

WestminsterResearch

http://www.westminster.ac.uk/westminsterresearch

Investigating AI: Reflections on Data Journalism in Practice Broussard, Meredith

This is a PhD thesis awarded by the University of Westminster.

© Ms Meredith Broussard, 2025.

The WestminsterResearch online digital archive at the University of Westminster aims to make the research output of the University available to a wider audience. Copyright and Moral Rights remain with the authors and/or copyright owners.

Investigating AI:

Reflections on Data Journalism in Practice

By

Meredith Broussard

A thesis submitted in partial fulfilment of the requirements of the University of Westminster for the degree of Doctor of Philosophy

April 2025

Abstract

This commentary includes analysis of and context for two books: *Artificial Unintelligence: How Computers Misunderstand the World* and *More Than a Glitch: Confronting Race, Gender, and Ability Bias in Tech.*

List of Contents

ACKNOW LEDGEMENTS	••••
AUTHOR'S DECLARATION	5
INTRODUCTION	6
OVERVIEW	8
LITERATURE REVIEW	8
PUBLISHED WORKS	11
Artificial Unintelligence	11
More Than a Glitch	15
EXTENDING THE ANALYSIS TO A NEW TECHNOLOGY	18
IMPACT	36
DISCUSSION AND FUTURE WORK	38
BIBLIOGRAPHY	42

Acknowledgements

I am grateful to my advisor, Dr. Pieter Verdegem, for his guidance and collegiality. I am grateful to my readers, Dr. Gina Neff and Dr. Maria Michalis, for their thoughtful comments on the work. I thank Dr. Natasha Whiteman and Dr. Alessandro D'Arma for chairing my committee. I thank my colleagues at CAMRI and at NYU, particularly those in the Arthur L. Carter Journalism Institute and in the Center for Faculty Advancement, as well as my editor and the staff at MIT Press. I am especially grateful to my family and friends for their support during this process.

Author's Declaration

I declare that all the material contained in this thesis is my own work, except for the AIgenerated text as indicated.

Introduction

I have been working as an academic, and later a data journalist, since 2006. In college, I studied computer science and English literature. I began my career as a professional computer scientist in the telecom industry, where I worked on 1990s-era AI and then transitioned to internet technology. After a few years, I quit tech to become a journalist. I began as a feature journalist because it was 2000, still early in the internet era, and I wasn't aware that it was possible to get a job doing what was then called computer-assisted reporting. Around 2011, I became aware of what people were calling a new kind of journalism, data journalism, which I define as the practice of finding stories in numbers and using numbers to tell stories. The new name encompassed the idea that data journalists were doing all the same things as computer-assisted reporters, like spreadsheets and databases, but were also adding web-based elements like data visualization and other affordances of commercial web browsers. I started doing freelance investigative data journalism. In 2013, I secured one of the first full-time tenure-track assistant professor jobs specifically focused on data journalism. I currently do investigative reporting with and about artificial intelligence.

From the beginning, it was clear that data journalism had a bumpy path inside the academy. At my first interview for my tenure-track job, a professor walked into the empty conference room where I was waiting, sat down with a huff, looked at me for a moment, and said, "You seem like a smart young lady. Why don't you have a Ph.D.?" I should mention that this colleague was much older and taught in a different sub-discipline of journalism; they had been promoted to full professor in the year I was born. I explained diplomatically that there was not at that time any Ph.D. in data or computational journalism. The highest degree available in related fields in the United States was a master-level degree, which I had. Nobody was offering a Ph.D. in data or computational journalism because there weren't any professors to teach in it. If got this job, I said, I'd be in a position to eventually help start a Ph.D. program. The professor seemed surprised but accepted my answer.

I tell this story because it's a useful example of the routine sexism and microaggressions that women, especially women of color like me, face in the academy. It's also a useful example of a particular kind of assumption that people make about computational work in journalism. Many people assume that because computers are ubiquitous in everyday life, and because consumer electronics are often so easy to use, a wide variety of high-level computational work must be happening inside every field. In my experience, this is rarely the case. As a new assistant professor, I started looking for tenured professors of data journalism senior to me and found only a few. These were full professors who occupied endowed chairs at a handful of universities. The Knight Foundation had established several chairs at American universities in order to accelerate data journalism within the academy. It was an effective strategy in that data journalism achieved a foothold in classrooms in journalism schools. However, it was less effective in changing academic norms. Because the Knight professors were not Ph.Ds., and were appointed at the rank of full professor without the requirement of any graduate degree, many of them had only bachelor's degrees and were essentially clinical or teaching professors. They were extremely effective at teaching data journalism skills to journalism students and normalizing quantitative analysis in the mainstream journalism curriculum. Their remit did not include changing the scholarly publishing norms inside journalism and mass communication; nor did it include updating the tenure and promotion standards at universities; nor did it include expanding quantitative inquiry in curriculum standards through the US accrediting body for journalism schools, the Accrediting Council on Education in Journalism and Mass Communications (ACEJMC).

This rocky path for computational faculty is in sharp contrast to the promises and prophecies that have been made about journalism in the digital age. Technology has often been portrayed as salvation inside academic journalism. Creech and Mendelsohn (2015) write: "At the center... often rests an idealized notion of the technologically adept journalist who creates compelling content by mastering digitally based techniques of reporting, producing, and distributing the news, thus offering a blueprint for managing journalism's current crisis." Just like the media industry, the academy has been eager to adopt new technologies. However, both spheres struggled (and still struggle) to absorb technologists into their institutional frameworks when it comes to management, promotion, and recognition of professional accomplishments. When I began submitting computational work to scholarly journals, I discovered that there were few peer reviewers available who could evaluate both technology and scholarship. I am grateful to the editors who went above and beyond to find reviewers for my work. I eventually wrote a paper about the need to integrate computational work into tenure and promotion standards in journalism and mass communication departments, as well as the need for reproducibility in scholarly code experiments (Broussard 2016).

All of this is to say: I have been working as a scholar and a journalist for a long time, I have published a lot of work in trade and scholarly journals, sometimes as the first scholar to do a particular kind of computational work, and I have followed an unconventional path to do so. I do not have a Ph.D. That is a situation I hope to rectify with this dissertation. The work in this portfolio is an analysis of original research in computational journalism and is presented as part of a Ph.D. by published work. In this commentary, I explore the strategy that animates my work: creating complex data journalism stories that serve as critical data studies interventions and help readers understand how longstanding social issues manifest inside technological systems.

Overview

This commentary focuses two of my books. *Artificial Unintelligence: How Computers Misunderstand the World* (2018) and *More Than a Glitch: Confronting Race, Gender, and Ability Bias in Tech* (2023). Both were both published by MIT Press and translated into several languages. These are the longest works I have produced in the past ten years. I have omitted from this portfolio many works of journalism, several scholarly articles, a handful of book chapters, and a number of coding projects. In this commentary, I situate my work in a scholarly context and summarize the argument that I make in each book critiquing artificial intelligence and tech culture generally. I then extend the critique to the latest group of AI tools, generative AI, by evaluating how a Google model summarizes my work. This latter exploration will likely form the basis of a future essay. The work in this portfolio reflects my interdisciplinary approach, my commitment to engaging multiple audiences, and my intellectual journey over the past several years.

Literature Review

In *The Elements of Journalism*, Kovach and Rosentiel (2014) write: "The primary purpose of journalism is to provide citizens with the information they need to be free and self-governing." With this goal in mind, I situate my investigative work within the field of algorithmic accountability reporting, which performs an important role in preserving democracy by shedding light on the ways that algorithms restrict freedom rather than enhancing it. I join journalists like Julia Angwin and Cathy O'Neil in exploring how many algorithms once thought to be "neutral" or "objective" in fact infringe on citizens' freedoms. (O'Neil 2016; J Angwin et al. 2016)

I follow Turner and Hamilton (2009) in defining computational journalism as "the combination of algorithms, data, and knowledge from the social sciences to supplement the accountability function of journalism." My work is interdisciplinary: I primarily engage with the field of journalism and mass communication, but I draw from ideas in sociology,

computer science, data science, and digital archiving. I write both prose and code, and I embrace the way that this allows me the flexibility to write for scholarly or general audiences on different platforms, choosing each platform based on what I want to communicate and which audience(s) I want to engage. In this way, I move between what Coddington (2015) calls "three quantitative forms of journalism—computer-assisted reporting, data journalism, and computational journalism."

My scholarly work comprises some of the first applied research in computational journalism. Because of this, there are few methodological precursors. I took this opportunity to develop my own interdisciplinary methods, informed by scholarship from journalism and mass communication as well as other fields. My investigative data journalism work is inspired by work on algorithmic accountability by Angwin (2016), Sweeney (2013), and Nelson (2024). Anderson's (2013) work on a sociology of computational and algorithmic journalism is helpful in situating my work in communication studies. The idea of making a prototype as a proof of concept is popular in design research (Wensveen and Matthews 2014) and software engineering. Reverse engineering, the technique I frequently use when writing about other people's software, comes from engineering (Friesinger and Herwig 2014). The idea of building technology for public benefit is part of public interest technology, a newer field (McGuinness and Schank 2021; Ward and Bruce 2022; Pahlka 2023). The idea of interrogating data-driven, AI, or algorithmic systems is considered part of critical data studies, critical algorithm studies, or critical AI studies. My work is also informed by algorithmic accountability, design justice, data feminism, and critical race perspectives on technology (Costanza-Chock 2020; D'Ignazio and Klein 2020; O'Neil 2016; Caplan et al. 2018; Hanna et al. 2019; Garcia, López, and Vélez 2018; McIlwain 2020; Verdegem 2021).

I endeavor to build knowledge in my work. When I approach a new project, whether it is in campaign finance or gender studies or critical race & digital studies, I first try to find out the general outline of the field. I determine what the big questions are that top scholars and thinkers are working on, and what problems have already been solved. Through the process of reading in the popular press and in scholarly journals (similar to a systematic review), plus interviewing a wide variety of amateurs and experts in the field, I develop a sense of what is known and what is unknown in the field. My interviews tend to start with informational interviews and then progress to semi-structured interviews (Weiss 1995; Lareau 2021). Then, I synthesize all of the gathered information in order to identify a highly specific research question that has the potential to slightly enlarge the boundary of what is known. Next, I pick the method of inquiry and the format for my research output. Sometimes I build a prototype in order to figure out if something is possible, as in when I built artificial intelligence software to investigate textbook shortages at Philadelphia public schools. In that case, I needed to build a very specific piece of software in order to answer my research question: "Do Philadelphia public schools have the books and other learning materials students would need in order to learn the material on the state-mandated standardized tests?" The software didn't exist, so I designed and built it. Using a solutions journalism approach (Wenzel, Gerson, and Moreno 2016), I added another research question: "If schools don't have the materials the students need, how much money would be required to remedy the situation at each school?"

I realized that the software I built could be useful to other reporters who were working on education stories, and I decided to publish my software online as a news app (Broussard 2015) accessible to others. In addition to designing the front end and back end code for the software system, I wrote a handful of journalistic pieces based on my findings (Broussard 2014b); I wrote an academic paper about using AI for investigative reporting (Broussard 2014a); and I published parts of the software on Github, a popular code-sharing platform. Usually, it is easier to publish only the front-end code rather than sharing the backend code as well. When I wrote Bailiwick, a subsequent AI system focused on finding stories in campaign finance data, I only published the front-end, or user-facing, system. If I am doing a story that involves reverse-engineering, I typically don't publish any code but I do explore the system's documentation; validate my findings via interview; and iterate in order to discern truth. I try to link to the documentation in whatever story I write.

The prototypes I make are entirely functional, and sometimes can be commercialized. The open-source school textbook management technology that I developed for a chapter of *Artificial Unintelligence* was adopted by a group of University of Pennsylvania students who read the book, and they tried to commercialize the technology and have it adopted by school districts.

Usher (2016) has written of the routine work practices of data journalists working in newsrooms. I've followed many of these processes and have adapted others to my own specific situation. I work as a freelance journalist in addition to my work as an academic. This means that I don't have the same constant group of colleagues or collaborators; I have to find them for each new project. For a code-intensive story, I typically do a lot of research, write some code, and then pitch the story to a news outlet in order to customize the code and the story for the publication outlet. I work with collaborators as appropriate. Funding is usually a challenge. I've used various strategies, including reporting grants and fellowships; contracted work with a news organization; funding from my research budget; or self-funding.

