Adapting to "AI"

How Will Generative AI Affect Work? How Should We Respond?

Joseph Cohen Department of Sociology City University of New York, Queens College 65-30 Kissena Boulevard Queens, New York 11367 Web: <u>incohen.commons.gc.cuny.edu</u> Email: <u>joseph.cohen@qc.cuny.edu</u> Mastodon: <u>@jncohen@fosstodon.org</u> Twitter: <u>@jncohen</u> LinkedIn: jncohen

DRAFT VERSION This is a preliminary draft being circulated for feedback. Please do not cite or circulate Comments, corrections, and advice are greatly appreciated.

Last Updated: August 30, 2023 © 2023 Joseph Nathan Cohen. All Rights Reserved.

Table of Contents

Adapting to "AI" 1
How Will Generative AI Affect Work? How Should We Respond?
Part 1: Background1
The Art Contest
Was It Really AI?
Image Generators Disrupt Content Creation Markets4
Chat-GPT is Released
Higher Education Disrupted5
General Economic Disruption
Part 2: How the Technology Works
Generative AI vs General AI 8
Overview of the Process
Training Data9
Converting Multimedia to Data 11
Pattern Identification 12
Prediction
The Result14
Automated Content Generators15
Part 3: Adapting to Generative AI15
Future Jobs Will Use AI15
Integrating the Technology into Workflows16
Retaining Human Mastery over the Process16
Advanced Technical Skills Remain Important18
Building Credibility
Leverage AI to Create New Opportunities18
Preparing for Technological Disruption19
Bibliography 19

Chat-GPT 3.5 was released to the public on <u>November 30, 2022</u>. Since then, there has been a scramble to develop and adopt *Generative AI*, computer programs that generate seemingly high-quality human-like communications, from text documents to computer code to multimedia files. Well-known examples of this kind of software include <u>Chat-GPT</u>, <u>Bard</u>, <u>LLaMA</u>, <u>Dall-E</u>, <u>Stable Diffusion</u>, or <u>Midjourney</u>.

There was an almost immediate economic fallout from very recent advances in Generative AI. Businesses and skill sets lost earning power almost immediately. It triggered <u>lawsuits</u>, <u>labor</u> <u>strikes</u>, <u>and government inquiries</u>. It was predicted to automate many people's jobs (Briggs, Kodani, and Pierdomenico 2023), and by some reports, the <u>process is already under way</u>.

This is not a development that can be ignored by universities, scholars, and their students. It has rendered teaching methods obsolete. It calls parts of university curricula into economic question. The technology will create automation pressures on long-insulated university-educated white-collar "knowledge" and "creative" workers. It creates a new source of information and knowledge to compete with educated and trained people.

Although the software's underlying methods has been developing over years (see Young et al. 2018), it has only recently come into widespread use. We do not yet have firm answers about Generative AI will change work and society, and how best to adapt to this new technology. Still, it is not an issue that the academy or the wider public can ignore until someone else develops an optimal response. It is worth engaging this technology mindfully, paying attention to what it does and how it changes society, and making thoughtful and reasoned ideas about how humans can adapt to a world in which Generative AI exists and improves constantly.

This essay considers some initial ideas about the academy's adaptation as we enter Year Two of the Generative AI era The paper proceeds in three parts. Part One introduces basic theoretical concepts and briefly summarizes Genarative AI's introduction into the popular consciousness and widespread usage during the academic year of 2022 – 2023. Part Two attempts to clarify the workings of this software in a non-technical but detailed way, with an eye towards discerning what it can and cannot do. Part Three shares some initial thoughts on how knowledge and creative workers – and the educators who train them – might adapt to this technological disruption.

Part 1: Background

Generative AI is a new technology in which computers generate seemingly-intellgent, seemingly-original communications content in response to a user-input prompt. It allows computers to create original, on-demand text, images, video, audio, or code whose quality is fast approaching what you would get from a skilled human analyst, creative, or communicator.

Prior to Fall 2022, there were several early attempts to develop and commercialize content generation software. Tech startups like <u>Jasper</u>, <u>copy.ai</u>, or <u>Writesonic</u> were marketing text-ondemand services to communications professionals (e.g., Figure 1 below). Chances are that, if you are a communications professional, ads for these services appeared on your Facebook timeline in 2021 and early 2022.



Figure 1: Example of Early Commercial Generative Al Applications, 2020. Source: Internet Archive Wayback Machine

What changed during the 2022 – 2023 academic year was that these technologies came into widespread use. A series of events caused societal attention to converge on the technology. This attention prompted cycles of adoption and investment, causing the technology to diffuse into practice.

The Art Contest

Generative AI's first big inroads into the popular consciousness occurred at the start of the academic year. In September 2022, an <u>"AI-generated" image won an art contest in Colorado</u>



Figure 2: "Theatre D'opera Spatial" [sic] by Jason Allen. Winner of the 2022 Digital Art Prize at the Colorado State Fair.

against human competitors. Jason Allen's "Theatre D'opera Spatial" [sic] (see Figure 2 below¹) won the Digital Art Prize at the Colorado State Fair.

The piece was produced using the software <u>Midjourney</u> and was initially framed and understood as "AI-generated" art, though Allen himself reported having spent 80 hours generating over 900 iterations of the piece (Harwell 2022). Strictly speaking, a computer working alone did not beat humans in an art contest, but rather a human using a new software tool beat other humans. Still, the moment was widely portrayed and understood as a meaningful moment in the advancement of artificial intelligence. For a moment, a computer was widely construed as having out-created humans in an art contest.

Was It Really AI?

Allen's award opened discussions about whether machines were getting close to artificial intelligence. People understood a computer to have beaten human competitors in an activity thought to be the exclusive province of humans: original artistic expression. It amounted to a first impression that a computer had satisfied a widely-referenced benchmark for assessing artificial intelligence: the *Turing Test*.

The Turing Test was proposed by the famed British mathematician and artificial intelligence research pioneer Allan Turing (pictured in Figure 3). Turing (1950) maintained that, regardless of whether machines could "think" in ways similar to humans, we might consider a computer to be "intelligent" if it could fool humans into thinking that it was also human. A machine might be deemed "intelligent" if humans could not discern that they were speaking to a machine.



Figure 3: Alan Turing. Source: United States National Security Agency Hall of Fame.

