses or simply misuse can bring great harm. In this paper, I employed the theory of social construction of technology to analyze the perception of GAI by various stakeholders, conflicts among them, and the beneficial and malicious applications of GAI. Furthermore, I explored possible mechanisms that will lead to conflict resolution, in which GAI gets widely adopted and some applications of GAI (e.g. beneficial use) prevail.

Social construction of technology

First proposed by Pinch and Bijker (1984), the theory of social construction of technology (SCOT) asks for a multi-directional view of the technological artifacts. The artifacts may be associated with different meanings within different social groups (known as interpretive flexibility), and the problems inflicted on some social groups and conflicts between these groups shape the process in which designs are proposed and adopted. Consequently, the interpretive flexibility may collapse through the closure process, and a more uniform view of the technology will be developed.

Despite being a recent topic, a few works that focus on the social construction of GAI or the broader topic of AI can be found in the literature. Using the SCOT framework, Eynon and Young (2021) investigated the characterization of AI for education within three different social groups – academia, industry, and government. They demonstrated that by developing different "technological frames" around AI, these social groups can advance their own agendas and influence how AI can be used for education. On the other hand, using DeepFake (Deepfakes, n.d.) as a case study, Kwok and Koh (2021) borrowed the idea from SCOT that technology itself is not simply good or bad. They argued that while DeepFake has malicious applications, beneficial uses of DeepFake can "preponderate its negative aspects" through the closure process and can avail industries such as tourism.

In the case of GAI, different social groups diverge in the way they perceive the nature and application of the technology. Fundamental differences in use cases and derived products result in conflicts that do not originate from the technology itself. In the following section, I will elaborate on the interpretations and applications of GAI within various social groups, their interaction with the GAI technology, and the conflicts between some of these groups.

Profit-driven organizations and conflict with end users

Companies and businesses see GAI as an opportunity to make profits as they see potential demand from the consumers. GAI brings groundbreaking technologies to domains such as entertainment and technology.

For entertainment, GAI gives rise to fundamentally different ways of interacting with images. A notable example of a GAI-based application built with GAN and its derivative works is FaceApp, which enabled users to edit facial attributes, and simulate and defy aging with a few clicks while producing realistic facial images. The App went viral in 2017 and again in 2019 (Rosenblatt & Thorp, 2019). In addition, there are applications such as Grammarly and GitHub Copilot that use GAI to promote productivity. Grammarly is an intelligent writing assistant that not only corrects your grammar but also helps you rephrase and adjust your delivery with its deep textual understanding, thanks to the recent breakthrough in language modeling. GitHub Copilot is an AI trained on public codebases hosted on GitHub and is able to intelligently suggest code completions for you (Ziegler, n.d.).

The conflicts arising from the use of GAI in commercial applications typically concern the collection and usage of data for inputs and training data. For instance, FaceApp needs human faces; Grammarly requires a piece of text. The way these inputs are processed and possibly stored on the cloud servers raises issues among users. A notable type of conflict between users and companies is the right and the extent to use user's uploaded images. A Chinese face swap App similar to FaceApp called ZAO was involved in such an incident (Doffman, 2019). ZAO's privacy agreement says ZAO "will try its best to only use content within a reasonable, necessary and

expressly stated extent". As the result, according to Lee (2019), China's Ministry of Industry and Information Technology requested the rectification of the app on September 4.

An example of a problem arising from the training data is the legal concern of GitHub Copilot, which is trained on publicly available codebases hosted on GitHub. However, the issue with it is that publicly available does not mean directly usable. Some of the codebases are protected by strong licenses that have restrictions. For example, the GNU general public license (GPL) is a family of licenses widely used by open-source code repositories. A notable example of a codebase that is licensed under GPL is the Linux kernel, which is an essential component for the majority of cloud servers. Important restrictions of GPL version 2 and onwards are that derivative works must also be open source and licensed under the same license as the parent work ("GNU General Public License, version 3", 2007). GitHub Copilot might produce similar or identical code snippets from these codebases protected by these strong licenses without the user knowing since it is common for a machine learning model to remember its training data. If it ever happened and was discovered by the original copyright holder, it can potentially put the user in legal trouble (FOSSA, 2021).

