

TARGETED ADS

🔍 The Infrastructure for Algorithmic Discrimination ↑ ZIP

Search related to your zipcode

- 🔍 Racial bias
- 🔍 Racial discrimination
- 🔍 Racial inequality

Matthew Bui

Ho-Chun Herbert Chang

Charlton McIlwain

Report racially biased predictions

ABOUT THE AUTHORS

Matthew Bui is an assistant professor at the University of Michigan School of Information. mattbui@umich.edu

Ho-Chun Herbert Chang is a doctoral student at the University of Southern California.

hochunhe@usc.edu

Charlton McIlwain is Vice Provost for Faculty Engagement & Development at New York University, Professor of Media, Culture, and Communication, and founder of the Center for Critical Race and Digital Studies.

cdm1@nyu.edu

ABOUT THE CENTER FOR CRITICAL RACE + DIGITAL STUDIES

The Center for Critical Race and Digital Studies (CR+DS) is a network of prominent, public scholars of color who produce research, distribute knowledge, and convene stakeholders at the intersections of race and technology. Our affiliates include social scientists, artists, computer scientists, cultural critics, lawyers, policymakers, ethicists, journalists who understand digital media technology manifest racialized power relationships throughout our societies. We produce cutting edge research that illuminate how race, ethnicity and identity shape and are shaped by the design, deployment and use of digital technology applications and platforms, and how these technologies structure power to impact communities of color. We mobilize our collective knowledge, expertise, and life experiences as scholars of color to inform and engage the public and influence the public agenda.

GRAPHIC DESIGN AND LAYOUT

[Selly Djap](#)

The authors gratefully acknowledge funding support from the Democracy Fund, without which this project would not have been possible.

First published November 2022

04 **SUMMARY**

06 **INTRODUCTION**

09 **SECTION 1: TOWARD A CONCEPT OF ALGORITHMIC DISCRIMINATION**

Research Problem and Context

Regulating AdTech: Emerging Cases and Challenges

Algorithmic Audits as an Intervention Strategy

Toward A Concept of Algorithmic Discrimination

Toward a Sociohistorical and Infrastructural Algorithmic Audit Approach

Challenging Racial Hierarchies through a Sociohistorical Approach to Technology

An Overview and Critique of Algorithmic Fairness

Research Questions

23 **SECTION 2: TOWARD A TOOLKIT FOR ALGORITHMIC DISCRIMINATION AUDITS**

Data Collection

Data Analysis Techniques

Geography-based similarity analysis

Publisher-level similarity analysis

Summary

28 **SECTION 3: PRELIMINARY RESULTS**

Education Case Study: New York City Ads for “College Scholarships”

Top Domains Analysis

Zip Codes Similarity Analysis

Employment and Housing Ads: “Jobs Near Me” and “Housing for Sale” in Los Angeles

“Jobs Near Me”

Housing for Sale Near Me

Ad-Level Content Analysis

38 **SECTION 4: KEY IMPLICATIONS AND NEXT DIRECTIONS**

41 **REFERENCES**

SUMMARY

Online targeted ads uphold, produce, and recreate racially discriminatory infrastructures within everyday life.



ONLINE TARGETED ADS **UPHOLD, PRODUCE, AND RECREATE** RACIALLY DISCRIMINATORY INFRASTRUCTURES WITHIN EVERYDAY LIFE

This report summarizes the findings of a one-year study of online targeted advertisements. It highlights important questions and considerations regarding how online targeted ads uphold, produce, and recreate racially discriminatory infrastructures within everyday life. First, we propose a novel framework—*algorithmic discrimination*—which purports targeted ads to be discriminatory infrastructures by design: namely, this conceptual and analytic tool situates the potential harms and risks of targeted ads in relation to a longer history of predatory processes, tactics, and classification schemas, especially within and against marginalized communities of color. Next, we discuss how this framework relates to our novel methodology for algorithmic discrimination audits, in light of ongoing discussions of algorithmic accountability and corporations’ seeming attempts to forestall such efforts. Focusing on third-party search data pertaining to queries for educational opportunities, employment, and housing, we use zip codes as a proxy for racial and sociodemographic data, to audit and assess trends in online ad targeting. We compare differences across and within neighborhoods in online targeting patterns; we also compare individual ad messaging content. In contrast, we argue that a sociohistorical and infrastructural approach to algorithmic audits elucidates the community-based harms and risks of targeted ad systems as well as the digital infrastructures targeted ads undergird and fuel. As such, this approach more aptly shows the longer-term impacts of targeted ads and how they re-instantiate—and amplify—legacies of racial inequality. We close with key questions and future directions for this exploratory framework and methodology, particularly considering ongoing concerns about tech regulation and policy, and the protection of vulnerable communities from further tech-driven exploitation and extraction.



INTRODUCTION

Targeted ads have only continued to come under scrutiny, especially as they have expanded and grown in the online format.



DISCRIMINATORY SYSTEMS

TARGETED ADS AS INFRASTRUCTURE FOR ALGORITHMIC DISCRIMINATION

In June 2016, in a “groundbreaking” ruling, a Brooklyn-based jury found Emigrant Mortgage Company guilty of targeting minoritized¹ homeowners for predatory home loans (Feuer, 2016; Lane, 2016; Relman Colfax, n.d.). In short, through their audience segmentation, micro-targeting, and messaging strategies (which deployed racial cues to appeal to communities), Emigrant Mortgage Company was found guilty of aggressively marketing itself toward African American and Latino homeowners through various multicultural and culturally-specific print media outlets. Their “reverse redlining” scheme was driven by high profit margins tied to defaulted, subprime, and high-interest mortgages from 1999 to 2008; and it was found to be deeply implicated in the 2008 financial crisis. The case was important for fair housing and other civil rights laws, but it also revealed: 1) how targeted advertising was weaponized to target communities of color for predatory purposes; and 2) a general skepticism over whether wealthy banks and other powerful actors would be held accountable (Henderson quoted in Lane, 2016).

Since this ruling, targeted ads have only continued to come under scrutiny, especially as they have expanded and grown in the online format. Indeed, online ad targeting has been implicated within multiple high-profile scandals such as Cambridge Analytica, the release of the “Facebook Files”, and misinformation campaigns tied to the 2016 U.S presidential election (Confessore, 2018). These concerns are rooted in seeing how the current advertising ecosystem both mirrors—and amplifies—a history of predatory advertising strategies: specifically, micro-targeting and market segmentation. That is, online advertising and the evolution of advertising technology (AdTech)² and programmatic

¹ Like Crooks and Currie (2021), we sometimes opt to use “minoritized” to draw attention to the sociopolitical order that oppresses and marginalizes individuals and communities for their racial, gender, and sexual identities.

² AdTech, or “advertising technology” refers broadly to technologies that automate digital advertising and marketing processes, from purchasing to content generation. For basic definitions see McStay (2018). Unless otherwise stated, we use the term AdTech in this broadest possible sense.

(automated) advertising have provided advertisers the ability to intensify their advertising and marketing efforts (Lau, 2020; Liu-Thompkins, 2019), exacerbating and amplifying targeting tactics with little to no regulatory accountability. While ongoing congressional hearings strive to discuss antitrust and related anti-discrimination issues pertaining to Big Tech—including a recent bill proposal to ban “surveillance advertising” (Kelly, 2022) and a settlement that excluded protected categories from being used within Facebook’s Ad Library (Gillum & Tobin, 2019)—many barriers remain to ensuring legal and policy protections against discriminatory forms of ad targeting (for sample coverage, see Dans, 2021; Feiner, 2021).

Considering this, we aim to explore and propose both a conceptual and methodological toolkit for analyzing and auditing the impacts of targeted ads from a lens of racial justice and anti-discrimination. The results of this work, generously supported by the Democracy Fund, are presented in this report. First, drawing from related work on technology-enabled racial discrimination, we articulate and develop a novel framework of *algorithmic discrimination* in Section 1. This concept is tied to a transdisciplinary and sociohistorical approach that places algorithmic ethics and auditing research in conversation with extant (critical) scholarship on issues of race and technology. Resultantly, this proffers a novel but important intellectual contribution: mainly, we signal the importance of foregrounding the structural (and meso- and macro-level) inequities that targeted ads and other automated decision systems exploit and amplify. As a response, we call for a paradigmatic shift away from dominant “objective” and ethical approaches, which generally reduce harms and interventions to focus on individuals and individual bias (and, resultantly, de-biasing interventions). Secondly, to further explore the need for and utility of a sociohistorical (and critical and structural) approach to algorithmic audits, we articulate a novel methodology for auditing, assessing, and exploring algorithmic harms in Section 2: that is, we use zip codes from extant third-party search data to infer and examine the racial and sociodemographic parameters that seemingly drive ad targeting and buying strategies. With this dataset, we use speculative design principles and draw from “dark patterns” research to delineate the potential harms and risks of ad targeting in the current digital ecosystem. The preliminary results of our ongoing study are presented in Section 3 and we close with a discussion about next directions and possible limitations in Section 4.

Online advertising and the evolution of AdTech and programmatic advertising have provided advertisers the ability to intensify their advertising and marketing efforts, exacerbating and amplifying targeting tactics with little to no regulatory accountability.

SECTION 1

TOWARD A CONCEPT OF ALGORITHMIC DISCRIMINATION

A key research problem in the area of algorithmic accountability is developing useful terms, concepts, and analytical toolkits (i.e., language) for engaging with emerging and persistent issues of data ethics, fairness, and justice, especially as they relate to advocacy by and for vulnerable communities. In this section, we articulate and explore a novel concept of *algorithmic discrimination*: this concept (and its associated sociohistorical and infrastructural approach to algorithmic auditing) aims to distinguish between more mainstream definitions of algorithmic fairness—which have been critiqued for reducing problems of algorithmic harms and risks to individual “bias”. Relatedly, interventions often focus on de-biasing systems and garnering corporate cooperation and/or data access. In place of such strategies, we propose and argue for the need for a community-based, “reparative” approach (Davis et al., 2021) that redresses algorithmic harms and ameliorates data-mediated power relations; such a framework foregrounds—and highlights—the unequal contexts in which algorithmic decision making systems are deployed and have been weaponized. It also moves beyond solely technical fixes and intervention strategies to address issues³.

RESEARCH PROBLEM AND CONTEXT

To begin, our entry point into the issue of targeted ads centers around multiple recent high-profile cases involving Big Tech companies (and digital advertising giants) such as Google and Facebook (now Meta). For one, Google’s AdSense has been shown to perpetuate racial stereotypes through its suggestive representations within online ads (Sweeney, 2013) while also blocking advertisements attempting to address issues of

³ Algorithmic discrimination as a concept is also developed in a forthcoming paper, currently under preparation.

racism (Mak, 2020). Meanwhile, in late 2021 Facebook was under scrutiny over the leaking of its internal research, which confirmed that the company was aware of how its content recommendations preyed upon and caused harm against children. Prior to the leaking of these documents, the company had already been under scrutiny for unethical research (Kramer et al., 2014) and the underregulation of misinformation, particularly within marginalized communities (Entous et al., 2017); most famously depicted in the Cambridge Analytica scandal (Confessore, 2018). More recently, amid an ongoing civil rights audit, *The Washington Post* reported that Meta's internal research team had detailed how hate speech within its Facebook platform targeted minorities (Dvoskin et al., 2021). Yet, findings were not shared with auditors nor fully addressed by the company (Baik & Sridharan, 2022; Dvoskin et al., 2021). Generally, tech companies have responded to these scandals and reports by issuing public apologies as well as attempting to scrub harmful content and racial categories from their advertising sites (Hall, 2020).

