

Algorithmic Misjudgement in Google Search Results: Evidence from Auditing the US Online Electoral Information Environment

Brooke Perreault
bp101@wellesley.edu
Wellesley College
Wellesley, Massachusetts, USA

Ropafadzo Shava
rs2@wellesley.edu
Wellesley College
Wellesley, Massachusetts, USA

Johanna Hoonsun Lee
hl105@wellesley.edu
Wellesley College
Wellesley, Massachusetts, USA

Eni Mustafaraj
emustafa@wellesley.edu
Wellesley College
Wellesley, Massachusetts, USA

ABSTRACT

Google Search is an important way that people seek information about politics [8], and Google states that it is “committed to providing timely and authoritative information on Google Search to help voters understand, navigate, and participate in democratic processes.”¹ This paper studies the extent to which government-maintained web domains are represented in the online electoral information environment, as captured through 3.45 Google Search result pages collected during the 2022 US midterm elections for 786 locations across the United States. Focusing on state, county, and local government domains that provide locality-specific information, we study not only the extent to which these sources appear in organic search results, but also the extent to which these sources are *correctly targeted* to their respective constituents. We label misalignment between the geographic area that non-federal domains serve and the locations for which they appear in search results as *algorithmic mistargeting*, a subtype of algorithmic misjudgement in which the search algorithm targets locality-specific information to users in different (incorrect) locations. In the context of the 2022 US midterm elections, we find that 71% of all occurrences of state, county, and local government sources were mistargeted, with some domains appearing disproportionately often among organic results despite providing locality-specific information that may not be relevant to all voters. However, we also find that mistargeting often occurs in low ranks. We conclude by considering the potential consequences of extensive mistargeting of non-federal government sources and argue that ensuring the correct targeting of these sources to their respective constituents is a critical part of Google’s role in facilitating access to authoritative and locally-relevant electoral information.

¹<https://elections.google/civics-in-search/>



This work is licensed under a Creative Commons Attribution International 4.0 License.

FACCT ’24, June 03–06, 2024, Rio de Janeiro, Brazil
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0450-5/24/06
<https://doi.org/10.1145/3630106.3658916>

CCS CONCEPTS

• Information systems → Web search engines; Personalization.

KEYWORDS

algorithm auditing, problematic behavior, google search audit, elections, algorithmic misjudgement

ACM Reference Format:

Brooke Perreault, Johanna Hoonsun Lee, Ropafadzo Shava, and Eni Mustafaraj. 2024. Algorithmic Misjudgement in Google Search Results: Evidence from Auditing the US Online Electoral Information Environment. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FACCT ’24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3630106.3658916>

1 INTRODUCTION

Electoral denialism—the belief that elections are rigged or that electoral fraud is rampant—in the United States (and other countries, such as Brazil and Peru) is emerging as a significant threat to democracy [5] and has roots in racial resentment and rising conspiracism [38]. Since the early 2000s, internet was regarded as a medium “uniquely receptive to conspiracism” [19], and in the past two decades, there is little doubt that it has contributed significantly to conspiracism going mainstream [34]. Following the 2020 US Presidential Election, “election rigging” conspiracies that came to be labeled as “The Big Lie” led to the violent attack to the US Capitol on January 6, 2021 [13]. While beliefs about electoral fraud are somewhat persistent [10], the 2020 electoral denialism movement was different because it was embraced by partisan state legislatures, as reported by the Brennan Center for Justice.² As a result, a large number of changes to electoral laws (some of them restricting and some of them expanding voting rights), became operational ahead of the 2022 US midterm election.³

In an environment of election denialism on one hand and changing electoral laws on the other, voters must seek accurate, up-to-date, and relevant information about elections and voting. A survey by the Bipartisan Policy Center in Fall 2022 found that voters are “most likely to look to their state and local election officials, and

²<https://www.brennancenter.org/our-work/research-reports/voting-laws-roundup-june-2023>

³<https://www.reuters.com/legal/how-new-us-laws-could-trip-up-voters-midterm-elections-2022-11-01/>

search engines” for election information,⁴ aligning with previous research [8]. Election officials similarly position themselves as authorities on election information. For example, the National Association of Secretaries of State launched a campaign ahead of the 2020 presidential elections called #TrustedInfo2020 to “promote election officials as the trusted sources of election information.” Operating as #TrustedInfo during non-federal election years and now #TrustedInfo2024 ahead of the 2024 presidential elections, NASS states, “By driving voters directly to election officials’ websites, social media pages, and materials, they will be able to receive credible, timely information on each step of the elections process.”⁵

Search engines also position themselves as trustworthy intermediaries for seeking election information. Google states that it is “committed to providing timely and authoritative information on Google Search to help voters understand, navigate, and participate in democratic processes.”⁶ In an article about their work to “support” the 2022 US midterm elections, Google’s Vice President for Trust and Safety writes, “Our work is centered around connecting voters to the latest election information ...”⁷ and highlights “features that show data from nonpartisan organizations” that work to “connect voters with accurate information about voter registration and how to vote,” referring to Google Search election information panels that contain aggregated information from *Democracy Works*⁸ about where and how to vote.

Study Motivation: Given this background, three assumptions motivate the current study: (1) voting-eligible US citizens who search for electoral information using Google’s search engine have the right to receive accurate and relevant search results, regardless of where they live; (2) similar to Ballatore et al. [2], we acknowledge a spatial dimension to searching and search results, in which users searching for information related to elections and voting deserve *results relevant to their geographic area*, and 3) given the size of the United States and the varying voting laws per state (most of them updated ahead of the 2022 elections), receiving locality-specific information from authoritative sources is a reasonable expectation. Following these assumptions, we proceed to audit Google Search to determine the extent to which government-maintained web domains, as authoritative sources on election information, are represented in the online environment of electoral information, during the 2022 US midterm elections. Specifically, we analyze where, how often, and how accurately state, county, and local government domains appear in Google’s organic results for election-related queries collected in 786 locations across the United States between October and November 2022, in a dataset of 3.45 million search engine result pages (SERPs). Drawing parallels to studies of geo-targeting errors in advertisements [4], we label the misalignment between the geographic area that a non-federal government web domain serves (such as a state or county) and the location for which it appears in search results as an instance of *mistargeting*, in which the search algorithm targets locality-specific information to a user

in a different (incorrect) location. In this paper, we answer two related research questions:

- **RQ1:** To what extent are government web domains represented in the online environment of electoral information about the 2022 US Midterm Elections, as mediated by Google Search?
- **RQ2:** To what extent are state, county, and local government web domains correctly targeted to their respective constituents?