Independent developer and conflicts with victims

Independent developers are individuals with advanced computing skills and can create derivative products on their own. Unlike for-profit companies, independent developers often create applications for fun without commercialization, and thus they are less restrained on the purpose of their application. The general-purpose nature of GAI leads to design flexibility, and some controversial applications are produced. For example, DeepFake is a complete pipeline for face swap that assembles several state-of-the-art works in the field of computer vision (Deepfakes, n.d.). It makes newly born technology in academia more accessible to regular programmers, significantly simplifying the steps required to setup. GitHub's insight of the DeepFake repository, which can reveal the contribution history of the codebase, shows that it is mainly a single developer who produced this work.

A more infamous example is DeepNude, an app based on similar technology as DeepFake but created to synthesize naked images of women by taking off clothes from their real pictures (Telford, 2019). Although the App has been voluntarily taken down by the developers, versions of the source code with various improvements are still being maintained on GitHub by individual programmers ("GitHub Topics – DeepNude", n.d.).

Conflict arises when such technologies are applied to people and create footage of something they did not say or do. This can cause embarrassment and spread misinformation. For example, a DeepNude bot is created and operated on Telegram channels with tens of thousands of subscribers, and it is used by "thousands of people every month who use it to create nude images of friends and family members, some of whom appear to be under the age of 18." (Burgess, 2020) Although these nude images are not real, they can still cause embarrassment if they go out to the public.

In addition, public figures such as celebrities and politicians are very common victims. For example, a video of actor Kit Harrington, dressed as Snow, apologizing for season 8 of Game of Thrones appeared online in 2019 (Kooser, 2019). It is apparent that someone made this video in response to viewers' dissatisfaction with the TV show. While this may seem harmless due to its nature for entertainment, celebrities whose faces are used in place of the face of pornography actors may disagree. Kristen Bell, an American actress, was shocked when she learned that her face is on some recent pornography videos circulating online. In the news report by DeSantis (2020), she expressed serious concern about consent: "We're having this gigantic conversation about consent and I don't consent — so that's why it's not okay".

Researchers and why they are conflict free

Researchers are the ones who innovate by improving upon previous works, tackling shortcomings, synthesizing each work's strength, and generalizing from specific use cases. Numerous papers that empower the current GAI technologies reveal that the authors are not particularly interested in any specific application of the proposed algorithms in their paper. For

example, the pix2pix paper (Isola et al., 2017), one of the foundational works of DeepFake, stated that its goal is to provide a general-purpose image to image translation architecture without the need to engineer anything specific to different image translation scenarios (sketches to drawings, labels to scenes, etc.). Similarly, the goal of the GPT-3 (Brown et al., 2020) that empowers modern chatbots, machine translation, and many other natural language-based tasks is to simply demonstrate to the academic community that scaling up language models can make them good few-shots learners, meaning they can adapt to new tasks quickly with very little learning just like humans. Discussions of how this technology can be used are absent from these papers.

Therefore, researchers are interested in the mere performance and technical details of the technology, and not any specific use cases. In fact, researchers commonly license their works with open source licenses such as BSD license (used by pix2pix), MIT license, and GPL (used by DeepFake) that give certain freedom to users of their source code. Most importantly, these licenses make the authors free of any liability that may arise from using their source code. Through these means, researchers are able to place themselves out of the main conflicts with profit-driven companies who use their products, victims of products based on their research, or policymakers who would like to regulate the technology. Consequently, they are free to either advance GAI research or develop GAI forensics, which will be discussed later.

Role of advocates

In the context of GAI, advocates refer to the group of people who help to raise the public's awareness of GAI technology and its issues. Often, by creating and popularizing generated content using GAI, advocates draw the public's attention to what the new technology can bring to society.

For instance, several video clips generated by DeepFake were created and played to millions of people to demonstrate how realistic and influential the technology can be and to warn people about its potential for misinformation. A fake speech and a dance by "Queen Elizabeth" aired on Channel 4 by a British broadcaster (Rahim, 2020). In the report, the broadcaster of the video said it intended to offer "a stark warning about the advanced technology that is enabling the

proliferation of misinformation and fake news in a digital age." Another famous example is a synthesized speech "In the Event of Moon failure" by President Richard Nixon (Fiscutean, 2021). It was created by a group of people at MIT and Respeecher, a voice-cloning startup. In the news report by Fiscutean (2021), the CEO of Respeecher says "the idea behind this seven-minute film was to show what online misinformation will look like in the future". Through these videos, advocates are trying to convey to the public that with GAI technologies, misinformation will be more indistinguishable from facts in real footage, and the public should be more aware of the possibility of synthesized content.