Yet, even after removing racial categories, Facebook developed "affinity" categories through which individuals were tagged and classified based on individual interests; these schemas seemingly preserved a system for using racial stereotypes and assumptions to inform audience segmentation and other ad targeting strategies (Keegan, 2021, McIlwain, 2017). Internal efforts to remediate racial discrimination have been minimally successful, hence a recent bill proposal to ban surveillance advertising (Kelly, 2022) and calls from community organizations such as Color of Change to address issues through their Stop Hate for Profit campaign. These anecdotes and cases altogether illustrate a common tactic from companies such as Facebook: being uncooperative or minimally cooperative with external audits and agencies, while continuing to find ways to conduct business as usual in ways that potentially harm people and communities of color.

To be clear, our critique resides not in the racial classification schemas of advertising technologies in and of themselves. Rather, we draw attention to the ways in which AdTech platforms (and their related tech companies) allow for the flourishing and expansion of racially predatory practices and processes within everyday life. Discussing the state of multicultural advertising prior to the emergence of AdTech, Oscar Gandy (2000), sums up the political economy of racialized audiences. Relatedly, drawing from this work, Rosa-Salas (2019) documents how the U.S. marketing industry is largely informed by the cultural politics of racial segregation; other scholars (e.g., Davila, 2001; Shankar, 2014; and Foster Davis, 2013) discuss how different racialized audiences have been constructed according to their perceived economic value, designated as new and

“opportune” markets. Thus, it is under the guise of racial diversity and inclusion that multicultural marketing—and now, AdTech—use racial categories to shape their business practices and strategies. In the process, these processes reify epistemologies and ideologies tied to the cultural politics of racial segregation and reproduce longstanding legacies of racial inequality, all the while the rich and powerful accrue more wealth and status through their continued presence and expansion.

Importantly, while defining themselves as publishers, Facebook and Google increasingly operate as advertising servers: technologies that store ads and serve them onto publisher sites. That is, Facebook and Google act as platforms because they bring together advertisers, publishers, exchanges, networks, and agencies. Ad Servers are often sold alongside Demand Side Platforms (DSPs) and Supply Side Platforms (SSPs) because they offer the technical capabilities to collect data and serve ads. Ad servers allow advertisers to collect their own data on creative performance and the audiences. Indeed, the Google Ad Manager is both a publishing platform and a server that automates the buying of inventory and quickly serves these ads to available inventory for an optimal price. Facebook at one point had its own ad server product, Atlas, before shutting it down in 2016 to focus on its DSP Facebook AdSense. Ad servers are especially unique: not only do they automate the facilitation of auctions but they are effective in collecting data about the auctions themselves. This data aids advertisers in re-targeting. This feature is important. Not only do servers facilitate auctions that are the backbone of real-time bidding but they are almost always associated with other technologies that enhance data collection and user tracking. This data surveils and sorts audiences in targetable segments.

It is under the guise of racial diversity and inclusion that multicultural marketing—and now, AdTech—use racial categories to shape their business practices and strategies.

REGULATING ADTECH: EMERGING CASES AND CHALLENGES

Several major pieces of legislation illustrate congressional pressure in relation to issues of targeting and surveillance. First, largely in response to consumer protection concerns, various senators have proposed a social media transparency bill, *The Platform Accountability and Transparency Act (PATA)*. It would require social media companies to share platform data, which is significant in light of Facebook’s recent removal of data access to New York University (NYU) researchers. Notably, these researchers at the NYU Ad Observatory were investigating the effects of targeting political advertising on the platform (Bobrowsky, 2021).

Relatedly, building on *The Kids Internet Design and Safety (KIDS) Act*—which both enacted protections for children and attended to the business model of online advertising—the *Justice against Algorithms Act* aims to reform Section 230 of *The Communications Decency Act*. Namely, it discourages Facebook and other companies from using personalized advertising for profit-based models for advertising, thereby opening up Facebook and other tech platforms to being sued over users' third-party content. More recently, Reps. Anna Eshoo (D-CA), Jan Schakowsky (D-IL), and Sen. Cory Booker (D-NJ) proposed the ambitious *Banning Surveillance Advertising Act*, which would prohibit digital advertisers from targeting any ads to users.

Despite these intentions to regulate Big Tech, it is difficult to regulate in practice, for a number of reasons (Caplan et al., 2018; Feathers, 2022; see also Gillum & Tobin, 2019, for exception). For one, information access and a lack of standardized auditing and documentation processes often magnify the obstacle of corporate influence in regulatory debates. Indeed, a recent working group for a Washington state-sponsored bill aiming to regulate automated decision systems (ADS) highlighted corporate influence as one of the most significant regulatory hurdles (Washington Technology Solutions, 2021). These arguments echo previous concerns about the evasive tactics involved in protecting “black-boxed” algorithms from independent audits and regulation, particularly in the name of competition and profit generation (Pasquale, 2016; see also Ananny & Crawford, 2016).

Issues in regulation only compound extant obstacles to fostering civil rights: i.e., protecting citizens from discrimination based on race, gender, sexual orientation, and other social identity categories. For one, [Petty \(2003\)](#) documents the various types and issues of civil rights harms related to racially-targeted advertising, particularly cases of deception and unfairness that warrant governmental oversight and intervention. That said, by and large, anti-discrimination remains an unrealized goal ever since the passage of the Civil Rights Act of 1964. In response, critical race theory (Delgado & Stefancic, 2012; West, 1995) and intersectionality (Crenshaw, 1990; Collins, 2019) offer an apt understanding for why and how current US socio-legal institutions are often (implicitly and explicitly) designed to erase and penalize communities of color. While an exploration of these legal studies frameworks is not within the scope of this work, we do consider their critiques insightful and informative. These frameworks caution about the limits of and challenges to seeking legal remedies for algorithmic accountability while simultaneously striving to ensure that legal and policy remedies for discrimination are protected and expanded in the public interest.

The dominance of Google and Facebook as AdTech technologies opens up discussion and potential regulation for anti-discrimination. Indeed, Facebook and Google are not just platforms where ads get placed. These companies control the means for ad distribution, dictate the parameters and terms of ad targeting, and design the algorithms that determine auction outcomes (and, consequently, the visibility of both display and search advertising). As such, they hold immense power to both discriminate and facilitate discrimination by industry bad actors in sectors that have histories of causing intentional and unintentional harms to marginalized groups. In all, these challenges to intervening in the space of targeted ads and tech regulation elucidate the difficulties in defining, conceptualizing—and therefore acting upon—discriminatory issues concretely. Internal company mechanisms for redress are a common—often ill-suited—avenue for intervention. We propose algorithmic discrimination as a concept that might provide language and clarity for, first, conceptualizing the harms of targeted ads in relation to a sociohistorical and contextualized perspective; and relatedly, how to intervene from a community-based, contextualized approach.

Algorithmic discrimination as a concept that might provide language and clarity for, first, conceptualizing the harms of targeted ads in relation to a sociohistorical and contextualized perspective; and relatedly, how to intervene from a community-based, contextualized approach.

ALGORITHMIC AUDITS AS AN INTERVENTION STRATEGY

As previously noted, the Platform Accountability and Transparency Act (PATA) seeks to ensure independent audits are protected from companies' attempts to prevent them. This bill is related to previous legal cases calling for the protection of independent audits of Big Tech's impact on society. That is, the ACLU filed a lawsuit—*Sandvig v. Barr*—in 2016, to challenge the *Computer Fraud and Abuse Act (CFAA)* and ensure that academics, journalists, and other researchers were protected in carrying out anti-discrimination audits (American Civil Liberties Union, 2019; Gilens & Williams, 2020). A 2020 landmark ruling issued protections for these investigations. The opinion, reaffirmed later on, encompasses auditors who seemingly violate a website's terms of service to carry out their inquiries for the public interest. While not focusing on targeted ads specifically, Senator Wyden (D-OR) proposed the *Algorithmic Accountability Act* in February 2022, as a means to empower the FTC to audit automated decision systems, particularly AI systems within the housing and employment sectors (Kaye, 2022).

This momentum creates the opportunity for watchdog projects such as the NYU Facebook Ad Observatory (previously noted) to carry out independent algorithmic audits. Unfortunately, Facebook's removal of

data access has hampered such efforts. The Ad Observatory was a particularly useful platform for enabling journalists to produce investigative pieces and conduct algorithmic audits. During the 2020 U.S. presidential election, Politico journalists leveraged it to identify the misuse of \$4.5 million of Covid-19 impact funds to target individuals in battleground states by an organization called WorkMoney (Scott & Montellaro, 2020). During this investigation, Politico reporters interviewed WorkMoney founder CJ Grimes, who explained that the organization’s Facebook spending targeted six battleground states to draw in politicians’ attention, although she previously described the organization as “nonpolitical” in a previous Politico interview.

Similarly, US-based nonprofit journalism outfit The Markup has carried out algorithmic audits and published data-driven inquiries into the impact of technology on society, while exploring a funding model not predicated on advertising revenue. They have documented political targeting in the oil industry and financial ads, detailing discrepant climate change stances in ads targeted for liberal versus conservative users (Merrill, 2021); they also found credit card ads targeting based on age, a violation of Facebook’s anti-discrimination policy (Faife & Ng, 2021).

Similarly, ProPublica found job ads on Facebook excluded individuals based on their race, age, and/or gender. Their reporters documented multiple instances in which companies such as Uber and Verizon placed Facebook job ads that excluded specific ages and genders. In other work, they gathered evidence for the exclusion of Jewish and African American individuals from housing ads placed on the Facebook site. This investigative work led to a settlement whereby Facebook advertisers could not use race, age, or zip code for ads related to housing, education, and credit (Gillum & Tobin, 2019).

In short, these examples demonstrate the promises and limits of algorithmic audits as an intervention strategy for addressing issues of discriminatory ad targeting. Indeed, they show the growing popularity—and utility—of research-driven efforts to assess the impacts, risks, and harms of various systems by journalists, researchers, and academics. Often, these investigative reports and audit findings inform legislation (i.e., congressional bills, bill proposals, and other legal documents). Yet, these cases also show immense challenges to this work: notably, the time- and resource-intensive processes of audits and the resultant, prolonged legal battles they shape—even to allow for algorithmic audit projects to merely persist. In more recent work, Constanza-Chock et al. (2022) draw attention to the growing—yet ad-hoc—nature of audits, which they argue warrants increased oversight and a more concerted effort to establish standards for evaluation (and possibly a certification program).

TOWARD A CONCEPT OF ALGORITHMIC DISCRIMINATION

To further articulate our concept of algorithmic discrimination, we first anchor it as complementary to the framework and concept of algorithmic reparation (Davis et al., 2021): that is, drawing from key tenets of undoing and revealing structural inequalities in society within critical race theory (Delgado & Stefancic, 2012; West, 1995) and intersectionality (Crenshaw, 1990; Collins, 2019), Davis and colleagues argue for the need to consider how algorithmic systems can be both evaluated and challenged from a *sociohistorical* approach. Algorithmic reparations centers on revealing structural inequities in society. This contextualized approach promotes questions about power and power relations as they are encoded by algorithmic systems. It also resists exclusively technical fairness interventions, which fail to engage with the structural problems of “biased” systems.