My scholarship takes different forms depending on whether I'm writing for a scholarly audience or the general public. In my journalism, I draw on literary journalism techniques and ethnographic participant-observation methods. My particular flavor of immersion journalism (Conover 2016) involves technological interventions where I build technology in order to illuminate an often-hidden world. Sometimes I conceptualize the project as a stunt, inspired by reporters such as Ida B. Wells or Nelly Bly (Todd 2021). For example: in one chapter of Artificial Unintelligence, I participate in a cross-country hackathon aboard what organizers call the "Startup Bus" in order to show readers what it is like to build technology, warts and all. I could have picked any hackathon to write about, but I picked the one that seemed at the time like the most provocative. In these kinds of projects, I am engaging in performative software development in order to demonstrate some underlying social issue. I am particularly invested in explaining complex technical topics in plain language in order to increase readers' computational literacy and include more people in democratic decision-making around algorithms. This strategy of detailed, audienceappropriate explanation has proved to be a meaningful intervention that has helped many people to feel more comfortable using and critiquing AI.

Published Works

Artificial Unintelligence

In *Artificial Unintelligence*, I offer a guide to understanding the inner workings and outer limits of technology—and warn that we should never assume that computers always get things right. Aimed at a general audience, the book explains artificial intelligence in plain language and takes the reader on a series of computational "adventures" designed to illuminate complex technical topics. The book developed out of a need I observed while working on applied research projects in computational journalism for which I built artificial intelligence tools for investigative reporting. In the classroom and the newsroom, I often had to explain what artificial intelligence is (and is not) in order to explain how I arrived at a particular conclusion in my reporting. I began to pay attention to the way that explanatory

journalism is essential to helping audiences understand the increasingly opaque technological world in which algorithms are used to make decisions on society's behalf.

Artificial Unintelligence couples explanatory reporting about artificial intelligence with original algorithmic accountability reporting, using a literary style derived from immersion journalism. An immersion journalist engages with a specific community as a participant-observer in order to convey an insider's perspective on what it is like to be a part of that unique world. The audience is taken along on the reporter's learning process. Since I have written code in at least a dozen different languages, I can take readers inside previously inaccessible technological worlds that require a certain amount of technical knowledge as the price of entry. I have developed a particular expertise in this kind of writing, which I think of as immersion reporting for technology. I usually put myself into the piece as a character in order to give the reader a window into an otherwise impenetrable world.

Consistent with the applied computational aspect of my research agenda, two chapters of the book focus on stories for which I developed investigative reporting software. The software is an original framework that I call a Story Discovery Engine. Two examples of the software, Bailiwick and Stacked Up, are accessible on Github upon request. The design is a variation on an expert system, a type of artificial intelligence system that was popular in the 1980s and 1990s. Expert systems are a sub-field of artificial intelligence. Machine learning is currently the most popular and well-known sub-field of artificial intelligence, but expert systems and natural language generation (NLG) are both AI sub-fields with useful applications for journalism.

The original idea for an expert system was that the computer scientist would collect knowledge about a domain, encode that knowledge in a set of rules, and embed the rules in a computational system. The system would then act like a kind of expert in a box, answering users' questions. This paradigm was touted as a way to replace experts like doctors or lawyers with computers. However, computer scientists realized after a few years that human knowledge is too complex to be captured in a simple rules-based system, and expert systems fell out of vogue.

I resurrected the idea of an expert system and hacked the idea, realizing that it is not necessary for the machine to autonomously come to a conclusion—that in journalism, it is perfectly sufficient to build a human in the loop system in which the system makes data visible in a way that allows the reporter to investigate a set of common stories on a public affairs beat. I wrote a scholarly paper about the Story Discovery Engine paradigm for the Computation + Journalism symposium, revising it for the journal *Digital Journalism*.

As I disseminated my findings from the project, the most common question I received was whether a Story Discovery Engine was an actual machine that spits out story ideas. (Notably, this was well before the current generative AI era, in which machines spit out mediocre story ideas as well as error-ridden stories.) I explained that was not my intention, that the engine was conceptual rather than physical. But after dozens of these conversations, I decided to explore whether it was possible to build such a machine. I had developed the first Story Discovery Engine, which I named Stacked Up, to help me report a story on textbook shortages in Philadelphia public schools. I had realized quickly that the system I built could be used by other reporters as well. I hypothesized that the Story Discovery Engine model I had developed for the textbook story would be generalizable to other public affairs domains that have lots of public data available. Thus was born Bailiwick, an AI engine to help investigative reporters quickly and efficiently uncover new story ideas in campaign finance data. I developed Bailiwick with the support of a research fellowship awarded by the Tow Center for Digital Journalism at Columbia Journalism School, and NYU generously supported its subsequent operating expenses. After its launch, it was used by editors and journalists, validating its effectiveness. I once gave a talk at the Computation + Journalism conference at Stanford about the software. An editor in the audience, who ran a high-end travel site, pulled out his laptop to look at Bailiwick. By the end of my talk's Q&A, he had used the software to find a story idea about politicians and travel expenses, and had assigned a writer to work on it.

As of January 1, 2017, Bailiwick included 94.2 million records in its database. The number of records increased every time new filings were added to the Federal Election Commission (FEC) website. The system gathered small updates from the FEC site every night, and then it completely refreshed all of its FEC data once a week. The system took a year to plan and build, which at the time was unusually fast for this kind of complex software and such a limited budget. I decommissioned Bailiwick when the FEC changed its data feed format in mid-2017.

To design the Bailiwick project, I identified the top campaign finance data reporters across the US and contacted them for interviews. I interviewed several of them at length at the annual National Institute for Computer Assisted Reporting (NICAR) conference, a data journalism conference that I regularly attend. I learned about the data journalists' writing process; I learned which tools they find useful; I learned which tools they thought would be useful and which didn't exist at that time. I asked about the stories they wrote, identifying the most common types of stories that are derived from structured data analysis. I did not use coding software such as NVIVO to analyze the interview data. Rather, I used basic reporting techniques, similar to content analysis, to identify themes and patterns from my interview notes. Then, I designed a system to surface the data that would help other reporters write those common stories. I devised a visual schema that represented all of the relevant entities identified by my interview subjects and mapped each of the entities to the data points found in multiple campaign datasets. Usually, data journalists work with a single dataset and mine it for story ideas. I took multiple datasets, joined them in a complex architecture similar to a knowledge map, and mined the resulting system for story ideas similar to the story tropes identified by my interview subjects.

Images of the Bailiwick front end may be found in the penultimate chapter of *Artificial Unintelligence*; a diagram of the back-end system is included. I created the system architecture and the majority of the code myself, hiring contractors for some of the visual design and coding work.

One important distinction between generative AI and this earlier work: a Story Discovery Engine does not explicitly suggest a story. Rather, it combines data elements and visualizations in a way that triggers the reporter's own creativity. A Story Discovery Engine provides a lead, then the reporter follows up and does some original reporting. Because the situation on the ground is different inside each school or each political campaign, the resulting story will be slightly different. Generative AI or NLG, given 25 different sets of similar-looking data, will generate 25 different almost-identical stories. A human reporter, given 25 sets of similar-looking data, will likely find an interesting angle on each because of their professional expertise and the inevitable surprises that arise during the reporting process. News requires interestingness, novelty, and serendipity; people ignore news if it is the same all the time. This is one of the major reasons that generative AI has not replaced journalists in newsrooms. Another reason, of course, is that generative AI generates untruths and journalism is a discipline of pursuing truth.

Artificial Unintelligence became part of an emergent scholarly field called critical algorithm studies, critical data studies, or critical race and digital studies. The book has become a core text in data ethics classes internationally and has been translated into 7 languages.

More Than a Glitch

In 2020, after completing several research projects focused on digital archiving, I transitioned back to exploring artificial intelligence, asking: how can people better understand the real-world biases embedded in AI systems? This appealed to me as both a technical challenge and a science communication issue.

The question emerged as a result of public engagements around *Artificial Unintelligence*. I had the great good fortune of being asked to speak about the book's topics to audiences around the world. People always asked more questions about Chapter 6, "People Problems." In it, I looked at the people who launched AI, and explored how their particular ideologies shaped the scientific field. I ended up discussing this in a documentary, "Coded Bias," that was nominated for an Emmy award. I was grateful, and also curious about why this idea landed with people. This curiosity led me to become even more interested in the history of science, and I began reading and thinking about the ways individuals shaped today's narratives about technology. Then, the Covid pandemic brought everything to a crashing halt, and I found myself asking questions about people's faith in technology—and the ways that technology lets us down.

I developed a new theme in my research that looks into the extremely human issues embedded in technology. *Artificial Unintelligence* explored institutional failures of technology in public education, campaign finance, and more. My next project built on that work in order to illuminate structural issues in society that are exacerbated by algorithmic technologies. In *More Than a Glitch: Confronting Race, Gender, and Ability Bias in Tech*, I argue that when technology reinforces inequality, it's not merely a glitch—it's a signal that we need to redesign our systems to create a more equitable world.

The word "glitch" implies an incidental error, as easy to patch up as it is to identify. The book asks the question: what if racism, sexism, and ableism aren't just bugs in mostly functional machinery—what if they're coded into the system itself? Inspired by the work of Safiya Umoja Noble, Cathy O'Neil, and Ruha Benjamin, I aspired to demonstrate how neutrality in tech is a myth and why algorithms need to be held accountable.

The first two chapters of the book orient the reader to the central argument and offer plain-language explanations of the math behind machine learning. I find that writing about complex technology requires a simple explanation of how the technology works as well as a detailed human example of how and when the technology might be used. Because so many people have math anxiety, I try to be slow, calm, and reassuring when I write about mathematical topics. There's no getting around the fact that AI is math, and the math is difficult to understand. However, most people can do it if they take it slow. Most tech writing uses examples from science fiction like Star Trek or Star Wars. I deliberately use domestic examples, as a feminist intervention. In these initial chapters, I link the technology to the larger social issues under discussion. I write:

Many people, when confronted with ableism, race, or gender bias in tech, tend to consider it a glitch. A glitch is something temporary, a mysterious blip that may or may not be repeated. A bug is a more serious matter that makes the software fail, and it is worth addressing. A glitch is ephemeral and can be dismissed; meanwhile, a bug is substantial, ongoing, and deserves attention. Developers use cognitive shortcuts to figure out which problems merit fixing. This is a normal strategy—it's how the human brain operates. We all use cognitive shortcuts, and in today's world where decision making is more complex than ever, shortcuts are essential. The problem is, shortcuts often contain bias. Shortcuts focus on race, gender, ability, and other superficial categories.

The book explores a range of examples. Chapter 3 looks at facial recognition technology trained only to recognize lighter skin tones. I use the example of Robert Julian-Borchak Williams, a Detroit man who was arrested for robbery based on a false match generated by a police facial recognition system. I use the Gender Shades project (Buolamwini and Gebru 2018), the subject of "Coded Bias," to illuminate the underlying technical and racial issues. Facial recognition systems are better at recognizing light skin than dark skin, better at recognizing men than women, and often don't consider trans, nonbinary, or gendernonconforming people at all. Together, the two elements of scholarship and lived experience vividly demonstrate real and potential algorithmic harms of facial recognition systems.

The facial recognition chapter is paired with another chapter on technology used in policing. All of the work on algorithmic accountability and disparate impact in policing flows from Latanya Sweeney's work on discrimination in online ad delivery (2013) and Julia Angwin's ProPublica work on the COMPAS algorithm (2016). I found Sarah Brayne's ethnographic work (2021) about how the Los Angeles Police Department uses technology to be especially useful in understanding CompStat, a statistical system used to collect and analyze crime statistics. I write:

CompStat was widely adopted by police forces in major American cities in the 1990s and 2000s. By relying heavily on crime statistics as a performance metric, the CompStat era trained police and bureaucrats to prioritize quantification over accountability. Additionally, the weekly meetings about crime statistics served as rituals of quantification that led the participants to believe in the numbers in a way that created collective solidarity and fostered

what organizational behaviorists Melissa Mazmanian and Christine Beckman call "an underlying belief in the objective authority of numbers to motivate action, assess success, and drive continuous organizational growth" (2018). In other words: technochauvinism became the culture inside departments that adopted CompStat and other such systems. Organizational processes and controls became oriented around numbers that were believed to be "objective" and "neutral." This paved the way for the adoption of AI and computer models to intensify policing—and intensify surveillance and harassment in communities that were already overpoliced.

After the policing chapters, I look at bias in educational technology using the example of an algorithm purchased by the International Baccalaureate, an organization that awards a prestigious international secondary school diploma. That system generated imaginary grades that were assigned to real students. The system was biased and resulted in lower-income students being denied admission to college or being denied college credit for their high school achievements. I next look at gender and databases, using the example of Jonathan Ferguson, a British civil servant who transitioned in 1958 (Hicks 2019). Most people don't realize that 1950s ideas about gender are embedded in today's software.

