



INTRO TO GENERATIVE A.I.

How to use this guide

This primer addresses the common questions that many people have about generative artificial intelligence (AI) systems. It contains information about generative AI capabilities, seven key issue areas where generative AI is having impact, and how generative AI may be used in the future. Although this primer was originally designed primarily for journalists, we have found that it is a valuable resource for anyone communicating about generative AI.

For a list of resources for where to find technical experts, please see [Finding Experts in AI](#). For information about AI more generally, please see [AI 101](#).

Visit techprimers.aspendigital.org to see all primers.

Intro to Generative AI © 2023 by Aspen Digital is licensed under [CC BY-NC 4.0](#).

CONTENTS

- 3 What is generative AI?**
What it is and how it's different from decision-making AI tools
- 3 What generative AI capabilities exist today?**
Text-to-text models like ChatGPT aren't the only thing out there
- 4 Poster: Tech that people call "artificial intelligence"**
- 5 How might generative AI be used?**
It will disrupt many aspects of society
- 6 Key Issues in Generative AI**
 - 7 Information ecosystems** Generative AI makes it much easier to mislead at scale
 - 8 Expectations & claims** Tools like ChatGPT are not designed to produce facts
 - 9 Intellectual property** Existing tools do not attribute to their source material
 - 10 Future of work** Automating creative work will impact peoples' livelihoods
 - 11 Environmental impacts** AI use is far from carbon-neutral
 - 12 Discriminatory effects** Generative AI systems can perpetuate societal harms
 - 13 Feedback loops** How will generative AI impact future AI development?
- 14 What comes next?**
- 15 Acknowledgements**

WHAT IS GENERATIVE AI?

Generative AI is a subset of artificial intelligence technologies that are used to create new content, such as images or text, based on patterns in large amounts of existing content. Generative AI differs from classification AI—like email spam filtering or tumor detection used in medical settings—because generative AI systems are designed to make content, not to make decisions.

WHAT GENERATIVE AI CAPABILITIES EXIST TODAY?

While [ChatGPT](#) captured the public's attention by allowing people to generate uncanny and seemingly confident responses to a vast array of written prompts, there is a diverse set of generative AI applications that have been made available to businesses and consumers. Increasingly, AI developers are creating multimodal tools, tools which incorporate multiple sources of input data (such as images, audio, or text) at once.

- Image-to-image [Canvas](#)
- Text-to-image [Midjourney](#), [DALL-E](#), [Stable Diffusion](#)
- Text-to-audio [MusicLM](#)
- Video-to-video [Project Morpheus](#), [AI video compression](#)
- Text-to-video [Make-A-Video](#)
- Image-to-text [Automatic image description](#)
- Text-to-text (including computer code) [Github Copilot](#), [Bard](#)

TECH THAT PEOPLE CALL “ARTIFICIAL INTELLIGENCE”

MACHINE LEARNING
Statistical (based on data)

AUTOMATED REASONING
Symbolic (based on rules)

Note: Machine learning is currently the most popular approach, but some believe that automated reasoning may make a comeback or that these two branches may be combined in new ways.

Deep Blue recommended chess moves by using the rules of the game and searching a parallelized tree of many possible combinations of moves.

The various ghosts in **Pac-Man** are programmed to chase you using different pathfinding methods. (Blinky, the red ghost, is most aggressive!)

EXPERT SYSTEMS

PLANNING

LABELED DATASETS

CLASSIFICATION

FaceID classifies images of faces using convolutional neural networks (CNNs).

MIXTURE OF LABELED & UNLABELED DATASETS

GENERATIVE

DALL-E produces images from text through an iterative process using a diffusion model.

ChatGPT generates text in response to typed prompts using a transformer deep learning model.

REINFORCEMENT

AlphaGo develops strategies for the game Go using self-play reinforcement learning.

UNLABELED DATASETS

CLUSTERING

Spotify clusters songs into music genres using hierarchical density-based spatial clustering with noise (HDBSCAN).

HOW MIGHT GENERATIVE AI BE USED?

