

Algorithmic Auditing

A rhetorical analysis of two organizations focused on investigating algorithmic inequity

Amy Toole
English, Technical Communication
University of Central Florida
Orlando, Florida, USA
amy.a.toole@knights.ucf.edu

ABSTRACT

Algorithms are embedded in daily life: they influence everything from what we see to what jobs we get to what our credit score is. Despite algorithm's ubiquity, there is no oversight, regulation, or system of appeals to address potential discrimination. Recent studies have found many algorithmic computations to be deeply flawed, and people are beginning to take action. This rhetorical analysis looks at two organizations on the front line of algorithmic auditing.

KEYWORDS

Algorithms, algorithmic auditing, social justice, rhetoric, technology, artificial intelligence, AI

ACM Reference:

Amy Toole. 2012. Algorithmic Auditing: A rhetorical analysis of two organizations focused on investigating algorithmic inequity. In *Proceedings of ACM Orlando conference (ORLANDO'22)*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/1234567890>

1 Introduction

Algorithms are used for a range of technological purposes. These technologies are often touted as ways to increase efficiency, enhance security, and reduce errors. However, upon closer inspection, it is apparent that many algorithmic computations are deeply flawed. Algorithmic processes play a large, if unrecognized, role in daily life, but there is no oversight, regulation, and system of appeals for a technology that is not well understood.

This case study addresses the rhetorical strategies used by two organizations: the O'Neil Risk Consulting and Algorithmic Auditing (ORCAA) and the Algorithmic Justice League (AJL). Each website offers a different solution for addressing algorithmic bias. ORCAA, founded by data scientist Cathy O'Neil, is a consulting firm that helps companies manage and audit algorithms. AJL, founded by researcher Joy Buolamwini, is a social justice organization focused on artificial intelligence (AI) awareness and advocacy. Assessing each organization's use of

rhetoric can lead to a better understanding of how ORCAA and AJL support algorithmic auditing.

1.1 Background

Society needs algorithms. During her tenure as New Jersey Attorney General, Anne Milgram helped improve the state's crime rates by implementing data-driven risk assessments to be used in conjunction with instinct and experience [9]. She concluded that issues as complex as criminal justice cannot be navigated by an individual's gut reactions or subjective judgment alone; technological processes such as algorithms must be included. Society is technologically dependent on an individual and collective level and algorithms are "indispensable" to this dependency [15]. James Brown Jr., a digital rhetorician and director of digital studies, posits that while algorithms provide solutions to complicated societal issues, they also generate complex ethical problems [2]. Algorithms can be biased or unfair based on how they are programmed [10] [16]. Even if a program is designed with good intentions, such as offering cheap loans or reducing prison sentences [13], it can have harmful implications.

Because algorithms are not well understood and not recognized in everyday life, it is difficult to assess when this technology perpetuates bias. In her 2017 TED Talk, Cathy O'Neil compares the obvious destruction of a plane crash to the invisible destruction of an algorithm, stating that "an algorithm designed badly can go on for a long time, silently wreaking havoc" [12]. Unlike the localized damage of a plane crash, biased algorithms can disrupt lives miles away from where they originated. Buolamwini equates the destruction of algorithmic bias to a virus, explaining that both can spread "on a massive scale at a rapid pace" [4]. Technology has progressed by leaps and bounds in recent history, but the assumption that "we've arrived" is far from the truth [3]. Questions about what is possible, who participates, and what is ethical should be explored regularly as technology advances [2].

I first heard of algorithmic profiling in Noble's book, *Algorithms of Oppression*. Noble makes the argument that the discriminations embedded in algorithmic code are a human rights issue [10]. Before reading Noble's book, I knew that algorithms were involved in my social media feed, but I was unaware of how embedded algorithms were in my daily functions. The Netflix

Commented [GP1]: According to CMS, in-text citations should follow one of two options, which can be found in detail [here](#). It's completely up to you on which style to use, however, I will be adding numbers (i.e., the "9" in this citation) as the journal requires citations to reference their corresponding number on the bibliography page(s).