I also look at mortgage-approval algorithms that encourage discriminatory lending and dangerous feedback loops that arise when medical diagnostic algorithms are trained on insufficiently diverse data. I became curious about medical racism as a result of thinking about disparities in care and impact during the Covid crisis. I was volunteering for the Covid Tracking Project, an effort to scrape or otherwise acquire Covid data from state websites and consolidate it in a single useful portal. The data showed that historically marginalized groups were dying from Covid at much higher rates. I was horrified. I was also recovering from surgery, the lead-up to which had meant an extended engagement with the consumer side of medical technology. I decided to investigate the ways that AI was being used in radiology. For one of the chapters in *Glitch*, I ran my own mammograms through an open-source cancer detection AI in order to write about the state of the art in AI-based cancer detection. The AI worked—but not as I expected. The difference between what I imagined and what was real is the same realization that we all face when confronted by technology that does not live up to our expectations.

Even when technologies are designed with good intentions, fallible humans develop programs that can result in devastating consequences. The solution isn't to make omnipresent tech more inclusive, but to root out the algorithms that target certain demographics as "other" to begin with. This book takes a more critical approach than *Artificial Unintelligence* and is often situated in the emerging field of critical data studies, critical AI studies, or critical

algorithm studies. I'm pleased to say that it cemented my reputation for explaining complex technical concepts in plain language, blending disciplines from computer science to sociology.

Extending the analysis to a new technology

Since publishing *Artificial Unintelligence*, the field of artificial intelligence has become both mainstream and hugely popular. Generative AI, in its infancy in 2018, became the thing everyone talked about. As I wrote this commentary, I became curious about how a generative AI system might summarize one of my books. Summarization is one of the muchtouted capabilities of generative AI. At the time of this writing, Google had just released a new generative AI product, NotebookLM, which it had developed in consultation with writer Stephen Johnson. NotebookLM was being touted as a useful tool specifically for writers. Having seen the launch (and frequent failures) of multiple digital products designed for writers, I was curious. I decided to load the text of my book into NotebookLM and compare its summaries to the ones I had written for this part of the dissertation. I selected *Artificial Unintelligence* because I knew it had been included (without my consent or permission) in one of the datasets commonly used to train AI systems. I theorized that having the text in the training data, and then uploading the text to re-train a smaller model on the book text, was likely to result in better summaries because NotebookLM would have been exposed to the text multiple times.

I took a PDF copy of the book and fed it into Google's NotebookLM product on October 8, 2024. I gave the instruction "summarize chapter X" 13 times, once for each chapter. It was interesting that the NotebookLM product offered a disclaimer below the text entry area that read, "NotebookLM may still sometimes give inaccurate responses, so you may want to confirm any facts independently." These are the summaries that resulted, along with some observations on their effectiveness.

Introduction Summary: Challenging Techno-Chauvinism

The introduction of "Artificial Unintelligence" sets forth the book's main argument: that technology, while often touted as superior, is not inherently better than other solutions and frequently falls short of its promises.

To illustrate this point, the author uses the example of writing "Hello, world." While a computer program can execute this task, it argues that there is no inherent value in the technological approach versus simply writing it down.

The introduction emphasizes that **understanding the social and technical contexts behind technology is crucial to recognizing its limitations.** It highlights the importance of data

journalism in holding decision-makers accountable by scrutinizing technology's flaws and challenging overly optimistic claims. The author aims to equip readers with a critical perspective to question assumptions about technology and work toward a future where it enhances, rather than threatens, our lives.

Initially, this looks like an adequate summary. In the introduction, I write: "This book is a guide for understanding the outer limits of what technology can do. It's about understanding the bleeding edge, where human achievement intersects with human nature. That edge is more like a cliff; beyond it lies danger... This book is about how to understand the outer limits of what technology can do. Understanding these limits will help us make better choices and have collective conversations as a society about what we can do with tech and what we ought to do to make the world truly better for everyone." The machine's summary that "understanding the social and technical contexts behind technology is crucial to recognizing its limitation" is reasonable.

The final sentence of the machine summary veers off course. Does technology "threaten" our lives? I don't think so. I doubt I have ever used this phrasing—it is unnecessarily dramatic and is precisely the kind of phrase I would edit out of my own writing or one of my student's in favor of more specificity. I usually go out of my way to help people understand that AI is not an existential threat, merely another technology. In chapter 4, in explaining the AI behind song recommendation engines, I write explicitly: "This process is procedural and unthreatening and shouldn't make anyone think that the machines are going to rise up and take over the world."

The machine summary neglects to mention that the introduction lays out the rhetorical path of the book. I write: "This isn't a technical manual or a textbook; it's a collection of stories with a purpose. I chose a handful of adventures in computer programming, each of which I undertook in order to understand something fundamental about technology and contemporary tech culture. All of those projects link together in a sort of chain, building an argument against technochauvinism. Along the way, I'll explain how some computer technology works and unpack the human systems that technology serves."

Technochauvinism, a term I coined, is a key concept in this book and in *Glitch*. I define it as the assumption that technological solutions are superior to other solutions. Instead, I argue that we should use the correct tool for the task. Sometimes this is a computer, sometimes it is something simpler like a book in the hands of a child sitting on a parent's lap. One is not inherently better than the other. This key theoretical concept should be mentioned in the first part of any summary of the book. The machine summary also leaves out an important road map. The book is organized into three parts of 3-5 chapters each. Part 1, How Computers Work, "covers a few basics about how computers work and how computer programs are constructed." Part 2, When Computers Don't Work, includes some AI history (chapter 6, "People Problems") and has four chapters devoted to four different computational adventures in which I build some technology in order to make a rhetorical point. Part 3, Working Together, is about collaboration; it includes two more computational adventures, and a final forward-looking chapter. It is conventional in academic writing to explain the path of the book in the first chapter. The machine is not aware of conventions (it is not "aware" at all) and thus omits this important conceptual element. The next summary is similarly close to adequate:

Chapter 1: My First "Hello, World" Program

In chapter 1, the author introduces the very foundation of computer programming – how to write a simple "Hello, world" program. The chapter walks the reader through three different ways to achieve this:

• **Typing it out:** Simply typing "Hello, world" into a word processing program and printing it.

• **Using Python:** This involves a slightly more technical approach using the Python programming language within a terminal window. The author provides step-by-step instructions on how to type and execute the command print("Hello, world!") to make the computer display the phrase.

• Using a loop in Python: The chapter then introduces the concept of a loop to automate the process of printing "Hello, world" multiple times. This demonstrates how even simple programs can be made more efficient and powerful using programming concepts. The author emphasizes that while these methods achieve the same result, none are inherently "better". The choice depends on the user's experience and the desired outcome. This example highlights the book's central theme: technology isn't inherently superior, but rather a tool that must be understood and used appropriately. The chapter then moves to introduce the concept of "data" and its importance in programming and, more broadly, in understanding the world. The author stresses a key principle: data is socially constructed, meaning it's generated and interpreted by people, even when it appears to be purely numerical. The example of the varying pocket money collected by kindergarteners illustrates this point, demonstrating how data can be influenced by social factors rather than objective truth.

This approaches a reasonable summary of the first chapter. I do indeed write about kindergarteners collecting something. However, I do not write about kindergarteners collecting *pocket money* as the machine summary suggests. My anecdote is about kindergarteners who do a weekly tally on the total number of pockets in their classroom. I find this classroom activity delightful because kids this age are fascinated by pockets. They keep all kinds of strange things in their pockets, and they are always excited to discuss whatever they are carrying. It's very cute to imagine. However, because the kids are just learning to count, the data they collect is always a bit off. It doesn't matter really; the point of the exercise is for the kids to practice their evolving numeracy. I make the point that at some

level, every data collection exercise is like the kindergarten pocket data collectors; there are always going to be human and mathematical mistakes in any dataset.

The transformation from "kindergarten pocket data collectors" (my text) to "pocket money collected by kindergarteners" (the machine's summary of my text) provides some insight into how generative AI works. This version of NotebookLM works by feeding the user's source input(s) into Gemini 1.5, a proprietary multimodal model ("8 Expert Tips for Getting Started with NotebookLM" 2024). A multimodal model can take in multiple kinds of inputs (such as text, images, audio, or video) and can create different types of outputs. NotebookLM can generate text or audio at the time of this writing. Gemini is a large scale foundational model, and the document(s) that the user uploads to it provide what is called the "context window" ("Long Context | Gemini API | Google AI for Developers," n.d.). The context window process essentially instructs Gemini to base its output on the material inputted by the user. The foundational model has gone through many mathematical processes, including one in which it is trained on the frequency and location in vector-space of different word tokens. The model's text output is basically next-token prediction, in which the model chooses the most statistically likely next word (token) in a sequence. It's impressive that this process works. However, such models produce inaccuracies and untruths (Palta, Angwin, and Nelson 2024). They produce mediocrity, not excellence.

The phrase "kindergarten pocket data collectors" is unique. Its omission is a useful illustration of how LLMs select for words that are popular, not good. I performed a frequency search on Google Trends on November 7, 2024, and found that I was the only person who had searched this phrase in the past day. "Pocket data," a similar phrase, had very few results for the time period 2004-2024. Most were related to a Malaysian company called PocketData (notice the missing space). On the other hand, the phrase "pocket money" is quite common according to Google Trends. It even has a Wikipedia page. I chose my phrasing because the juxtaposition of these terms struck me as absurd and delightful. The model chose its phrase "pocket money collected" because it is predicting the statistically most likely word to come after "pocket" based on words that appear in its training data. In this case, we can tell that the phrase "pocket money" is far more common than "pocket data" because of the Google Trends search results. This explains why the model selected "money" instead of "data" to follow "pocket."

Following are the summaries of subsequent chapters. These summaries are very similar to each other in structure and tone. Each has a bolded title and most feature bullet points. Many of them end with a two-sentence paragraph that begins "Overall, Chapter X..."

This indicates that there is a kind of template programmed into NotebookLM, likely in the reinforcement learning part of the development process. This strategy of generating text based on templates is one of the earliest popular ways of using AI in journalism (Fassler 2012; Haim and Graefe 2017; Carlson 2014; Montal and Reich 2017; Clerwall 2014). Most of the summaries included footnotes and hyperlinks that seemed like they referred to parts of the text. However, when I clicked on each hyperlinked footnote, it took me to a part of the text that was unrelated to the summary section. The footnotes were hallucinations. The length of each chapter summary is sub-optimal, but all are presented here in full in order to demonstrate the repetitive nature of the AI output.

Chapter 2: The Anatomy of a Computer and the Social Life of Data

Chapter 2 expands on the concepts introduced in chapter 1, providing a deeper understanding of how computers work and further emphasizing the social construction of data.

Building a "Turkey Club" Computer:

The chapter utilizes the analogy of a turkey club sandwich to explain the different layers of a computer system:

• **Hardware:** This is the foundation of the sandwich, akin to the physical components of a computer like the motherboard, processor, and memory.

• **Machine Language:** Just as the bread holds the sandwich together, machine language acts as the interface between hardware and software. It uses binary code (0s and 1s) to communicate instructions.

• **Assembly Language:** This layer is like the condiments, making the sandwich more palatable. Assembly language uses symbols to represent machine code, making it easier for humans to write programs.

• **Operating System:** This is a crucial software layer, similar to the main ingredients of the sandwich. The operating system (like Windows or macOS) manages hardware resources and provides a platform for other programs to run.

• **Programs:** These are the toppings of the sandwich, representing the applications we use, like web browsers, word processors, or games.

• **Data and Content:** Just as the sandwich's flavor comes from its ingredients, data and content are the heart of what we create and consume on computers. This layer represents everything from text documents to images to videos.

• **Content Management Systems (CMS):** For managing large amounts of data, like in media organizations, a CMS acts like a plate for the sandwich, providing organization and structure.

• **Client-Server Model:** This final layer, representing the act of eating the sandwich, involves the exchange of data between a client (e.g., web browser) and a server (another computer hosting websites or services).

The Human Element in Data:

The chapter reiterates the idea that data, despite its seemingly objective nature, is inherently shaped by human decisions and actions.

• **Data as a Social Construct:** The sources emphasize that all data originates from people counting and recording information. Even data generated by computers ultimately stems from human-created algorithms and datasets.

• **Case Study: Giffords' Recovery:** The chapter uses the example of Congresswoman Gabby Giffords' recovery from a traumatic brain injury to illustrate how understanding the human brain is crucial even when using technology like fMRI scans. The scans themselves are just data; interpreting them requires human expertise and understanding of the brain's complexity.

By drawing parallels between a familiar object like a sandwich and the complex workings of computers, chapter 2 aims to demystify technology and highlight its limitations. The chapter emphasizes that while computers excel at processing information, they lack the nuanced understanding and critical thinking abilities inherent to humans.

The chapter concludes by urging readers to approach technology with a critical eye, acknowledging its potential while remaining mindful of its limitations and the importance of human judgment.