Although generative AI tools are still in the early stages of development, they are already being used to produce content at surprisingly high speeds, low costs, and with relative ease for end users. This newfound accessibility has labor and operational implications for software engineering, media production, education, the commercial art market, and more. No one knows exactly what will emerge from the explosion of generative AI tools hitting the market, but early experiments point toward the potential for larger scale disruption to business, security, and society at large:

Hyper-personalized Content

Traditionally, the cost of making individually personalized content, such as ads with your face in them or movie trailers narrated in the sound of a loved one's voice, was prohibitively high. People may now easily use generative AI to realize this level of [extreme personalization](#), either for their own fulfillment or to manipulate others.

The Rise of "No-Code" Application Development

Historically, in order to develop websites or computer applications, you needed to know programming languages. Now, it is becoming possible to use conversational language to prompt an AI tool to produce computer code for you (even if today's systems are still imperfect). These tools may lower costs by expanding the number of people who are able to create and contribute to software development and make a wide range of products and services more accessible. However, they could also negatively impact how much people are paid for these skills and change the nature of work to make it more tedious and less collaborative.

HOW MIGHT GENERATIVE AI BE USED?

Better Augmented Reality

Real-time rendering of believable digital environments is computationally intensive and expensive. These graphical requirements have been a pernicious issue for augmented or virtual reality applications because time delays in rendering can create a jarring and unnatural user experience. Generative AI systems could be used to approximate (if not perfectly replicate) complex physical phenomena, like lighting and shadows, making these virtual scenes feel more immersive.

There are still many unknowns and opportunities for discovery. We have only scratched the surface on possible uses of these tools.

KEY ISSUES IN GENERATIVE AI

There are a number of promising applications of generative AI systems, a subset of artificial intelligence technologies that are used to create new content based on patterns in large amounts of existing content. These uses are not without [their risks](#), however. The following sections highlight a number of the most pressing issues associated with generative AI, with links to a number of illustrative articles exploring perspectives on each of these issues.

7 Information ecosystems

8 Expectations & claims

9 Intellectual property

10 Future of work

11 Environmental impacts

12 Discriminatory effects

13 Feedback loops

INFORMATION ECOSYSTEMS

How will generated content affect the trustworthiness of media?

Media created to mislead is not a new problem, but generative AI makes it much easier to create mis- and disinformation at scale and to create convincing human-like AI interactions that could be used to exploit users with scams or security attacks. As generated content improves, it will be harder to detect inauthentic content in the wild, and it will be easier for smaller actors to manufacture large-scale disinformation campaigns.

- **Deepfakes:** Humans may find deepfake faces more trustworthy than real ones

[scientificamerican.com/article/humans-find-ai-generated-faces-more-trustworthy-than-the-real-thing](https://www.scientificamerican.com/article/humans-find-ai-generated-faces-more-trustworthy-than-the-real-thing)

- **Confident nonsense:** “While the answers which ChatGPT produces have a high rate of being incorrect, they typically *look like they might be good*,” burdening human moderators

meta.stackoverflow.com/questions/421831/temporary-policy-chatgpt-is-banned

- **Democratizing malware:** ChatGPT makes it possible for anyone to produce malware and phishing emails

arstechnica.com/information-technology/2023/01/chatgpt-is-enabling-script-kiddies-to-write-functional-malware

- **Impact on elections:** Generative AI is already being used to influence voters and reduce trust in the electoral system

www.brennancenter.org/our-work/analysis-opinion/how-ai-puts-elections-risk-and-needs-safeguards

EXPECTATIONS & CLAIMS

What are the limitations of generative AI systems?

People selling AI products and services benefit from systems being perceived to be [more reliable and capable than they are](#), from the anthropomorphization of “smart assistants” to the typing animations of text-generators like ChatGPT. Peeking behind the curtain reveals that the AI tools on the market are specialized in scope, not general-purpose “intelligences,” and still have crucial vulnerabilities and flaws. Although it might be tempting to ascribe greater power to these systems, there are still many questions about whether they are appropriately effective for the widespread adoption we are already seeing, let alone that we are on the verge of “Artificial General Intelligence” that surpasses humans in a broad range of capabilities. Systems like ChatGPT and Bard were designed to produce confident sounding text, not factual statements.