To many, news that an AI-generated art piece beat human artists in an art contest sounded like a computer had passed something like a Turing test. Media outlets rushed to publish pieces about computers turning on humans. Over time, the public came to understand that there are clear differences between this specific software and a weaponizable general artificial

¹ Figure 2 is a reproduction from Roose (2022). The graphic is reproduced for the reader to view the cultural object being discussed in this history, as media outlets like the *New York Times* – at that time one of America's foremost media outlets – conveyed it to the public.

intelligence system that can destroy humanity. Still, many people were left jarred by the experience that the basic elements of artificial intelligence were much further along than most of us realized.

Image Generators Disrupt Content Creation Markets

Within the content creation markets that I research, Allen's art prize was big news for more practical, and less philosophical, reasons. The story widely informed the public that computers could create reasonably good-quality, copyright-free images – including highly photorealistic ones – with very little money or technical knowledge. It was possible to visit a range of easy, inexpensive online apps (e.g., <u>OpenAl's Dall-E</u>) to generate original, purportedly free visuals for pennies.

This was an economically-relevant development, because image generation or curation is an important function in content creation that requires labor or money. People started to substitute human-produced images with computer-generated ones. Video game developers started using generated art instead of hiring artists to decorate <u>their video environments</u> (Needleman 2023; Robertson 2023; Zwiezen 2022). In blogging and podcasting, generated visuals were an inexpensive alternative to stock photo or clip art subscriptions, while offering far more possibilities with far less time investment than any attempt to find and verify the usability of public domain or Creative Commons materials.

Concerns soon arose about those whose livelihoods were built on the creation and distribution of images: the photographers, graphic artists, logo designers, and image vendors. There were entire businesses and markets dedicated to providing people with this type of visual content. They were all confronting a new competitor who offered instantly-delivered, highly-customized visual content that was unencumbered by intellectual property and priced in the fractions of pennies.

Chat-GPT is Released

Then, in November 2022, OpenAI released Chat-GPT. The software was like Midjourney, but it generated textual as opposed to visual content. As mentioned earlier, it was not the first of its type under development. Earlier generations of this software had been released publicly, though reviewers argued that those earlier iterations were "<u>usually easily identifiable as non-human</u>" (Vincent 2019). Earlier that summer, a Google engineer went public with <u>claims that the company's private text generator (named "LaMDA") was sentient</u> (Tiku 2022).

Though there were others available, Chat-GPT was unique because (1) the model became quite good at generating human-like content and (2) it was distributed publicly through an inexpensive, stable, expandable, and easy-to-use deployent model. Unlike earlier generations, Chat-GPT allowed people to experiment with the technology, even if they did not have much money or the technical skill to do so otherwise. It was not just that the technology was useful, but that a mechanism had been created to let people try the technology and figure out practical uses for it. The result is reported to be the biggest software launch in history, taking three months to reach 100 million monthly active users (Hu 2023; Wodecki 2023).

Initial assessments were that Chat-GPT could instantly generate passable university-level essays, even by standards that prevailed in prestigious university programs (Kelly 2023; Scott 2023; Westfall 2023). By March 2023, OpenAI released the next iteration of their text generator – GPT-4 – which bested most people on highly advanced tests, like law or graduate school admissions tests. Figure 4 (below) reproduces OpenAI's product performance claims, as posted on their website in April 2023:

Simulated exams	GPT-4	GPT-4 (no vision)	GPT-3.5
	estimated percentile	estimated percentile	estimated percentile
Uniform Bar Exam (MBE+MEE+MPT) ¹	298/400	298/400	213/400
	-90th	-90th	-10th
LSAT	163	161	149
	-88th	-83rd	-40th
SAT Evidence-Based Reading & Writing	710/800	710/800	670/800
	-93rd	-93rd	-67th
SAT Math	700/800	690/800	590/800
	-39th	-89th	-70th
Graduate Record Examination (GRE) Quantitative	163/170	157/170	147/170
	-90th	-62nd	-25th
Graduate Record Examination (GRE) Verbal	169/170	165/170	154/170
	-99th	-96th	-63rd
Graduate Record Examination (GRE) Writing	4/6	4/6	4/6
	-54th	-54th	-54th
JSABO Semifinal Exam 2020	87/150	87/150	43/150
	99th-100th	99th-100th	3tst-33rd
USNCO Local Section Exam 2022	36/60	38/60	24/60
Medical Knowledge Self-Assessment Program	75%	75%	53%
Codeforces Rating	392	392	260
	below 5th	below 5th	below 5th
AP Art History	5	5	5
	88th-100th	95th-100th	86th-100th
AP Biology	5	5	4
	856h-100kh	85th-100th	62nd-85th
AP Calculus BC	4	4	1
	43ed_58th	43rd-50th	0th-7th

Figure 4: OpenAI's Product Performance Claims for Chat-GPT 4. Source: Screen capture of OpenAI web site, April 2023.

GPT was now outperforming most high-performing high school and undergraduate students in the creation of text-based content like essay-writing or test-taking. It also seemed able to perform the world of highly-trained people. It was <u>diagnosing patients</u> (Hughes 2023), <u>analyzing case law</u> (Perlman 2023), or <u>writing a TV script</u> (Coyle 2023).

Higher Education Disrupted

Over the 2022 – 2023 academic year, the basic fact of highly-available, very powerful automated content generators pressed itself in all corners of academia. Many were concerned about cheating (Barnett 2023). People were pitching that we bring back oral exams (Dobson 2023). School systems banned and filtered the technology outright, including New York City's public primary and secondary education systems (Yang 2023) and many university professors. Throughout the semester, people desperately waited for someone to develop effective AI-detection programs, not wholly unlike the way people waited for pharmaceutical firms to develop a vaccine during COVID.

In my view, concerns about cheating are missing a deeper and more important problem: the coming obsolescence of the "B" student, at least as we have been training them. Chat-GPT replicates the work of dogged rule- and example-followers. It produces the work of students who could do difficult jobs *if* given detailed instructions, guidance, and oversight. Either will give you a basic and unreliable version of what you request. The difference is that Generative AI does the job in seconds and for pennies. It won't matter if everyone can pose as a "B" student

in university if "Bs" aren't needed to do things because the computer outperforms them at the job. Many educational curricula teach and test skills whose economic value is falling.