Algorithmic discrimination expands upon the algorithmic reparations approach in two ways. First, it promotes both a sociohistorical and *infrastructural* approach to inquiry. Second, While Davis et al. (2021) identify data curation and redistributed AI power as two tangible methods for fostering algorithmic reparation, this approach encompasses algorithmic audits as another aligned avenue. That is, although not always, algorithmic audits are similarly attuned to revealing injustices and documenting the harms and risks of automated decision making systems.

In our concept of algorithmic discrimination, we draw inspiration from ongoing algorithmic audit projects to highlight the need to attend to an *infrastructural* approach to inquiry. To fill out this sociohistorical and infrastructural approach, we delineate key arguments from Gandy (2021 & 1993) and Chun (2021), who both offer prescient insights into how algorithmic systems surveil and control communities of color, encoding and re-encoding racial hierarchies by and through data-driven systems.

TOWARD A SOCIOHISTORICAL AND INFRASTRUCTURAL ALGORITHMIC AUDIT APPROACH

Our work draws inspiration from extant algorithmic accountability efforts but we also take a broader view of such projects—viewing targeted ads as infrastructures for racial targeting and discrimination. By infrastructure, we mean that targeted ads act as, and are part of, sociotechnical systems that facilitate various digital resources. In fact, targeted ads are an integral vehicle through which digital resources flow and are distributed that they have significant material impacts within and on society. For example, to

understand the impacts of Facebook and Google, Sandvig⁴ and colleagues call attention to the need to investigate digital technologies from an integrative perspective that merges both infrastructure studies and platforms studies (Plantin et al., 2018). In this view, the impacts of a singular commercial platform are scrutinized in relation to a broader ecosystem of multiple actors and platforms, illustrating the multiple benefactors and underlying mechanisms that fuel and undergird digital technologies. The platform view focuses on specific platforms and APIs whereas the infrastructural approach posits the importance of analyzing how such sites act in tandem with one another—their interoperability (See Table 1). Therefore, this framework can push algorithmic audits away from a platform-centric approach, i.e., one that sees Facebook and Google as singular platforms that only have specific impacts pertaining to their commercial products. Rather, it is important to view Facebook and Google as powerful actors within a larger sociotechnical system. Their actions have implications on the larger ecosystem, a viewpoint limited by a platform-centric approach.

The infrastructural approach demonstrates three easily obscured realities about the interconnectedness of platforms. An integrative perspective will also push for such analysis alongside a broader view of the interconnected impacts of platforms: the ubiquitous scale of systems, their monopolistic control, and thus the need to consider questions of regulation in the public interest. This shift to a broader view of algorithms beyond them as singular platforms is important because it allows for definitions and concepts to be attuned to the very large scale of multiple AdTech networks, for example, working in tandem; and the seeming inescapability of their impacts. Most importantly, the infrastructure view shows how “opting out” of such infrastructures—for example, Facebook or Google-driven Ads—is difficult, as they are seemingly ubiquitous in everyday life, hard to escape without becoming reclusive.

Coupled with the concept of *algorithmic reparation* (Davis et al., 2021, discussed earlier), the algorithmic discrimination concept encourages both a sociohistorical and infrastructural view of algorithmic audits. It promotes an approach to tech regulation that pushes us to continually foreground minoritized community-based harms and risks of targeted ads, to view them not as siloed networks and with singular impacts. In fact, it is important to foreground—and audit—the oft-invisible impacts of the interconnected data-driven systems that increasingly govern everyday life.

⁴ Notably, Sandvig was one of the key researchers involved in the legal protection of algorithmic audits from tech companies’ lawsuits in the 2016 ACLU lawsuit

Table 1: Summary Table from Plantin et al. (2018)

Table summarizing Infrastructure and platform properties.

	Infrastructure	Platform
Architecture	Heterogeneous systems and networks connected via sociotechnical gateways	Programmable, stable core system; modular, variable complementary components
Relation between components	Interoperability through standards	Programmability within affordances, APIs
Market structure	Administratively regulated in public interest; sometimes private or public monopoly	Private, competitive, sometimes regulated via antitrust and intellectual property
Focal interest	Public value; essential services	Private profit, user benefits
Standardization	Negotiated or de facto	Unilaterally imposed by platforms
Temporality	Long-term sustainability, reliability	Frequent updating for competitive environment
Scale	Large to very large; ubiquitous, widely accessible	Small to very large; may grow to become ubiquitous
Funding	Government, subscription, lifeline services for indigent customers, pay-per-use (e.g. tickets)	Platform purchase (device), subscription (online), pay-per-use (e.g. TV shows), advertising
Agency of users	"Opt out," for example, going off the grid	"Opt in," for example, choosing one platform instead of another; creating mashups

API: application programming interface

This novel approach of AD is attuned to institutionalized histories of racial inequality and the various barriers erected to prevent marginalized groups from accessing a myriad of opportunities. As McIlwain (2019) argues, constructs such as fairness fail to identify core concerns with respect to the impacts of technology in and on communities of color. Explaining the concept *Black Software*, McIlwain states: “in between these two versions of Black software... lies a most significant question, not about recently popularized concepts like computer bias or fair algorithms, or platform inequality, or digital ethics. No, the question goes to the heart of the matter that these concepts merely skirt around. Will our current or future technological tools ever enable us to outrun white supremacy? After all, this is not just our country’s founding principle. It is also the core programming that preceded and animated the birth, development, and first uses of our computational systems” (p.8, emphasis ours). As such, algorithmic discrimination extends, and is connected to, ongoing efforts to audit algorithms and pushes for accountability from the perspective of impacted communities and their concerns—concerns that are fundamentally associated with race and the future of democracy. AD historicizes and contextualizes current and ongoing debates about the amplification of predatory targeting tactics through AdTech and other algorithmically-enabled processes—from a broader societal view.

Will our current or future technological tools ever enable us to outrun white supremacy?

CHALLENGING RACIAL HIERARCHIES THROUGH A SOCIOHISTORICAL APPROACH TO TECHNOLOGY

To further guide the formulation of our framework, we draw largely from scholars of technology attenuated to questions of racial difference and exploitation. Notably, among others, we propose the work of Gandy (2021 [1993]) and Chun (2021) as particularly formative contributions, to guide and propose new directions for a body of work that is more equipped to engage with algorithms from a sociohistorical and infrastructural approach.

First published in 1993, Gandy's (2021) *The Panoptic Sort* provides new and updated evidence for how newly developed technologies have reinforced a "panoptic sort" whereby individuals and communities are monitored through various commercial and government data-driven systems. He demonstrates how it is the most marginalized communities—specifically, Black Americans—who are often disparately surveilled, penalized, and controlled by and through data collection and analytics efforts:

While classifications are also parts of identificatory processes, their application is not primarily orientated toward particular individuals but toward particular **types of individuals**

— Gandy, 2021, p.6, emphasis ours

Thus, Gandy's (2021 & 1993) discussion of the social construction of data-driven systems is pertinent to discussions of algorithmic fairness because it elucidates how technologies, under justifications of commercial and governmental utility and necessity, are often expanded to increasingly monitor *specific* individuals and communities; and how they discriminate between, and sort, populations to reproduce extant racial hierarchies in society. In this work, Gandy elaborates how targeted advertisements are an example of predatory processes, especially as they sort and surveil Black people and other communities for profit. In related work, Gandy (2000) provides an overview of the political economy of race-targeted marketing, detailing various developments in African American, Latino, and Asian markets: collectively, these efforts show the linkages between the growth of race-targeted marketing and the perceived economic "value" of a community. Similarly, Nakamura (2002) documents the dangers of racial stereotypes, as they are reproduced and amplified within digital spaces; and as they are leveraged and performed by different types of actors.

Applying an infrastructural approach, we argue that it is through the interconnectedness of targeted advertising processes (and other surveillant online data-driven collection efforts) that technological infrastructures are built and expanded for racial capitalism (Robinson, 2012). Considering past critiques of non-digital-multicultural marketing (e.g., [Davila, 2001](#); [Foster Davis, 2013](#); [Rosa-Salas, 2019](#); [Shankar, 2014](#)), it was through racial targeting strategies that ideologies of racial difference—and, even, racial segregation—were propagated, reproduced, and amplified. Thus, contemporary AdTech has only expanded the reach and scope of these extractive and exploitative processes—and ideologies—inextricably linking racism and racial hierarchies with the buttressing of the targeted advertising business model: a model that entails the use and interlinking of multiple data-driven platforms and systems to surveil, track, and profile communities for profit. As much as tech companies and the online advertising industry try to hide the negative impacts and far reach of these processes, racial targeting is a foundational, networked business practice for growing and expanding their profit margins.

Relatedly, to call attention to the racialized harms and risks of data-driven ideologies, Chun (2021) scrutinizes common assumptions from the fields of data science and statistics. In *Discriminating Data*, Chun traces the histories and interconnections between data, statistics, and various—problematic—uses and applications of data, such as eugenics and redlining, and more recently, predictive policing and facial recognition technology. She denaturalizes concepts such as *homophily*—instead asserting a need to consider *segregation* and *difference* within social networks—to demonstrate how these racialized metaphors bleed into our contemporary thinking. Indeed, considering the deep ties of these concepts with eugenics and other racial projects, there is a need to revisit and reflect on the implications of such taken-for-granted concepts, especially considering how such dominant concepts and frameworks have erased or obscured ideas and notions about race, racism, and racial capitalism. Therefore, we draw from Chun’s (2021) work to further reflect upon and reconsider how targeted ads magnify and make use of the “default settings of whiteness” to perpetuate discrimination; and how technologies amplify data-driven processes of sorting, classification, and racial segregation and segmentation—that is, racial capitalism—in everyday life. Indeed, as we have previously mentioned, these reflections might lead us to interrogate and more deeply examine how racial capitalism and other systems of racial extraction undergird our everyday lives, with little to no scrutiny and/or under various guises of individual benefits.

AN OVERVIEW AND CRITIQUE OF ALGORITHMIC FAIRNESS

This approach exists in contrast to the algorithmic fairness approach dominating literature on audits and algorithms. In their development of algorithmic reparation as an alternative approach to algorithmic inquiry, Davis and colleagues (2021) draw largely from Corbett-Davis and Goel's (2018) critique of three definitions of fairness to discuss the limits of technical attempts to de-bias systems and ultimately address issues of social inequality: that is, the most common definitions of fairness—anti-classification, classification parity, and calibration—often fail to account for the social contexts in which these fixes occur. In this section, we outline similar critiques, to highlight the limits of individualized, technically-oriented definitions of algorithmic fairness and harm issues; and to assert a reparative or justice-minded—and community-based—notions of harms and risks for algorithmic accountability efforts.

First, when synthesizing the literature on algorithmic fairness, we often observe a distinct difference between how algorithm designers (computer scientists) and regulatory observers (legal experts, critical race scholars) view them—as theoretical constructs and instances of these constructs, respectively. Fairness, by definition, is “absence of prejudice or favoritism toward an individual or a group based on inherent or acquired characteristics” (Mehrabi et al., 2021). Issues in algorithmic fairness were particularly brought to light when discussing COMPAS—the legal system that assisted judges in passing legal sentences and found bias toward African Americans—although earlier work has substantively discussed these issues (e.g., Friedman & Nissenbaum, 1996). The fairness paradigm has since been extended to cases such as the medical field, childhood welfare systems, and autonomous vehicles.