Findings and Contributions: We find that 40.6% (1,848/4,556) of domains that appear among organic results are government sources, and that these government sources constitute nearly 40% of all organic results occurrences. Despite this proportionality, we find that 71% of the occurrences of state, county, and local government sources were mistargeted, with a handful of locality-specific government domains appearing disproportionately often among organic results, despite providing locality-specific information that may not be relevant to all voters. Despite extensive mistargeting, we find that correctly targeted government sources usually appear in higher ranks, potentially mitigating possible harm inflicted by mistargeting. We conclude by considering how mistargeting factors into Google’s role in facilitating access to accurate, authoritative, and locally-relevant electoral information to voters.⁹

2 LITERATURE REVIEW

2.1 Algorithm Auditing and Problematic Behavior

Online platforms and other sociotechnical and algorithmic systems are neither neutral nor unbiased entities. As Gillespie [15] argues, the presentation of platforms as open or neutral parties obscures the consequences of deliberate choices platforms make as to what content is published and how content is shown. Through a systematic literature review, Bandy organizes potential problematic behavior of algorithmic systems into four main categories (a) discrimination based on identity or socioeconomic status; (b) distortion or obscuration of reality, (c) exploitation of sensitive user information, and (d) misjudgement and inaccurate classifications [3]. Potential consequences of algorithmic systems become more wide-reaching when considering the interplay between the technical components of algorithmic systems and the social systems in which they exist; one such theme of sociotechnical harm is social systems harms, including how algorithmic systems contribute to the spread of misinformation, disinformation, and malinformation (informational harms) and how algorithmic systems impact governance, democracy, and civil liberties (political and civic harms) [37].

Sandvig et al. [35] propose algorithm auditing as a way of studying problematic behavior in Internet platforms and algorithmic systems, particularly problematic behavior that is subtle or unobservable via one-off instances. To date, search engines are among the most audited algorithmic systems, often for **discrimination** or **distortion** problematic behavior. Prior research has focused on personalization of search results, including personalization based

⁴<https://bipartisanpolicy.org/blog/new-survey-data-election-information/>

⁵<https://www.nass.org/initiatives/trustedinfo>

⁶<https://elections.google/civics-in-search/>

⁷<https://blog.google/outreach-initiatives/civics/our-ongoing-work-to-support-the-2022-us-midterm-elections/>

⁸Democracy Works describes itself as “a nonpartisan, nonprofit organization that collaborates with election officials, leading tech platforms, and world-class partners to drive voter access and participation”. Source: <https://www.democracy.works/about>.

⁹Audit and analysis materials, including query phrases, locations, and scripts, are available on GitHub: <https://github.com/credlab/facct24>

on logged-in status [17], browsing history [16, 23], and location [22]; composition of search results, including the presence and content of snippets [18], Top stories [25, 39], and other non-organic panels [33]; and prominence of different types of sources among organic results [6, 14, 26]. The pervasive concern in such studies is that the algorithmic choices are leading to a distorted view of reality. Meanwhile, studied to a lesser extent is problematic behavior related to **algorithmic misjudgement**. As Bandy defines it [3], misjudgement occurs when an algorithm makes incorrect predictions or classifications, and misjudgement often leads to other types of problematic behavior. Algorithmic misjudgement has been studied in the criminal justice context (eg. recidivism prediction [9]), as well as the advertising context, where algorithms incorrectly infer information about users [3]. One implication of advertising algorithm misjudgement is that purchased advertisements do not reach intended audiences. For example, Bandy and Hecht [4] find that geotargeting errors are common in Google Display Network, with users often seeing advertisements not intended for their ZIP code. However, they find that severity of such geotargeting errors were minor in the context of advertising goals, with most users living in the intended county, designated market area, or state [4]. With this study, we contribute to the literature on algorithmic misjudgement in the context of location-sensitive electoral information.

2.2 Politics and Search Engines

Much of previous work on auditing search engines is rooted in a political context, motivated by the potential informational, political, and civic harms that could result from biased search. Epstein and Robertson [11] hypothesized, for example, that biased search results can impact electoral outcomes by shifting voter preferences based on manipulated rankings, as studied in controlled, lab-based experiments [11, 12]. However, algorithm audits in the wild find little evidence for this concern [32, 33]. Findings that support the presence of filter bubbles in Google Search and Google News are also limited [6, 16, 30, 32, 40]. Trielli and Diakopoulos [40] in particular find that while members of different ideological groups use different queries to search for political information, such partisan queries do not lead to a filter bubble effect; rather, they observe a “mainstreaming effect”, where highly similar results appear even for diverging partisan queries. Moreover, Robertson et al. [31] find that user choice plays a larger role in engagement with partisan or unreliable news than algorithmic curation, further evidence against the idea of partisan filter bubbles.

Given the focus on potential bias or filter bubbles, the query phrases used in most political information audits are intended to capture such phenomena with respect to either the candidates or issues central to an election. Examples include queries related to Donald Trump’s inauguration [33], candidate names [26], partisan terms about immigration [23], socio-political themes, [6], and crowd-sourced queries from users about senate candidates [40]. Most of the audits that focus on candidate names are conducted during election periods, including the 2016 US Congressional Elections [27], the 2018 US Midterms [26, 40], and 2020 US Presidential Primaries [21, 41]. Diakopoulos et al. [7] and Lurie and Mulligan [24] similarly study various information panels that appear in

Google search results when searching for candidate or representative names, with the latter finding that featured snippets are often likely to mislead searchers by presenting incorrect or incomplete information about a representative. Lurie and Mulligan contextualize this behavior as “algorithmic breakdown”, and identify place name ambiguity as an important contributor to such breakdown, wherein the algorithmic system is unable to discern a particular location (and the respective locality-relevant information, like the representative of a congressional district) from an ambiguous place name, such as a county and city that share the same name. Diakopoulos et al. [7] additionally study the composition of organic results for queries of Democrat and Republican candidates ahead of the 2016 US presidential election, finding that official sources, such as campaign websites, were clustered at the top of organic results in high ranking positions. Most of these studies focus on query phrases selected by researchers in various ways. Mustafaraj et al. [28] discuss that voters have information needs not previously captured in such audits. Building upon prior literature, we create an election-relevant query phrases dataset composed of popular queries from Google Trends and queries formulated by users as reported in [28].

2.3 Source Geoprovenance

As Ballatore et al. [2] describe, search engines “increasingly mediate not just information but also spatial knowledges and experiences,” as users rely on search engines to inform them about destinations, where to shop, or other ways to interact with their “lived everyday geographies” - including participating in civic processes, such as voting. Moreover, *where* information originates from, particularly if information originates from local or non-local sources, can impact how people learn about and interact with places near them, to the extent that local community members and non-community members view their locality differently [36]. The geographic origin of information, or geoprovenance [36], has largely been studied in places where volunteered geographic information (e.g. geotagging) is accessible, such as Wikipedia [36] and Twitter [20]. Ballatore et al. [2] study the extent to which different country variants of Google Search direct users to locally produced information about places, the first such audit study of search engines. Using a geoprovenance inference algorithm developed by Sen et al. [36], they identify the geoprovenance of a URL at the country level and compare URL geoprovenance to the country of search, finding a greater degree of local content for wealthy and well-connected countries. This study uses the concept of geoprovenance of a URL to study algorithmic misjudgement in search results. We audit not only the extent to which government sources are present in search results, but also the extent to which state, county, and local government sources are shown to the geographic areas which they serve.