Amid so much concern about preventing students from using Chat-GPT, teachers themselves found the software to be rather useful in performing their own job tasks. A Walton Family Foundation-commissioned study by Impact Research reportedly found that a higher proportion high school teachers were using Chat-GPT than students (Impact Research 2023). Faculty reportedly started having Chat-GPT write recommendation letters (Bogost 2023). Researchers started using Chat-GPT to write research papers, and <u>listing the content generator as a co-author</u> (Stokel-Walker 2023).

General Economic Disruption

A popular narrative holds that automation was the scourge of those who worked with their hands, but not with their minds. Economists argue that this insulation from automation helped bolster the economic gains and material privilege of those with advanced education (e.g., Brynjolfsson and McAfee 2012).

Generative AI has triggered widespread concern about job loss and impoverishment among university-educated, white-collared workers. A recent report by Goldman Sachs estimates that the technology exposes 300 million full-time jobs to substantial task automation, and that these automation pressures will be strongest in white-collar jobs: office and administrative jobs, law, architecture and engineering, the sciences, business management, finance, sales, computer programming and systems administration (Briggs et al. 2023). Elondou and colleagues (2023) estimates that 80% of workers could see at least one-tenth of their work tasks automated by GPT, and just under one-fifth might see half of their work tasks automated. It's probably too early to pin down firm numbers, but there is much confidence that the technology will disrupt jobs, and many of them will be in work that had been sought by university students and for which universities trained students.

While there is good reason to be concerned about people's jobs, it is also worth retaining a critical attitude apocalyptic expectations as well. Those concerns have been a mainstay of modernity. They motivated the original Luddites' war on the automation of early 19th-century England's cloth-spinning industry. The combustion engine wrecked an entire economic system built around people's reliance on horses. Electrification likely hurt the gas lighting or ice delivery business. Such framings implies that automation pushes humans into obsolescence, causes long-term unemployment, and results in our material immiseration.

This is one narrative, but it does not describe the necessary consequence of automation. Decades ago, Queens College employed professional typists who would type up manuscripts. Their office might have resembled Figure 6 (below), which is an old photo of a typing pool <u>published by the Scottish Government</u>.² It was a room full of people typing other people's

² Scottish Government (2009) "One of the typing pools" Published to Flickr <<u>https://www.flickr.com/photos/scottishgovernment/3829002585</u>> on a CC BY 2.0 license

manuscripts. This job became obsolete with the introduction of personal computers, as we came to expect that people type up their own written documents. These particular jobs disappeared from colleges, but college jobs grew on the whole since the personal computer's introduction.



Figure 5: Undated photo published by Scottish Government on Flickr. "One of the typing pools"

In the grand scheme of things, decades of automation have made humans materially richer. Even today's poorest Americans are – on the whole – far wealthier than the most privileged of the pre-Industrial era (Pinker 2019). Although automation allows employers to replace workers with machines in the short-term, people seem to devise new products, tasks, and roles in the economy (Acemoglu and Restrepo 2019; Autor and Salomons 2018). So even though the technology may disrupt workers and benefit wealthier people in the short-term (Autor and Salomons 2018), some economists expect that humanity as a whole will be made wealthier by this kinds of technology-induced social changes (Briggs et al. 2023; Manyika 2017).

It can be comforting to think that things will all sort themselves out in the long-run if your employability does not feel threatened in the present. There are more immediate, practical questions facing people who work in (or were aspiring to work in) professions that will be affected by Generative AI. It is one thing to declare blithely that everyone will be fine in the end, and an altogether different matter to develop practical strategies that help today's affected workers navigate their way towards this future utopia. The next two sections focus on this question. We begin by clarifying what this technology actually does, and then discuss concrete ways that people might adapt to a coming, technology-altered workplace.

Part 2: How the Technology Works

In May 2023, "AI" industry leaders were exhorting Congress to regulate AI due to the threats posed by the technology (Roose 2023), warning that "Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war." It is odd to see industry leaders asking to be regulated. Part of the confusion stems from

there being two "AI's" – Generative AI (i.e., automated content generators) and General AI (programs with humanlike mental abilities). There are real concerns about the advancement of general artificial intelligence. Concretely, the "AI" of the software currently diffusing and integrating into practice (like Chat-GPT) is the far more limited Generative AI.

Before formulating responses to the technology, it is worth moving past flashy labels to develop a firmer grasp of what it actually does. This section attempts to develop a detailed but nontechnical explanation of Generative AI's inner workings. It is written for fellow social science and humanities professors, and focuses on the details of these methods in order to clarify the technology's basic character and limits for the purposes of contemplating future changes in workflow, job design, and vocational training systems. Young et al. (2018) offers a good, specialist-written reviews of the underlying work developed in the computer programming and statistics foundations of these fields. For introductory discussions of machine learning, neural networks, and deep learning, see Kulkarni and Harmon (2011), Suk (2017), Bi *et al.* (2019), and Sergihou and Rough (2023).

Generative AI vs General AI

The term "Generative AI" intimates that this technology is "artificial intelligence." There is a considerable dose of marketing in describing these technologies as unqualified "AI". Describing this software as such tacitly overstates its capacities and distorts our response to its introduction.

Generative AI lacks many of the abilities or functions that are possessed by humans and would be expected parts of an artificial mind that could meaningfully substitute for that of a human. For example, we would expect a synthetic human mind to show self-awareness, introspection, internal reasoning, or self-direction (McCarthy 1995, 2007a; Schank 1987). The programs do not develop new functionalities or adapt new roles. Generative AI does not create or reason, so much as it synthesizes human-provided words, pixels, or samples that recur in a set of training data. It is, in the words of Judea Pearl (2019), "curve fitting", generating communication in an "almost exclusively in a statistical, or model-blind, mode, which is analogous in many ways to fitting a function to a cloud of data points." It is more like an incomprehensibly complex calculator in which you enter phrases instead of numbers and operators as input, and from which you get an estimate of a media file described by that prompt instead of a math answer. The methods by which these predictions are calculated are described next.

Overview of the Process

Generative AI predicts the content of a digital media file based on a user-input prompt. You give the software some text, and the software will return digital content (in text, images, or audio) that pertains to that input. This feat is produced by a process described below in Figure 6:



Figure 6: Generative AI Development Process Depicted

Broadly, the process is built on the digitalization of human communications, in which we have converted text, images, audio, video, and code into digital computer files (Point 1 in Figure 6). The digital files are compiled into large libraries, and parsed to dissect these holistic acts of human communication (like a speech or painting) into constituent elements (like word fragments or pixels). Then, algorithms are deployed to find patterns in these dissected communications acts (2). Once these regularities are encoded into a predictive system, we can make guesses about the contents of hypothetical digital files (3). These predictions can then be reverse-coded into instances of what seems to be meaningful human communications (4).