Relatedly, Mehrabi and colleagues (2021) provide a comprehensive survey on the state of current algorithmic bias and fairness research, broken down into broad categories of data biases and algorithmic design. From their data, they note 23 ways that data can be biased. Each of these definitions of bias share a common word: i.e., *predictor*. Put plainly, the field of fairness research often considers how algorithms can be designed to address specific criteria of—and measure—bias. Indeed, Jacobs' work calls attention to the varying theories of fairness as a key source of debates in the field; Jacobs' later work (2022) proposes that constructing fairness measurements is an important site for intervention and critical reflection, particularly for their impacts on governance processes. Namely, it is at this level—of operationalizing and defining constructs—that governance issues manifest: specific interventions can be

uplifted or forestalled based on how fairness is defined and constructed at this stage. Meanwhile, Chouldechova (2017) and Kleinberg (2016) allude to the “impossibility of fairness” due to an inability to satisfy its competing definitions.

We echo calls for an expansion of the terms and scope of fairness interventions to address its constraints in understanding and addressing community- and society-level harms. For one, Hoffmann (2021) discusses the predatory nature of inclusion within datasets: namely, being included within datasets to maximize fairness fails to address core issues and concerns of racial and social inequality. In fact, it often enables inequalities to be reproduced, as problematic applications are made more accurate, not stopped. Meanwhile, Green and Viljoen (2020) express the need for “realist” approaches to reduce harms, as opposed to idealized notions of their impacts. Relatedly, Bui and Noble (2020) note the idealized assumptions of fairness as an objective and neutral approach to data ethics, which only serves to reinscribe rather than redistribute power relations. In all, these works foreground a longer, more grounded view of vulnerable populations and their experiences of algorithmic harms, rather than romanticizing debiasing interventions that shallowly rectify individual instances of decision making bias.

As Davis et al. (2021) emphasize, algorithmic reparation is “geared towards building better systems and holding existing ones to account.” It is through this focus on systems, structures, and sociohistorical context that the predatory manifestations of racial capitalism within the targeted advertising environment become especially apparent. Moreover, rather than reducing concerns and issues to an individual level (and thus promote de-biasing algorithmic interventions), a sociohistorical and infrastructural approach generates a greater understanding of the longer standing, interconnected, and dynamic–sociotechnical–nature of the issues and concerns as play. As such, the importance of marginalized users and communities is further heightened within calls for accountability, as a longer view of history shows their disparate, greater bearing of burdens. In addition, a lack of intervention and regulation within algorithmic accountability discussions can thus be viewed as an attempt to maintain the racially unequal status quo in society, whether intentional or unintentional, by design or not.

Algorithmic reparation is “geared towards building better systems and holding existing ones to account.”

RESEARCH QUESTIONS

Amid an intensive moment of reflection over technology-enabled injustices and discrimination, our project scrutinizes whether targeted ads—and the companies publishing and profiting from them—should segment and presuppose their target audiences, especially within the realms of employment, education, and housing—opportunities explicitly protected by law from anti-discrimination. We ask broad questions such as: How do online ads discriminate within and across racialized communities? Why do we need to deconstruct how targeted ads are, by nature, racially discriminatory? How do we track, measure, and audit the loss of opportunities—i.e., the harms and risks of technology-enabled racial discrimination—across racial categories?

While this first section of the report developed our framework to address and explore these questions, the next section of our report details a novel methodological toolkit for, first, measuring the construct of algorithmic discrimination; and secondly, discussing and understanding these measures in accordance with a sociohistorical, infrastructural view. Using zip-codes as the basis of our analysis of third-party search data, we analyze the differences in ad coverage across multiple market sectors. The goal of investigating multiple sectors is to illustrate how different products can diverge greatly, the types of targeting that occurs, and the primary actors in each sector responsible for biased behavior. Additionally, we build a context-based model that allows us to assess whether evidence merits the classification of racial discrimination. Finally, we use content analysis methods to disentangle what is being marketed to different racial markets. To guide this inquiry, the report focuses on the following research questions:

- ▶ **RQ1: Where are the top domains targeting their ads (in terms of zip code)?**
- ▶ **RQ2: Which zip codes are targeted for the best and worst (employment, education, and housing) opportunities? (Is there a bias in the “hotspots” for online ads?)**
- ▶ **RQ3: To what degree do distributions of online ads reflect current and historical racial-spatial inequalities (i.e., segregation)? (Are race and/or class strong predictors of ad targeting?)**

These research questions discuss how targeting distributions—as measured by differential targeting behaviors by domains (RQ1) and zip codes (RQ2)—relate to, and amplify, extant racial inequality (RQ3). Domains refer to companies’ registered internet addresses so RQ1 will allow for an inquiry into the individual-level dynamics of targeted ads whereas zip codes (RQ2) will show the community and meso-level relations, and their connection with broader sociohistorical trends and relations (RQ3).

SECTION 2

TOWARD A TOOLKIT FOR ALGORITHMIC DISCRIMINATION AUDITS

In this section we propose (and explore) a methodological toolkit tied to our novel framework for algorithmic inquiry. First, we outline our data collection process, considering research issues of data access, limited resources, and tech companies' general resistance to external algorithmic audits. Next, using a mixed-methods approach, two sets of preliminary findings are presented: first, visualizations and mappings of trends in advertisements' targeted populations (as evidenced by sociodemographic data tied to more vs. lesser targeted zip codes); and secondly, we are developing a content analysis of individual messages, utilizing natural language processing techniques and methods⁵.

DATA COLLECTION

Table 2: Important keywords we tracked through third-party ad aggregators.

Keyword	Volume	CPC (USD)	Competitive Density	Sector
covid-19	3,350,000	\$0.00	0	Public Health
jobs near me	1,830,000	\$1.08	0.32	Employment
houses for sale near me	673,000	\$0.48	0.45	Housing
houses for rent near me	823,000	\$0.33	0.46	Housing
college scholarships	40,500	\$1.9	0.63	Education
bad credit mortgage	4,400	\$1.51	0.71	Housing
online degree programs	4,400	\$27.14	0.94	Education

⁵ Earlier forms of this methodology and empirical work have been presented and published within the archived conference proceedings for the 2022 Hawaii International Conference on System Sciences (Chang et al., 2022).

Data was primarily procured through SEMRush's Competitor Discovery platform, with a focus on sectors explicitly named within extant state and federal anti-discrimination protections (i.e., education, employment, and housing). By inputting a zip code and keyword, we tracked the top 80 to 120 domains vying for each ad keyword, including their **rank** (within Google Search Results), relative **visibility** and estimated **traffic**. As an example, if the sector of interest was education, we first used *education* as a seed keyword, and then selected the top keyword by search volume. Table 2 below shows a partial list of the most important keywords we tracked.

We primarily focused New York City and Los Angeles to be able to examine how racial discrimination manifests within and across large, urban metropolitan areas long known for a wide range of racial and socioeconomic backgrounds and issues of inequality. Together, this yielded 248,884 url-zip code pairs for New York and 191,697 url-zip code pairs for Los Angeles. This also resulted in 22,063,726 unique text ads. Upon scraping the ads, we then scraped the racial, socioeconomic (income), and educational profile of each zip code, using a publicly available Python wrapper for the Census API⁶ to retrieve 2015 U.S. Census data. This allows us to understand how independent ad publishers target specific audiences based on covariates, and whether race is a significant predictor. While race and ethnicity is a protected category and thus often prohibited from inclusion within explicit ad targeting strategies, the triangulation of ad targeting zip codes and U.S. census data allowed for an inquiry into the implied—and more subtle—forms of targeting behaviors that result from zip code based segmentation. Due to the project's launch around the start of the Covid-19 pandemic, we also collected data for "covid-19" as a keyword. This keyword often provided a baseline from which we could determine different targeting patterns, due to its equal targeting for all populations.

DATA ANALYSIS TECHNIQUES

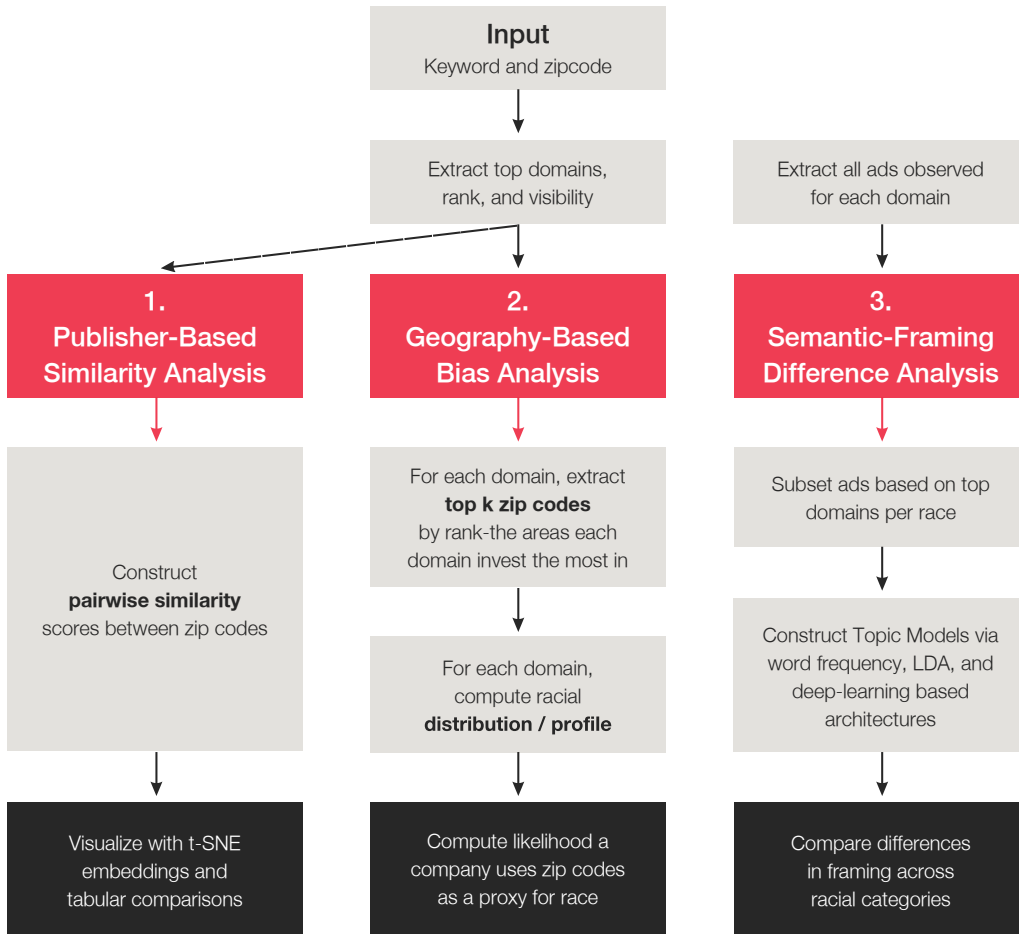
The general strategy of our data analysis was to identify patterns of difference between and across zip codes. To explore such trends, we conducted a series of different types of analysis to determine whether and how ad publishers differed and varied in their zip code targeting.