Government and policymakers' view of GAI

While government and policymakers tend to not involve in direct conflicts with any social groups, their role in society requires them to protect user rights and pay attention to social stability, and therefore need to combat harmful usage of GAI. However, their main weapon, regulation, often falls short when it comes to new technology with diverse derived applications. As the result, GAI often lands in a legal grey area. For example, an analysis conducted by Gieseke (2020) showed that existing laws are insufficient to protect victims of GAI-derived pornography.

In addition, since GAI can be utilized to spread misinformation, policymakers need to worry about its potential in destabilizing society and affecting national security. For example, DeepFake can be used to create fake speeches of famous politicians to influence political campaigns and shape public opinion. By analyzing the potential influence of DeepFake from a securitization standpoint, Taylor (2021) argued that DeepFake makes it difficult for ordinary citizens to distinguish between the truth and rumor, and the state cannot effectively control the repercussions of synthesized information originating from DeepFake.

From rampant spread to controlled growth: the future of GAI

Despite the fact that social groups derive different use cases from GAI, hold different views towards GAI, or undertake different actions when it comes to GAI, we can draw some common trends regarding the future direction of GAI. We have companies making use of GAI and creating

useful products, even though some issues, such as those with regard to data usage, need to be addressed. We have creative developers making derivative products, even though some are controversial. We have researchers whose sole goal is to advance the technology. We have advocates who aim to raise the public's awareness of what GAI can bring. We have government and policymakers who, despite lagging behind in terms of regulation, are actively passing new laws regulating GAI. From these different groups, we can see the forward driving force that brings this novel technology into daily life, into the vision of more people.

On the other hand, the divergent nature of the driving force, originated from the design and interpretive flexibility of the technology, results in differing goals and conflicts among those groups. However, most of the conflicts arise from malicious, unethical, or irresponsible use of the technology. Therefore, I envisioned that given the current effort of advocates, regulators, and researchers, the growth of GAI will be controllable in the future converge towards its beneficial side, leading to conflict resolution. In the theory of SCOT, such a converging process in which interpretive and design flexibility collapse is referred to as closure.

In order to achieve closure, various social groups should achieve a basic consensus on the use cases of GAI technology. Two types of feasible consensuses can be achieved: either the beneficial use of GAI prevails or the technology is outright banned and cease to exist in daily life. It is extremely unlikely that the latter case will occur because an immense amount of effort has been invested into the technology by profit-driven organizations, independent developers, and researchers. In addition, as pointed out by Toyama (2020), GAI as a kind of digital technology is not as easy to prohibit as others such as the nuclear weapon, due to GAI's low cost to use and high utility or profits for many social groups. Hence, given the current widespread applications of the technology and the efforts of some social groups to bring it onto the right track, I will focus on the case in which GAI becomes another useful technology.

Road towards closure

The development of regulation serves as a great starting point to guide the public toward the beneficial use of GAI. A straightforward approach is to ban the uses of GAI that would cause harm. Several laws have already been passed, addressing DeepFake directly. For example, Texas and California have passed laws prohibiting the use of DeepFake that may harm public candidates and influence elections (Greengard, 2019). China passed a law that prohibits the use of DeepFake in creating fake news (Jing, 2019). While these approaches are effective against the type of crimes they are designed to eliminate, they are by no means comprehensive and there can be loopholes. A comprehensive approach would be to standardize GAI technologies, defining the scope of GAI technologies and what constitutes good and responsible use of GAI technologies.

Of course, just having regulations will be insufficient to prevent or reduce the harmful use of GAI. The reason is that law and punishment can only prevent harm to some extent. Society needs defense mechanisms against the harmful use of GAI so that either the cost of misuse is high or utility to the persons conducting misuse is minimized. We can borrow the success of the internet to illustrate this point. Internet is a technology that almost become a necessity for the modern world. Similar to GAI, the internet can be used for good and for bad. While you can talk to your friends anywhere with your phone through the internet, cybercriminals can cut off oil pipelines for ransom (Salam, 2021). Other than existing laws prohibiting exploits of the internet, the stability of the internet is maintained by cybersecurity experts who build infrastructure and study the techniques to defend against cyberattacks, so attacking any internet facilities will incur a high cost, in terms of both knowledge required and resources needed. Therefore, large-scale cyberattacks are relatively rare, rendering the internet a useful technology that should not be discontinued due to a few harmful cases.