Training Data

The process begins by compiling databases of human-produced communications. To train a text model, this might involve downloading and databasing people's written work in published books and articles, or in people's social media activity on platforms like <u>Twitter</u>, <u>Reddit</u>, <u>Github</u>, or <u>Stack Overflow</u>. To train an image-generation model, data can also be taken from online galleries, image-focused social media (e.g., <u>Instagram</u> or <u>Flickr</u>) or images indexed by <u>Google Images</u>.

For example, Figure 7 (below) is an entry from a <u>database of 300,000 articles and highlights</u> <u>from CNN and the Daily Mail</u> (from See, Liu, and Manning 2017). This is the kind of data that

article (string)	highlights (string)	id (str
(CNW) At least 14 people were killed and 60 others wounded Thursday when a bomb…	visit . Blast went off 15 minutes before…	"bf0cd4
"(CNN) Football superstar, celebrity, fashion icon, multimillion-dollar…	"Beckham has agreed to a five-year contract with Los Angeles Galaxy . New…	"2f43e9
"(CNN) A virus found in healthy Australian honey bees may be playing a role…	"Colony collapse disorder has killed millions of bees . Scientists suspect a…	"eb68bc
"LONDON, England Savers at a leading UK mortgage bank lined up for a second day to…	"Savers at leading UK mortgage bank lined up to empty their accounts . Northern Roc…	"ad021a
"LAS VEGAS, Nevada (CNN) Former football star O.J. Simpson will be held without bail…	"No bail for ex-NFL star accused of directing men in alleged armed robbery	"9d7fc7
"LAGOS, Nigeria (Reuters) Nigeria's television survival show has been suspended…	"Anthony Ogadje, 25, reportedly drowned in Shere Hills Lake . He was preparing for…	"30cb67
"(CNN) A former government contract employee was indicted on charges of stealin	"NEW: Indictment: Man tried to pass nuclear filters to foreign agent . NEW:…	"d41dc7
"LONDON, England Chelsea are waiting on	"Chelsea are still waiting on the fitness	

Figure 7: Instance of Databased Communication

would be used to build a text engine, like ChatGPT. The leftmost column ("articles") contains the text of an article, and the second column ("highlights") gives a short article summary that was composed by the news outlet. The task is to find patterns in "article" column that match patterns in "highlights."

Legal Conflicts. The legalities of using these data sources have become a sticking point in the development of this technology. Those in content creation-oriented occupations and organizations rightly see this technology as a threat to their earnings, and a technology that has been built on their intellectual property. Individual creators have begun to <u>initiate lawsuits</u> claiming that those who built Generative AI models did not have their permission to use their content when training models (Brittain 2023; Setty 2023). <u>Major media are reported to be following suit</u> (Allyn 2023).

Generative AI's involvement in legal disputes could redefine copyright boundaries. Traditionally, copyright prevents direct reproductions, but it does not confer ownership of style, ideas, concepts, or other abstract elements of speech that can be used across individual concerete speech acts (US Copyright Office 2021). These legal fights will evoke questions of whether people can claim ownership over elements of style or works that were inspired by others. For example, human visual artists often adopt stylistic elements from each other, as the artist Mickalane Thomas does from Andy Warhol in the left panel in Figure 8 (below left). This work, "Sweet and Out Front", was characterized in the *New York Times* as an "homage" to Andy Warhol (Sargent 2018). The right panel, which I generated using the prompt "ham sandwich in the style of Andy Warhol" on Dall-E, does the same. It has recognized and reprod A computer has recognized regularities in Warhol's images, and creates an adapted original work that picks up on those regularities. How is the process different? We will have to wrestle with whether these two acts are in fact different, and, if so, how to restrict creation via Generative Ai without stifling the larger creation of ideas and culture [which is the ultimate social goal of allowing people to privatize speech (Reid 2019)].





Figure 8: Two Works Inspired by Andy Warhol

Converting Multimedia to Data

Once the database is assembled and cleaned, the process continues with *tokenization*, in which the individual items of a training data set are broken down into constituent elements and organizes them as number sequences. A *token* is a continuent element of a multimedia file's constituent elements, like words in text, pixels in images, or samples in audio.

So, in the creation of a text generator like Chat-GPT, we might take pieces of writing like those depicted in Figure 7 (above). To start, we would convert the words, word fragments, or short phrases in that set into numbers, and store them along vectors (or sequences) that mark each encoded word's position in sentences and the overall document, relative to other words. Likewise, images can be encoded as sequences of pixels, as in Figure 9 (below) [reproduced from Iglesias *et al.* (2021)]. Typically, image pixels are encoded with 255-step intensity scales of red, green, and blue (RGB) light, with different combinations of these rendering the spectrum of colors displayed on device screens.



Figure 9: Image encoded in numeric matrix.

Pattern Identification

Generative AI uses machine learning, and specifically a machine learning method called "neural networks", to find patterns among tokens and between tokens and media file descriptors (like the story summaries in Figure 7 above). *Machine learning* blends developments in statistics and computer science to allow machines to "learn" from experience and without programming (Samuel 1959). Here, *learning* involves some kind of process in which a computer improves its performance in some task based on "experience" (i.e., data). There are many machine learning methods other than neural networks, like decision trees, support vector machines, or clustering (for a brief and lucid overview, see Bi et al. 2019).

Neural networks are simplified methods of pattern recognition modeled after the human brain (Krogh 2008; Suk 2017; see Warner and Misra 1996). Their task here is to identify configurations of words, pixels, or sound frequencies that recur in the training data. The process is often described in a way similar to Figure 10 (below), reprinted from Warner and Misra (1996).



Figure 10: Depiction of a Neural Network. From Warner and Misra (1996).

The process might begin by discerning observed combinations of tokens (individually represented as x's in the figure). This processing is done by artificial "neurons" (depicted by circles), a programming concept that is taken to represent the mechanism by which these networks identify and register the existence of patterns in training data. In each "layer" or step in the process (depicted in the figure as rows of circles), these neurons identify "weights", or parameters that determine the strength of the relationship between lower-order tokens or neurons. These weights are adjusted during the training process based on the network's performance, with the goal of minimizing the difference between the network's predictions and the actual data.