Commented [AT2R1]: Changed in-text citations to numbers.

Commented [GP3]: In tandem with your professor's feedback, a brief introduction to Cathy O'Neil, similar what you did with James Brown Jr., would be beneficial for reader awareness.

Commented [AT4R3]: I explain who O'Neil and Buolamwini are in the introduction and in their respective sections.

Commented [GP5]: Also in conjunction with the introduction comment, another one here for Buolamwini would provide context to who exactly this is.

Commented [AT6R5]: See comment response above.

documentary *Coded Bias* [7] deepened my understanding of the myriad of ways that algorithms affect daily life. *Coded Bias* features interviews with researchers, authors, politicians, and others to explore issues with algorithmic discrimination. The documentary covered solutions ranging from government oversight to third-party auditing. There are many ways that algorithmic oppression can be reduced and prevented. Noble emphasizes that algorithms are “worthy of our interrogation” [10], which is what ORCAA and AJL claim to do.

1.2 Assumptions

This case study accepts the following assumptions: that algorithms can perpetuate bigotry; that this bigotry stems from the biases of the programmers and the companies they work for; that those algorithms are used to make or support decisions that affect a myriad of activities; that the role and design of algorithms are opaque; and that companies and governments have dual responsibilities to develop unbiased, fair algorithms and to correct existing algorithms that are biased or unfair.

2 Rhetorical Analysis

2.1 ORCAA

The O’Neil Risk Consulting and Algorithmic Auditing (ORCAA) is a consulting company that works with companies to “manage and audit their algorithmic risks” [14]. The ORCAA team includes data scientists, strategists, and researchers. There is currently little regulation surrounding algorithmic fairness, but as more people understand the role that algorithms play in daily life, an increasing number of companies begin to explore ways to mitigate algorithmic risks. The ORCAA website uses various rhetorical strategies to make an argument for its services, establish credibility, and engage its audience.

2.1.1 Positioning ORCAA as the Solution. Upon arriving at the ORCAA website, the reader is faced with this sentence in bold, white text: “It’s the age of the algorithm and we have arrived unprepared” (Figure 1).




Figure 1: The ORCAA homepage, warning readers that they are unprepared for the “age of the algorithm.”

There is little else on the screen to distract from this sentence. By making this claim, ORCAA places the reader in a moment in history. It suggests that computers – not humans – hold the power in society, and we have passively allowed this to happen. The suggestion is that while the reader is unprepared, ORCAA is ready with the solution. The final word on the screen is “unprepared,” a cliffhanger that draws the reader into the site. By scrolling down, the reader hopes they will find out how to prepare themselves for this ominous sounding “age of the algorithm.”

A short scroll down the page gives the reader more context about the problem. A brief, two-paragraph “About” section explains how algorithms negatively affect daily life in ways such as offering women less credit than men and making it harder for people with mental health statuses to get jobs. Figure 2 shows examples of common experiences, which demonstrate the scope of the issue: algorithmic inequality is a widespread issue that affects many populations.




Figure 2: The ORCAA website goes on to describe how algorithms impact daily life. After outlining the issues, the “What We Do” section is visible, guiding the reader to continue scrolling to see how ORCAA can help with the problem of algorithmic discrimination.

The “About” section is so brief that the next section, “What We Do,” is already visible. While algorithmic discrimination may be a pervasive problem, ORCAA is guiding the reader to seek its solution.

ORCAA continues its focused argument by dedicating several screens to answering how it provides the solution. The “What We Do” section varies from the first two sections in its length and amount of content. This section uses concise and direct language to inform the reader that ORCAA offers a range of services, including auditing, vetting, procurement, and workshops. These services indicate that ORCAA supplies a comprehensive solution to companies across industries. Specific words such as “due diligence,” “operationalize,” and “strategic communications” target corporate decision-makers who are responsible for ensuring their businesses are successful. The reader may be concerned about the problem of the “algorithmic age,” but the detailed information in this section lets the reader know that ORCAA can provide an answer.

Commented [GP11]: I’m unsure why this is capitalized. Perhaps this a term relating to the class in which you wrote this paper for?