⁶ <https://github.com/datamade/census>

GEOGRAPHY-BASED SIMILARITY ANALYSIS

First, a natural question that arises when considering the distribution of ads is whether certain publishers—such as Zillow for housing ads—target specific zip codes, or even clusters of zip codes within their audience and market segmentation. To identify clusters, we created a network of zip codes where two zip codes are closely connected if they featured the same advertisers; and if these advertisers bid in the same way (i.e., targeted the same zip codes for their ad targeting strategies). Once the zip codes were connected as a network, a variety of state-of-the-art network algorithms were used to assess patterns in clustering, on a sector basis.

Figure 2: Overview of Data Analysis Process



⁶ <https://github.com/datamade/census>

PUBLISHER-LEVEL SIMILARITY ANALYSIS

While the network and clustering analysis was useful to understand the distribution of ads from a geographic (i.e., zip code-based) point of view, it failed to reveal details about individual publishers' targeting strategies. To further examine this, we measured and compared the racial and sociodemographics of each individual advertiser's most highly targeted zip codes, to compare whether and how these highly targeted zip codes were similar or different. To guide this, we mainly focused on zip codes' racial and ethnic demographics, based on the 2015 U.S. Census. We then: 1) ranked individual publishers based on the "diversity" of zip codes; and 2) compared them to the city-wide baseline. In this case, diversity was defined as the discrepancy between the city average, based on the census data for New York City overall. The results shed light on whether specific publishers' targeting strategies were seemingly tied to racial and ethnic criteria beyond a random level—i.e., potentially racially discriminatory. In the end, we selected the top 10 zip codes each ad publisher (domain) targeted, then compared it with the city-average to see certain domains target zip codes on the basis of race. There are multitudes of ways this data can be cross-sectioned, given its granularity.

A useful framework can be borrowed from current research in algorithmic fairness and bias. These are essentially mathematical formulations of different types of fairness (also known as parity measures), with a non-exhaustive list included below:

- ▶ **Demographic parity:** Given a treatment—such as the distribution of scholarships or in our case, ad serving—the ad serving rate between all groups should be equal. This was what proved the COMPAS program's bail denied rate as biased toward Black individuals $P(C=1|A=a)=P(C=1|A=b)$.
- ▶ **Equal Opportunity:** The rate of true positives between groups must be the same.
- ▶ **Equal Odds:** The rate of true positives and false positives between groups must be the same.

These differing definitions are immensely useful in clarifying algorithmic decision-making, where a ground-truth can be obtained. However, this framework is more compatible with cases where judgment must be made, rather than assessing the general state (or bias) of a system. In the results, we further elaborate and justify how such racial biases in ad targeting are arguably tied to more discriminatory behaviors.

SEMANTIC-LEVEL DIFFERENCES

Beyond auditing racial discrimination at the distribution level of online targeted advertisements, we investigate discrepancies in how language—and the individual ad messaging—is used within the sectors of interest. We first employed a classic topic-modeling technique (Latent Dirichlet Allocation) to identify key themes and topics subset on different sectors. Given the scope of ads, this introduced too much noise. Moving forward, we will employ BERT, a state-of-the-art neural network on a regression task—in essence, to see if the algorithm can predict the targeted racial composition based on the language alone. This specific type of analysis is being integrated and used to expand our preliminary findings and results, discussed in the next section. We will also consider how extant language processing tools and models are in need of refinement and development to assess racial fairness, bias, and discrimination. In particular, translating extant tools for different racialized—and semantically divergent—contexts and domains—education, employment, housing—is critical to this approach.

SUMMARY

Figure 2 illustrates our novel methodological process and the overall pipeline of data collection and data analysis for this long-term study of targeted ads. It is different in that it merges a network approach with tests. This is in an effort to extend the notion of demographic parity to a system-level analysis, using canonical tests of statistical significance. The contributions are thus two-fold: first, we theoretically and conceptually expand the notion of fairness to justice by analyzing system-wide phenomena rather than collections of judgements where ground-truths may be provided; secondly, we include a modular framework that is compatible in different statistical scenarios.

SECTION 3

PRELIMINARY RESULTS

To address our RQs, the first part of our empirical work largely draws from Taneja, Wu, and Edgerly's (2018) work on an integrative and infrastructural perspective of online use and engagement. That is, in contrast to the preference view, an infrastructural view foregrounds the format and materiality of digital systems: it is attendant to systems architecture and design, and the role of structural factors in mediating users' experiences, regardless of user content. More generally, an infrastructural perspective (Plantin et al., 2018) foregrounds the interconnectedness and "interoperability" of digital ecosystems, as explained above.

This perspective is particularly relevant in the case of search queries and targeted ads, as individuals attempt to access various opportunities and resources through their online search queries. Yet, as McIlwain (2017) shows, the political economy of web traffic operates in such a way that race-based hierarchies are shown to implicitly and explicitly emerge within online environments through systems architecture design mechanisms. Due to various web-based ranking systems and online targeting systems in place, a searcher—in this case, an individual from a minoritized community—might not be able to view and/or access the opportunities they seek, based on criteria used to track and surveil individual consumers, what Xian and Jacobs (2022) call "opportunity harms". Even if racial and ethnic background are not explicitly used to determine what ads they do, and can, view, other proxies and inferred attributes foreclose their access to specific ads while showing them others (Keegan, 2021, McIlwain, 2017). Therefore, an infrastructural approach to online environments and algorithmic audits reveals how automated systems are interconnected—with one another and with societal structures. In turn, this leads us to question whether and how specific individuals and communities are prioritized and privileged over others based on domains' strategies and their locations (RQ1 and RQ2); and how the form and function of these systems might predetermine such negative impacts upon communities of color.

An infrastructural approach to online environments and algorithmic audits reveals how automated systems are interconnected—with one another and with societal structures.

To answer this, we utilize data from the SEMRush platform, which tracks various web traffic metrics, including the relative position of domains within Google's search engine results pages (SERPs). Of the web traffic data brokers contacted (SEMRush, ComScore, Similar Web), SEMRush was the only vendor to provide data at the zip code level, thereby allowing for the use of zip codes to approximate the racial and sociodemographic characteristics of targeted audiences. For every keyword search, the SEMRush web traffic data allowed us to gain a sense of the "top" domains, or the company websites and advertisements that were relatively higher in position, vs. other domains (i.e., competitors). It also allowed us to consider questions of dis/advantage as they pertained to what targeted ads and domains were shown to users from different communities, as they attempted to use Google Search for various housing, employment, education and/or political information inquiries.

Put simply, if someone were to search for "houses for rent near me" or "jobs near me", we seek to ask whether and how communities are differentially exposed, and given access to, better opportunities: namely, more affordable and/or high-quality rental units or higher-paying and more stable jobs with benefits. Whether and how internet infrastructures encode and re-encode extant racial hierarchies through targeted advertising that allows for the privileged to continually accrue, and benefit from, their extant capital and advantage, based on assumptions and hierarchies tied to racial stereotypes. Indeed, Sweeney (2013) discusses how the advertising industry is often fraught with tenuous connections between racial identity and assumed consumer interests, using racial stereotypes to generate ill-informed ad targeting and segmentation strategies.

With this process, our study explores four of the noted components of Internet traffic systems: websites, platforms, ad sources, web and advertising traffic metrics. Namely, using web traffic metrics, it explores the key ad sources (i.e., top domains) within the Google search results page (i.e., one of the top search engines). We focus on areas covered under non-discrimination protections—education, employment, and housing—to extrapolate how and why these targeting processes contribute to disparate gatekeeping of resources (i.e., differential access to various domains and their offerings) within online infrastructures. We also include political ads as an area of inquiry, due to recent scandals involving target. It is important to reiterate that algorithmic audits into these issues are often limited by data access, hampered by tech companies' general resistance to external and independent audits. Thus, data sourcing and analysis is often a creative, adaptive process driven by available resources and access.

We begin by providing a full overview of our methods and analysis of a single keyword for the education sector (i.e., “college scholarship” ads), and then provide brief summaries of our results for ads within the employment and housing sectors (i.e., ads for “jobs near me” and “houses for sale near me”, respectively). In all, these results present evidence that targeted ads are driven to target and find specific racial and ethnic demographics beyond a random level of chance. This warrants continued and further exploration—and action and intervention, such as policy and regulation reform.

EDUCATION CASE STUDY: NEW YORK CITY ADS FOR “COLLEGE SCHOLARSHIPS”

TOP DOMAINS ANALYSIS

The first step in our analytical process was to provide some descriptive summaries of the dataset. The top domains for each sector observed are documented in Table 3. In the education sector, the top domain for White, Black, and Asian populations are *landmark.edu*, *collegeboard.org*, and *myscholly.com*.

Table 3: Top domains for each racial category for the city of New York for College Scholarships

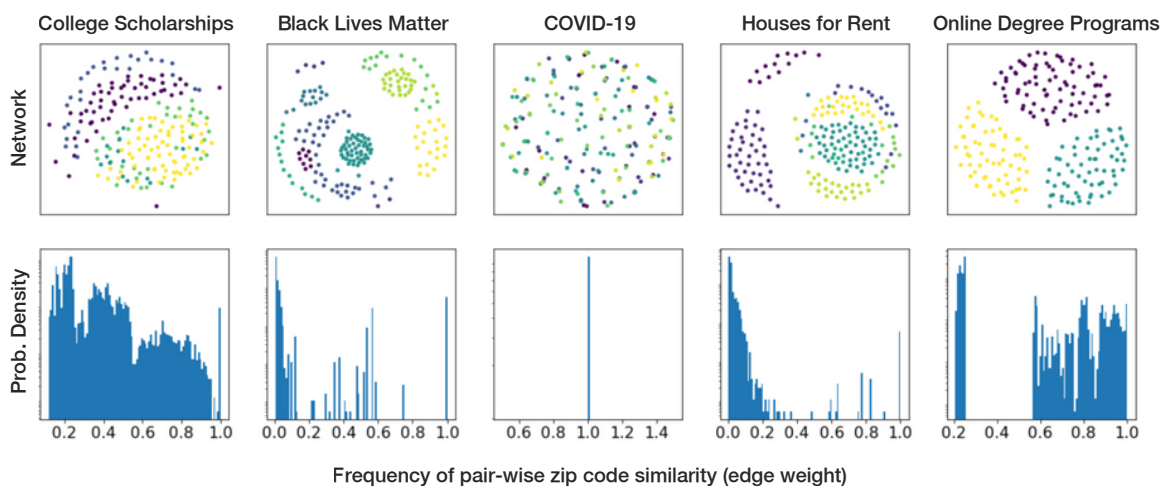
Rank	White	Infrastructure	Platform
1	"landmark.edu"	"collegeboard.org"	"myscholly.com"
2	"studentscholarships.org"	"phoenixpubliclibrary.org"	"hope.edu"
3	"firstinspires.org"	"fastweb.com"	"contracosta.edu"
4	"collegesofdistinction.com"	"cuny.edu"	"wvu.edu"
5	"coca-colascholarsfoundation.org"	"collegescholarships.com"	"macomb.edu"
6	"compostfoundation.org"	"unigo.com"	"gocolumbia.edu"
7	"mometrix.com"	"jumpstart-scholarship.net"	"alpenacc.edu"
8	"pinterest.com"	"collegegreenlight.com"	"palmbeachstate.edu"
9	"ed.gov"	"sfcollege.edu"	"ccis.edu"

Upon identifying these top domains, we also examined a) their racial composition and b) geographic distribution. We turn our attention to the university level for our education query dataset, considering the many different “.edu” domains present. Figure # shows the domains that generated the greatest biases toward specific racial groups as a result of their bidding strategy, based on their relative target demographics. To recap, the relative demographics—normalized on White, Black, and Asian racial categories—better compares across these three groups because census data contains multiple demographic fields. Relative demographics allows direct comparisons with each other and to the city-wide values (also normalized), and thus also more advanced statistical testing. The next section details these results, using zip codes as a proxy for racial and ethnic demographics.