The concrete programming mechanism by which these patterns are extracted depends on the type of content being modeled. Figure 11 (below) reproduces an excellent animation of how this process might work in image scanning, as described by <u>Amidi and Amidi</u> (2018) (for an alternative depiction, see Suk 2017).



Figure 11: Depicting of Process by Which Image Features are Mapped. Reproduced from Amidi and Amidi (2018)

In this example, each "neuron" (in blue) specifies a pattern in the training data (in red), whose resemblance to pixel clusters is recorded to a "feature map" (in purple). This map captures the intensity of image features (like an edge or texture) in the images used to train the model. The process will continue, with the computer finding higher-order relationships among lower-order feature maps. For example, the model to "recognize" a hand, perhaps as a flat cube with five tubes emanating out of it, after seeing many such combinations of cubes and tubes, alongside image descriptions containing the word "hand." Further along the process, it might come to recognize color schemes or visual elements routinely associated with artist's names. So, it may come to recognize "Warhol" as being associated with four, neon-colored panels of the same image, as so many images with his name attached to it deliver images in that scheme (e.g., Figure 8 above).

Eventually, this web of recognized features will form a "fully connected", which means a large agglomeration of multiple layers of neurons that have picked out an extremely large set of patterns that prevail in the training data. The process can generate tens of billions of parameters, or relationships between different features or feature combinations observed in the data.

Prediction

Once the model is "trained", it is capable of generating responses based on patterns that it learned from its training data. The process begins when a user devises a *prompt* that describes the content that they want to generate. This prompt is tokenized, encoded to numbers, and processed by our trained model. So, if I ask Chat-GPT to "Write a romance story about an

apple", it will convert the prompt into a numerically-encoded ordered vector of ["Write", "a", "romance", "story", "about", "an", "apple"] for processing.

In response to the prompt, the model generates a sequence of tokens based on the associations learned in the training data. For instance, it may recognize from the pattern "Write a _____ story" that the output should be a structured narrative, while drawing on documents that contain a "romance story" or "story about an apple" to populate the characters, setting, and plot. These higher-order abstractions – like romance or apple stories, or the concept of stories themselves – are themselves built on patterns of words arranged into phrases arranged into sentences arranged into paragraphs. This complex web of associations or parameters generates predictions about the image, document, program, or other content described in prompt. When recoded into natural human communications, we have been surprised by how life-like and sophisticated the results have become.

The Result

The essence of Generative AI, and its shortcomings as a general-purpose artificial mind, are illustrated when we as an image generator to create a "Map of New York City, circa the year 500" (as I did using Dall-E 2 in Figure 12 below). The map does not show an actual map of New York in the year 500, but rather a predicted, typical pixel configuration found among the many images by "map", "New York" and "circa year 500" in the model's training data.



Figure 12: Generated Image: "Map of New York City, circa Year 500."

It has the general contours of a New York City map, and its notations and fonts look like one of those old-time maps that you see depicted as exhibits in books about history. It's a reverse-coded statistical prediction of pixels described with those three phrases used in combination, and based on a huge library of images. The configuration of lines – for example, of the islands and text – are good predictions for a computer with no cognizance of the City and its history

circa 500, but they clearly do not convey hard or reasoned information. It only loosely looks like that.

Automated Content Generators

After contemplating the inner workings of this technology, a few things should be clear. To start, the program has a far narrower functionality than one might presume of general AI. The software does not have a conscious mind or independent will. It does not reprogram itself to perform new functions. It does not understand the concepts or real-world referents encoded in its word, pixel, or frequency predictions. So this technology only automates part of the content creation process. It cannot generate a meaningful decision, analysis, or communication. However, it does allow the person who is orchestrating the creation of information, culture, or decisions to do more of the groundwork through computer generation instead of human labor.

In essence, Generative AI is a content generation algorithm. It models and generates a statistical prediction of what humans might express in response to a prompt. It is the multidimensionally average textual passage, image, or sound that appears in association with the prompt. The machine delivers inescapably average content. Yes, people can use it for creative purposes, but that creativity comes through the human effort of entering the prompt, or from the meaning-making that occurs as the human makes sense of the output. The algorithm is not guiding humans towards meaningful or socially valuable innovation on its own. It is a tool wielded by people in larger practical operations. It is still up to humans to figure out how to use this input-output box to practical, useful, or meaningful ends, and it is up to them to figure out how to feed and tweak that box to perform in a way that delivers them maximally useful output.

Part 3: Adapting to Generative AI

At this point, no one really knows how Generative AI will affect work, let alone the optimal response to it. However, some initial ideas have emerged after a half-year of wrestling with the technology.

Future Jobs Will Use AI

It seems quite likely that Generative AI will become a widely-used tool in many occupational tasks involving "mental" or "creative" work. Early studies by economists suggest that the technology helps workers become more productive, and renders considerable performance gains to low-performing workers (Brynjolfsson, Li, and Raymond 2023). Surveys of today's business leaders suggest that businesses plan to use the technology to automate, particularly once issues related to data quality, usage rights, data security, and regulatory compliance issues are cleared up (IBM 2023). The technology is coming to the workplace, people will need to learn how to use it, and educational institutions will likely need to help them.

Educators need to rethink their policies of prohibiting the technology from classes or school computers, or broadly conceiving of all uses of Generative AI as cheating. There is a clear rationale to learning to perform certain skills "by hand", much like we teach students to hand-

calculate as opposed to using a calculator. However, blanket restrictions on student use of Generative AI runs the risk of wholesale purging the technology from school. Students will have to learn to seize upon the technology's affordances outside of class, and teachers will not be pressed to adapt lessons to the new technological environment. Eventually, you will be left with the equivalent of a statistics program that teaches everyone to hand-calculate chi-square tests, while the rest of the world is using R and Python to scrape terabytes of data and simulate complex models. Worse yet, prohibition does not remove the technology from the classroom. Students with the family resources to purchase non-school computers will have access to the tech, and clever students will find a way to use Generative AI in a way that makes formal charges of plagiarism difficult to impossible (Terry 2023).

Integrating the Technology into Workflows

A recent, student-written op-ed in the *Chronicle of Higher Education* detailed the many ways that students could use Chat-GPT to write college essays without detection (Terry 2023). He described a strategy of using Chat-GPT to automate parts of the essay-writing process (topic brainstorming, essay outlining, advice in items to research for inclusion in individual subsections) to generate more ideas and come to a final product with less effort. To my mind, this describes an effective way to use the technology.