Commented [AT12R11]: This does not need to be capitalized.

Commented [GP7]: Is there a reference for this quote? Or is this from ORCAA’s website?

Commented [AT8R7]: Added this reference.

Commented [GP9]: Per CMS guidelines, companies should be referred to in the singular, so copy edits will be made throughout to adhere to this.

Commented [AT10R9]: Noted. Accepted these suggestions.

Commented [GP13]: This is a good area to provide some quoted examples, just like you do later on in the paragraph.

Commented [AT14R13]: I agree, but I opted not to add examples in this specific instance.

2.1.2 Establishing Credibility. For ORCAA to transform the reader into clients, it must establish credibility. There is no accreditation for companies that perform algorithmic auditing, so ORCAA must use other methods to demonstrate its competency. It does this by aligning itself with other reputable companies and highlighting its professional achievements.

Establishing credibility is critical to acquiring new clients. In a relatively new field such as algorithmic auditing, companies cannot depend on a long history to establish expertise. Word of mouth and referrals are effective ways of building trust, and ORCAA accomplishes this by naming past partners in the "Industries + Clients" section. The branding logos of Airbnb, the State of Illinois, and Consumer Reports, to name a few, are displayed directly following ORCAA's list of services. Displaying these popular brands seems to ask the reader: *If Airbnb, a popular and successful company, has hired ORCAA, then why shouldn't you also use ORCAA for your business?* For those readers who do not see their industry represented in the company brands, ORCAA explains in a brief paragraph that it works with everyone: "private and public companies of all sizes and public agencies, in the US and internationally." [The website does not offer specific details about the services provided to each of these past partners.] There are no testimonials or recommendations directly from these past partners. ORCAA maintains a focused and streamlined website, and these details could distract readers from how ORCAA could help their business. Rather than getting caught up in the scale, scope, or success of the services ORCAA provides, the reader might only see that the two organizations are connected. The branding logos are displayed as evidence of ORCAA's credible and quality solutions.

ORCAA dedicates most of its site to show the reader why they need to hire ORCAA. The "Industries + Clients" and "In the News" sections showcase the companies that partner with ORCAA. It's not until the reader scrolls to the bottom of the page that they learn more about the individuals on the ORCAA team. Details about ORCAA team member accomplishments and experience is crucial to establishing credibility. The "People" section introduces the consultants with the phrase "We are ORCAA" followed by biographies and headshots for each of the five individuals who make up the ORCAA team. This is the only place on the website with photographs of people. The biographical information of professional experience, awards, and accomplishments, combined with a photograph of each person, contributes to a feeling of trustworthiness. At this point on the website, the reader has read about why algorithmic auditing is important, know the services that ORCAA provides, have seen who has partnered with it in the past, and now can put a face to the organization. The "People" section serves to personalize the company as well as establish another level of credibility, one that the reader will be more receptive to, now that they know the purpose and accomplishments of the company.

2.1.3 Engaging the Audience. ORCAA clients come from various sectors including finance, education, and hospitality. The primary audience is business owners and leaders. The purpose of the website is to inform potential clients about how ORCAA can help their company be successful by identifying potential algorithmic risks. A streamlined website enables busy corporate executives to quickly locate information and then contact ORCAA for their services. The website design is focused, easy to navigate, and has few distractions.

Algorithms may seem complicated, but this website is not. Outside the world of computer programmers and data scientists, algorithms are not widely understood. Many people have a vague awareness of their existence but would be hard-pressed to describe what an algorithm is or how it functions. ORCAA understands this. For a company that seeks to demystify complex computer processes, the simplicity of the website seems extremely intentional. The entire site is one long continuous page with a top menu bar that allows the reader to navigate between eight sections. There are no submenus, no drop-down boxes, and no sidebars. [The layout is so simple that it borders on boring, but there is just enough variety in type size, text length, and imagery that the site appears professional, not neglected.] This visual focus extends to the imagery used and beyond, to the content that is – and isn't – included.