3 DATA COLLECTION

An audit comprises a set of inputs and the corresponding outputs. Our inputs are the query phrases and the geographical locations where Google searches were performed. The outputs are the search engine results pages (SERPs) generated for each query-location pair. We describe how we selected the queries and locations and how we performed the audit.

3.1 Selection of Queries

In line with previous research about centering voters in algorithm audit methodology [28], we developed a list of queries aimed at capturing the diverse information needs of voters who use Google to search for election information.

In early October 2022, we collected Top and Rising related terms from Google Trends for the seed queries “midterm elections”, “voting near me”, and “ballot”, as well as from the Google Trends category “United States election, 2022 - General election”. The location of Google Trends was the United States. From the received results, we excluded terms that were irrelevant to the topic of elections; terms that included the names of political figures or candidates; and any terms that included specific states, state abbreviations, or localities so that queries generalize across locations.

A review of the query list after these steps revealed a lack of general queries about candidates and key voter issues, as well as a lack of biased and/or partisan queries. To fill these gaps, we augmented our query list by manually selecting queries formulated by eligible US voters included in the Voter Searches dataset in [28]. Lastly, a few additional seed queries were added related to the topic of abortion, which was one of the top voting issues in the 2022 election.¹⁰ These processes resulted in a final list of 161 queries, of which 74 were from Google Trends.

Queries were reviewed qualitatively by one author to develop three major themes of a) *Voting Logistics* (in which queries are primarily about logistical information related to voting); b) *Candidates and Issues* (in which queries are primarily about candidates and candidates’ opinions, stances, and/or leanings)¹¹ and c) *General Election Topics*, (where queries are primarily general inquiries about the midterm election, such as what midterms are and when they occur). Three labelers then independently labeled all queries into these three categories. The majority label was used as the final label per query. Table 1 shows the query distribution by category and some examples for each category.

3.2 Selection of Locations

We selected locations for data collection to correspond to U.S states and congressional districts whose House or Senate midterm races were classified as toss-ups, leaning Democrat, and leaning Republican, according to the Cook Political Report.¹² We identified 30 congressional districts with competitive House races across 23 states, as well as 10 states with competitive Senate races; 3 states had only competitive Senate races, making for a total of 26 unique states.

For house races, because congressional district borders do not align with county borders, we first manually compiled a list of counties within the relevant congressional district, according to

¹⁰<https://www.pewresearch.org/politics/2022/08/23/abortion-rises-in-importance-as-a-voting-issue-driven-by-democrats/>

¹¹The category *Candidates and Issues* includes 9 queries about abortion, a key voting issue in the 2022 US midterms. While these 9 queries could constitute their own category, we found that separating “Candidates” and “Issues” queries did not lead to meaningful changes in any results with respect to mistargeting rates, as government sources appear very minimally for these 9 queries. Thus, it was sufficient to lump these queries in with queries about candidates and candidate stances. Future work may study mistargeting for a more expansive set of queries representing different key voting issues.

¹²<https://www.cookpolitical.com/ratings/senate-race-ratings> and <https://www.cookpolitical.com/ratings/house-race-ratings>

Wikipedia,¹³ including only the counties that were described as having most or all of their localities within the congressional district. Next, we identified the associated Federal Information Processing System (FIPS) Codes for each county¹⁴ and used these codes to filter 2020 U.S. Census data¹⁵ to include localities from the relevant counties only. Additionally, we filtered locations to only include cities, towns, and townships, as identified on the Census.

To sample locations, we identified the median population of possible locations, and selected (up to) 15 locations below the median population and (up to) 15 locations above the median population, as well as the location with the highest population. Some states (for example, Nevada, which had only a competitive Senate race) had fewer than 31 cities, towns, and townships from the relevant counties. This process resulted in 872 locations across 26 states. Google Places API was used to get latitude and longitude coordinates for each of the locations. This information is needed to change the location of the Chrome browser to the desired location for the audit, as shown next.

3.3 Scraping Audit and SERP Validation

We use the Python-scripted Selenium web driver to automate the process of querying Google Search, changing the geolocation, and saving the SERPs. For each query, our custom script opens a blank-slate of the Chrome browser, enters the query, and scrolls to the bottom of the page to update the location using latitude and longitude coordinates (Figure 1). The page reloads, displaying the results for the new location, and the entire HTML page is stored.




Figure 1: The script automatically updates the location with latitude and longitude coordinates so that the search location is “from your device” rather than IP address. Previous audits [22] have tested and confirmed that search results are personalized based on provided coordinates, rather than the IP address.

SERPs were collected for 30 days, starting on October 18, 2022 and ending November 18, 2022. Data were not collected on November 14 and 15. Data were collected using 12 desktop computers, with 2 to 3 states collecting per computer. Because of the large amount of locations, data collection started each day at 2am EST and continued until the evening. The order of locations was not changed daily, so SERPs per query were collected at approximately the same time within each location, and locations were completed sequentially. Due to random hardware or software failure, some SERPs failed to collect. 3,838,128 SERPs were collected in total (92.3% collection rate).

To validate that locations were changed correctly, we extracted the text of the location box on each SERP and compared it to the intended location. The validation process revealed two types of

¹³Eg. https://en.wikipedia.org/wiki/California's_27th_congressional_district

¹⁴<https://transition.fcc.gov/oet/info/maps/census/fips/fips.txt>

¹⁵Dataset titled “Incorporated Places and Minor Civil Divisions Datasets: Sub-county Resident Population Estimates: April 1, 2020 to July 1, 2021 (SUB-EST2021)” from <https://www.census.gov/data/tables/time-series/demo/popest/2020s-total-cities-and-towns.html>

Query Category	Example Queries	n
Voting Logistics	voting early near me, sample ballot for my county, what district am i in	68
Candidates and Issues	candidate midterm info, candidates for my district, candidate opinions	50
General Election Topics	2022 midterm, mid term elections, should i vote,	43

Table 1: The audit used 161 queries that cover three major themes related to logistics of voting, candidates, and general inquiries about midterms. The query categories were assigned by three independent labelers.

errors: (1) the location of a SERP failed to update properly for unknown reasons, and (2) the latitude and longitude coordinates from the Google Places API did not exactly match the actual intended location. To resolve these issues, we remove all 86 locations that had errors and conduct analysis on 3,455,822 SERPs from the remaining 786 locations across 25 states.

4 METHODS

Since our research questions revolve around the presence of government-maintained websites in SERPs, we describe in this section how we extracted their occurrences, labeled them according to the government type and geographic area, and defined when they are mistargeted with respect to the search location.