To create and process an informational or cultural product – like an essay – the larger project requires that several steps be taken. A person's *workflow* is whole sequence of steps or tasks required to perform a job. Some of these steps may be amenable to automation through Generative AI, which might allow people to do the job more quickly or better. Many of the tasks required to complete such projects are formulaic and lower-stakes, and could be done more efficiently (and perhaps more effectively) with computer assistance.

Generative AI works well in mental tasks that are formulaic, and for which low-quality or quality lapses are acceptable (Brynjolfsson and Mitchell 2017). Many information, analytical, or communicative tasks meet these criteria. The technology is useful for idea generation and brainstorming activities, and can find links between ideas that users might not anticipate. It is useful for quickly generating up the fallible first draft of a basic computer program or short story outline. It can synopsize the basic ideas present in a text document. It is excellent for synopsizing technical information, and for helping solve practical problems when using technology.

This suggests that human attention will be best reserved for tasks to which Generative AI is not well-suited, and for which quality requires some sort of guarantee. Humans will have to program the tool, parse out the nonsense, and fashion it into something that other people will value.

Retaining Human Mastery over the Process

Students often try to pass off Generative AI output as their own work. In effect, this is a situation in which someone is wholesale offloading the mental tasks of their role to a statistical

model. There are many reasons why it is a bad idea to pass off unvetted or lightly-vetted Generative AI output as your own work. It is important to convey these reasons to early-career people, as I believe that the technology can just as easily stunt their professional development as enhance their abilities:

The Need for Human Quality Guarantors. The algorithm is fallible, and an enterprise's failure to safeguard against that fallibility can cause problems. In May 2023, a New York lawyer threw himself on the mercy of the court after he was discovered to have <u>filed a Chat-GPT generated</u> <u>legal brief replete with citations of cases that do not exist</u> (Weiser 2023). A lawyer who is caught submitting fictitious claims might damage their client's case, and ultimately leave their employers or teams worse off for having contracted them. Enterprises rely on people to act as a layer of security between raw Generative AI output and a final product to be introduced to market. When a person fails to perform that function, they leave the larger enterprise of which they are a part vulnerable. Skilled professionals are expected to use their training, attention, and mental energies to direct, vet, correct, and improve upon raw Generative AI output.

By many indications, society is moving towards a model in which people in certain occupations are expected to act as guarantors of Generative AI output-derived content. The fallout of the "Chat-GPT lawyer" signal the expectations that lawyers guarantee the court brief is true, even when created with computers. A similar logic is at work in the aforementioned decision by the major journal *Nature*'s decision to forbid listing Chat-GPT as co-authors as a meaningful (and in my opinion positive) choice; it makes <u>scientists co-sign whatever content emerges from a process involving Generative AI</u> (Nature Editorial Board 2023). It is also at work when a company hires a professional illustrator to create graphic products, in hopes of avoiding embarrassing situations like printing <u>coloring book pages with three-legged chickens or foxes with the bird mouths</u> (Tiffany 2023). Somewhere in the real-world production chains that integrate Generative AI, there will be people held accountable. One strategy is to develop knowledge and skills to act as this checkpoint. It is to become the local expert who writes customized prompts built on a body of basic field-specific knowledge, and has the capacity to evaluate and fix the Generative AI output.

Communications is an Artifact of Human Attention. In much knowledge and culture work, there is an implicit understanding that signed content (e.g., reports, articles, books, artwork) represents meaningful communication, built on a person's meaningful mental processing. When someone sends us an email or text message, we assume that its contents represent the contents of the sender's mind. Meaning is drawn from the fact that a communication's contents reflects the subjectivity of a fellow person.

You do not hire an analyst to deliver *a* report, but rather to assess a case and convey the contents of that assessment in a report as an act of communication. The report is a marker and summary of an interpretation, judgment, or other line of thought from a fellow human with specialized knowledge and training. The same can be said for a painting or cartoon. There is little to no underlying exercise of human mental processing in a job that delivers an unvetted or

lightly vetted Generative AI output, or really anything that we would consider deliberate mental processing. Generative AI is just a statistics-based simulator of what a human might say.

A Path to Personal Obsolescence. Passing off lightly vetted and lightly-touched Generative AI as one's work is a path to personal obsolescence. As a mental exercise, consider the economic decisions involved in hiring research assistants. Prior to GPT, this was a job for which trained students were routinely employed. Currently, the Chat-GPT API generates text at the cost of \$0.02 per thousand tokens, or about \$0.00002 cents per word. That's about \$1.30 for as much text as a fairly weighty, 65 thousand-word tome. This is the substitution cost of a knowledge or creative worker who submits raw Generative AI output.

Advanced Technical Skills Remain Important

A recent survey of CEOs suggest that business leaders are deprioritizing technical skills (IBM 2023). It is my sense that advanced technical skills will retain their importance, though the specific skills involved in playing the role of expert may change.

Deep expertise allows people to vet the quality or applicability of the formulaic or generic content encoded by the models. For example, it takes expertise in advanced statistics to know if Generative AI is implementing or interpreting an advanced statistical operation correctly. Likewise, you need some understanding of visual design best practices before judging whether a computer-generated image conforms to them. In order to vet and fix raw Generative AI output, you need to know how to judge output quality and what to do when that output does not meet quality standards.

Building Credibility

Generative AI is creating a deluge of content. Amazon's ebook store was <u>floded with AI-generated books</u> (Baker-Whitelaw 2023; Terech 2023). AI-generated junk is flodding Etsy (Tiffany 2023). People will be facing a deluge of content, and the question of how to separate quality or useful content from the coming flood of low-quality stuff. Moreover, Generative AI can be used to game search engines to make sure that the low-quality stuff is what regular users find at the top of their query returns. The technology will be used to media with synthetic people who will offer advice. In that type of environment, an individual's reputation and credibility may become a valuable asset. Personal reputation and affiliations with well-known enterprises may be more valuable in establishing the reputation and credibility to become a trusted alternative to AI.

Leverage AI to Create New Opportunities

People often imagine job layoffs when they think about how automation sows efficiency and helps productivity. They do often do not think about how efficiency and productivity can be enabling by making new products and projects feasible. By cutting the costs, Generative AI may make it viable to create information, communications, or culture available to new audiences. Niches that were once too small to serve may not become more viable.