Readers come to the ORCAA website from a variety of backgrounds, and they each have unique needs and goals. The limited use of imagery and color prevents distraction and focuses attention on the text. Algorithmic discrimination has a human cost, which ORCAA mentions in the "About" section. Rather than emphasize the impact of discrimination in images, ORCAA chose instead to remove any visual reference to humanity. Upon arriving at the website, the reader sees an abstract image of a dark background with light dots connected with short lines (Figure 1). The image is impersonal and suggestive of a computer network. This forms the background beneath the message about the Age of the Algorithm and sets the tone for the site. The muted color pallet, alternating blocks of solid backgrounds, and limited use of photos or other imagery focuses the reader attentions to the text.

Considering that algorithms affect everything from credit approvals to insurance to hiring, there is an opportunity to direct reader attention to a myriad of topics. Instead, ORCAA's website severely reduces distractions by limiting the number of links, details, or opportunities for engagement. There are few external links, which means that the reader has limited opportunities to leave the page. Content is presented to keep the focus on ORCAA and the services it provides. Rather than confuse readers with details about how algorithms inform daily decisions or make a case about whether algorithms should or shouldn't be regulated, ORCAA informs the reader that it is equipped to solve the problem of algorithmic bias in your company. The focus of the site is on the reader and what ORCAA can do for the reader. It is one-sided, with only a single opportunity for reader participation:

Commented [GP15]: Interesting. I think this says a lot and could be explored/analyzed further apart it being "a distraction."

Commented [AT16R15]: Included additional sentence exploring this idea.

Commented [GP19]: It may be worth considering incorporating screenshots/visual aids if you're considering submitting this to the journal for added context.

Commented [AT20R19]: Added screenshots of the two websites throughout the article.

Commented [GP21]: What isn't included?

Commented [AT22R21]: The following paragraph discusses what isn't included.

Commented [GP23]: I really love this sentence, the latter half feels so dark - adds a bit of good sting to your writing.

Commented [AT24R23]: Thanks!

Commented [GP17]: There being only five people on the team is another intriguing point that could be explored.

Commented [AT18R17]: Agree, but chose not to explore this in the paper.

the “Contact” section offers the sole request for reader participation by allowing the reader to send an email. There are no requests for donations, no opportunities to volunteer, and no links to social media accounts. The reader needs only to partner with ORCAA, and ORCAA will do the rest. ORCAA will determine how your company’s algorithms are being used and they will assess who these algorithms affect.

2.2 Algorithmic Justice League

The Algorithmic Justice League (AJL) was founded by Joy Buolamwini after she uncovered racial bias in facial recognition technology. She and a team of academics, authors, activists, and artists pursue algorithmic justice by raising awareness, sharing research, and influencing policy change. AJL has a multi-pronged mission, and the website reflects this complexity. The website uses various rhetorical strategies to explain the problem of algorithmic injustice, establish AJL’s credibility, and persuade the reader to participate.

2.2.1 Solution First, Then the Problem. The problems of artificial intelligence (AI), specifically, and algorithmic justice, generally, are addressed indirectly at first. The home page features action-oriented words such as “justice,” “movement,” and “towards.” This positions AJL as the solution to a problem without immediately describing the problem. This screen distinguishes between two groups: the “privileged few” and “all of us.” The reader is invited to participate by clicking a link to donate or entering their email address to “join the movement.” Even without knowing exactly what the problem is, the reader feels compelled to act based on the use of participatory language.

The home page goes on to make the argument that AI intensifies “racism, sexism ableism, and other forms of discrimination.” This discrimination has significant implications, including “harm[ing] vulnerable and marginalized people, and threaten[ing] civil rights.” While the mechanisms of algorithms are opaque and complex, AJL promotes a message of engagement and empowerment to its readers. The lines are drawn: to join AJL is to join a movement against not only computers but against the “privileged few” who benefit from discrimination. The “us versus them” framing simultaneously invites newcomers in while setting parameters on who the opposition is. AJL members are “agents of change” who “spread awareness,” “report bias,” host a workshop,” “grow the movement” and “request an audit.” The message of AJL is clear: you can help fight algorithmic injustice.