Chapter 3: Unpacking Artificial Intelligence (AI): Dreams vs. Reality

Chapter 3 tackles the concept of artificial intelligence (AI), a term often shrouded in misconception and hype. The author distinguishes between the **fantastical portrayal of AI in popular culture and the more grounded reality of its current capabilities.** Debunking the "Real AI" Myth:

The chapter begins with an anecdote about the author presenting an AI-powered journalism tool and encountering the question, "Is it a *real* AI?" This interaction highlights a common misunderstanding: people often conflate AI with human-like intelligence, as seen in science fiction.

The chapter clarifies that **true artificial general intelligence (general AI), capable of replicating human cognitive abilities, remains a distant aspiration.** What we currently have is **narrow AI**, which excels at specific tasks but lacks the broader understanding and adaptability of human intelligence.

How Narrow AI Works: Building Models from Data:

To explain narrow AI, the chapter focuses on **machine learning**, a popular AI approach. It describes how machine learning involves:

• **Training a model:** Feeding a computer algorithm with data to identify patterns and relationships.

• **Creating a prediction model:** Using the trained model to analyze new data and make predictions or classifications.

The chapter uses the example of Google's AlphaGo program, which defeated a professional Go player, to further illustrate narrow Al. While an impressive feat, the author emphasizes that AlphaGo's success is limited to the specific rules and patterns of the game. It doesn't signify the dawn of sentient machines, but rather the power of specialized algorithms within defined domains.

Data Journalism: A Real-World Application of AI:

The chapter then connects these AI concepts to the field of data journalism. It provides examples of how journalists use data analysis and visualization techniques, powered by AI tools, to uncover stories and hold institutions accountable. These examples demonstrate how **narrow AI can be a powerful tool for investigative journalism and promoting transparency**.

The Importance of Context in AI:

The chapter concludes by reiterating the importance of **understanding the social and technical context** surrounding AI. It argues that while AI systems can process vast amounts of data, they lack the nuanced judgment and understanding of human experts. In essence, chapter 3 aims to provide a more realistic and grounded understanding of AI, moving beyond the hype to demonstrate its potential and limitations. It emphasizes that AI, **particularly in its current narrow form, is a tool that must be used responsibly and ethically, always with a critical awareness of its inherent biases and the importance of human oversight.**

Chapter 4: Data Journalism - More Than Just Numbers

Chapter 4 introduces the concept of data journalism, exploring its various forms and highlighting its significance in investigative reporting. While previous chapters focused on the technical aspects of computers and AI, this chapter shifts the focus to how these technologies are applied in a real-world profession like journalism.

Different Flavors of Data Journalism:

The chapter emphasizes that data journalism encompasses a wide range of approaches, going beyond simply using numbers in news stories:

• **Data Visualization:** This involves presenting data in visually compelling ways, such as charts, graphs, and interactive maps. This approach helps make complex data more accessible and understandable to the public. The source gives the example of Amanda Cox, editor of "The Upshot" section in *The New York Times*, who won an award for her visual representation of the consumer price index.

• **Computational Journalism:** This branch of data journalism utilizes computer programming and algorithms to gather, analyze, and visualize data. The sources highlight this approach in the context of investigating large datasets, citing the example of *The Atlanta Journal-Constitution*'s investigation into sexual abuse by doctors.

• Algorithmic Accountability Reporting: This crucial aspect of data journalism involves scrutinizing the algorithms used by governments and corporations,

particularly those that impact people's lives in areas like criminal justice and finance. The source highlights *ProPublica*'s "Machine Bias" investigation, which uncovered racial bias in an algorithm used for sentencing.

The Evolution of Data-Driven Reporting:

The chapter traces the historical development of data journalism:

• **Precision Reporting:** This early form of data journalism, pioneered by Philip Meyer, emphasized using social science methods and statistical analysis in reporting. The sources cite Meyer's Pulitzer Prize-winning work on the 1967 Detroit riot as a key example.

• **Computer-Assisted Reporting (CAR):** With the advent of personal computers, journalists began using spreadsheets and databases to organize and analyze data, leading to the rise of CAR. The sources cite the movie *Spotlight*, which portrays *The Boston Globe*'s investigation into child sex abuse in the Catholic Church, as a prominent example of CAR.

• **Data Journalism in the Digital Age:** The emergence of the internet and digital tools led to the evolution of CAR into what we now know as data journalism. This encompasses a broader range of techniques, including those mentioned earlier, along with bot-making and online data analysis.

Key Figures and Institutions in Data Journalism: The chapter acknowledges the contributions of key figures and organizations in advancing data journalism:

• **Nicholas Diakopoulos:** A professor at Northwestern University, Diakopoulos is recognized for his work on algorithmic accountability reporting and his contributions to the field's academic development.

• **Nikki Usher:** An ethnographer and author of *Interactive Journalism: Hackers, Data, and Code*, Usher has provided valuable insights into the practices and culture of data journalism through her fieldwork and interviews with practitioners.

• **Cindy Royal:** Known for her work on journalists' use of code, Royal has contributed to understanding how these skills are integrated into newsrooms and journalism education.

• **James T. Hamilton:** Author of *Democracy's Detectives*, Hamilton has highlighted the crucial role of data-driven investigative journalism in serving the public interest.

• **Nate Silver:** Founder of FiveThirtyEight.com and author of *The Signal and the Noise*, Silver is a prominent figure in data journalism, particularly for his work on election forecasting and statistical analysis.

The Enduring Value of Data Journalism:

The chapter concludes by emphasizing the enduring importance of data journalism in an era of increasing data availability and algorithmic decision-making.

By providing a comprehensive overview of data journalism, chapter 4 highlights the crucial role that journalists play in analyzing complex datasets, uncovering hidden patterns, and holding powerful institutions accountable.

The chapter connects with the broader themes of the book by demonstrating how technology, when applied with a critical and ethical lens, can be a powerful tool for promoting transparency and social good. Just as previous chapters cautioned against blindly trusting technology, chapter 4 reminds readers that data itself requires careful interpretation and contextualization, with journalists playing a vital role in bridging the gap between raw data and meaningful narratives.

Chapter 5: Unmasking the Technochauvinism in Education Reform

Chapter 5, "Why Poor Schools Can't Win at Standardized Tests," examines the flawed logic of using standardized testing as the primary measure of educational success, particularly in underfunded schools. Broussard argues that this approach exemplifies technochauvinism, as it assumes that technology (in this case, standardized tests) offers a simple and objective solution to complex social problems.

• **The Problem with Standardized Tests:** The chapter challenges the notion that standardized tests are neutral measures of academic achievement. Broussard argues that these tests are often biased towards specific knowledge found in expensive textbooks, putting low-income schools at a disadvantage.

• **Philadelphia's Textbook Shortage:** To illustrate her point, Broussard investigates the state of textbooks in Philadelphia public schools, uncovering a severe shortage of up-to-date materials. Her findings reveal a direct correlation between students' access to textbooks and their performance on standardized tests. For instance, she cites Tilden Middle School, where eighth-grade students scored far

below the district average in reading, likely due to having access to less than half the required reading textbooks.

• **The Human Element in Data Systems:** While acknowledging the potential benefits of data-driven approaches in education, Broussard highlights how even the most sophisticated systems are ultimately reliant on humans for data collection and implementation. She points to Philadelphia's electronic textbook inventory system, which proved ineffective due to understaffing and a lack of training among school personnel.

• **Exposing the "Inventory is Destiny" Myth:** Broussard criticizes the technochauvinist assumption that simply providing computers and software will automatically lead to improved educational outcomes. She argues that this approach ignores the crucial role of human factors like teacher training, curriculum development, and adequate infrastructure.

The chapter concludes by emphasizing that true educational equity requires a multifaceted approach that addresses the underlying social and economic factors contributing to the achievement gap. Broussard cautions against simplistic, technocentric solutions that fail to account for the complexities of the education system and the needs of diverse student populations.

Tracing the Roots of Technochauvinism: From Minsky's AI to Cyberculture

Chapter 6, "People Problems," explores how a small group of influential individuals shaped the trajectory of technology and its impact on society, often by promoting an overly optimistic view of its capabilities.

The Enduring Influence of Marvin Minsky: The chapter focuses on Marvin Minsky, a pioneering figure in artificial intelligence (AI) who co-founded the MIT AI Lab. Broussard highlights Minsky's significant contributions to the field, including his work on neural networks and his influence on early AI research. For example, Minsky consulted with Stanley Kubrick on the creation of HAL 9000, the iconic AI computer in the film "2001: A Space Odyssey".

The Intersection of Counterculture and Technology: Broussard argues that Minsky's vision of AI was deeply intertwined with the counterculture movement of the 1960s. She connects figures like Stewart Brand, publisher of the "Whole Earth Catalog," to Minsky's circle, showing how they shared a belief in technology's potential to empower individuals and create a more utopian society. The chapter cites Brand's famous line, "We are as gods and might as well get good at it," as evidence of this techno-utopian mindset.

From Abacus to AI: A Historical Perspective: The chapter provides a brief historical overview of how humans have used tools and machines to augment their capabilities, from the abacus to the astrolabe to early computers. This historical context emphasizes that the desire to create machines that can perform complex tasks has been a long-standing human aspiration. **The Perils of Unchecked Optimism:** While acknowledging the positive aspects of

technological progress, Broussard cautions against blindly embracing technology without considering its potential downsides. She argues that the counterculture's emphasis on individual freedom and its skepticism towards authority may have inadvertently contributed to a "move fast and break things" mentality in the tech industry, where ethical considerations and social responsibility are often sidelined in the pursuit of innovation and profit.

The chapter concludes by suggesting that understanding the historical and cultural forces that shaped our current technological landscape is crucial for developing a more critical and nuanced perspective on AI and its role in society. It sets the stage for the subsequent chapters, which examine specific instances where technochauvinism has led to unrealistic expectations and potentially harmful outcomes.

Chapter 7: Demystifying Machine Learning and Its Limitations

Chapter 7, "Machine Learning: The DL on ML," aims to clarify what machine learning is, how it works, and its limitations. Broussard argues that the term "machine learning" is often misunderstood, leading to inflated expectations and a failure to recognize the crucial role of humans in shaping these systems.

• **Defining Machine Learning**: The chapter acknowledges the difficulty in defining machine learning, as it encompasses a range of techniques and approaches. Broussard presents several definitions from different sources, highlighting their strengths and weaknesses. She ultimately favors Tom Mitchell's definition:

"We say that a machine learns with respect to a particular task T, performance metric P, and type of experience E, if the system reliably improves its performance P at task T, following experience E."

• **Machine Learning in Practice: Predicting Titanic Survivors**: To illustrate the concepts, Broussard walks the reader through a practical example: predicting which

passengers survived the Titanic disaster using a machine learning model. She provides a step-by-step explanation of the process:

Data Preparation: Importing necessary libraries. loading the dataset. understanding the data dictionary, and performing basic exploratory data analysis (EDA).

Feature Selection: Identifying potentially relevant features 0 (passenger class, sex, age, fare) for predicting survival.

Model Training: Splitting the data into training and test sets, 0 handling missing values, converting categorical variables, and training a decision tree classifier.

Model Evaluation: Assessing the model's performance by applying it \sim to the test data and comparing predictions to actual outcomes.

Interpreting Results: Analyzing the model's feature importance scores to understand which factors were most influential in predicting survival.

The Importance of Context: Broussard emphasizes that while machine learning models can identify patterns and make predictions based on available data, they often fail to account for the broader social, historical, and human factors that shape real-world events.

The Human Element in the Titanic Disaster: For example, the chapter discusses how the actions of individual officers on the Titanic, such as their interpretation of evacuation orders, significantly impacted survival rates. These nuances are not captured in the dataset and highlight the limitations of relving solely on data-driven models.

Narrow AI vs. General AI: Broussard reiterates the distinction between narrow AI (task-specific systems like the Titanic survivor predictor) and general AI (hypothetical machines with human-like intelligence). She emphasizes that current machine learning techniques fall under the category of narrow AI, cautioning against attributing human-like understanding or sentience to these systems.

The Allure and Limitations of Data: While acknowledging the potential benefits of machine learning, Broussard stresses the importance of critical thinking and a healthy skepticism towards data-driven approaches. She warns against blindly accepting the results of machine learning models without considering their limitations and potential biases.

Overall, Chapter 7 provides a grounded introduction to machine learning, demystifying the technology by demonstrating its inner workings and emphasizing its reliance on human choices and interpretations. This chapter lays the foundation for subsequent chapters by highlighting the need to approach AI with a balanced perspective, recognizing both its capabilities and its limitations.

The Illusion of Self-Driving Cars: Exposing the Limits of Computing Chapter 8, "This Car Won't Drive Itself," examines the challenges and limitations of autonomous vehicle technology, challenging the hype surrounding self-driving cars and emphasizing the crucial differences between human and machine intelligence. Broussard argues that while driver-assistance features have improved, truly autonomous vehicles are still a distant reality, and the push for their rapid development poses significant risks.