4.1 Extracting domains

SERPs are composed of organic results (the blue links) and non-organic results that include top stories, videos, knowledge panels, and, during election season, election information panels. This study focuses on organic results only, and in particular the domains that occur in organic results. 22,750 unique links were parsed from 31,310,263 organic results. Once all such links were extracted from all SERPs, they were cleaned to their root domain; however, given our focus on various government sources, in some cases we maintain the subdomain and avoid aggregating to the root, particularly if the subdomain is a county or town website. For example, while `clintonco.illinois.gov` is a subdomain of the state government website `illinois.gov`, we maintain the county government website. 4,556 unique domains were identified. For each SERP, we additionally record the rank of the organic results. In determining rank, we do not account for the presence of non-organic panels that may be on the page; rank is based only on the organic results.

4.2 Identifying Government Websites

To identify government websites, we make use of two repositories maintained by the US government: (1) a list of nearly 10,000 .gov domains maintained by the Cybersecurity and Infrastructure Security Agency (CISA),¹⁶ and (2) a list of 9,200 government-maintained websites outside of the .gov and .mil domains, maintained by the US General Services Administration (GSA).¹⁷ Both datasets include information about the type of government domain, such as state, county, or local, the agency that maintains the domain, and some geographic information about the organization that the domain serves (such as the state or county). 726 domains were found in the CISA repository and 716 domains from the GSA repository. Because

of how links were cleaned and aggregated, some .gov websites, particularly those corresponding to towns, cities, or counties, were not included in these repositories because they were subdomains of state root domains. We manually examined remaining websites ending with .gov and included any that correspond to a town, city, or county government, manually filling in relevant information about government type and geographic area.

Additionally, we identified some county and city websites using the .us domain that were not included in the GSA repository; because the .us domain is not strictly governmental, we manually reviewed all websites with the .us domain and included only those which corresponded to a county or city. 362 domains were added by manual review. In total, 1,848 websites were identified, which we call “government domains” or “government sources”. The distribution of domains across government types is shown in Table 2.

4.3 Measuring Mistargeting

We label misalignment between the geographic area that non-federal domains serve and the locations for which they appear in search results as *mistargeting*, a type of algorithmic misjudgement in which the search algorithm targets locality-specific information to users in different (incorrect) locations. Mistargeting can be studied at varying degrees of granularity.

4.3.1 Domain level. First, we examine mistargeting with respect to government domains and where they appear. To do so, we first identify how many government domains that appeared in organic results are from states in which we did not collect data; these are domains which are always mistargeted. Next, for domains originating in states for which we collected data, we calculate the proportion of times the domain appeared in the appropriate state and the proportion of times the domain was mistargeted (at the state-level). This gives insight into whether there are some locality-specific government domains that Google shows often across locations and are frequently mistargeted.

4.3.2 Aggregate Organic Results. We also study mistargeting by looking at all organic results in aggregate and by parsing differences in mistargeting by query categories, domain type, and state. Every organic result is labeled as to whether it is a government domain, and if so, we compare the domain’s geoprovenance (the geographic area it serves, such as a state, county, or city) to the search location. This gives us a more precise understanding of mistargeting than at the domain level, in which we only compare states. Organic results are labeled as either correct or mistargeted. Mistargeted results for county and city domains are also labeled as either in-state, if they serve a different location in the same state as the search location, or out-of-state, if they do not, which functions

¹⁶<https://github.com/cisagov/dotgov-data/blob/main/current-full.csv>

¹⁷https://github.com/GSA/govt-urls/blob/main/1_govt_urls_full.csv

Government Domain Type	Examples (geoprovenance)	n
Federal	whitehouse.gov (United States)	44
State	az.gov (Arizona)	160
County	cookcountyclerkil.gov (Cook County, Illinois)	697
Local	cityofmaywood.com (Maywood, California)	806
Representative	bishop.house.gov (Georgia, 2nd District)	137
Senator	portman.senate.gov (Ohio)	3
Native	nativefederation.org (Alaska)	1

Table 2: 1,848 government sources were identified from the data set of 4,556 domains (40.6%). Government sources were labeled for type and geoprovenance. We focus our analysis on state, county, and local government sources, as senator, representative, and Native sources show up minimally, composing less than 0.1% of all organic results.

as an indicator for the severity of mistargeting. For each query category, and across all queries, we calculate the proportion of organic results composed of *all* non-federal government sources and the proportion of results composed of *correctly targeted* government results. Mistargeting rate is calculated by subtracting the proportion of correctly targeted from the total proportion, and dividing by the total proportion. We similarly calculate the proportion of results that are mistargeted but in-state, which gives insight into the nature of mistargeted results.

4.3.3 SERP Level. Lastly, we quantify mistargeting at the SERP level by calculating descriptive statistics about the number of government sources that appear per SERP, as well as descriptive statistics about the rank of correctly and incorrectly targeted government results. Since users heavily rely on ranking of search results [29], this informs the severity of misjudgement with respect to a user.

5 RESULTS

5.1 RQ1: To what extent are government sources represented in the online environment of electoral information?

Of the 4,556 domains that appeared in organic results, 1,848 were identified as a government domain (40.6%). Overall, links belonging to these 1,848 government domains compose 39.97% of all organic results captured in the 3.45 million SERPs (31.3 million organic results) analyzed. Non-federal government domains compose 32.89% of all organic results. However, the extent to which government domains appear depends on the type of query, as well as the type of government source as shown in Table 3. In particular, government sources constitute 60.55% of organic results for *Voting Logistics* queries, with state government sources composing half of such appearances. In contrast, government sources are underrepresented among organic results for *General Election Topics* and *Candidates and Issues* queries, composing just 26.36% and 23.15% of organic results, respectively.

Across all queries, local government sources are minimally represented among organic results. The prominence of state-level government sources compared to other government sources may be explained by the key roles that state-level actors play in elections: while the composition of election administrations vary by state, elections in each state are overseen by a chief election officer (such

as a secretary of state, a governor, or a board), and their responsibilities include enforcing election laws, administering a statewide voter registration database, and other duties related to election logistics.¹⁸ However, their dominance is striking. While state government websites make up 3.5% of all websites (160 out of 4,556), they compose 30.06% of *Voting Logistics* results and 19.67% of overall results.

On a SERP level, *Voting Logistics* queries had a median of 6 government results, with medians of 3 state-level and 2 county-level government sources, aligning with the aggregate proportions in Table 3. In contrast, *General Election Topics* queries had a median of 2 government source per page and *Candidates and Issues* queries had a median of 1.

5.2 RQ2: To what extent are state, county, and local government web domains correctly targeted to their respective constituents?