Preparing for Technological Disruption

There is much that we do not know about Generative AI. The technology is fairly new, and its integration into real-world enterprise is even more recent. Still, it is a development that we cannot ignore. This essay was an attempt to bring more people into discussions about how we can adapt. Part 1 described the ways that the technology is disrupting markets. Part 2 sought to break past the hype to explain the nuts and bolts of the technology's workings. Part 3 offered some initial thoughts on practical strategies that might help workers adapt.

If there is any core lesson to be gleaned, it is that this technology is worth our attention. Like the personal computer or the cellular phone, it will become a commonplace tool. The technology promises many benefits, and its threat to humanity seems quite overblown (at least this specific incarnation of "AI" in its present form). If workers are to stay relevant, and if the universities that train them are to adapt to a changing technological environment, we will have to engage the technology mindfully, rather than ignore it.

Bibliography

- Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33(2):3–30. doi: 10.1257/jep.33.2.3.
- Allyn, Bobby. 2023. "'New York Times' Considers Legal Action against OpenAI as Copyright Tensions Swirl." NPR, August 16.
- Amidi, Afshine, and Shervine Amidi. 2018. "Convolutional Neural Networks Cheatsheet." Retrieved May 24, 2023 (https://github.com/afshinea/stanford-cs-230-deep-

learning/blob/4653bc01297b269edb19e844b01127ba13de59df/en/cheatsheet-convolutional-neural-networks.pdf).

- Autor, David, and Anna Salomons. 2018. "Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share."
- Baker-Whitelaw, Gavia. 2023. "Amazon's Ebook Charts Are Full of Al-Generated Spam." *The Daily Dot*. Retrieved August 27, 2023 (https://www.dailydot.com/unclick/amazon-kindle-ai-ebooks-spam/).
- Barnett, Sofia. 2023. "ChatGPT Is Making Universities Rethink Plagiarism." Wired, January 30.
- Bi, Qifang, Katherine E. Goodman, Joshua Kaminsky, and Justin Lessler. 2019. "What Is Machine Learning? A Primer for the Epidemiologist." *American Journal of Epidemiology* 188(12):2222–39. doi: 10.1093/aje/kwz189.
- Bogost, Ian. 2023. "The End of Recommendation Letters." *The Atlantic*. Retrieved May 15, 2023 (https://www.theatlantic.com/technology/archive/2023/04/chatgpt-ai-college-professors/673796/).
- Briggs, Joseph, Davesh Kodani, and Giovanni Pierdomenico. 2023. *The Potentially Large Effects of Artificial Intellgience on Economic Growth (Briggs/Kodnani)*. Goldman Sachs.
- Brittain, Blake. 2023. "Lawsuit Says OpenAl Violated US Authors' Copyrights to Train Al Chatbot." *Reuters*, June 29. Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2023. "Generative Al at Work."
- Brynjolfsson, Erik, and Andrew McAfee. 2012. *Race Against the Machine: How the Digital Revolution Is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*. 58895th edition. Lexington, Mass: Digital Frontier Press.
- Brynjolfsson, Erik, and Tom Mitchell. 2017. "What Can Machine Learning Do? Workforce Implications." Science, December 22.
- Coyle, Jake. 2023. "ChatGPT Is the 'terrifying' Subtext of the Writers' Strike That Is Reshaping Hollywood." Fortune. Retrieved May 15, 2023 (https://fortune.com/2023/05/05/hollywood-writers-strike-wga-chatgpt-ai-terrifying-replace-workers/).
- Dobson, Stephen. 2023. "Why Universities Should Return to Oral Exams in the AI and ChatGPT Era." *Gizmodo*. Retrieved May 14, 2023 (https://gizmodo.com/chatgpt-ai-universities-colleges-should-give-oral-exams-1850363808).

- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. "GPTs Are GPTs: An Early Look at the Labor Market Impact of Potential Large Language Models."
- Harwell, Drew. 2022. "He Used AI to Win a Fine-Arts Competition. Was It Cheating?" *Washington Post*, September 2.
- Hu, Krystal. 2023. "ChatGPT Sets Record for Fastest-Growing User Base Analyst Note | Reuters." *Reuters*. Retrieved June 12, 2023 (https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/).
- Hughes, Stephen. 2023. "How Good Is ChatGPT at Diagnosing Disease? A Doctor Puts It through Its Paces." *The Conversation*. Retrieved May 15, 2023 (http://theconversation.com/how-good-is-chatgpt-at-diagnosingdisease-a-doctor-puts-it-through-its-paces-203281).
- IBM. 2023. 2023 Chief Executive Officer Study: Decision-Making in the Age of Al.
- Iglesias, Lara Lloret, Pablo Sanz Bellón, Amaia Pérez del Barrio, Pablo Menéndez Fernández-Miranda, David Rodríguez González, José A. Vega, Andrés A. González Mandly, and José A. Parra Blanco. 2021. "A Primer on Deep Learning and Convolutional Neural Networks for Clinicians." *Insights into Imaging* 12(1):117. doi: 10.1186/s13244-021-01052-z.
- Impact Research. 2023. Teachers and Students Embrace ChatGPT for Education.
- Kelly, Samantha Murphy. 2023. "ChatGPT Passes Exams from Law and Business Schools | CNN Business." CNN. Retrieved May 14, 2023 (https://www.cnn.com/2023/01/26/tech/chatgpt-passes-exams/index.html).
- Klein, Gary, Roberta Calderwood, and Anne Clinton-Cirocco. 2010. "Rapid Decision Making on the Fire Ground: The Original Study plus a Postscript." *Journal of Cognitive Engineering and Decision-Making* 4(3):186–209.
- Krogh, Anders. 2008. "What Are Artificial Neural Networks?" Nature Biotechnology 26(2):195-97.
- Kulkarni, Sanjeev, and Gilbert Harman. 2011. An Elementary Introduction to Statistical Learning Theory. 1st ed. Wiley.
- Manyika, James. 2017. "Technology, Jobs, and the Future of Work." *McKinsey Global Institute Executive Briefing*. Retrieved May 16, 2023 (https://www.mckinsey.com/featured-insights/employment-andgrowth/technology-jobs-and-the-future-of-work).
- McCarthy, John. 1995. ""Making Robots Conscious of Their Mental States.""
- McCarthy, John. 2007a. "From Here to Human-Level AI." Artificial Intelligence 171(18):1174–82.
- McCarthy, John. 2007b. ""What Is Artificial Intelligence?" Stanford University, Stanford, CA.
- Nature Editorial Board. 2023. "Tools Such as ChatGPT Threaten Transparent Science; Here Are Our Ground Rules for Their Use." *Nature* 613(7945):612–612. doi: 10.1038/d41586-023-00191-1.
- Needleman, Sarah. 2023. "How AI Is Building the Next Blockbuster Videogames." Wall Street Hournal, April 17.
- Pearl, Judea. 2019. "The Limitations of Opaque Learning Machines." in *Possible Minds: 25 Way of Looking at AI*, edited by J. Brockman. Penguin Books.
- Perlman, Andrew. 2023. "The Implications of ChatGPT for Legal Services and Society." *Harvard Law School Center* on the Legal Profession. Retrieved May 15, 2023 (https://clp.law.harvard.edu/article/the-implications-ofchatgpt-for-legal-services-and-society/).
- Pinker, Steven. 2019. Enlightenment Now: The Case for Reason, Science, Humanism, and Progress. Reprint edition. New York: Penguin Books.
- Reid, Amanda. 2019. "Copyright Policy as Catalyst and Barrier to Innovation and Free Expression." *Catholic University Law Review* 68(1):33–86.
- Robertson, Joshua. 2023. "System Shock Remake Gets Dragged On Twitter For Using Al Art, Doubles Down." *TheGamer*. Retrieved May 13, 2023 (https://www.thegamer.com/system-shock-remake-criticised-ai-art/).
- Roose, Kevin. 2022. "An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy." *The New York Times,* September 2.
- Roose, Kevin. 2023. "A.I. Poses 'Risk of Extinction,' Industry Leaders Warn." The New York Times, May 30.
- Samuel, A. L. 1959. "Some Studies in Machine Learning Using the Game of Checkers." *IBM Journal of Research and Development* 3(3):210–29. doi: 10.1147/rd.33.0210.
- Sargent, Antwaun. 2018. "Seven Artists on the Warhol Influence." *The New York Times*, November 1. Schank, Roger C. 1987. "What Is AI, Anyway?" *AI Magazine* 8(4):59–59.
- Scott, Inara. 2023. "Yes, We Are in a (ChatGPT) Crisis." Inside Higher Ed. Retrieved May 14, 2023
 - (https://www.insidehighered.com/opinion/views/2023/04/18/yes-we-are-chatgpt-crisis).