2.2.2 Establishing Credibility. To transform readers into participants, AJL must establish the need for algorithmic justice and demonstrate its credibility in this field. The fight against algorithmic injustice is new, so AJL establishes its credibility by aligning itself with well-known brands, organizations, images, and movements. The most prominent connection is with the Netflix documentary *Coded Bias* [7]. This documentary follows Buolamwini through her discovery and investigation of AI discrimination and ends with her founding the Algorithmic Justice

League. Navigating to the AJL home page prompts a popup that reads “You’ve seen the film” with an image from *Coded Bias* (Figure 3).

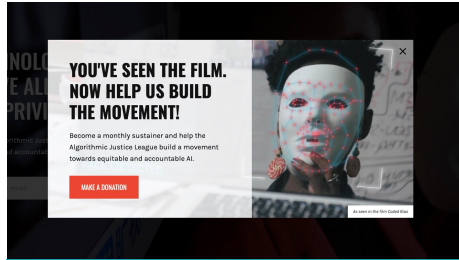


Figure 3: The AJL homepage with the popup that simultaneously aligns the organization with popular streaming service Netflix and asks visitors to donate to the “movement.”

This promotes a sense of community by assuming that every visitor to the AJL site has seen the same documentary and is now seeking more information. Visitors who have not seen the documentary might feel left out, and therefore might be compelled to watch *Coded Bias* in order to fully understand the reference. Building a community is integral to social justice, and AJL attempts to draw the reader into the movement from the first moment.

AJL must convince readers that it is not a fringe movement based on esoteric and cerebral ideals. Positioning AJL alongside Netflix suggests that AJL is already accepted into mainstream culture. The site continues this positioning in other ways. There is a link to Buolamwini’s TED talk. The company logos of organizations such as *Forbes*, *Wired*, and *The New York Times* are on the home page. There is a link to the documentary *Coded Bias* and bios for featured cast members. The implication is that if TED, *Forbes*, and *The New York Times* are credible sources, then readers are more likely to view AJL as credible after seeing these connections.

According to the home page, the algorithmic justice movement is for “all of us,” but the individuals that AJL mentions specifically are the “enthusiast, engineer, journalist, or policymaker.” The companies and faces that AJL uses to establish credibility suggest that AJL readers hold progressive views, are educated, and are aware that algorithms affect decision-making processes (whether they understand exactly how algorithms work or not). The reader may not agree with all the views promoted by each source, but they are familiar with the general ethos of each company.

AJL is not simply a repository for information, it is a social justice organization. Algorithmic justice may be a new movement, but racial justice is not. AJL equates its movement with the larger racial justice movement in a big banner notification at the top of

Commented [GP25]: What would this mean for those who haven't seen the film, however?

Commented [AT26R25]: Included sentence to address this question.

the page. Words such as “equitable” and “accountable” are familiar to those involved with social justice. These word choices indicate that the goals of algorithmic justice are similar to other activist movements. Drawing the reader’s attention to familiar images, brands, and phrases helps bridge the gap between the reader’s existing knowledge and the new realm of algorithmic injustice.

2.2.3 Engaging the Audience. Every social justice movement depends on participants, and AJL is no different. The site has numerous opportunities for engagement, including making a financial donation, entering an email address, sharing a personal “Coded Bias” story, or following AJL on social media. Appeals for reader participation are repeated often on the website, sometimes multiple times on one page. The abundance of participation opportunities contributes to the culture of inclusivity and engagement that AJL promotes. Not every reader can contribute financially, and not every reader is interested in signing up for an email newsletter. AJL is persistent in its solicitation, but it offers a variety of ways to get involved.

In addition to participation opportunities, AJL uses layout and imagery strategies to engage the reader. The layout boasts clean lines, few animations, and consistent formatting. The content is easy to read, but this simple format belies the quantity of content. The site is organized into six pages, and each page has multiple sections with numerous internal and external links. Information is duplicated from page to page (there are at least four links to donate). AJL addresses injustice on multiple fronts using a range of strategies, and its website reflects this complex approach. Rather than guiding the reader to focused action, the quantity of information combined with redundant content could lead readers to feel overwhelmed or distracted.