5.2.1 Domains. At the domain level, 174 of the 1,804 (9.6%) non-federal government sources represent a state for which we did not collect data; in other words, these sources were always mistargeted. This includes 71 local-level, 53 county-level and 48 state-level government domains, along with 2 senator websites. Appearances from these domains constitute 13.65% of appearances of non-federal government sources, or about 4.4% of all organic results. The remaining 1,630 non-federal government sources match, at the state level, the states for which we collected SERPs. 1,268 of these domains always show up in the appropriate state (although not necessarily in the appropriate county or locality, for county and local sources). While this is the majority of non-federal government domains (70.3%), only 4.96% of all organic appearances by non-federal government sources can be attributed to these domains.

Table 4 shows the 10 non-federal government domains which make up the highest proportion of organic results, and together compose 35.96% of appearances for non-federal domains, or 11.8% of organic results overall. Because the number of locations for which data was collected varied per state, it is not appropriate to directly compare the frequency with which these domains occur to each other. However, it is clear that these domains suffer from high rates of mistargeting, with many appearing more often in other states than the states for which they serve. Notably,

¹⁸<https://www.ncsl.org/elections-and-campaigns/election-administration-at-state-and-local-levels>

	Voting Logistics	General Election Topics	Candidates and Issues	All Queries
% Federal	4.85	9.1	8.44	7.08
% State	30.06	11.99	11.87	19.67
% County	22.91	4.75	2.66	11.86
% Local	2.62	0.52	0.18	1.31
% Total excl. Federal	55.7	17.26	14.71	32.89
% Total	60.55	26.36	23.15	39.97

Table 3: Government domains compose 39.97% of organic results overall, with state and county government domains constituting the majority of government domain appearances, while local government domains are minimally represented in organic results. The degree to which government domains appear depends largely on the query type: government sources compose 60% of organic results for *Voting Logistics* queries but only around a quarter for *General Election Topics* and *Candidates and Issues*.

Domain	Type	State Domain Serves	% Appearances Correct State	% Appearances Other States	Collected data in state?
pa.gov	State	Pennsylvania	55.25	44.75	Yes
ca.gov	State	California	61.98	38.02	Yes
sos.state.tx.us	State	Texas	18.86	81.14	Yes
wi.gov	State	Wisconsin	26.30	73.69	Yes
wakegov.com	County	North Carolina	5.78	94.52	Yes
lavote.gov	County	California	25.79	74.21	Yes
ncsbe.gov	State	North Carolina	34.45	65.55	Yes
georgia.gov	State	Georgia	28.28	71.71	Yes
maryland.gov	State	Maryland	0	100.00	No
votebrevard.gov	County	Florida	5.83	94.17	Yes

Table 4: These 10 domains constitute 35.96% of all non-federal government organic results, which is 11.8% of all organic results. While collecting from more locations in certain states (such as Pennsylvania and California) increases the number of times these domains occur, all of these domains appear often in states for which they do not serve, indicating frequent mistargeting. maryland.gov is one of 174 non-federal sources that appear despite no data being collected for its respective state.

wakegov.com, a county government site for Wake County, North Carolina, and votebrevard.gov, the Supervisor of Elections site for Brevard County, Florida, are mistargeted over 94% of the time. Overall, 84 domains were mistargeted more often than they were correctly targeted, and they compose 44% of non-federal organic results (14.45% of all organic results). In contrast, 245 domains are correctly targeted more often than they are mistargeted, and they compose 41% of non-federal organic results (13.5% of all organic results). Thus, while the majority of non-federal government sources are always correctly targeted at the state-level, these sources make up only a minority of all organic results; a relatively small subset of domains are mistargeted more often than they are correctly targeted, and these mistargeted results show *frequently* among organic results.

5.2.2 Mistargeting by Query Category and State. Table 5 shows the percent of all organic results that are correctly targeted to the appropriate locality, and Table 6 shows the mistargeting rates for each state, county, and local government sources for different query categories. Across all SERPs, 71% of all appearances of locality-specific government websites in organic results were mistargeted. Mistargeting is highest among *Candidates and Issues* queries at 78.61%. *Voting Logistics* queries, which see the highest proportion of government results by far, have the lowest mistargeting rate of the 3 query categories, albeit still 69.4%. Across all queries, county-level govern-

ment sources suffer from the highest levels of mistargeting, at 91.9%. State-level government domains have the lowest mistargeting rate of the locality-specific sources at about 58.3%. Mistargeting of state-level government sources is lowest for *Voting Logistics* queries at 51.75%, notable since state-level sources are the most commonly appearing government source in this query category. Lastly, local government sources suffer from a mistargeting rate of 76%, unfortunate considering the minimal appearances of local government sources in organic results.

5.2.3 Mistargeting Severity. While mistargeting rates are high, particularly for county and local government domains, it is useful to distinguish between error rate and error severity. Error severity can be assessed in two ways: 1) the rank of mistargeted organic results, 2) how close mistargeted results are to the correct location.

With respect to rank, while only 28.8% of non-federal government sources are correctly targeted, these correctly targeted results generally appear in higher ranks among organic results than mistargeted results (Figure 2). The median rank of a correctly-targeted government result was 3, while the median rank of a mistargeted government result was 7. At the SERP level, 5.6% of SERPs analyzed had only correctly targeted results, 19.9% of SERPs had only mistargeted results, 43.7% of SERPs had both correctly-targeted and mistargeted government results, and the remaining 30.7% of SERPs had no subnational government source. Thus, just under half of

	Voting Logistics	General Election Topics	Candidates and Issues	All Queries
% State Correct	14.5	4.1	3.0	8.2
% County Correct	1.89	0.42	0.11	0.96
% Local Correct	0.63	0.13	0.03	0.31
% Total Correct	17.02	4.65	3.14	9.47

Table 5: Percent of all organic results that are correctly targeted to the appropriate locality for state, county, and local sources; the denominator here is the number of all organic results (about 31 million). While 32.89% of all organic results can be attributed to locality-specific government sources, only 9.47% of organic results are composed of correctly targeted government sources.

	Voting Logistics	General Election Topics	Candidates and Issues	All Queries
% State Mistargeted	51.75	65.8	74.73	58.29
% County Mistargeted	91.73	91.12	95.7	91.94
% Local Mistargeted	75.92	75.06	81.79	76.08
% Total Mistargeted	69.4	73.05	78.61	71.18

Table 6: Percent of government results that were mistargeted. Mistargeting rate is calculated by subtracting the proportion of correctly targeted (Table 5) from the total proportion (Table 3), and dividing by the total proportion. Overall, 71% of locality-specific government organic results were mistargeted, with the mistargeting rate being highest for county government sources at nearly 92%.

SERPs (49.3%) had at least one correctly targeted result. Moreover, for the 43.7% of SERPs that contained both correctly-targeted and mistargeted government results, the first correctly-targeted result appeared above the first mistargeted result 83.9% of the time. Lastly, a mistargeted government result appeared in the first organic position for only 3.3% of SERPs analyzed, while a correctly-targeted government result appeared in the first organic position for 19.3% of SERPs.