ZIP CODES SIMILARITY ANALYSIS

The next research question seeks to explore whether we find differences in terms of zip codes, across keywords (RQ2). Figure 3 shows the network topology of select keywords from Table 2. We observe some immediate differences across sectors. For instance, there are clear clusters with *Black Lives Matter*, *Houses for Rent*, and *Online Degree Programs*.

Figure 3: Zip code-similarity networks (by keyword)



On the other hand, keywords such as *COVID-19* seem to be more well-mixed (i.e., little to no clustering). This can be explained by observing the distribution of pairwise similarity weights, shown in the second row.

COVID-19, for instance, shows only one value of similarity. This is because across all zip codes, ads about the pandemic--primarily from the CDC and state of New York--targeting all zip codes equally and thus act as a control. Importantly, this result suggests there was nondifferential targeting for pandemic-related queries.

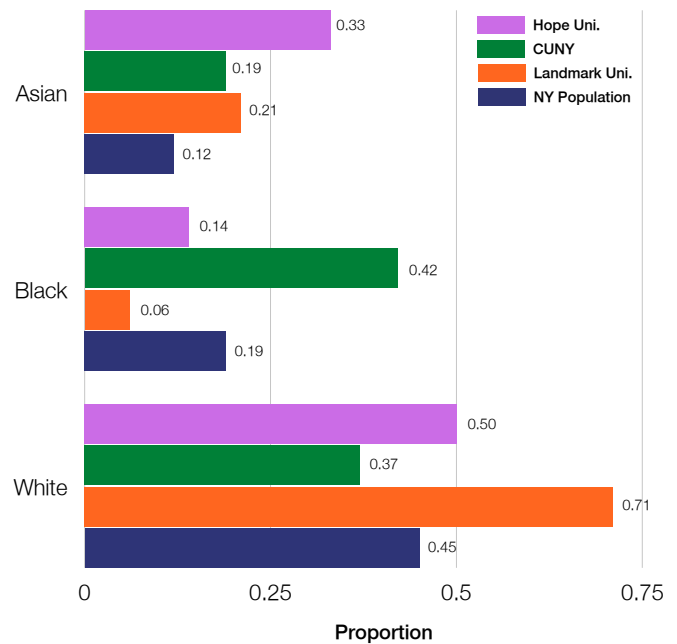
Our next step is to establish *how* and *to what degree* these zip codes are differentially targeted. We focus our attention on the education sector, with the keyword *college scholarship*. As a brief reminder, for every domain, we consider their top *k* zip codes for which they compete for, in terms of ranking and visibility. Table 3 shows the top domains based on three racial categories for the city of New York.

Based on Table 3, a few observations can be made. First, the top contenders for the keyword often consist of universities, external scholarship providers (e.g., coca-colascholarsfoundation.org) and general tertiary education websites (e.g., studentscholarships.org, finaid.org), which aggregate scholarship information. This includes government-sponsored websites such as studentaid.gov.

As a closer look, we consider how the top domains for “college scholarships” differ in terms of the marketing strategy by race. First, note the relative demographics for New York City at large are as follows: 12% Asian, 19% Black, and 45% White. Landmark University (in orange), which is the top domain in terms of targeting White audiences, serves 71% of its ads to White-heavy zip codes, 6% to Black-heavy zip codes, and 21% to Asian-heavy zip codes.

Meanwhile, CUNY (in green) serves ads mostly to predominantly Black zip codes (42%); and a proportion to White and Asian demographics: 37% and 19%, respectively. In this case, Black audiences encounter ads at a rate of more than 2 times the expected rate, based on the city average. Finally, Hope University (in purple) serves its ads accordingly: 33% to Asian zipcodes, 50% to White zipcodes, and 14% to Black zip codes; this rate for predominantly Asian zipcodes is almost three times the city average.

Figure 4: Top "College Scholarship" Domains by ZIP-Weighted Racial Composition



The different rates of targeting may be from a variety of reasons, due to allocation of marketing budget and to deliberate, race-based choices. In the best case scenario, we observe distinctions in who different universities bid for. Additionally, these differences in levels may certainly be attributed to characteristics of specific zip codes, such as income. Here, we offer a map of New York, its five boroughs and the top 20 zip codes that Hope, CUNY, and Landmark dedicate their bidding efforts to.

Landmark's (orange) audience can be seen predominantly around the Manhattan area and parts of Brooklyn. In contrast, Hope University's dominant audience is found in the Queens area and Long Island, with some bidding in South Brooklyn. CUNY, in contrast, bids the most across all five boroughs, with top bids in Staten Island, Manhattan, Brooklyn, and notably the Bronx, which is the largest departure compared to the other two schools.

Lastly, we offer a way to visualize the deviation from distributional fairness through a ternary diagram. Ternary diagrams are useful for analyzing interactions across three categories. This ternary diagram indicates the relative bias of each school, based on the city-average. That is, the midpoint of a triangle illustrates where NYC targeted ads are expected to lie, when no racial biases are detected. However, these clear clusters show that targeted ads for "QUERY" searches do not distribute according to expected ratios, even adjusting and accounting for varying levels of racial/ethnic representation.

Figure 5: Map of Attention from Top "College Scholarship" Domains


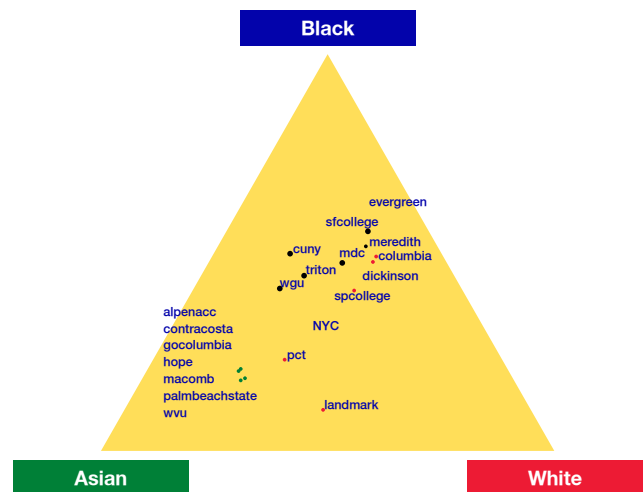


Figure 6: Tertiary Diagram of Top "College Scholarships" Domains by Zip-weighted Racial Composition





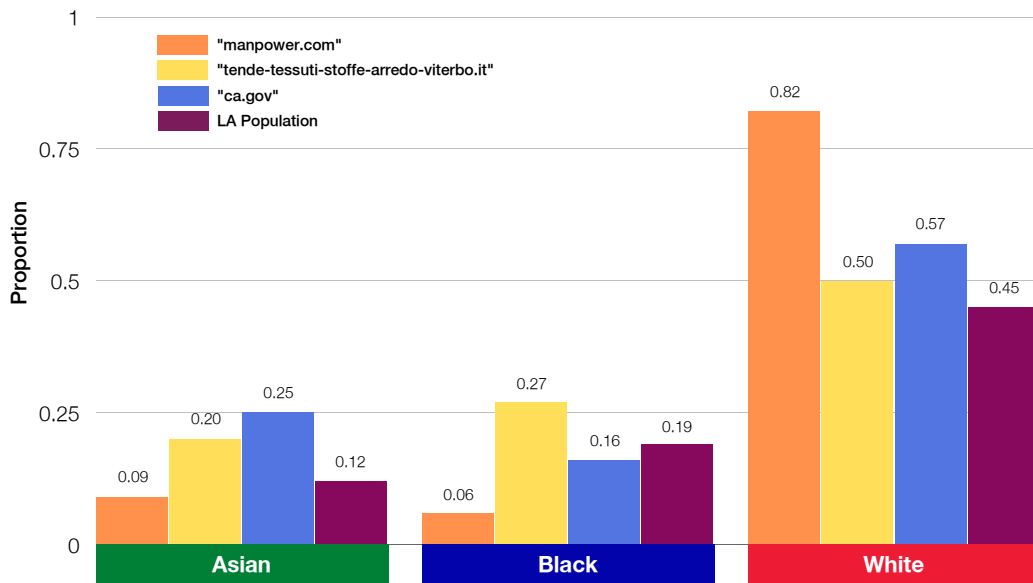
EMPLOYMENT AND HOUSING ADS: “JOBS NEAR ME” AND “HOUSING FOR SALE” IN LOS ANGELES

Upon providing the previous full run-through of a single sector’s ads (education ads for “college scholarships” in New York City), we now present two key findings in the employment and housing sector, using Los Angeles as an example.

“JOBS NEAR ME”

In Los Angeles, we find manpower.com, tende-...it, and ca.gov spent the most advertising revenue on White, Black, and Asian populations. Manpower.com is a professional job aggregator with a leaning toward engineering, tending to service white neighborhoods. On the other hand, we find Black neighborhoods get smaller local-shops looking for manual labor, such as “dulairoofing,” the second most popular domain looking to hire roof-tilers. The Californian Government (ca.gov) advertises job-postings to racial minority groups–Asian and Black populations much more frequently.

Figure 7: Top-domains competing for the employment sector in the city of Los Angeles



Notably, an interesting result arises from the large proportion of the Italian ads for “jobs near me”. Upon further investigation, this originates from postings targeting Little Italy’s high levels of Black, Hispanic, and Albanian populations. Indeed, Little Italy is a well-known working class “melting pot” district in Los Angeles, and the ads from local jobs directly reflect this population.

Table 4: Top Domains for "Jobs Near Me" in Los Angeles (by Zip Code weighted Racial Composition)

Domain_A	Domain_B	Domain_W
"ca.gov"	"tende- stoffe-arredo-viterbo.it"	"manpower.com"
"providence-california.jobs"	"dulairoofing.com"	"linkedin.com"
"craigslist.org"	"usps.com"	"target.com"
"backstage.com"	"dollargeneral.com"	"aldi.us"
"elipal.com.br"	"5strada.it"	"cub.com"
"bestbuyshop.online"	"usps.com"	"kaiserpermanentejobs.org"

We note here that these results seem to illustrate how the definition of “near” varies in different algorithmic contexts. Whereas affluent, whiter communities prescribe a greater notion of mobility, minority communities, in particular the Black population, have a greater percentage of local labor. As a note, this data reflects the period prior to COVID-19 when there were higher levels of working-from-home. Preliminary analyses of education ads (for “online degree programs”) also confirms this: as Coursera tended to market itself toward affluent white communities in Manhattan where as it is absent in Brooklyn and Queens, which skew toward more working class and racially and ethnically diverse.