Buolamwini founded AJL in response to discriminatory facial recognition technology. The abundance of faces on the website seems to be a direct response to Buolamwini’s investigations. These faces remind the reader that algorithmic injustice is personal. Algorithmic discrimination is not a random technological accident but a widespread issue that affects everyone, everywhere, in every industry and walk of life. This bias disproportionately affects women and people of color, and these are the faces displayed the most. Buolamwini’s face appears frequently in various forms, which seems fitting as she is the founder of AJL and the primary subject of *Coded Bias*. One of the first photos the reader sees on the site is of Joy holding a white mask (Figure 4).




Figure 4: Buolamwini holding a white mask, illustrating that present day technology supports white faces but not Black ones.

This mask was the key to Buolamwini’s discovery of algorithmic bias: the computer would not recognize her Black face, but it would recognize the white mask as a face. The photo of Joy and the white mask consumes most of the first page and illustrates AJL’s overarching goal of centering individuals who have previously been masked by white uniformity.

There are other faces that the reader might recognize. Alexandria Ocasio-Cortez, Trevor Noah, and Michelle Obama are shown throughout the website. These images do not have a caption or explanation, suggesting that AJL expects the reader to know who they are. These individuals have a large audience of younger, progressive, socially active individuals. Using these images is another way that AJL aligns itself with existing successful and credible movements.

3 Conclusion

3.1 Implications

People make decisions every day. Some decisions are based on human thought while other decisions are aided by technological processes known as algorithms. Automating decision-making processes help society improve efficiency, but the underlying assumptions built into these processes must be carefully assessed [2]. Algorithms are “opinions embedded in code” [12], and if the opinions of algorithmic coders are racist, sexist, classist, or ableist, then the code will reflect these opinions. Algorithms should be audited [4] [12] to assess the data and measures of success. Algorithms provide the underpinning to crucial decisions in all areas of life, and when these mechanisms are not evaluated and revised, then they pose a threat to job security, financial independence, criminal justice, and more. If algorithms are left unmonitored, there could be catastrophic consequences.

Algorithmic decisions should be put into a larger context. O’Neil describes algorithms as proxies for difficult decisions that society wants to avoid [11]. Algorithmic bias affects everyone and is one example of a “rhetorical ecology” that scholars in all fields should work to understand [2]. Ensuring non-discrimination is just one element of algorithmic justice. Buolamwini points out that if

Commented [GP27]: This might also be a good area to add a screenshot/picture in your submission (while keeping in mind the maximum length the journal allows).
Commented [AT28R27]: Agree. Added screenshot.

individual freedoms are not protected, then accurate facial recognition software merely enables better surveillance for systems and governments to use [7]. The field of technical communication is positioned to witness how algorithmic technologies, ethics, and implementations are playing out across an array of industries.

3.2 Future Work

Scholars and computer scientists need to focus on algorithmic auditing and assessment. In a 2018 interview, Buolamwini explains that claims regarding algorithmic accuracy should be closely examined and questioned [3]. Latanya Sweeney, professor of the Practice of Government and Technology, is confident that we can build technology that “distinguishes between desirable and undesirable discrimination” [16]. Cynthia Dwork, professor of Computer Science [5], presents a framework for “fair classification,” an algorithm that emphasizes that “similar individuals are treated similarly” [6]. This algorithm can detect and certify unfairness [6] but it must be reassessed periodically. Ideally, frameworks such as this one are made available to the public, open to scrutiny, and continually refined [6]. Algorithmic influence will continue to increase in every aspect of daily life, so it is of utmost importance to examine these calculations to ensure that no group is being neglected, abused, or misrepresented.