Given that users tend to focus on search results in the top positions [29], these findings suggest that while mistargeting rates are high, their severity, within the context of impact on a user, may be low. Additionally, these findings may indicate that the search algorithm prioritizes non-federal government sources in ranked results only when there is a high likelihood that the source geoprovenance matches the user’s locations; otherwise, non-federal government




Figure 2: While only 28.8% of non-federal government domains are correctly targeted, these results usually appear in higher ranks among organic results. On the 43.7% of SERPs that have both correctly-targeted and mistargeted results, the correctly targeted result appeared above the mistargeted result 83.9% of the time. Thus, while mistargeting rates are high, their severity, within the context of impact on a user, may be low, given that users often rely on top ranking results.

sources appear in lower ranked positions, and other content (such as a news or non-profit organizations) is given the top ranks.

With respect to the nature of mistargeted results, a high percentage of mistargeting for local government sources can be explained by the appearance of sources that serve other localities in the same state as the search location. While the overall mistargeting rate for local government sources is 76%, 64% of local government results are for other localities in the correct state. Figure 3 visualizes this phenomenon; in most states, nearly all of the mistargeted local government results are from other local government sources in the respective state.

A similar phenomenon with targeting county-level government sources is observed, though to a lesser degree and with more variation between states. These findings imply a *lack of precision* with targeting locality-specific government sources to the appropriate audience when using general queries that do not specify location information. While a similar lack of precision has been observed with




Figure 3: Much of the mistargeting of local government sources (and county government sources, to a lesser extent) can be explained by the appearance of sources that serve other localities in the same state as the search location. This indicates a *lack of precision* with targeting locality-specific government domains to their specific audience when relying only on geocoordinates to represent user location.

geotargeting of Ads on Google [4], the consequences of this lack of precision is greater in the context of “low information voters”¹⁹ searching for election information as it relates to their locale, given that other county or locality information is likely not be applicable to them.

6 DISCUSSION, LIMITATIONS, AND FUTURE WORK

6.1 Discussion

This paper studies the extent to which local, county, and state government sources appear in election-related search results, as well as the extent to which these sources are correctly targeted to the locations which they serve. We label misalignment between the geographic area that non-federal domains serve and the locations for which they appear in search results as *mistargeting*, a subtype of algorithmic misjudgement in which the search algorithm targets locality-specific information to users in different, incorrect locations.

6.1.1 Overrepresentation of Some Locality-Specific Government Sources. With respect to our first research question, we find that 40.6% (1,848/4,556) of websites that appeared in organic results are government sources, and combined, these sources compose 39.97% of all organic results. This might look like an excellent outcome, given the almost perfect matching between the proportion of domains and the proportion of their occurrences in the search results. However, we find that that some locality-specific government domains are over represented and frequently mistargeted in organic results. About 1/3 of all occurrences of non-federal government sources can be attributed to ten domains, with these domains frequently appearing in locations for which they do not serve (see Table 4). This phenomenon is reminiscent of source concentration and mainstreaming effects found in previous work for news sources [39, 40]. While it might not be harmful to an individual to read national news from New York Times more than from Baltimore Sun (although, collectively, source concentration has led to the so-called ‘news desert’ phenomenon [1]), for elections, all out-of-state web domains might be irrelevant to a voter, when it comes to their voting information.

6.1.2 Extensive Mistargeting in Low Ranks. Furthermore, we find extensive mistargeting of locality-specific government sources: 71.18% of all occurrences of non-federal government sources were mistargeted, with county-level government sources having the highest mistargeting rate at 91.94% (Table 6). However, we also find that correctly targeted results generally appear in higher ranks among organic results than mistargeted results (Figure 2), and for the 43.7% of SERPs that contained both correctly targeted and mistargeted government results, the first correctly targeted result appeared above the first mistargeted result 83.9% of the time. In the absence of user behavior data, these findings are encouraging that the severity of harm to a user due to mistargeted results may be low, since low-ranking mistargeted results are usually accompanied by high-ranking correctly targeted results.

6.1.3 Government Sources, SERP Composition, and Consequences of Mistargeting. The study of mistargeted government results in election SERPs is motivated by how voters and election officials position election officials and government sources as important, trusted conveyors of election information.²⁰ It is not our aim to claim that government sources are the most (or only) relevant type of source for electoral information. Indeed, it is conceivable that higher quality, non-government sources containing relevant information may be more beneficial to a user than lower quality, correctly-targeted government sources. Exploring the extent to which government sources are quality sources for electoral information is beyond the scope of this paper. However, regardless of whether government sources should be considered as the best and highest quality, our audit shows not only that locality-specific government sources *are* showing in SERPs for different types of queries (even if minimally), but also that these results are often not showing up *in the correct places*. Questioning whether a locality-specific government source is correctly targeted to its constituents perhaps precedes questions about whether this source type is most favorable for a user. At best, mistargeting of government sources among organic search results constitutes wasted space on a SERP that could otherwise be used to provide voters with relevant electoral information. At worst, mistargeting constitutes a failure in Google’s ability to facilitate (easy) access to accurate, authoritative, and locality-specific electoral information to voters who turn to the search engine with their election-related information needs.

Google states that it is committed to providing voters with authoritative information about US elections. Our focus on mistargeting of government sources among organic results reveals room for improvement in this regard. Showing locality-specific government sources in organic results *and* ensuring that these sources are correctly targeted to their respective constituents are necessary criteria for ensuring that Google Search can appropriately provide users across different locations access to authoritative *and* locally-relevant election information.

6.2 Limitations

There are limitations to our study. First, we focus exclusively on government web domains in search results, but in reality there are other types of information that users encounter in search results, including results from other types of sources and curated information panels on the SERP. During the 2022 US midterm elections in particular, Google populated election-related SERPs with election information panels, many of which contained curated, state-specific information; an analysis of these panels is necessary to get a complete understanding of how locality-specific electoral information is surfaced among search results, but this is left to future work. Additionally, our audit was conducted when there were around 10 organic results per SERP, though since December 2022 (after our data collection), Google has moved to “continuous scrolling”, which loads up to 60 results at once.²¹

¹⁹https://en.wikipedia.org/wiki/Low_information_voter

²⁰See Introduction; <https://bipartisanpolicy.org/blog/new-survey-data-election-information/> and <https://www.nass.org/initiatives/trustedinfo>

²¹<https://www.theverge.com/2022/12/6/23495973/google-search-desktop-continuous-scrolling>

Second, we rely on geocoordinates to serve as a proxy for locations of real users and do not study how including location information in the search query (e.g. “voting near me Boston, MA”) impacts the composition and mistargeting of government search results. We also do not study whether targeting based on geocoordinates differs from IP-based targeting, nor do we consider how user profiles impact targeting. It is conceivable that IP-based targeting and/or user profiles lead to more location precision and thus less mistargeting among organic results. Furthermore, while the queries come from Google Trends and real users, we do not study how users actually interact with government sources in search results.