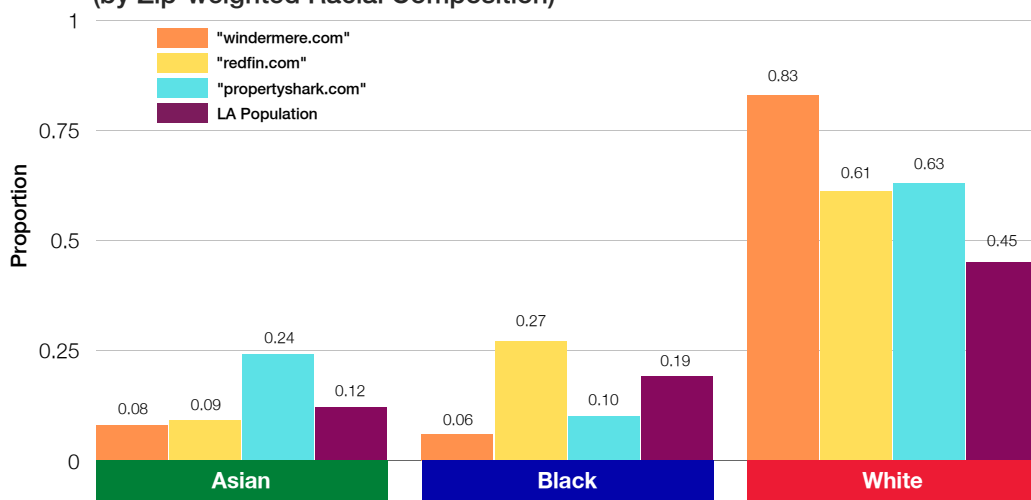
Table 5: Top Domains for "Jobs Near Me" in New York (by Zip Code Racial Composition)

Domains Targeting Asian Communities	Domains Targeting Black Communities	Domains Targeting White Communities
"usatoday.com"	"careerbuilder.com"	"joblinkapply.com"
"adeccousa.com"	"jobs-ups.com"	"glassdoor.com"
"nychealthandhospitals.org"	"linkedin.com"	"ny.gov"
"atriumstaff.com"	"atriumstaff.com"	"forbes.com"
"cub.com"	"craigslist.org"	"dropbox.com"
"usatoday.com"	"careerbuilder.com"	"joblinkapply.com"

HOUSING FOR SALE NEAR ME

Again focusing on Los Angeles ads, we find interesting results for the housing sector in terms of racially differentiated targeting strategies. In housing advertising, the network effect is very strong, thus generating a few primary players. This is in stark contrast with the employment sector where we observe a spectrum of locality in terms of ad distribution. Due to the low number of players, we observe the emergence of “turf-wars” across these aggregators. This can be observed further in Table 4, where overlaps between top domains does not exist.

Figure 7: “Houses for Sale Near Me” Domains in Los Angeles Ads (by Zip-weighted Racial Composition)



For ads delivered to Black populations, auction.com stands out as uniquely differentiated, as it deals with a particular type of property—i.e., ones that have been foreclosed. Notably, since the cost of advertising across different platforms is low, this is even more striking as this implies these platforms—Zillow, Windermere, Redfin, and Propertyshark—appear to have their own geographical agendas, therein warranting further research.

Table 6: Top “Houses for Sale” Domains in Los Angeles (by Zip-weighted Racial Composition)


Domains Targeting Asian Communities	Domains Targeting Black Communities	Domains Targeting White Communities
"propertyshark.com"	"redfin.com"	"windermere.com"
"zillow.com"	"trulia.com"	"homesandland.com"
"coldwellbanker.com"	"century21.com"	"homefinder.com"
"tollbrothers.com"	"auction.com"	"bhhs.com"
"reali.com"	"lakehomes.com"	"homesnap.com"

AD-LEVEL CONTENT ANALYSIS

The last part of our study builds on the “infrastructural” analysis by conducting a content analysis of various advertisements. This part of the project touches on 2 components of Internet traffic systems: websites and advertising messages. Beyond distributional discrimination, we also consider the actual messages as a site of discrimination. To do so, we implement natural language processing algorithms to assist our framing analysis.

Unfortunately, canonical topic modeling techniques such as Latent Dirichlet Allocation (LDA) generate too much noise. However, simple word clouds reveal how top text ads diverge in terms of their framing. We observe a few notable preliminary results:

- ▶ To Asian neighborhoods, top keywords stress ratings and service.
- ▶ To Black neighborhoods (subset on auctions), ads focus on foreclosed properties.
- ▶ To White neighborhoods, top keywords stress design, location, and luxury features.



Comparatively, the top words encode for socio-economic status, with the most stark contrast between auctioned foreclosed properties in black neighborhoods and “luxurious” amenities found in white neighborhoods. Relatedly, [Boeing \(2019\)](#) shows that the language with rental ads within Craigslist similarly focuses on neighborhood amenities for predominantly White communities whereas rental ads for heavily Black neighborhoods focus on restrictions and specific renter requirements (e.g., foreclosure and eviction histories). Interestingly, Asian neighborhoods feature ads that focus on new listings, there is a greater emphasis on their intermediary—in other word, trust in an agent.

SECTION 4

KEY IMPLICATIONS AND NEXT DIRECTIONS

In summary, as much as designers and computer scientists attempt to remove or ignore issues of race within technological design, race often operates as a technology in and of itself. From eugenics and redlining to facial recognition and targeted advertising, technologies have gleaned information based on racial identity categories, making predictions and institutionalizing inequalities tied to racial hierarchies (even if racial variables are not present). Thus, scholars of race and technology such as Gandy and Chun push us to think beyond the present moment, to consider the historical legacies of racial inequality, and the continued oppression and subjugation of racialized individuals and communities. Their work calls attention to how racial capitalism often undergirds and justifies the development of technologies to exist, to expand data collection and analytics efforts, and to foreclose other imaginaries outside of late capitalism. Yet, as Benjamin (2019) urges in her documentation of the New Jim Code, there is a need to contest, deconstruct, and reimagine new worlds outside of these logics; arguably, to push for what Davis et al. (2021) term as a “reparative” approach to algorithms. In different ways, their interventions highlight the harms and disparate risks borne by communities of color in the development of both public and privately funded data-driven systems. It complements the sociohistorical and infrastructural approach that we develop in this project and report.

In our work, using such sociohistorical approaches to algorithmics as a guiding lens, we present a novel concept and framework of algorithmic discrimination. We also extend and apply this concept to provide and explore a related methodology for algorithmic discrimination audits. This algorithmic audit approach integrates an infrastructural and socio-

Technologies have gleaned information based on racial identity categories, making predictions and institutionalizing inequalities tied to racial hierarchies.

historical perspective on issues and potential interventions for redressing algorithmic harms. First, this entailed collecting data and tracking keywords from sectors explicitly implicated within and seemingly protected by anti-discrimination policy in the United States; therein aiding in the documentation and tracking of issues and concerns. Next, we compare and examine differences within various sectors—healthcare, housing, education, and employment—to explore the different logics of ad distribution. In education, stark differences can be found in which schools target zip codes but we hope to expand our analysis by analyzing trends and differences in the continuing education industry and (predatory) for-profit university sector.

In all, this report presents a novel research framework and methodology that calls attention to how AdTech platforms (and their related digital platforms) contribute to an ecosystem that reproduces social and economic inequalities and hierarchies. To better illustrate the racist harms of and logics of racial capitalism within AdTech—and thus pave the way for more effective interventions and regulation—we echo calls for research concepts and designs that are attendant to the sociohistorical contexts and longstanding concerns of marginalized users and communities (e.g., Bui & Noble, 2020; Birhane et al., 2022; [Constanza-Chock et al., 2022](#); Davis et al., 2021). Within the algorithmic audit space specifically, our approach uses a sociohistorical approach to more aptly show the longer-term impacts of targeted ads and how they re-instantiate—and amplify—legacies of racial inequality—upon and within marginalized communities. We challenge dominant AI ethics paradigms and interventions to better grapple with the need for a sociohistorical—and thus, an inherently structural and political economic—understanding of algorithmic accountability issues and interventions. Indeed, given past difficulties in regulating the advertising industry for racial unfairness ([Petty, 2003](#)), it is of paramount importance to consider how contemporary advertising technologies extend, magnify, and obscure such predatory and racist outcomes—at a greater scale and speed—and amid growing concerns over a lack of algorithmic accountability writ large: how and why the digital distribution of resources is tied to ongoing power struggles and processes of racial control.

Importantly, we hope this novel concept, framework, and methodology push for an expansion of fairness interventions to more aptly and reflexively address societal issues; to move beyond invidiating and debiasing interventions to meso- and macro-level campaigns for justice



and redress. We hope to refine and develop this work through more in-depth data analyses and future publications related interdisciplinary venues for communication technology, urban/social policy, and data ethics/policy research. In the immediate future, we are actively expanding our analysis to assess directly socioeconomic and educational covariates comparatively between Los Angeles and New York City, in addition to our presented city-level analysis. We are also developing an interview study that engages with AdTech employees about how they construct and reconstruct different categories in their everyday work, even though they cannot explicitly discuss topics of race, age, and gender.

REFERENCES

American Civil Liberties Union. (2019, May 22). *Sandvig v. Barr: Challenge to CFAA prohibition on uncovering racial discrimination online* (Court Cases). <https://www.aclu.org/cases/sandvig-v-barr-challenge-cfaa-prohibition-uncovering-racial-discrimination-online>

Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3), 973–989. <https://doi.org/10.1177/1461444816676645>

Baik, J. S., & Sridharan, H. (2022, January 10). *Big Tech civil rights audits should be mandated by law*. Tech policy press. <https://techpolicy.press/big-tech-civil-rights-audits-should-be-mandated-by-law/>

Benjamin, R. (2019). *Race after technology: Abolitionist tools for the New Jim Code*. Polity Press.

Birhane, A., Kalluri, P., Card, D., Agnew, W., Dotan, R., & Bao, M. (2022). The values encoded in machine learning research. *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, 173–184. <https://doi.org/10.1145/3531146.3533083>

Bobrowsky, M. (2021, August 4). Facebook disables access for NYU research into political-ad targeting. *Wall Street Journal*. <https://www.wsj.com/articles/facebook-cuts-off-access-for-nyu-research-into-political-ad-targeting-11628052204>

Boeing, G. (2020). Online rental housing market representation and the digital reproduction of urban inequality. *Environment and Planning A: Economy and Space*, 52(2), 449–468. <https://doi.org/10.1177/0308518X19869678>

Bui, M. L., & Noble, S. U. (2020). We're missing a moral framework of justice in artificial intelligence: On the limits, failings, and ethics of fairness. In M. D. Dubber, F. Pasquale, & S. Das (Eds.), *The Oxford handbook of ethics of AI* (pp. 162–179). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780190067397.013.9>



- Caplan, R., Donovan, J., Hanson, L., & Matthews, J. (2018, April 18). *Algorithmic accountability: A primer*. Data & Society.
https://datasociety.net/wp-content/uploads/2019/09/DandS_Algorithmic_Accountability.pdf
- Carbaugh, J., Gregg, F., Knauss, D., Garland, H., LLamado, R., Tumminaro, A., Giuffrida, J., Culver, J., Hand, J., Ghenis, M., Welsh, B., Bond, D., Thomas, K., Menezes, R., Taylor, A., Riffer Xia-Reiner, B., Smith, B., Davis, M., Spence, M., . . . Sweger, B. (2022). *Datamade/Census* (Version 0.8.19) [Computer software]. GitHub.
<https://github.com/datamade/census>
- Chang, H.-C. H., Bui, M., & McIlwain, C. (2022). Targeted ads and/as racial discrimination: Exploring trends in New York City ads for college scholarships. In T. X. Bui (Ed.), *Proceedings of the 55th Hawaii International Conference on System Sciences* (pp. 2806–2815). University of Hawaii at Manoa. <http://hdl.handle.net/10125/79682>
- Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2), 153–163.
<http://doi.org/10.1089/big.2016.0047>
- Chun, W. H. K. (2021). *Discriminating data: Correlation, neighborhoods, and the new politics of recognition*. MIT Press.
- Collins, P. H. (2019). *Intersectionality as critical social theory*. Duke University Press.
- Confessore, N. (2018, April 4). Cambridge Analytica and Facebook: The scandal and the fallout so far. *New York Times*.
<https://www.nytimes.com/2018/04/04/us/politics/cambridge-analytica-scandal-fallout.html>
- Constanza-Chock, S., Raji, I. D., & Buolamwini, J. (2022). Who audits the auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, 1571–1583.
<https://doi.org/10.1145/3531146.3533213>
- Corbett-Davis, S., & Goel, S. (2018, August 14). *The measure and mismeasure of fairness: A critical review of fair machine learning*. arXiv.
<https://doi.org/10.48550/arXiv.1808.00023>