The 2007 Urban Challenge: The chapter recounts Broussard's experience at the 2007 DARPA Urban Challenge, where she witnessed firsthand the struggles of autonomous vehicles to navigate complex environments.

Little Ben's Near-Disaster: She describes her ride in "Little Ben." a self-driving car developed by University of Pennsylvania students. The experience involved several near-misses, highlighting the car's difficulty in accurately perceiving and responding to real-world situations like unexpected obstacles and sudden movements. For example, the car once veered towards a snowplow and had to be manually overridden.

The Limitations of Pre-Programming: Broussard contrasts the challenges faced by the Urban Challenge cars with the relative simplicity of earlier autonomous vehicle competitions, such as the 2005 DARPA Grand Challenge, where cars primarily navigated open desert terrain. She explains that in the controlled environment of a desert, obstacles could be easily avoided with basic pre-programmed rules. However, urban environments require a much higher level of perception, decision-making, and adaptability that current AI systems struggle to achieve.

• **The "Karel the Robot" Analogy**: To illustrate the fundamental limits of computing, Broussard uses the analogy of "Karel the Robot," a programming exercise used to teach computational thinking.

• Karel operates in a simplified grid world with known obstacles and pre-defined rules. This predictability allows programmers to anticipate and code for all possible scenarios. In contrast, real-world driving involves countless unpredictable variables, making it impossible to pre-program every situation a car might encounter.

• **The SAE Levels of Automation**: The chapter discusses the SAE Levels of Automation, a standardized system for classifying the capabilities of autonomous vehicles.

• Broussard emphasizes that while lower levels of automation (e.g., lane-keeping assistance, adaptive cruise control) can be beneficial, higher levels (e.g., Level 3 and above, where the car assumes full control in certain situations) introduce significant risks due to the potential for system failures and the difficulty of seamlessly transitioning control between human and machine.

• **The "Tesla Autopilot" Fallacy**: Broussard challenges the misconception that Tesla's "Autopilot" feature represents true self-driving capability.

• She argues that the term "Autopilot" itself is misleading, creating a false sense of security and potentially encouraging drivers to become overly reliant on the technology, even in situations where it may not be appropriate or safe.

• **Questioning the Safety Imperative**: Broussard critically examines the claim that self-driving cars will significantly reduce traffic fatalities, highlighting the lack of robust data to support this assertion.

• She points out that while some statistics suggest a decrease in accidents involving autonomous vehicles, these numbers are often based on limited testing environments and may not accurately reflect real-world driving conditions.

• **The Importance of Human Judgment**: The chapter concludes by emphasizing the irreplaceable role of human judgment and intuition in navigating complex and unpredictable environments.

• Broussard argues that while autonomous vehicle technology has made significant strides, it still falls short of replicating the nuanced perception, adaptability, and ethical reasoning capabilities of human drivers. She cautions against overestimating the capabilities of AI and stresses the need for continued human oversight and intervention in the development and deployment of self-driving cars.

Unveiling the "Popular Doesn't Mean Good" Fallacy

Chapter 9, titled "Popular Doesn't Mean Good," critiques the tendency to equate popularity with quality, particularly in the realm of computational decision-making. Broussard argues that this conflation can lead to flawed outcomes and perpetuate existing biases, especially when dealing with subjective judgments and complex social issues.

• **Challenging the Selfie Experiment**: Broussard opens the chapter by dissecting a 2015 media experiment that claimed to identify the characteristics of a "good" selfie using data science. She criticizes the experiment for perpetuating harmful stereotypes by associating specific facial expressions and poses with attractiveness, ultimately reinforcing societal biases instead of challenging them. The chapter notes that a human would be able to understand that this type of analysis is problematic.

• **The Limits of Ranking and the Rise of Internet Fraud**: Broussard examines the societal obsession with measurement and ranking, fueled in part by the internet's ability to quantify and order everything. She argues that while ranking can be useful in some contexts, it often fails to capture the nuances of quality and can be easily manipulated.

• **The "Hot or Not" Phenomenon**: Broussard cites the example of early social media platforms like "Hot or Not," where users rated each other's attractiveness based on photos. She argues that such platforms reduced complex human qualities to superficial metrics, contributing to a culture of objectification and superficiality. Broussard connects this ranking phenomenon to the problem of internet fraud, as illustrated by the example of a Facebook scam involving the fraudulent inflation of "likes" for a

commemorative coin. She notes that this type of fraudulent activity is pervasive online, highlighting how the pursuit of popularity can be exploited for malicious purposes.

• **The COMPAS Case: Algorithmic Bias in Criminal Justice**: The chapter examines the controversial COMPAS algorithm, used in the US criminal justice system to predict recidivism risk. Drawing on ProPublica's investigative reporting, Broussard discusses how the algorithm, despite its claims of objectivity, exhibited racial bias, assigning higher risk scores to Black defendants compared to white defendants with similar criminal histories. She highlights the inherent difficulty of creating algorithms that are both predictive and fair, especially when dealing with complex social factors and historical inequalities.

• **The Importance of Context and Critical Thinking**: Broussard concludes by emphasizing the importance of considering the broader social and ethical implications of computational decision-making. She argues that instead of blindly accepting popularity as a proxy for quality, we need to ask critical questions about *who* benefits from these systems and *how* they might perpetuate or exacerbate existing inequalities.

• **The Silicon Valley Drug Culture**: As an example, she cites the normalization of illicit drug use in Silicon Valley, driven in part by the pressure to work long hours and meet demanding deadlines. She argues that this culture, while seemingly popular within a specific industry, can have negative consequences for individuals and society as a whole. Broussard suggests that instead of equating popularity with "good," it's crucial to interrogate the underlying values and assumptions driving these trends and consider their wider impact.

Overall, Chapter 9 serves as a call for greater critical thinking and a more nuanced understanding of the relationship between popularity and quality in the digital age. Broussard encourages readers to question the metrics used to define success, challenge the algorithms that shape our lives, and consider the broader ethical implications of our technological choices.

Riding the Startup Bus: A Critical Look at Tech Innovation Culture

Chapter 10, "On the Startup Bus," chronicles Broussard's experience participating in Startup Bus, a cross-country hackathon where teams compete to develop a viable business idea and pitch it to investors. Through her firsthand account, Broussard provides a critical examination of startup culture, questioning the romanticized narrative of "disruptive innovation" while highlighting the importance of collaboration and communication in software development.

• **Startup Bus: A Microcosm of Tech Culture**: Broussard uses her experience on Startup Bus to illustrate both the excitement and the challenges of the tech industry. She describes the intense pressure to develop a working product and a compelling pitch within a short timeframe, mirroring the demanding pace and high-stakes environment of Silicon Valley.

• **Pizzafy: A Case Study in Rapid Prototyping**: Broussard's team on Startup Bus developed "Pizzafy," an app designed to help users calculate the amount of pizza they need for a group. While the concept itself may seem trivial, the process of building the app from scratch, designing a user interface, and testing it with real users provides insight into the iterative nature of software development.

• **Collaboration and Communication: Essential Ingredients for Success**: Broussard emphasizes the importance of effective communication and collaboration in the success of any software development project, particularly within the fast-paced and often chaotic environment of a hackathon.

• **The Value of Face-to-Face Interaction**: She contrasts the efficiency of face-to-face communication with the limitations of digital communication tools, noting that in-person interactions allow for richer information exchange, faster problem-solving, and stronger team dynamics. This observation aligns with her earlier critique of technochauvinism, as she highlights the enduring value of human connection and collaboration in a technology-driven world.

• **Questioning the "Disruptive Innovation" Myth**: Broussard uses her experience on Startup Bus to challenge the romanticized notion of "disruptive innovation," often associated with the tech industry.

• **The Allure of the Overnight Success Story**: She notes how the media often glorifies overnight success stories, creating a misleading perception of how innovation actually happens. While acknowledging that

some startups do achieve rapid growth and success, she emphasizes that these cases are often exceptions rather than the rule.

• **The Importance of a Critical Lens**: The sources do not describe the outcome of the Startup Bus competition or the long-term fate of any projects, including "Pizzafy." This lack of closure reinforces Broussard's point that the true value of such experiences lies not necessarily in creating a commercially successful product, but in the learning process, the collaboration, and the critical examination of technology's role in society.

Overall, Chapter 10 encourages readers to approach the world of tech startups with a discerning eye, recognizing both the potential benefits and the potential pitfalls of this fast-paced and often hype-driven industry. Broussard's account underscores the importance of human connection, ethical considerations, and a healthy skepticism toward the "move fast and break things" mentality often associated with Silicon Valley.

Building "Bailiwick": A Case Study in Third-Wave AI for Investigative Reporting

Chapter 11, "Third-Wave AI," focuses on Broussard's development of "Bailiwick," an AI-powered tool designed to assist investigative journalists in uncovering campaign finance irregularities. Through this project, Broussard advocates for a more nuanced and ethical approach to AI development, one that prioritizes human expertise, contextual understanding, and social impact.

• **Beyond Automation: Augmenting Human Intelligence**: Broussard emphasizes that the goal of "Bailiwick" is not to replace journalists but to enhance their investigative capabilities. She argues that while AI can automate certain tasks, such as data collection and analysis, it's crucial to retain human judgment and expertise in interpreting findings, identifying patterns, and pursuing leads.

• **The Limits of First and Second-Wave AI**: Broussard distinguishes between three waves of AI. First-wave AI focused on symbolic reasoning and rule-based systems, while second-wave AI emphasized machine learning and statistical pattern recognition. She argues that both waves have limitations, particularly when dealing with complex, context-dependent domains like journalism.

• **Toward a Third-Wave AI**: Broussard proposes a "third wave" of AI that combines the strengths of machine learning with human domain knowledge and ethical considerations. She argues that this approach is essential for developing AI tools that are not only technically sophisticated but also socially responsible and aligned with human values.

• **Designing "Bailiwick": From Data to Insights**: The chapter details the process of designing and building "Bailiwick," highlighting the importance of collaboration, domain expertise, and iterative development.

• **Understanding the Problem**: Broussard consulted with experienced campaign finance reporters to understand their workflow, the challenges they faced, and the types of anomalies they looked for in financial disclosures. This human-centered design approach ensured that "Bailiwick" addressed real-world needs and augmented existing investigative practices.

• **Data Collection and Preprocessing**: The chapter discusses the technical aspects of gathering campaign finance data from the Federal Election Commission (FEC) and structuring it in a way that could be analyzed by the Al system.

• **Identifying Patterns and Generating Story Ideas**: "Bailiwick" was designed to identify unusual patterns in campaign finance data, such as large donations from obscure sources, sudden spikes in spending, or discrepancies between reported income and expenditures. The system then presented these anomalies to reporters as potential story leads.

• "Bailiwick" in Action: A Case Study in Algorithmic Accountability: Broussard showcases how "Bailiwick" can be used to investigate campaign finance issues by examining spending patterns related to Donald Trump's 2016 presidential campaign.

• **Visualizing Spending Patterns**: The system's treemap visualization allowed reporters to quickly identify the largest donors and expenditures, revealing, for example, that a significant portion of Trump's campaign funds went towards "Make America Great Again" hats.

• **Uncovering Potential Irregularities**: While "Bailiwick" did not make accusations of wrongdoing, it highlighted potential areas of interest for further investigation, such as the high proportion of Trump's campaign spending attributed to "administrative expenses."

• The Importance of Transparency and Human Oversight: Broussard emphasizes the importance of transparency and human oversight in the development and deployment of AI systems like "Bailiwick."

• **Avoiding "Black Box" Algorithms**: She argues against relying on "black box" algorithms, where the decision-making process is opaque and difficult to understand. Instead, she advocates for systems that provide clear explanations for their findings, allowing journalists to critically evaluate the results and avoid drawing unwarranted conclusions.

• **The Need for Continuous Evaluation and Refinement**: Broussard acknowledges that AI systems are not perfect and require continuous evaluation and refinement. She stresses the importance of ongoing collaboration between journalists, computer scientists, and ethicists to ensure that AI tools are used responsibly and do not perpetuate existing biases or harm individuals.

In conclusion, Chapter 11 presents "Bailiwick" not just as a successful AI project but as a model for a more ethical and human-centered approach to AI development. By prioritizing collaboration, transparency, and social impact, Broussard demonstrates how AI can be a powerful tool for investigative journalism and a force for greater accountability in society.

"Aging Computers": Reflecting on the Past and Shaping the Future of AI

Chapter 12, "Aging Computers," concludes Broussard's exploration of artificial intelligence by reflecting on the field's historical trajectory and outlining a path toward a more ethical and human-centered future for AI. Drawing on her experiences investigating algorithmic bias, developing AI tools for journalism, and studying the history of computing, Broussard advocates for a critical and nuanced understanding of AI's capabilities and limitations.