We analyze data from 786 locations across 25 states, a size which is sufficient for statistical analysis but does not capture all of the United States. Additionally, we do not analyze how mistargeting varies by search location features (including the extent to which the congressional election for a location is competitive) nor by date. Finally, this study only focuses on the United States, given that electoral laws and election authorities are different state by state, making this kind of analysis meaningful. In other countries, where there is a single election authority, the needs and assumptions might be different. Indeed, it is unclear if these results can be generalized to other countries with similar political systems to the United States. In addition to the similarity of political systems, one would have to consider the strength of a country’s online government and electoral information infrastructure when generalizing results.

6.3 Future work

One challenge that we encountered when analyzing our results was finding accurate information about government websites (as discussed in Methods). However, government websites were only 40.6% of the entire dataset of websites. The rest was composed of news organizations, non-profit civic organizations, companies, universities, individual blogs, and more. There are no high-quality and freely available databases that provide detailed information on web domains (their nature and what audiences they serve). Nuanced audits of search engine results need such information; therefore, methods to create and maintain such databases need to be developed. (We tried to use some news related curated lists for other studies, such as [14], but their overlap with our dataset was not high.)

A scraping audit is useful for controlling the inputs to the algorithmic system, so that we can systematically compare results, but it does not involve real voters with real information needs, that is, it lacks ecological validity. Currently, there are efforts underway like the National Internet Observatory²² that will enable researchers to carry out crowdsourced audits, similar to what was used in [31]. Future audits can then address both issues that are central to our audit: what queries do users formulate (and reformulate) to find information about elections and with what results do they engage? For example, do they click on Google’s election boxes, do they click on their own local/county/state websites, or do they visit non-governmental sources? Additionally, such audits might allow us to find out if Google Search is bypassed entirely and users go directly to their trusted sources.

²²<https://nationalinternetobservatory.org/>

6.4 Ethical Considerations

Our study does not include any personal data, as Google Search result pages are public data. Google handles circa 8.5 billion of searches per day,²³ so our scraping audit of 140 K daily searches provided a minimal burden. Finally, a US Federal Court has ruled that audits such as ours are legal and do not violate the Computer Fraud and Abuse Act.²⁴

ACKNOWLEDGMENTS

We are grateful to members of Wellesley Cred Lab and acknowledge funding from NSF grant ISS 1751087.

REFERENCES

- [1] Penelope Muse Abernathy. 2018. *The expanding news desert*. Center for Innovation and Sustainability in Local Media, School of Media and ...
- [2] Mark Graham Andrea Ballatore and Shilad Sen. 2017. Digital Hegemonies: The Localness of Search Engine Results. *Annals of the American Association of Geographers* 107, 5 (2017), 1194–1215. <https://doi.org/10.1080/24694452.2017.1308240> arXiv:<https://doi.org/10.1080/24694452.2017.1308240>
- [3] Jack Bandy. 2021. Problematic Machine Behavior: A Systematic Literature Review of Algorithm Audits. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 74 (apr 2021), 34 pages. <https://doi.org/10.1145/3449168>
- [4] Jack Bandy and Brent Hecht. 2021. Errors in Geotargeted Display Advertising: Good News for Local Journalism? *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 92 (apr 2021), 19 pages. <https://doi.org/10.1145/3449166>
- [5] Maxwell A Cameron. 2023. Electoral Denialism in American Democracies. (2023).
- [6] Cédric Courtois, Laura Slechten, and Lennert Coenen. 2018. Challenging Google Search filter bubbles in social and political information: Disconforming evidence from a digital methods case study. *Telematics and Informatics* 35, 7 (2018), 2006–2015. <https://doi.org/10.1016/j.tele.2018.07.004>
- [7] Nicholas Diakopoulos, Daniel Trielli, Jennifer Stark, and Sean Mussen. 2018. I vote for—how search informs our choice of candidate. *Digital Dominance: The Power of Google, Amazon, Facebook, and Apple*, M. Moore and D. Tambini (Eds.) 22 (2018).
- [8] William Dutton, Bianca Reisdorf, Elizabeth Dubois, and Grant Blank. 2017. Search and Politics: The Uses and Impacts of Search in Britain, France, Germany, Italy, Poland, Spain, and the United States. *SSRN Electronic Journal* (01 2017). <https://doi.org/10.2139/ssrn.2960697>
- [9] Grant Duwe and KiDeuk Kim. 2017. Out With the Old and in With the New? An Empirical Comparison of Supervised Learning Algorithms to Predict Recidivism. *Criminal Justice Policy Review* 28, 6 (2017), 570–600. <https://doi.org/10.1177/0887403415604899> arXiv:<https://doi.org/10.1177/0887403415604899>
- [10] Adam M Enders, Joseph E Uscinski, Casey A Klofstad, Kamal Premaratne, Michelle I Seelig, Stefan Wuchty, Manohar N Murthi, and John R Funcheon. 2021. The 2020 presidential election and beliefs about fraud: Continuity or change? *Electoral studies* 72 (2021), 102366.
- [11] Robert Epstein and Ronald E. Robertson. 2015. The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences* 112, 33 (2015), E4512–E4521. <https://doi.org/10.1073/pnas.1419828112> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.1419828112>
- [12] Robert Epstein, Ronald E. Robertson, David Lazer, and Christo Wilson. 2017. Suppressing the Search Engine Manipulation Effect (SEME). *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 42 (dec 2017), 22 pages. <https://doi.org/10.1145/3134677>
- [13] James J Fahey. 2023. The big lie: Expressive responding and misperceptions in the United States. *Journal of Experimental Political Science* 10, 2 (2023), 267–278.
- [14] Sean Fischer, Kokil Jaidka, and Yphtach Lelkes. 2020. Auditing local news presence on Google News. *Nature human behaviour* 4, 12 (December 2020), 1236–1244. <https://doi.org/10.1038/s41562-020-00954-0>
- [15] Tarleton Gillespie. 2010. The politics of ‘platforms’. *New Media & Society* 12, 3 (2010), 347–364. <https://doi.org/10.1177/1461444809342738> arXiv:<https://doi.org/10.1177/1461444809342738>
- [16] Mario Haim, Andreas Graefe, and Hans-Bernd Brosius. 2018. Burst of the Filter Bubble?: Effects of personalization on the diversity of Google News. *Digital Journalism* 6 (03 2018), 330–343. <https://doi.org/10.1080/21670811.2017.1338145>
- [17] Aniko Hannak, Piotr Sapiezynski, Arash Molavi Kakhki, Balachander Krishnamurthy, David Lazer, Alan Mislove, and Christo Wilson. 2013. Measuring

²³<https://blog.hubspot.com/marketing/google-search-statistics>

²⁴<https://www.aclu.org/press-releases/federal-court-rules-big-data-discrimination-studies-do-not-violate-federal-anti>