See, Abigail, Peter J. Liu, and Christopher D. Manning. 2017. "Get To The Point: Summarization with Pointer-Generator Networks."

- Serghiou, Stylianos, and Kathryn Rough. 2023. "Deep Learning for Epidemiologists: An Introduction to Neural Networks." *American Journal of Epidemiology* kwad107. doi: 10.1093/aje/kwad107.
- Setty, Riddhi. 2023. "Sarah Silverman, Authors Hit OpenAI, Meta With Copyright Suits." *Bloomberg Law*. Retrieved August 19, 2023 (https://news.bloomberglaw.com/ip-law/sarah-silverman-authors-hit-openai-meta-with-copyright-suits).
- Stokel-Walker, Chris. 2023. "ChatGPT Listed as Author on Research Papers: Many Scientists Disapprove." *Nature* 613(7945):620–21. doi: 10.1038/d41586-023-00107-z.
- Suk, Heung-II. 2017. "An Introduction to Neural Networks and Deep Learning." Pp. 3–24 in *Deep Learning for Medical Image Analysis*, edited by S. K. Zhou, H. Greenspan, and D. Shen. San Diego, UNITED STATES: Elsevier Science & Technology.
- Terech, Kristtina. 2023. "Amazon Has a Big Problem as AI-Generated Books Flood Kindle Unlimited | TechRadar." *Tech Radar*. Retrieved August 27, 2023 (https://www.techradar.com/computing/artificialintelligence/amazon-has-a-big-problem-as-ai-generated-books-flood-kindle-unlimited).
- Terry, Owen Kichizo. 2023. "I'm a Student. You Have No Idea How Much We're Using ChatGPT." *The Chronicle of Higher Education*, May 12.
- Tiffany, Kaitlyn. 2023. "Al-Generated Junk Is Flooding Etsy." The Atlantic, June 15.
- Tiku, Natasha. 2022. "The Google Engineer Who Thinks the Company's AI Has Come to Life." *Washington Post*, June 11.
- Turing, Allan M. 1950. "Computing Machinery and Intelligence." *Mind* LIX(236):433–60. doi: 10.1093/mind/LIX.236.433.
- US Copyright Office. 2021. *Circular 1: Copyright Basics*.
- Vincent, James. 2019. "OpenAI's New Multitalented AI Writes, Translates, and Slanders." *The Verge*. Retrieved August 17, 2023 (https://www.theverge.com/2019/2/14/18224704/ai-machine-learning-language-models-read-write-openai-gpt2).
- Warner, Brad, and Manavendra Misra. 1996. "Understanding Neural Networks as Statistical Tools." *The American Statistician* 50(4):284–93. doi: 10.2307/2684922.
- Weiser, Benjamin. 2023. "Here's What Happens When Your Lawyer Uses ChatGPT." The New York Times, May 27.
- Westfall, Chris. 2023. "Educators Battle Plagiarism As 89% Of Students Admit To Using OpenAl's ChatGPT For Homework." *Forbes*. Retrieved May 14, 2023

(https://www.forbes.com/sites/chriswestfall/2023/01/28/educators-battle-plagiarism-as-89-of-students-admit-to-using-open-ais-chatgpt-for-homework/).

- Wodecki, Ben. 2023. "UBS: ChatGPT Is the Fastest Growing App of All Time." *AI Business*. Retrieved June 12, 2023 (https://aibusiness.com/nlp/ubs-chatgpt-is-the-fastest-growing-app-of-all-time).
- Yang, Maya. 2023. "New York City Schools Ban AI Chatbot That Writes Essays and Answers Prompts." *The Guardian*, January 6.
- Young, Tom, Devamanyu Hazarika, Soujanya Poria, and Erik Cambria. 2018. "Recent Trends in Deep Learning Based Natural Language Processing." *Ieee Computational IntelligenCe Magazine* 13(3):55–75.
- Zwiezen, Zack. 2022. "Rick And Morty Creator Used Controversial AI Art, Voice Acting In New Shooter." *Kotaku*. Retrieved May 13, 2023 (https://kotaku.com/high-on-life-justin-roiland-ai-art-rick-morty-1849900835).