• **Reconsidering the Narrative of Progress**: Broussard challenges the prevailing narrative of technological progress as inherently linear and inevitable. She argues that the development of AI has been marked by both advances and setbacks, successes, and failures.

• **The ENIAC: A Reminder of Al's Origins**: Broussard's visit to the ENIAC, one of the first digital computers, serves as a reminder that even the most sophisticated technologies have humble beginnings. She notes that the ENIAC, once considered a marvel of engineering, now appears outdated and limited in its capabilities.

• **The Importance of Historical Context**: The sources do not elaborate on specific examples of setbacks or failures in Al development. However, Broussard's broader argument suggests that understanding the historical context of Al is crucial for avoiding the pitfalls of techno-utopianism and hype. By recognizing that Al has not always lived up to its promises, we can approach current and future developments with a healthy dose of skepticism.

• **Centering Human Values in AI Development**: Broussard emphasizes the importance of embedding human values, such as fairness, accountability, and transparency, into the design and deployment of AI systems. She argues that AI should serve human needs and aspirations, rather than dictating or undermining them.

• Al Now Institute's Call for Algorithmic Accountability: Broussard highlights the Al Now Institute's recommendation that public institutions cease using "black box" Al systems that lack transparency and accountability. She echoes their call for greater oversight and public scrutiny of algorithms used in high-stakes domains such as criminal justice, healthcare, and education.

• **The Role of Human Oversight in Content Moderation**: The author cites content moderation on social media platforms as an example of the importance of human judgment in Al systems. While algorithms can flag potentially inappropriate content, human reviewers play a critical role in making nuanced decisions about what constitutes a violation of platform policies and what should be removed.

• **Toward a More Inclusive and Equitable AI**: Broussard argues that AI has the potential to perpetuate or even exacerbate existing social inequalities if developed and deployed without careful consideration of its impact on marginalized communities.

• Addressing Algorithmic Bias: Broussard emphasizes the importance of addressing algorithmic bias, drawing on her previous investigation of racial disparities in risk assessment tools used in the criminal justice system

(Chapter 9). She argues that AI systems can inherit and amplify biases present in the data they are trained on, leading to unfair or discriminatory outcomes.

• **The "Good Selfie" Experiment Revisited**: The sources do not offer specific solutions for mitigating algorithmic bias. However, Broussard's earlier discussion of the "good selfie" experiment (Chapter 9) suggests that one approach is to critically examine the datasets used to train AI models and to consider the potential for biases to influence the results.

• **Journalism as a Force for Algorithmic Accountability**: Broussard reiterates her belief that investigative journalism has a vital role to play in holding powerful institutions accountable for the ways in which they develop and use Al.

• **Data Journalism's Evolving Role**: She argues that data journalists are uniquely positioned to investigate the societal impacts of AI, uncover algorithmic bias, and shed light on the decision-making processes behind these systems.

• **Recognizing the Importance of Investigative Data Projects**: The author cites the emergence of awards specifically recognizing impactful investigative data projects as a positive sign, indicating growing awareness and support for this critical form of journalism.

• A Call for Collaboration and Critical Engagement: In closing, Broussard calls for greater collaboration and critical engagement between computer scientists, journalists, ethicists, and the public at large. She argues that shaping the future of Al is not the sole responsibility of technologists; rather, it requires a collective effort to ensure that these powerful tools are used responsibly and for the benefit of society as a whole.

Overall, Chapter 12 serves as a call to action, urging readers to approach AI not with blind optimism or unfounded fear, but with a critical and discerning eye. Broussard's work emphasizes that shaping the future of AI requires a commitment to transparency, accountability, and a deep understanding of the complex interplay between technology and society.

I was bored reading the NotebookLM summaries. The language was flat and the

repetitive structure failed to keep my interest. Because each chapter summary was about 500-

800 words (which is currently considered the sweet spot for the length of an online piece

according to metrics about readers' attention), the summaries took up ten manuscript pages.

Ten pages of summary is far too much. Here's the way I summarized chapters 5-10 in my

own introduction, clocking in at 1,153 words, just over two manuscript pages, and without

frequent bolding or bullet points:

In the first computational adventure (chapter 5), I investigate why, after two decades of education reform, schools still can't get students to pass standardized tests. It's not the students' or the teachers' fault. The problem is far bigger: the companies that create the most important state and local exams also publish textbooks that contain many of the answers, but low-income school districts can't afford to buy the books.

I discovered this thorny situation by building artificial intelligence software to enable my reporting. Robot reporters have been in the news in recent years because the Associated Press (AP) is using bots to write routine business and sports stories. My software wasn't inside a robot (it didn't need to be, although I'm not averse to the idea), nor did it write any stories for me (ditto). Instead, it was a brand-new application of old-school artificial intelligence that helped reveal some fascinating insights. One of the most surprising findings of this computational investigation was that, even in our high-tech world, the simplest solution—a book in the hands of a child—was quite effective. It made me wonder why we are spending so much money to put technology into classrooms when we already have a cheap, effective solution that works well.

The next chapter (chapter 6) is a whirlwind tour through the history of machines, specifically focused on Marvin Minsky—commonly known as the father of artificial

intelligence—and the enormous role that 1960s counterculture played in developing the beliefs about the Internet that exist in 2017, the time this book was written. My goal here is to show you how the dreams and goals of specific individuals have shaped scientific knowledge, culture, business rhetoric, and even the legal framework of today's technology through deliberate choices. The reason we don't have national territories on the Internet, for example, is that many of the people who made the Internet believed they could make a new world beyond government—much like they tried (and failed) to make new worlds in communes.

In thinking about tech, it's important to keep another cultural touchstone in mind: Hollywood. A great deal of what people dream about making in tech is shaped by the images they see in movies, TV programs, and books. (Remember my childhood robot?) When computer scientists refer to artificial intelligence, we make a distinction between general AI and narrow AI. General AI is the Hollywood version. This is the kind of AI that would power the robot butler, might theoretically become sentient and take over the government, could result in a real-life Arnold Schwarzenegger as the Terminator, and all of the other dread possibilities. Most computer scientists have a thorough grounding in science fiction literature and movies, and we're almost always happy to talk through the hypothetical possibilities of general AI.

Inside the computer science community, people gave up on general AI in the 1990s. General AI is now called Good Old-Fashioned Artificial Intelligence (GOFAI). Narrow AI is what we actually have. Narrow AI is purely mathematical. It's less exciting than GOFAI, but it works surprisingly well and we can do a variety of interesting things with it. However, the linguistic confusion is significant. Machine learning, a popular form of AI, is not GOFAI. Machine learning is narrow AI. The name is confusing. Even to me, the phrase machine learning still suggests there is a sentient being in the computer.

The important distinction is this: general AI is what we want, what we hope for, and what we imagine (minus the evil robot overlords of golden age science fiction). Narrow AI is what we have. It's the difference between dreams and reality.

Next, in chapter 7, I define machine learning and demonstrate how to "do" machine learning by predicting which passengers survived the Titanic crash. This definition is necessary for understanding the fourth project (chapter 8), in which I ride in a self-driving car and explain why a self-driving school bus is guaranteed to crash. The first time I rode in a self-driving car was in 2007, and the computerized "driver" almost killed me in a Boeing parking lot. The technology has come a long way since then, but it still fundamentally doesn't work as well as a human brain. The cyborg future is not coming anytime soon. I look at our fantasies about technology replacing humans and explore why it's so hard to admit when technology isn't as effective as we want it to be.

Chapter 9 is a springboard for exploring why *popular* is not the same as *good* and how this confusion—which is perpetuated by machine-learning techniques—is potentially dangerous. Chapters 10 and 11 are also programming adventures, in which I start a pizza-calculating company on a cross-country hackathon bus trip (it's popular but not good) and try to repair the US campaign finance system by building AI software for the 2016 presidential election (it's good but not popular). In both cases, I build software that works—but it doesn't work as expected. Its demise is instructive.

My goal in this book is to empower people around technology. I want people to understand how computers work so that they don't have to be intimidated by software. We've all been in that position at one time or another. We've all felt helpless and frustrated in the face of a simple task that should be easy, but somehow isn't because of the technological interface.

Even my students, who grew up being called digital natives, often find the digital world confusing, intimidating, and poorly designed.

When we rely exclusively on computation for answers to complex social issues, we are relying on artificial unintelligence. To be clear: it's the computer that's artificially unintelligent, not the person. The computer doesn't give a flying fig about what it does or what you do. It executes commands to the best of its abilities, then it waits for the next command. It has no sentience, and it has no soul.

People are always intelligent. However, smart and well-intentioned people act like technochauvinists when they are blind to the faults of computational decision-making or they are excessively attached to the idea of using computers to the point at which they want to use computers for everything—including things for which the computer is not suited.

I think we can do better. Once we understand how computers work, we can begin to demand better quality in technology. We can demand systems that truly make things cheaper, faster, and better instead of putting up with systems that promise improvement but in fact make things unnecessarily complicated. We can learn to make better decisions about the downstream effects of technology so that we don't cause unintentional harm inside complex social systems. And we can feel empowered to say "no" to technology when it's not necessary so that we can live better, more connected lives and enjoy the many ways tech can and does enhance our world.

Obviously, I prefer my own work because I'm the one who wrote it. However, the writing is shorter and snappier than that produced by the AI, with some unexpected juxtapositions of concepts. Readers tend to appreciate this as well.

After I finished creating individual summaries and saving them as "notes" in the application, I noticed a button in the NotebookLM interface labeled "Summarize." I clicked the button and got the following response:

Summary: This text explores the limitations of technology, arguing that it's often overhyped and doesn't always live up to its promises. The author emphasizes the importance of understanding how technology works, its potential biases, and the need for human oversight. **Explanation**: The passage provides an overview of several chapters from a book about artificial intelligence (AI) and technology. The author challenges the idea that technology is always the best solution by highlighting its limitations and potential biases. The author argues that AI, particularly "narrow AI" which is designed for specific tasks, is often mistaken for "general AI" that has human-like intelligence. They use examples like self-driving cars and educational technologies to demonstrate how technology can fall short of its promises and even perpetuate social inequalities.

The text emphasizes the importance of data journalism in holding powerful institutions accountable for how they develop and use AI. By investigating algorithmic bias and promoting transparency, journalists can help ensure that technology is used responsibly and ethically. **Key Terms**:

- **Technochauvinism**: The belief that technology is inherently superior to other solutions.
- Narrow AI: AI systems designed for specific tasks, lacking the broader understanding of general AI.
- General AI: Hypothetical AI with human-like cognitive abilities.
- Algorithmic Bias: When algorithms produce unfair or discriminatory outcomes due to biases in the data they were trained on.
- **Data Journalism**: Using data analysis and visualization techniques to uncover stories and hold institutions accountable.

This version got technochauvinism as a key term, which is good. The bullet points and the strategic bold text is again clearly a template. The summary is accurate, but it leaves out all of the things that make the book interesting and different and easy to understand. This AI-generated text is probably not compelling enough to persuade a reader to buy the book.

It is notable that the summary used the pronoun "they." I personally use she/her pronouns. However, usually when I create a machine summary of my work or use a machine translation program on my work, my personal pronouns are translated as or assumed to be male. I suspect this is because of gender bias baked into the training data. There are far more male tech writers than female or nonbinary tech writers, and the machine chooses pronouns based on popularity, not social appropriateness. I am curious whether the choice to use "they" is also a templated decision based on current social norms. In any case, I appreciate not being misgendered by a machine. It is progress.

I then selected all of the individual chapter summaries and selected the NotebookLM button marked "Critique" in order to get a machine critique of the machine-generated summaries. I received the following:

This is a well-structured and comprehensive summary of several chapters in a book about technology, likely Broussard's "Artificial Unintelligence." The use of headings and bullet points makes the text easy to follow, and the summaries accurately capture the main arguments and examples from each chapter.

Suggestions for improving grammar and style:

- Minimize bolded text. While bolding can be useful for emphasis, using it too frequently diminishes its impact. For example, in the Chapter 3 summary, "narrow Al" and "general AI" could be bolded once for initial emphasis, then left unbolded in subsequent mentions.
- Vary sentence structure. The text relies heavily on simple sentences. Introducing more complex sentences with varied punctuation (semicolons, em dashes) would create a more engaging rhythm.
- Reduce repetition of "The chapter." Try starting sentences differently: "Broussard argues..." "In this example..." "This highlights..." Suggestions for improving content:
- Provide more context/background. Who is Broussard? What are her credentials? Briefly mentioning the book's target audience and overall purpose would be helpful.
- Connect the chapters more explicitly. The summaries stand somewhat independently. Adding a sentence or two at the end of each summary to link it to the next chapter would improve flow and demonstrate the book's overarching narrative.
- **Consider adding specific quotes.** While paraphrasing is fine, including a few striking quotes from the book would add authenticity and make the summaries more compelling.
- Address the intended audience. The summaries are guite detailed; consider whether this level of detail is appropriate for the intended purpose. If this is for a book review or blurb, some summaries could be condensed.
- Expand on the critique of technochauvinism. While the summaries mention the term, they could more directly engage with its implications and why it's problematic. This would strengthen the book's central argument.