- Personalization of Web Search. In *Proceedings of the 22nd International Conference on World Wide Web* (Rio de Janeiro, Brazil) (WWW '13). Association for Computing Machinery, New York, NY, USA, 527–538. <https://doi.org/10.1145/2488388.2488435>
- [18] Desheng Hu, Shan Jiang, Ronald E. Robertson, and Christo Wilson. 2019. Auditing the Partisanship of Google Search Snippets. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 693–704. <https://doi.org/10.1145/3308558.3313654>
- [19] Nigel James. 2000. Militias, the Patriot movement, and the internet: the ideology of conspiracism. *The Sociological Review* 48, 2_suppl (2000), 63–92.
- [20] Isaac L. Johnson, Subhasree Sengupta, Johannes Schönig, and Brent Hecht. 2016. The Geography and Importance of Localness in Geotagged Social Media. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 515–526. <https://doi.org/10.1145/2858036.2858122>
- [21] Anna Kawakami, Khonzodakh Umarova, and Eni Mustafaraj. 2020. The Media Coverage of the 2020 US Presidential Election Candidates through the Lens of Google's Top Stories. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 868–877.
- [22] Chloe Kliman-Silver, Aniko Hannak, David Lazer, Christo Wilson, and Alan Mislove. 2015. Location, Location, Location: The Impact of Geolocation on Web Search Personalization. In *Proceedings of the 2015 Internet Measurement Conference* (Tokyo, Japan) (IMC '15). Association for Computing Machinery, New York, NY, USA, 121–127. <https://doi.org/10.1145/2815675.2815714>
- [23] Huyen Le, Raven Maragh, Brian Ekdale, Andrew High, Timothy Havens, and Zubair Shafiq. 2019. Measuring Political Personalization of Google News Search. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 2957–2963. <https://doi.org/10.1145/3308558.3313682>
- [24] Emma Lurie and Deirdre K. Mulligan. 2021. Searching for Representation: A sociotechnical audit of googling for members of U.S. Congress. *CoRR* abs/2109.07012 (2021). arXiv:2109.07012 <https://arxiv.org/abs/2109.07012>
- [25] Emma Lurie and Eni Mustafaraj. [n. d.]. Opening Up the Black Box: Auditing Google's Top Stories Algorithm. *Proceedings of the ... International Florida Artificial Intelligence Research Society Conference* 32 ([n. d.]). <https://par.nsf.gov/biblio/10101277>
- [26] Danaë Metaxa, Joon Sung Park, James A. Landay, and Jeff Hancock. 2019. Search Media and Elections: A Longitudinal Investigation of Political Search Results. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 129 (nov 2019), 17 pages. <https://doi.org/10.1145/3359231>
- [27] Panagiotis Takis Metaxas and Yada Pruksachatkun. 2017. Manipulation of Search Engine Results during the 2016 US Congressional Elections. <https://api.semanticscholar.org/CorpusID:11867830>
- [28] Eni Mustafaraj, Emma Lurie, and Claire Devine. 2020. The Case for Voter-Centered Audits of Search Engines during Political Elections. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT* '20). Association for Computing Machinery, New York, NY, USA, 559–569. <https://doi.org/10.1145/3351095.3372835>
- [29] Bing Pan, Helene Hembrooke, Thorsten Joachims, Lori Lorigo, Geri Gay, and Laura Granka. 2007. In Google We Trust: Users' Decisions on Rank, Position, and Relevance. *Journal of Computer-Mediated Communication* 12, 3 (2007), 801–823. <https://doi.org/10.1111/j.1083-6101.2007.00351.x> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1083-6101.2007.00351.x>
- [30] Cornelius Puschmann. 2019. Beyond the Bubble: Assessing the Diversity of Political Search Results. *Digital Journalism* 7, 6 (2019), 824–843. <https://doi.org/10.1080/21670811.2018.1539626> arXiv:<https://doi.org/10.1080/21670811.2018.1539626>
- [31] Ronald Robertson, Jon Green, Damian Ruck, Katherine Ognyanova, Christo Wilson, and David Lazer. 2023. Users choose to engage with more partisan news than they are exposed to on Google Search. *Nature* 618 (05 2023), 1–7. <https://doi.org/10.1038/s41586-023-06078-5>
- [32] Ronald E. Robertson, Shan Jiang, Kenneth Joseph, Lisa Friedland, David Lazer, and Christo Wilson. 2018. Auditing Partisan Audience Bias within Google Search. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 148 (nov 2018), 22 pages. <https://doi.org/10.1145/3274417>
- [33] Ronald E. Robertson, David Lazer, and Christo Wilson. 2018. Auditing the Personalization and Composition of Politically-Related Search Engine Results Pages. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) (WWW '18). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 955–965. <https://doi.org/10.1145/3178876.3186143>
- [34] Nancy L Rosenblum and Russell Muirhead. 2019. *A lot of people are saying: The new conspiracism and the assault on democracy*. Princeton University Press.
- [35] Christian Sandvig, Kevin Hamilton, Karrie Karahalios, and Cédric Langbort. 2014. Auditing Algorithms : Research Methods for Detecting Discrimination on Internet Platforms. <https://api.semanticscholar.org/CorpusID:15686114>
- [36] Shilad W. Sen, Heather Ford, David R. Musicant, Mark Graham, Os Keyes, and Brent Hecht. 2015. Barriers to the Localness of Volunteered Geographic Information. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 197–206. <https://doi.org/10.1145/2702123.2702170>
- [37] Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Roshtamzadeh, Paul Nicholas, N'Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, and Gurleen Virk. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (Montreal, Canada) (AIIES '23). ACM, New York, NY, USA, 723–741. <https://doi.org/10.1145/3600211.3604673>
- [38] Charles Stewart III. 2023. Public Opinion Roots of Election Denialism. *SSRN* (2023). <http://dx.doi.org/10.2139/ssrn.4318153>
- [39] Daniel Trielli and Nicholas Diakopoulos. 2019. Search as News Curator: The Role of Google in Shaping Attention to News Information. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3290605.3300683>
- [40] Daniel Trielli and Nicholas Diakopoulos. 2022. Partisan search behavior and Google results in the 2018 US midterm elections. *Information, Communication & Society* 25, 1 (2022), 145–161. <https://doi.org/10.1080/1369118X.2020.1764605> arXiv:<https://doi.org/10.1080/1369118X.2020.1764605>
- [41] Aleksandra Urman, Mykola Makhortyk, and Roberto Ulloa. 2022. The Matter of Chance: Auditing Web Search Results Related to the 2020 U.S. Presidential Primary Elections Across Six Search Engines. *Social Science Computer Review* 40, 5 (2022), 1323–1339. <https://doi.org/10.1177/08944393211006863> arXiv:<https://doi.org/10.1177/08944393211006863>