- **AI hype:** An AI model “passing” an exam designed for humans is not indicative of intelligence
technologyreview.com/2023/08/30/1078670/large-language-models-arent-people-lets-s-top-testing-them-like-they-were
- **Confident but inaccurate:** CNET used ChatGPT to generate articles that ended up riddled with errors
washingtonpost.com/media/2023/01/17/cnet-ai-articles-journalism-corrections
- **False equivalency:** Equating GPT’s text prediction capabilities with consciousness is misleading
karawynn.substack.com/p/language-is-a-poor-heuristic-for
- **Unintuitive failures:** Certain strange prompts (like usernames) elicit nonsense from ChatGPT
vice.com/en/article/epzyva/ai-chatgpt-tokens-words-break-reddit

INTELLECTUAL PROPERTY

Who owns what?

Many datasets used to build today's generative AI have been compiled by scraping, or extracting information, from the web. For example, to make a generative AI tool that can output digital images, developers scraped millions of existing images hosted on large art platforms, like Flickr and DeviantArt. Content collected online in this manner is often used without the consent or knowledge of the original creator. Even if the original content is not reproduced by the system ([although in some cases it can be](#)), this process leads to thorny questions around attribution, intellectual property, monetization of generative AI tools, and economic harms to creative industries.

- **Sensitive data:** Scraping isn't flawless—personal health data was found in a popular image dataset
arstechnica.com/information-technology/2022/09/artist-finds-private-medical-record-photos-in-popular-ai-training-data-set
- **IP protections at work:** Google won't release its text-to-music generator because there is a chance it will reproduce copyrighted music it was trained on
techcrunch.com/2023/01/27/google-created-an-ai-that-can-generate-music-from-text-descriptions-but-wont-release-it
- **Fights for attribution:** A class action lawsuit is filed against Microsoft for lacking attribution for code used to train Github Copilot
theverge.com/2022/11/8/23446821/microsoft-openai-github-copilot-class-action-lawsuit-ai-copyright-violation-training-data
- **Copyright ramifications:** AI-generated works should not be permitted copyright protection
lawfaremedia.org/article/ai-generated-works-artists-and-intellectual-property

FUTURE OF WORK

How will generative algorithms impact peoples' livelihoods?

Generative tools can be used as assistants, augmenting human creativity, but they can also be used to automate certain types of work, from writing copy to creating spot art for articles to coding. There are many open questions about what tasks will be most easily automated and whether that automation will result in a reduction in total jobs, a profound change in how certain work is valued, or a restructuring of labor as new jobs are created. For example, a software developer that once created website templates could either (1) lose their job because someone else can use a tool to do themselves what they would have hired the developer to do, (2) get a reduction in salary as they face more competition in the market or, (3) no longer code as much manually, but instead be in charge of generating outputs using the AI.

- **Expanded creativity:** Generative AI could make it much easier for people to create websites and apps without needing to know how to code

venturebeat.com/ai/how-code-generating-ai-code-llms-is-creating-new-challenges-as-it-matures

- **Impact on creative industries:** The use of generative AI in Hollywood was a key sticking point in labor negotiations

wsj.com/business/media/hollywood-studios-can-train-ai-models-on-writers-work-under-entative-deal-aedae589

- **Augmenting work:** How professionals can use ChatGPT today

linkedin.com/pulse/chatgpt-professionals-guide-using-ai-allie-k-miller

- **Invisible labor:** Automating some tasks just creates a different kind of work

logicmag.io/failure/the-automation-charade

ENVIRONMENTAL IMPACTS

How does generative AI contribute to climate change and consume natural resources?

Generative AI systems are more resource-intensive than many similar technologies. The large data centers (collections of connected computers) used to develop and deploy these tools require power and cooling, using electricity and water. While the development of AI models typically requires a large amount of energy, generative AI is unique in that people's ongoing use of these systems makes up the bulk of its consumption. Additionally, although today the water usage of generative AI is dwarfed by the amount used for other purposes, like growing almonds or making potato chips, water and energy usage would rise further if proposals such as [incorporating generative AI into every Google search](#) are implemented.