It was interesting to read the LLM's critique of its own summaries. The bolded text criticism was particularly interesting because the bold text was the machine's own choice. Similarly, the suggestion to "provide more context/background" was also ludicrous. As the author, I obviously don't need additional background on my own life or my own credentials. I also made sure to include appropriate personal details in the book so that the reader would have an idea of my background and perspective-a technique that is considered useful in

first-person narrative nonfiction writing. This points to a major weakness of generative AI. Generative AI has no "awareness" of its audience and must be specifically instructed to write for a particular level. The default reader programmed into generative AI is the software developer. Demographically speaking, this is a white American upper middle class coastal man with some college experience. A subtle bias against the humanities and toward STEM is obvious in the stilted prose suggestions. The machine's suggestion to state the "overall purpose" of the book is naïve. As early as second grade, teachers instruct students to state the purpose of a book in order to build literacy. It is a beginner step, not an advanced interpretation. I agreed with the machine critique that the machine text "relies heavily on simple sentences." Notably, the simple sentences are a design choice made by the NotebookLM or the Gemini programmers. I also agreed with the machine that the summaries are repetitive.

I was curious about what the product would say about the text of the entire book, not just the summaries. I had already uploaded a PDF copy of the entire manuscript. I gave NotebookLM the instruction "This is a book. Critique it." I received the following critique:

A Critical Look at "Artificial Unintelligence"

While the provided source excerpts don't encompass the entirety of "Artificial Unintelligence," they offer valuable insights into the book's core arguments and themes. Based on these excerpts, it's clear that Broussard critiques the pervasive belief in technology's inherent superiority, a concept she terms "technochauvinism." She argues that blindly embracing technology as the solution to all problems is not only misguided but potentially harmful. Broussard advocates for a more nuanced understanding of technology, recognizing its limitations alongside its capabilities. This involves challenging the prevailing narrative of technology as inherently progressive and acknowledging that its development is marked by both successes and failures.

One of the book's strengths lies in its use of personal anecdotes and real-world examples to illustrate its arguments. Broussard's firsthand experience developing AI software for journalism and investigating algorithmic bias provides a unique perspective, adding weight to her critiques.

The excerpts suggest that Broussard's writing style is accessible and engaging, making complex technical concepts understandable for a broader audience. Her use of humor and vivid imagery, such as comparing the implementation of change in a public school classroom to "turning a cruise ship that's running at high speed," [1] helps to demystify technology and make her arguments more relatable.

However, the excerpts also reveal some areas where the book could be strengthened:

- **Explicitly Addressing the Target Audience:** While the writing style suggests a general audience, clearly stating the intended readership within the book itself would be beneficial. This could involve tailoring the level of technical detail and providing additional context for readers unfamiliar with specific concepts or industries.
- Strengthening the Critique of Technochauvinism: While the excerpts introduce the term and provide examples of its manifestations, a more direct and comprehensive critique would enhance the book's central argument. This could involve exploring the historical and cultural factors contributing to technochauvinism, its ethical implications, and its impact on different sectors of society.
- Offering Concrete Solutions: While the book effectively highlights the problems with current approaches to technology, it could benefit from a more developed discussion of potential solutions. This could involve exploring alternative models for technological development, proposing policy recommendations, or showcasing

examples of organizations or individuals successfully challenging technochauvinistic practices.

Overall, the source excerpts suggest that "Artificial Unintelligence" offers a timely and important critique of our relationship with technology. By challenging assumptions, exposing biases, and advocating for human-centered design, Broussard encourages readers to engage with technology more critically and thoughtfully. To further enhance the book's impact, considerations for strengthening the critique of technochauvinism, explicitly addressing the target audience, and offering concrete solutions could prove valuable.

Obviously, the critique starts out incorrect. I fed in the *entire* text of the book, not an excerpt. This could be a hallucination, or it could be a design choice to allow the AI's creators to dodge some kind of copyright infringement allegation. I noticed that this critique and the one before it both mention target audience, suggesting that explicitly naming the target audience would be an improvement. I strongly disagree, both as a writer and as an instructor. Writers of literary nonfiction do not explicitly name their target audience.

It's clear that generative AI is nifty, but it has a long way to go in generating anything but mediocre prose.

Impact

My work has had impact on policy, scholarship, and popular culture. In *Artificial Unintelligence*, I first articulated the theoretical concept of technochauvinism, which has since been adopted by other scholars to explain a specific kind of pro-technology bias. As of this writing, Google Scholar shows 375 scholarly publications that use the term.

In May 2019, I was invited to speak at United Nations headquarters as part of an event called "Media for Democracy: Journalism and Elections in Times of Disinformation." In October 2019, I was invited to testify in front of the U.S. House of Representatives Task Force on Artificial Intelligence, a part of the Committee on Financial Services. The hearing was entitled "AI and the Evolution of Cloud Computing: Evaluating How Financial Data is Stored, Protected, and Maintained by Cloud Providers." I spoke about the realities of AI and cloud computing as a way of thinking through the human-scale issues with running bank operations in the cloud.

In October 2021, I was invited again to testify before the Task Force on Artificial Intelligence of the Committee of Financial Services. I presented my testimony before the US House of Representatives during the first session of the 117th Congress. The session was called "Beyond I, Robot: Ethics, Artificial Intelligence, and the Digital Age." I spoke about a practical vision for recognizing AI, and spoke about AI generally as well as discrimination, algorithmic auditing, and regulatory sandboxes.

Again in October 2021, I was invited to participate in a workshop sponsored by the National Institute of Standards and Technology (NIST) to kick off its draft Artificial Intelligence Risk Management Framework. The goal of the Framework is to manage risks to individuals, organizations, and society associated with AI. I presented my work as part of a small panel and engaged in dialogue with NIST about its future plans for the AI RMF, exploring and contributing to the overarching project goals. Several NIST staffers told me that Artificial Unintelligence was helpful to them in getting up to speed on AI in order to engage in policy conversations and design policy interventions. I have presented at numerous small government events, like speaking at the US State Department's privacy conference or giving a keynote at the National Library of Medicine at the National Institute of Health campus. I've been fortunate to be asked to speak with audiences all over the world about AI and bias. I've given 200+ talks and interviews in the past six years, in front of audiences ranging from 10 to 10,000. I'm pleased to say that I regularly hear from readers inside and outside academia who say the work has influenced their thinking. The Digital, Governance and Sovereignty Chair at SciencesPo created an entire report based on technochauvinism (Cobbe 2022).

One of the most delightful pop culture impacts showed up on social media. There is a Twitter/Mastodon bot called @NYT first said that announces the first time a term is used in the New York Times. "Being in the New York Times is this standard of notability," the bot's developer, Max Bittker, told the New Yorker's Max Norman. Norman writes: "The bot's rigorous simplicity isn't so much trivializing as philological: it's built on the presumption that the language the Times uses is worthy of our attention, and of our scrutiny. That, at the end of the day, the *Times* actually is a kind of language, a shared, increasingly global English, aspiring, for better or for worse, to the almost encyclopedic, universal function of helping people 'understand and engage with the world.' @NYT first said matters then, if only as the historical record of a paper expanding to include ever more of that world-including @NYT first said itself, which the Times reported on in 2019, prompting the bot to tweet its own name. It's a flashing beacon from the far reaches of the paper's ever-widening cultural empire" (Norman 2023). I was pleased to say that "technochauvinism" was mentioned by @NYT first said after I did an interview with a New York Times reporter. Artificial Unintelligence has won awards, including a PROSE award and the Hacker Prize from the Society for History of Technology.

More Than a Glitch, similarly, reached audiences all over the world. Glitch was a PROSE award finalist and received the getAbstract Business Impact book award. The Glitch

chapter on breast cancer was excerpted in WIRED, and has been noted as a completely novel experiment in self-quantification—in part because very few people working in open-source AI cancer detection get their own mammograms. WIRED has a monthly readership of 30 million across its platforms (Condé Nast 2024). Another chapter, on why the next frontier in gender equity is in databases, was excerpted in TIME. TIME reaches more than 120 million readers each month across its brands (TIME USA, LLC 2024). A third chapter, on algorithmic auditing, was excerpted in *Issues in Science and Technology*, the magazine of the National Academies of Sciences, Engineering, and Medicine. In 2023 the issues.org website reached about 75,000 people per month. The book received substantial media coverage, with favorable mentions in *Cosmopolitan, Glamour, FT, The Next Big Idea Club*, National Public Radio (NPR), and more.

The empirical contribution of the work in this portfolio is that it manifests many of the theoretical arguments that communication scholars have been making about data journalism for several years. Scholars like Anderson, Lewis, Hamilton, Usher, and McIlwain have theorized about the possibilities of accelerating different types of journalistic and media inquiry using computational tools. My work shows that yes, such things are possible, and such work is made more robust by drawing from and contributing to scholarly discourse.

Discussion and Future Work

The work in this portfolio is grounded in the literature of media studies and critical data studies as well as applied computer science. In the future, I plan to continue this path and do additional investigative journalism with and about artificial intelligence. My near future plans include three projects related to my work in AI. One is a scholarly article on medical racism for Oxford University Press, in which I will write about algorithmic audits of medical systems such as the systems used to evaluate patients' eligibility for kidney transplants. Another is a project with a civil society group, Data & Society, on algorithms used in education that target members of historically marginalized groups. I am interviewing educators and evaluating an algorithm relative to the 2023 executive order on AI, with an eye toward developing a standardized set of review criteria for algorithms in education. I'm also adapting a previous paper, on AI in investigative reporting, to discover if machine learning methods can be used to detect or predict grand corruption in cross-border investigations. This work is informed by open-source investigative techniques sourced from journalism and from law enforcement. I have been doing this corruption investigation work in collaboration with

the International Anti-Corruption Court, a new international organization that is drafting its founding treaty and looking to use state-of-the-art technologies in achieving justice.

I also plan to explore ways that computational journalism can be integrated into journalism curricula. Since the time I began writing this commentary, a new Ph.D. in AI in Journalism began at the University of Surrey, led by Bahareh Heravi of the Institute for People-Centred AI (2024). I look forward to learning more about the program and seeing what can be adapted for US universities. Does it make sense to start a computational journalism Ph.D. program at NYU? This is a question I would like to explore further.

An ideal curriculum for a computational journalism Ph.D. might include mandatory technology classes that cover spreadsheets, data visualization, data scraping, data cleaning, and basic programming in either Python or R. Writing-intensive classes would include the history of data journalism; the economics of investigative journalism; practical work strategies for fulltime and freelance journalists; and the basics of writing with and about numbers. After these introductory concepts, the program might include a larger investigative project that results in new findings in a field.

When designing a computational journalism Ph.D. program, one of the major challenges involves methods for funding student investigative work. It's hard to talk about high-end computational journalism projects without talking about money. One of the things I've learned on my academic journey is that money is always a limiting factor for investigative journalism. In-depth investigations, especially computational ones, take a lot of people and a lot of time. In practical terms, this translates to expense.

In Spring 2024, I ran a collaboration with Proof News, a new investigative startup, that illustrates the scale of what's required for these types of projects. I worked with three Proof staffers to design two 3.5-hour class sessions for my 22 data journalism students. Together, we concluded that the 5 major AI chatbots continued to generate inaccurate information about the 2024 US presidential election, more than half the time, despite the AI companies promising to do better. Twenty-six people put in a combined 320 hours of work to draw this one conclusion. This time does not include the many hours it took to develop and test the bespoke software needed to test the AI chatbots, nor does it include the many hours of developing the project and fundraising and developing the original methodology and piloting the benchmarking method that Proof developed in collaboration with a lab at the Institute for Advanced Study. All in all, ours was an innovative and expensive conclusion. Is this level of effort practical for every computational journalism class, especially given the

adjunct-ification of American universities and the accompanying labor issues? This is a complicated issue.

Computational journalism is a team sport, and the more people who work on a project, the more expensive the project becomes. Hamilton (2016) analyzed the cost of a major ProPublica investigation. It took about a million US dollars and several years for a single investigation. Foundation models like Gemini and ChatGPT mask the need for funding, making it seem like computational tools are widely available and free. The tools are widely available, and some of their features are free, but the more advanced features (the ones that are more useful for high-end investigative work) are anything but free. The same is true for cloud computing, which is portrayed as less expensive than running one's own server. It is less expensive for small short-term projects, yes. For larger projects, the cloud bills add up quickly—as do the ongoing costs of maintaining digital projects online indefinitely. Unfortunately, talking about money is often considered indecorous inside (and outside) the academy. This will have to change, somehow, if we truly want to train and encourage more data journalists.