- Crenshaw, K. (1991). Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stanford Law Review*, 43(6), 1241–1299. <https://doi.org/10.2307/1229039>
- Crenshaw, K., Gotanda, N., Peller, G., & Thomas, K. (Eds.). (1995). *Critical Race Theory: The key writings that formed the movement*. The New Press.
- Crooks, R., & Currie, M. (2021). Numbers will not save us: Agnostic data practices. *The Information Society*, 37(4), 201–213. <https://doi.org/10.1080/01972243.2021.1920081>
- Dans, E. (2021, June 30). Congress rolls out some tough regulatory proposals for Big Tech. *Forbes*. <https://www.forbes.com/sites/enriquedans/2021/06/12/congress-rolls-out-some-tough-regulatory-proposals-for-big-tech/?sh=73da5ccf43c4>
- Dávila, A. (2001). *Latinos, Inc.: The marketing and making of a people*. University of California Press.
- Davis, J. L., Williams, A., & Yang, M. W. (2021). Algorithmic reparation. *Big Data & Society*, 8(2), 1–12. <https://doi.org/10.1177/20539517211044808>
- Delgado, R., & Stefancic, J. (2012). *Critical Race Theory: An introduction* (2nd ed.). New York University Press.
- Dwoskin, E., Tiku, N., & Timberg, C. (2021, November 21). Facebook's race-blind practices around hate speech came at the expense of Black users, new documents show. *Washington Post*. <https://www.washingtonpost.com/technology/2021/11/21/facebook-algorithm-biased-race/>
- Entous, A., Timberg, C., & Dwosin, E. (2017, September 25). Russian operatives used Facebook ads to exploit America's racial and religious divisions. *Washington Post*. https://www.washingtonpost.com/business/technology/russian-operatives-used-facebook-ads-to-exploit-divisions-over-black-political-activism-and-muslims/2017/09/25/4a011242-a21b-11e7-ade1-76d061d56efa_story.html
- Faife, C., & Ng, A. (2021, April 29). *Credit card ads were targeted by age, violating Facebook's anti-discrimination policy*. The Markup. <https://themarkup.org/citizen-browser/2021/04/29/credit-card-ads-were-targeted-by-age-violating-facebooks-anti-discrimination-policy>

Feathers, T. (2022, January 4). *Why it's so hard to regulate algorithms*. The Markup. <https://themarkup.org/news/2022/01/04/why-its-so-hard-to-regulate-algorithms>

Feiner, L. (2021, December 31). *2022 will be the "do or die" moment for Congress to take action against Big Tech*. CNBC. <https://www.cnbc.com/2021/12/31/2022-will-be-the-do-or-die-moment-for-congress-to-take-action-against-big-tech.html>

Feuer, A. (2016, June 27). Emigrant Savings Bank discriminated against minorities, Brooklyn jury says. *New York Times*. <https://www.nytimes.com/2016/06/28/nyregion/emigrant-savings-bank-discriminated-against-minorities-brooklyn-jury-says.html>

Foster Davis, J. (2013). Realizing marketplace opportunity: How research on the Black consumer market influenced mainstream marketers, 1920–1970. *Journal of Historical Research in Marketing*, 5(4), 471–493. <https://doi.org/10.1108/JHRM-02-2013-0006>

Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. *ACM Transactions on Information Systems*, 14(3), 330–347. <https://doi.org/10.1145/230538.230561>

Gandy, O. H. (1993). *The panoptic sort: A political economy of personal information*. Westview Publishing.

Gandy, O. H. (2000, June 7). *Audience construction: Race, ethnicity, and segmentation in popular media* [Paper]. 50th Annual Conference of the International Communication Association, Acapulco, Mexico. <https://web.asc.upenn.edu/usr/ogandy/targeting.pdf>

Gandy, O. H. (2021). *The panoptic sort: A political economy of personal information* (2nd ed.). Oxford University Press.

Gilens, N., & Williams, J. (2020, April 6). *Federal judge rules it is not a crime to violate a website's terms of service*. Electronic Frontier Foundation. <https://www.eff.org/deeplinks/2020/04/federal-judge-rules-it-not-crime-violate-websites-terms-service>

Gillum, J., & Tobin, A. (2019, March 19). *Facebook won't let employers, landlords, or lenders discriminate in ads anymore*. ProPublica. <https://www.propublica.org/article/facebook-ads-discrimination-settlement-housing-employment-credit>

- Green, B., & Viljoen, S. (2020). Algorithmic realism: Expanding the boundaries of algorithmic thought. *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAT* '20)*, 19–31. <https://doi.org/10.1145/3351095.3372840>
- Hall, K. (2020). Public penitence: Facebook and the performance of apology. *Social Media + Society*, 6(2), 1–10. <https://doi.org/10.1177/2056305120907945>
- Hoffmann, A. L. (2021). Terms of inclusion: Data, discourse, violence. *New Media & Society*, 23(12), 3539–3556. <https://doi.org/10.1177/1461444820958725>
- Jacobs, A. Z. (2021). *Measurement as governance in and for responsible AI*. arXiv. <https://doi.org/10.48550/arXiv.2109.05658>
- Kaye, K. (2022, February 8). *This Senate bill would force companies to audit AI used for housing and loans*. Protocol. <https://www.protocol.com/enterprise/revised-algorithmic-accountability-bill-ai>
- Keegan, J. (2021). *Facebook got rid of racial ad categories. Or did it?* The markup. <https://themarkup.org/citizen-browser/2021/07/09/facebook-got-rid-of-racial-ad-categories-or-did-it>
- Kelly, M. (2022, January 18). *Democrats unveil bill to ban online “surveillance advertising.”* The verge. <https://www.theverge.com/2022/1/18/22889903/democrats-targeted-advertising-facebook-google-surveillance>
- Kleinberg, J., Mullainathan, S., & Raghavan, M. (2016). *Inherent trade-offs in the fair determination of risk scores*. arXiv. <https://doi.org/10.48550/arXiv.1609.05807>
- Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences of the United States of America*, 111(24), 8788–8790. <https://doi.org/10.1073/pnas.1320040111>
- Lane, B. (2016). *Groundbreaking ruling? Federal jury finds Emigrant Bank liable for predatory lending*. HousingWire. <https://www.housingwire.com/articles/37419-groundbreaking-ruling-federal-jury-finds-emigrant-bank-liable-for-predatory-lending/>

Lau, Y. (2020). *A brief primer on the economics of targeted advertising* (Economic Issues). Bureau of Economics, Federal Trade Commission. https://www.ftc.gov/system/files/documents/reports/brief-primer-economics-targeted-advertising/economic_issues_paper_-_economics_of_targeted_advertising.pdf

Liu-Thompkins, Y. (2019). A decade of online advertising research: What we learned and what we need to know. *Journal of Advertising*, 48(1), 1–13. <https://doi.org/10.1080/00913367.2018.1556138>

Mak, A. (2020, August 13). *Google's advertising platform is blocking articles about racism*. Slate. <https://slate.com/technology/2020/08/googles-ad-exchange-blocking-articles-about-racism.html>

McIlwain, C. (2017). Racial formation, inequality, and the political economy of web traffic. *Information, Communication & Society*, 20(7), 1073–1089. <https://doi.org/10.1080/1369118X.2016.1206137>

McIlwain, C. (2019). *Black software: The internet and racial justice, from the AfroNet to Black Lives Matter*. Oxford University Press.

McStay, A. (2018). Digital advertising and adtech: Programmatic platforms, identity, and moments. In J. Hardy, I. Macrury, & H. Powell (Eds.), *The advertising handbook*. Taylor and Francis. <https://doi.org/10.4324/9781315558646>

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 115:1–115:35. <https://doi.org/10.1145/3457607>

Merrill, J. B. (2021, April 13). *How Facebook's ad system lets companies talk out of both sides of their mouths*. The Markup. <https://themarkup.org/citizen-browser/2021/04/13/how-facebooks-ad-system-lets-companies-talk-out-of-both-sides-of-their-mouths>

Nakamura, L. (2002). *Cybertypes: Race, ethnicity, and identity on the internet*. Routledge.

Pasquale, F. (2016). *The black box society: The secret algorithms that control money and information*. Harvard University Press.



- Petty, R. D., Harris, A. G., Broaddus, T., & Boyd, W. M., III. (2003). Regulating target marketing and other race-based advertising practices. *Michigan Journal of Race & Law*, 8(2), 335–394. <https://repository.law.umich.edu/mjrl/vol8/iss2/1>
- Plantin, J.-C., Lagoze, C., Edwards, P. N., & Sandvig, C. (2018). Infrastructure studies meet platform studies in the age of Google and Facebook. *New Media & Society*, 20(1), 293–310. <https://doi.org/10.1177/1461444816661553>
- Relman Colfax PLLC. (n.d.). *Saint-Jean v. Emigrant Mortgage Co.* (Case Profiles). <https://www.relmanlaw.com/cases-emigrant>
- Robinson, C. J. (2020). *Black Marxism: The making of the Black radical tradition* (3rd ed.). University of North Carolina Press.
- Rosa-Salas, M. (2019). Making the mass White: How racial segregation shaped consumer segmentation. In G. D. Johnson, K. D. Thomas, A. K. Harrison, & S. A. Grier (Eds.), *Race in the marketplace: Crossing critical boundaries* (pp. 22–38). Palgrave Macmillan.
- Scott, M., & Montellaro, Z. (2020, October 28). *Union-linked group spending big to target swing states on Facebook*. Politico. <https://www.politico.com/news/2020/10/28/union-group-spending-facebook-swingstates-433457>
- Shankar, S. (2014). *Advertising diversity: Ad agencies and the creation of Asian American consumers*. Duke University Press.
- Sweeney, L. (2013). Discrimination in online ad delivery: Google ads, Black names and White names, racial discrimination, and click advertising. *ACM Queue*, 11(3), 10–29. <https://doi.org/10.1145/2460276.2460278>
- Taneja, H., Wu, A. X., & Edgerly, S. (2018). Rethinking the generational gap in online news use: An infrastructural perspective. *New Media & Society*, 20(5), 1792–1812. <https://doi.org/10.1177/1461444817707348>
- Washington Technology Solutions. (2021). *Automated Decision-Making Systems Workgroup report*. <https://watech.wa.gov/sites/default/files/public/privacy/Automated%20Decision%20Systems%20Workgroup%20Report.pdf>
- Xian, L. & Jacobs, A. (2022). Harms as process: Opportunity harms in the US mortgage market. Working paper.