- **Carbon footprint:** Generative AI use eats up resources needed to mitigate the climate crisis
slate.com/technology/2023/08/chatgpt-ai-arms-race-sustainability.html
- **Energy usage versus other technologies:** ChatGPT uses five to nine times as much electricity as a Google search
towardsdatascience.com/chatgpts-energy-use-per-query-9383b8654487
- **Freshwater usage:** Generative AI systems contribute to stress on local water infrastructure in drought-prone areas
themarkup.org/hello-world/2023/04/15/the-secret-water-footprint-of-ai-technology
- **Reducing climate impacts:** Some argue that companies can make AI greener by fine-tuning existing models instead of training new ones
hbr.org/2023/07/how-to-make-generative-ai-greener

DISCRIMINATORY EFFECTS

How do generative systems perpetuate societal harms?

Unless the data ingested into AI models is carefully curated—which datasets scraped from the web rarely are—tools built using that dataset will reflect the biases of the unfiltered internet. Even with careful dataset curation, however, AI tools need to be fine-tuned by human content moderators to mitigate systemic biases. In some cases, creators or deployers of a system will manually override the AI system to limit output of harmful material, but these sorts of interventions are necessarily brittle and imperfect.

- **Biased assumptions:** The "magic avatars" created using Lensa AI sexualize women and whiten people of color regardless of their users' wishes
polygon.com/23513386/ai-art-lensa-magic-avatars-artificial-intelligence-explained-stable-diffusion
- **Mitigating harms:** GPT-3 will make biased statements against certain groups, but it may be possible to mitigate this with extra training focused on fairness
vox.com/future-perfect/22916602/ai-bias-fairness-tradeoffs-artificial-intelligence
- **Content moderation:** In protecting the world from biased and discriminatory outputs in these systems, content moderators suffer the consequences
time.com/6247678/openai-chatgpt-kenya-workers

FEEDBACK LOOPS

How will generative AI impact future AI development?

Future datasets scraped from the web will be impacted as everyday people, content farms, and disinformation campaigns saturate the internet with generated content. New AI models that are trained using these datasets may perpetuate existing biases documented in large language models like GPT-3 or image generation models like Stable Diffusion. Detecting generated content to exclude it from datasets is an active field of research but is by no means a solved problem. This feedback loop of using content produced by machines to train machines to produce more content could reduce the quality and performance of future AI systems.

WHAT COMES NEXT?

While these are the early days, many experts agree that generative AI systems will have far-reaching implications for society. Unlike blockchain and other emerging technologies that have caught the tech industry's eye, generative AI tools have sparked the public's interest and imagination with creatives and business leaders alike identifying ready applications. Most immediately, here are some things that could come next:

- Tech companies vying to both define and control new markets carved out by generative AI tools, deploying work-in-progress tools into an unregulated space
- The establishment of legal frameworks and precedent to better define both intellectual property rights and consumer protection with regards to generative AI and the AI space more broadly
- The immediate disruption of some existing labor markets while new areas of work are still being defined

ACKNOWLEDGEMENTS

This work was produced by Eleanor Tursman, B Cavello, and Tom Latkowski, and was made possible thanks to generous support from Siegel Family Endowment, the Patrick J. McGovern Foundation, and the John S. and James L. Knight Foundation.



[Aspen Digital](#) is a technology and information-focused nonpartisan organization that brings together thinkers and doers to uncover new ideas and spark policies, processes, and procedures that empower communities and strengthen democracy. This future-focused [Aspen Institute](#) program inspires collaboration among diverse voices from industry, government, and civil society to ensure our interconnected world is accessible, safe, and inclusive—both on and offline. Across its initiatives, Aspen Digital develops methods for elevating promising solutions and turning thought into networked impact.

[Contact us](#) with questions or corrections regarding this primer. Please note that Aspen Digital cannot guarantee access to experts or expert contact information, but we are happy to serve as a resource. To find experts, please refer to **Finding Experts in AI**.

Visit techprimers.aspendigital.org to see all primers.