3.3 Conclusion

O’Neil and Buolamwini have identified issues surrounding the accuracy and fairness of algorithms, and they have produced two solutions to the problem: work with companies directly to audit algorithmic risks, and advocate for better regulations and more accurate algorithms. Algorithms are crucial to automating systems and processes. As society continues to improve efficiency, increase output, and expand automated systems, it is more important than ever to evaluate the programming and impact of algorithms. Humans are increasingly connected to each other and to technology. As the interdependence between human processes and technology advances, the ethical implication of this dependence should be explored in new and robust ways [2]. This is a complex undertaking, as algorithmic technology can both “foster” and “thwart” discrimination [16]. In a push toward technological advancement, it may be tempting to try to automate everything, but the public should be wary of completely removing human judgment. Addressing social problems requires a “solution-driven, people-centered” approach” [15]. No computer can consider nuance, utilize empathy, or employ compassionate responses to the extent that a human can.

REFERENCES

[1] Algorithmic Justice League. <https://www.ajl.org>.
 [2] James J. Brown, Jr. 2015. *Ethical Programs: Hospitality and the Rhetorics of Software*. University of Michigan Press, 19, 87, 98, 177-178.
 [3] Joy Buolamwini. “Can We Trust The Numbers?” January 2018, in TED Radio Hour, produced by Guy Raz, NPR, podcast, www.npr.org/programs/tes-radio-hour/580617765/can-we-trust-the-numbers.
 [4] Joy Buolamwini. “How I’m fighting bias in algorithms,” November 2016, TEDxBeaconStreet, video, 08:35,

www.ted.com/talks/joy_buolamwini_how_i_m_fighting_bias_in_algorithms?referrer=playlist-the_inherent_bias_in_our_tech&t=9674.
 [5] “Cynthia Dwork,” Harvard University, 2021, dwork.seas.harvard.edu.
 [6] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel, 2012. *Fairness Through Awareness*, ITCS.
 [7] Coded Bias, directed by Shalini Kantayya (2020; 7th Empire Media), www.netflix.com/watch/81328723.
 [8] “James J. Brown Jr.” Rutgers University, 2021, www.camden.rutgers.edu/expert/brown.
 [9] Anne Milgram. “Why smart statistics are the key to fighting crime,” October 2013, TED@ BCG San Francisco, video, 12:29, www.ted.com/talks/anne_milgram_why_smart_statistics_are_the_key_to_fighting_crime.
 [10] Safiya Umoja Noble. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press, 1-2; 171.
 [11] Cathy O’Neil, “Can We Trust The Numbers?” January 2018, in TED Radio Hour, produced by Guy Raz, NPR, podcast, www.npr.org/programs/tes-radio-hour/580617765/can-we-trust-the-numbers.
 [12] Cathy O’Neil. “The era of blind faith in big data must end.” April 2017, TED2017, video, 13:10, www.ted.com/talks/cathy_o_neil_the_era_of_blind_faith_in_big_data_must_end.
 [13] Cathy O’Neil. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York, Crown Publishers, 32.
 [14] ORCAA. <https://orcaarisk.com>.
 [15] Roxana Radu. 2019. *Negotiating Internet Governance*. Oxford, Oxford University Press, 164-165.
 [16] Latanya Sweeney, 2013. *Discrimination in Online Ad Delivery*. *Communications of the ACM*, 56, 5 (May 2013), 44-54. DOI:10.1145/2447976.2447990.
 [17] Latanya Sweeney. “Latanya Sweeney, Ph.D.” 2019, <http://www.latanyasweeney.org/index.html>.

Commented [GP29]: This paragraph, apart from this first sentence, is strictly citations. Consider interweaving some of your own thoughts on the sources you’re citing to balance the content out.

Commented [AT30R29]: Good point. Added my own thought at the end of the paragraph.

Commented [GP31]: Looking at the feedback your professor provided here specifically, I think this sentence actually serves a purpose here; it ties into your previous paragraph/topic sentence.

Commented [GP32]: I love this ending sentence, it’s probably one of the best I’ve read in a paper of this nature. It’s impactful and leaves the reader thinking.

Commented [AT33R32]: Thank you!

Commented [GP34]: Ensure that the bibliography references are formatted to CMS conventions, which can be found [here](#).

Commented [AT35R34]: Adjusted bibliography to match CMS guidelines.