For GAI, in my opinion, the former objective (raising the cost of malicious use) involves tremendous difficulty, as GAI technologies are just as accessible as the internet, with the source code of numerous applications hosted publicly on platforms such as GitHub. On the other hand, the latter objective is more feasible. Since malicious use of GAI utilizes its deceptive nature, it

can be countered through recognition, by either human perception and awareness or counter technologies developed to identify and label AI-generated content. Advocates are already on the track to educating the public about the technology. With their help, more people can be aware of the possibility of content being forged or manipulated, thus becoming more likely to classify them by looking for traces of AI or by fact-checking with other sources.

In terms of counter technologies, there are already ongoing works towards GAI forensics, with the help of the very group of researchers that propose GAI. As per a survey of deep-learning-based DeepFake detection by P and Sk (2021), "United States Defense Advanced Research Projects Agency (DARPA) initiated a research scheme in media forensics (named Media Forensics or MediFor) to accelerate the development of fake digital visual media detection methods". Numerous works that propose new methods to effectively detect AI-generated content can be found in academia. For example, Bonettini et al. (2021) used an ensemble of neural networks to detect potentially swapped faces in videos. At the same time, some researchers like Perov et al. (2020) chose to advance the DeepFake pipeline and improve DeepFake quality so high-quality AI-generated images can be used as datasets for GAI forensics. While not a lot of algorithms have been proposed to identify AI-generated passages, with enough time and funding, it will be a reality soon. Although researchers diverge in the ways they deem appropriate to counter GAI technologies, it can be reasonably concluded that, like the relationship between cybersecurity experts and cybercriminals, GAI technologies and GAI forensics will establish an equilibrium that makes GAI prosper and eliminate most cases of misuse.

Conclusion

Introduced during the AI bloom, GAI technologies are groundbreaking and bring many interesting applications. The application of GAI can be seen in entertainment, productivity, and many other domains. Even though GAI can be immensely useful, AI-generated content can also be used for fake news, scams, and influencing political outcomes. As with many technologies, the irresponsible and unethical use of GAI can bring harm to society. The popularity that comes with

GAI means that GAI has been influential to many social groups.

From the perspective of the social construction of technology, I discussed the interpretations of GAI within several social groups. I demonstrated that variability in interpretations and design flexibility of GAI leads to alternative designs and use cases that may be threats to other social groups. However, I proposed that through regulation, education, and development of parallel and counter technologies with the cooperation across various social groups, alternative designs and interpretations of GAI can collapse and stabilize to something beneficial to society. Using the internet as an analog, I illustrated that GAI can be regulated and made into a dynamic equilibrium with its counter technologies. In this equilibrium, the generative power of AI thrives and brings new ideas and joy to people, while the counterforce maintained by regulators and researchers balances the power of GAI and makes it under control.

References

Bonettini, N., Cannas, E. D., Mandelli, S., Bondi, L., Bestagini, P., & Tubaro, S. (2021). Video face manipulation detection through ensemble of cnns. *2020 25th International Conference on Pattern Recognition (ICPR)*, 5012–5019. https://doi.org/10.1109/ICPR48806.2021.9412711

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., . . . Amodei, D. (2020). Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems* (pp. 1877–1901). Curran Associates, Inc. https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf

Burgess, M. (2020). A deepfake porn bot is being used to abuse thousands of women. https://www.wired.co.uk/article/telegram-deepfakes-deepnude-ai

Caporusso, N., Zhang, K., Carlson, G., Jachetta, D., Patchin, D., Romeiser, S., Vaughn, N., & Walters, A. (2020). User discrimination of content produced by generative adversarial networks. *Human interaction and emerging technologies* (pp. 725–730). https://doi.org/10.1007/978-3-030-25629-6_113

Deepfakes. (n.d.). Deepfakes/faceswap: Deepfakes software for all. https://github.com/deepfakes/faceswap

DeSantis, R. (2020). Kristen bell recalls shock of learning her face was used in pornographic deepfake: 'it's not ok'. https://people.com/human-interest/kristen-bell-shock-face-used-pornographic-deepfake/

Doffman, Z. (2019). Chinese deepfake app zao confirms privacy issue for millions of users. https://www.forbes.com/sites/zakdoffman/2019/09/03/chinese-viral-app-zao-forced-to-backtrack-after-wechat-ban-and-privacy-backlash