In addition to securing new funding streams for innovative computational work, the academy might adapt by embracing computational journalism outputs as artifacts that "count" for tenure and promotion. As I wrote earlier, tenure and promotion standards should be updated so that truly novel code can be recognized as a creative work. (Not all code is novel, of course.) Impact might also be considered in evaluating computational journalism work. Getting a policy changed or seeing staffing changes as a result of a story has long been considered substantial impact among journalists. However, this type of impact is hard to track and measure. Other impact measures are awards, online views, social sharing statistics, citations, and so on. Tracking all of these measures is challenging for news organizations, and even more challenging for individual journalists or academics. Expanding existing alt-metric tools to account for journalistic impact might make a big difference.

These are only some of the many challenges of the field of computational journalism. Figuring out ways to meet these challenges will require collective effort and a certain amount of re-imagining the field. I am particularly excited to work on some of these computational journalism problems in the context of investigating AI. AI is advancing at a rapid pace right now in 2025, and the pre-existing social problems embedded in AI systems and their training data will only become more entrenched as time goes on. I look forward to continuing with accountability work that brings such issues to light. I also look forward to the next time someone inquires why I don't have a Ph.D., because I will be able to give a very different answer.

Bibliography

- "8 Expert Tips for Getting Started with NotebookLM." 2024. Google. October 18, 2024. https://blog.google/technology/ai/notebooklm-beginner-tips/.
- Anderson, CW. 2013. "Towards a Sociology of Computational and Algorithmic Journalism." *New Media & Society* 15 (7): 1005–21. https://doi.org/10.1177/1461444812465137.
- Angwin, J, J Larson, S Mattu, and L Kirchner. 2016. "Machine Bias ProPublica." ProPublica. 2016. https://www.propublica.org/article/machine-bias-risk-assessmentsin-criminal-sentencing.
- Angwin, Julia. 2024. "The Future of Trustworthy Information: Learning from Online Content Creators." Harvard Kennedy School Shorenstein Center on Media, Politics and Public Policy. https://shorensteincenter.org/future-trustworthy-information-learning-onlinecontent-creators/.
- Brayne, Sarah. 2021. *Predict and Surveil: Data, Discretion, and the Future of Policing*. New York, NY: Oxford University Press.
- Broussard, Meredith. 2014a. "Artificial Intelligence for Investigative Reporting: Using an Expert System to Enhance Journalists' Ability to Discover Original Public Affairs Stories." *Digital Journalism*, November, 1–18. https://doi.org/10.1080/21670811.2014.985497.
 - ——. 2014b. "Why Poor Schools Can't Win at Standardized Testing." *The Atlantic*, July 15, 2014. http://www.theatlantic.com/features/archive/2014/07/why-poor-schools-cant-win-at-standardized-testing/374287/.
- ———. 2015. "Preserving News Apps Presents Huge Challenges." Newspaper Research Journal 36 (3): 299–313. https://doi.org/10.1177/0739532915600742.
- 2016. "Big Data in Practice: Enabling Computational Journalism through Code-Sharing and Reproducible Research Methods." *Digital Journalism* 4 (2): 266–79. https://doi.org/10.1080/21670811.2015.1074863.
 - —. 2018. Artificial Unintelligence: How Computers Misunderstand the World. MIT Press.
- ——. 2023. *More than a Glitch: Confronting Race, Gender, and Ability Bias in Tech.* Cambridge, Massachusetts: The MIT Press.

- Buolamwini, Joy, and Timnit Gebru. 2018. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." In *Proceedings of Machine Learning Research*, 81:1–15.
- Bureau of Labor Statistics, U.S. Department of Labor. 2024. "News Analysts, Reporters, and Journalists." Occupational Outlook Handbook. August 29, 2024. https://www.bls.gov/ooh/media-and-communication/reporters-correspondents-andbroadcast-news-analysts.htm.
- Caplan, R., J. Donovan, L. Hanson, and J. Matthews. 2018. "Algorithmic Accountability: A Primer." Data & Society. https://datasociety.net/output/algorithmic-accountability-aprimer/.
- Carlson, Matt. 2014. "The Robotic Reporter: Automated Journalism and the Redefinition of Labor, Compositional Forms, and Journalistic Authority." *Digital Journalism*, November, 1–16. https://doi.org/10.1080/21670811.2014.976412.
- Clerwall, Christer. 2014. "Enter the Robot Journalist: Users' Perceptions of Automated Content." *Journalism Practice* 8 (5): 519–31. https://doi.org/10.1080/17512786.2014.883116.
- Cobbe, Jennifer. 2022. "Technochauvinism." SciencesPo Digital, Governance and Sovereignty Chair. https://www.sciencespo.fr/public/chaire-numerique/wpcontent/uploads/2022/06/website-.pdf.
- Coddington, Mark. 2015. "Clarifying Journalism's Quantitative Turn." *Digital Journalism* 3 (3): 331–48. https://doi.org/10.1080/21670811.2014.976400.
- Condé Nast. 2024. "WIRED Press Center." WIRED. November 7, 2024. https://www.wired.com/about/press/.
- Conover, Ted. 2016. *Immersion: A Writer's Guide to Going Deep*. Chicago Guides to Writing, Editing, and Publishing. Chicago ; London: The University of Chicago Press.
- Costanza-Chock, Sasha. 2020. Design Justice: Community-Led Practices to Build the Worlds We Need. Information Policy. Cambridge, MA: The MIT Press.
- Creech, Brian, and Andrew L. Mendelson. 2015. "Imagining the Journalist of the Future: Technological Visions of Journalism Education and Newswork." *The Communication Review* 18 (2): 142–65. https://doi.org/10.1080/10714421.2015.1031998.
- D'Ignazio, Catherine, and Lauren Klein. 2020. *Data Feminism*. https://data-feminism.mitpress.mit.edu/pub/h1w0nbqp/release/3.
- Fassler, Joe. 2012. "Can the Computers at Narrative Science Replace Paid Writers?" *The Atlantic*, April 12, 2012.

http://www.theatlantic.com/entertainment/archive/2012/04/can-the-computers-at-narrative-science-replace-paid-writers/255631/.

- Friesinger, Günther, and Jana Herwig, eds. 2014. The Art of Reverse Engineering: Open, Dissect, Rebuild; [... Is Based on the Symposium "Open. Dissect. Rebuild.", Which Took Place 2013 in the Context of the Paraflows Festival in Vienna]. Cultural and Media Studies. Bielefeld: Transcript-Verl.
- Fu, Angela. 2022. "Study by Six Gannett Unions Finds Racial and Gender Pay Inequities." *Poynter*, October 18, 2022. https://www.poynter.org/business-work/2022/study-bysix-gannett-unions-finds-racial-and-gender-pay-inequities/.
- Garcia, Nichole M., Nancy López, and Verónica N. Vélez. 2018. "QuantCrit: Rectifying Quantitative Methods through Critical Race Theory." *Race Ethnicity and Education* 21 (2): 149–57. https://doi.org/10.1080/13613324.2017.1377675.
- Haim, Mario, and Andreas Graefe. 2017. "Automated News: Better than Expected?" Digital Journalism 5 (8): 1044–59. https://doi.org/10.1080/21670811.2017.1345643.
- Hamilton, James. 2016. *Democracy's Detectives: The Economics of Investigative Journalism*. Cambridge, Massachusetts: Harvard University Press.
- Hamilton, James T., and Fred Turner. 2009. "Accountability Through Algorithm: Developing the Field of Computational Journalism." Developing the Field of Computational Journalism. Center For Advanced Study in the Behavioral Sciences Summer Workshop: Stanford University.

http://www.stanford.edu/~fturner/Hamilton%20Turner%20Acc%20by%20Alg%20Fi nal.pdf.

- Hanna, Alex, Emily Denton, Andrew Smart, and Jamila Smith-Loud. 2019. "Towards a Critical Race Methodology in Algorithmic Fairness," December. https://doi.org/10.1145/3351095.3372826.
- Heravi, Bahareh. 2024. "Fully Funded PhD Studentship in AI and Journalism | University of Surrey." September 10, 2024. https://www.surrey.ac.uk/fees-andfunding/studentships/fully-funded-phd-studentship-ai-and-journalism.
- Hicks, Mar. 2019. "Hacking the Cis-Tem." *IEEE Annals of the History of Computing* 41 (1): 20–33. https://doi.org/10.1109/MAHC.2019.2897667.
- Kovach, Bill, and Tom Rosenstiel. 2014. *The Elements of Journalism: What Newspeople Should Know and the Public Should Expect.* Revised and Updated third edition. New York: Three Rivers Press.

 Lareau, Annette. 2021. Listening to People: A Practical Guide to Interviewing, Participant Observation, Data Analysis, and Writing It All Up. Chicago Guides to Writing, Editing, and Publishing. Chicago ; London: The University of Chicago Press.

"Long Context | Gemini API | Google AI for Developers." n.d. Accessed November 7, 2024. https://ai.google.dev/gemini-api/docs/long-context.

- Mazmanian, Melissa, and Christine M. Beckman. 2018. "'Making' Your Numbers: Engendering Organizational Control Through a Ritual of Quantification." Organization Science 29 (3): 357–79. https://doi.org/10.1287/orsc.2017.1185.
- McGuinness, Tara Dawson, and Hana Schank. 2021. *Power to the Public: The Promise of Public Interest Technology*. 1st ed. Princeton: Princeton University Press.

McIlwain, Charlton D. 2020. *Black Software: The Internet and Racial Justice, from the AfroNet to Black Lives Matter.* New York, NY: Oxford University Press.

Mercedes, Dilcia. 2023. "How Newsrooms Pay Journalist-Coders Today." April 19, 2023. https://source.opennews.org/articles/how-newsrooms-pay-journalist-coders-2023/.

Montal, Tal, and Zvi Reich. 2017. "I, Robot. You, Journalist. Who Is the Author?:
Authorship, Bylines and Full Disclosure in Automated Journalism." *Digital Journalism* 5 (7): 829–49. https://doi.org/10.1080/21670811.2016.1209083.

- Norman, Max. 2023. "Do You Speak New York Times?" *The New Yorker*, March 7, 2023. https://www.newyorker.com/culture/rabbit-holes/do-you-speak-new-york-times.
- O'Neil, Cathy. 2016. Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. First edition. New York: Crown.

Pahlka, Jennifer. 2023. *Recoding America: Why Government Is Failing in the Digital Age and How We Can Do Better*. New York: Metropolitan Books.

- Palta, Rina, Julia Angwin, and Alondra Nelson. 2024. "How We Tested Leading AI Models Performance on Election Queries." The AI Democracy Projects. https://www.proofnews.org/how-we-tested-leading-ai-models-performance-onelection-queries/.
- Sweeney, Latanya. 2013. "Discrimination in Online Ad Delivery." *Communications of the ACM* 56 (5): 44. https://doi.org/10.1145/2447976.2447990.
- TIME USA, LLC. 2024. "TIME Media Kit." November 7, 2024. https://mediakit.time.com.
- Todd, Kim. 2021. Sensational: The Hidden History of America's Girl Stunt Reporters. New York: HarperCollins Publishers.
- US Bureau of Labor Statistics. 2017. "Newspaper Publishers Lose over Half Their Employment from January 2001 to September 2016." TED: The Economics Daily.

https://www.bls.gov/opub/ted/2017/newspaper-publishers-lose-over-half-their-employment-from-january-2001-to-september-2016.htm.

- Usher, Nikki. 2016. Interactive Journalism: Hackers, Data, and Code. Urbana: University of Illinois Press.
- Verdegem, Pieter. 2021. "Introduction: Why We Need Critical Perspectives on AI." In AI for Everyone? Critical Perspectives, edited by Pieter Verdegem, 1–18. University of Westminster Press. https://doi.org/10.16997/book55.a.
- Ward, Amy Sample, and Afua Bruce. 2022. The Tech That Comes next: How Changemakers, Philanthropists, and Technologists Can Build an Equitable World. Hoboken, New Jersey: John Wiley & Sons, Inc.
- Weiss, Robert S. 1995. Learning from Strangers: The Art and Method of Qualitative Interview Studies. Riverside: Free Press.
- Wensveen, Stephan, and Ben Matthews. 2014. "Prototypes and Prototyping in Design Research." In *The Routledge Companion to Design Research*, edited by Paul A. Rodgers and Joyce Yee, 1st ed., 262–76. Routledge. https://doi.org/10.4324/9781315758466-25.
- Wenzel, Andrea, Daniela Gerson, and Evelyn Moreno. 2016. *Engaging Communities through Solutions Journalism*. New York: Tow Center at Columbia Journalism School.