Eynon, R., & Young, E. (2021). Methodology, legend, and rhetoric: The constructions of ai by academia, industry, and policy groups for lifelong learning. *Science, Technology, & Human Values*, *46*(1), 166–191. https://doi.org/10.1177/0162243920906475

Fiscutean, A. (2021). Nixon's unheard moon-disaster speech is now a warning about the deepfake future. https://www.zdnet.com/article/nixons-grim-moon-disaster-speech-is-a-now-a-warning-about-the-deepfake-future/

FOSSA. (2021). Analyzing the legal implications of github copilot. https://fossa.com/blog/analyzing-legal-implications-github-copilot/

Gieseke, A. P. (2020). "the new weapon of choice": Law's current inability to properly address deepfake pornography notes. *Vanderbilt Law Review*, *73*, 1479.

Github topics – deepnude. (n.d.). https://github.com/topics/deepnude

Gnu general public license, version 3 [Last retrieved 2020-01-01]. (2007).

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*. Curran Associates, Inc. https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf

Greengard, S. (2019). Will deepfakes do deep damage? *Commun. ACM*, *63*(1), 17–19. https://doi.org/10.1145/3371409

Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. *CVPR*.

Jing, M. (2019). China cracks down on deepfake technology in news production. https://www.scmp.com/tech/apps-social/article/3039978/china-issues-new-rules-clamp-down-deepfake-technologies-used

Kirubarajan, A., & Dugan, L. (2020). Learning to trick humans: Evaluation criteria for human-written and machine-generated text.

Kooser, A. (2019). Game of thrones deepfake makes jon snow apologize for season 8. https://www.cnet.com/culture/entertainment/game-of-thrones-deepfake-makes-jon-snow-apologize-for-season-8/

Kwok, A. O. J., & Koh, S. G. M. (2021). Deepfake: A social construction of technology perspective. *Current Issues in Tourism*, *24*(13), 1798–1802. https://doi.org/10.1080/13683500.2020.1738357

Lee, C. (2019). 'deepfake' app zao sparks major privacy concerns in china. https://www.zdnet.com/article/deepfake-app-zao-sparks-major-privacy-concerns-in-china/

P, S., & Sk, S. (2021). Deepfake creation and detection:a survey. *2021 Third International Conference on Inventive Research in Computing Applications (ICIRCA)*, 584–588. https://doi.org/10.1109/ICIRCA51532.2021.9544522

Perov, I., Gao, D., Chervoniy, N., Liu, K., Marangonda, S., Umé, C., Dpfks, M., Facenheim, C. S., RP, L., Jiang, J., Zhang, S., Wu, P., Zhou, B., & Zhang, W. (2020). Deepfacelab: A simple, flexible and extensible face swapping framework. *CoRR*, *abs/2005.05535*. https://arxiv.org/abs/2005.05535

Pinch, T. J., & Bijker, W. E. (1984). The social construction of facts and artefacts: Or how the sociology of science and the sociology of technology might benefit each other. *Social Studies of Science*, *14*(3), 399–441. https://doi.org/10.1177/030631284014003004

Rahim, Z. (2020). 'deepfake' queen delivers alternative christmas speech in warning about misinformation. https://www.cnn.com/2020/12/25/uk/deepfake-queen-speech-christmas-intl-gbr/index.html

Rosenblatt, K., & Thorp, F. (2019). Photo editor faceapp goes viral again, prompting security concerns. https://www.nbcnews.com/tech/tech-news/photo-editor-faceapp-goes-viral-again-prompting-security-concerns-n1030776

Salam, E. (2021). Cyber-attack forces shutdown of one of the us's largest pipelines. https://www.theguardian.com/technology/2021/may/08/colonial-pipeline-cyber-attack-shutdown

Taylor, B. C. (2021). Defending the state from digital deceit: The reflexive securitization of deepfake. *Critical Studies in Media Communication*, *38*(1), 1–17. https://doi.org/10.1080/15295036.2020.1833058

Telford, T. (2019). 'the world is not yet ready for deepnude': Creator kills app that uses ai to fake naked images of women. https://www.washingtonpost.com/business/2019/06/28/the-world-is-not-yet-ready-deepnude-creator-kills-app-that-uses-ai-fake-naked-images-women/

Toyama, K. (2020). Bring on the technology bans! https://theconversation.com/bring-on-the-technology-bans-120209/

Ziegler, A. (n.d.). Research recitation. https://docs.github.com/en/github/copilot/research-recitation