Errors in Geotargeted Display Advertising: Good News for Local Journalism?

JACK BANDY and BRENT HECHT, Northwestern University, USA

The rise of geotargeted online advertising has disrupted the business model of local journalism, but it remains ambiguous whether online advertising platforms can effectively reach local audiences. To address this ambiguity, we present a focused study auditing the positional accuracy of geotargeted display advertisements on Google. We measure the frequency and severity of geotargeting errors by targeting display ads to random ZIP codes across the United States, collecting self-reported location information from users who click on the advertisement. We find evidence that geotargeting errors are common, but minor in terms of advertising goals. While 41% of respondents lived outside the target ZIP code, only 11% lived outside the target county, and only 2% lived outside the target state. We also present details regarding a high volume of suspicious clicks in our data, which made the cost per sample extremely expensive. The paper concludes by discussing implications for advertisers, the business of local journalism, and future research.

CCS Concepts: Information systems~Display advertising; Geographic information systems

KEYWORDS: Geopositioning, Advertising, Algorithm Auditing, Google

ACM Reference format:
Jack Bandy and Brent Hecht. 2021. Errors in Geotargeted Display Advertising: Good News for Local Journalism? *Proc. ACM on Hum.-Comput. Interact. Vol. 5,* CSCW1, Article 92 (April 2021), 19 pages, https://doi.org/10.1145/3449166

1 INTRODUCTION

Geotargeted online advertising has driven rapid growth for large technology companies while creating a crisis for local journalism. As part of ongoing antitrust investigations in the United States, one congressman recently alleged that Google's advertising dominance is "a key factor in crushing local and regional print news" [57], since advertisers now purchase geotargeted ads from platforms rather than local newspapers. Ad revenue for newspapers has dropped from over $50 billion in 2004 to an estimated $14.3 billion in 2018 [65]; in the same time period, revenue for online advertising platforms skyrocketed from less than $10 billion to over $107 billion, dominated by Facebook and Google [46]. The resulting crisis is vividly demonstrated by "news deserts" expanding across the U.S., as plummeting revenue forces local newspapers to close [1]. Newspaper closures have demonstrable negative impacts on communities, such as decreased civic engagement [72], polarized voting behavior [21], and lack of public accountability [31].

HCI and social computing researchers have shown an increasing interest in local journalism as a means of "supporting cities, neighborhoods, and local communities" [20]. This line of work

2573-0142/2021/04 - Art92 $15.00. https://doi.org/10.1145/3449166

includes CSCW scholarship in the production of local news [81], the role of local journalists in crisis communication [18], and the distribution of journalism on new platforms [28,53,70]. But research in this area has yet to examine geotargeted advertising, the very system that disrupted newspapers' core business model and is driving the economic crisis in local journalism [57].

This study seeks to clarify the effectiveness of online geotargeted advertising, asking whether it provides accurate local advertising that warrants its disruption to local newspaper advertising. Our work is motivated by early anecdotal [14,37] and more formal [48] evidence that geotargeted advertising on platforms often fails to reach local audiences. A 2012 study by Jones et al. [48] found that Google's geotargeted search advertisements "targeted correctly in just half the cases." Since the Jones et al. study, Google has made improvements to its location tracking technologies (e.g. [56]), and there have also been numerous advancements in positioning technology (e.g. [32]). However, these new positioning technologies may become less effective due to recent policy changes – namely the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) – as well as changes to Apple iOS and Google Android that require more active consent when collecting user information. Early evidence suggests these privacy policies and software changes make geotargeted advertising less accurate [26,33,49].

More generally, ad platforms currently face heightened scrutiny for monopoly power, potential acts of misconduct, and negative societal impacts, with some journalists and researchers suggesting that "ad tech could be the next internet bubble" [24,45]. High-profile examples include Facebook inflating audience estimates in video engagement metrics [63], as well as Google selling fruitless paid search ads (where organic results produced the same returns) [10], collecting users' locations for geotargeting even when users turned off location services [15], and illegally including children in audience targeting on YouTube [76]. An advertising watchdog for the industry recently warned Facebook that "it could be denied accreditation due to deficiencies in how its reports on the effectiveness of advertising on its products," and Google has faced accreditation challenges from the same watchdog [91,93]. Many advertisers are also scrutinizing online ad platforms and reckoning with their societal impact, as illustrated by the "Stop Hate for Profit" boycott organized in 2020 [6].

To address the effectiveness of online geotargeted advertising, we draw on methods from platform auditing (e.g., [28,53,70]) and geographic positioning (e.g., [30,39,50,68]) to conduct a focused audit study of geotargeting accuracy on Google's advertising platform. We deployed Google display ads to random ZIP codes throughout the United States, linking to a survey that collected location information. The survey asked whether participants lived in the target location or had some other familiarity with it (due to work, previous residence, travel, or otherwise), which allowed us to analyze different types of geotargeting errors. More formally, we address the following research questions about geographically targeted advertisements in the United States:

- **RQ1**: How *frequent* are geotargeting errors in the Google Display Network?
- **RQ2**: How *severe* are geotargeting errors in the Google Display Network?

The results of our survey suggest that geotargeting errors on Google Display ads are quite common, but most are insignificant for advertising purposes. Despite well-known challenges related to collecting user location information [44,51,75], as well as additional challenges recruiting participants via Google Display ads, we received 111 responses from our survey, which came from 99 unique ZIP codes and 39 unique states. 41% of respondents lived outside the target ZIP code, but only 17% lived outside the target county, and just 2% lived outside the target state. Of all errors that occurred, 47% reached a neighboring ZIP code, and the median centroid-to-centroid distance from the target ZIP code to the respondent's reported ZIP code was 17 kilometers. Thus, for the purposes of advertising, geotargeted ads generally appear to reach local audiences effectively. However, we found hints of other potential deficiencies that could create barriers for local advertising, including suspicious and fraudulent clicks which made data collection extremely expensive. We conclude by discussing implications for advertisers, the business of local journalism, and future research exploring these areas.

2 RELATED WORK

2.1 CSCW and Local Journalism

For social computing researchers aiming to support groups and communities, the task of supporting journalism has become an important research topic [2]. The production and distribution of local journalism is of particular interest, as the social computing community has expressed a desire to better understand "how information technologies can be used to assist local communities" [20]. This has led CSCW scholars to study the production of local journalism and "the role of technologies in supporting cities, neighborhoods, and local communities" (as articulated by Daly et al. [20]). For example, in 2009, Hossjer and Eklundh [43] showed how the use of electronic mail was affecting news production at a local newsroom in Sweden. Vaataja and Egglestone [81] also studied the local news production process, and specifically how mobile technologies could help coordinate news reporting by delivering assignments to journalists. Also, as suggested by Daily and Starbird in their research into crisis reporting [18,19], journalism labor often involves crowdsourcing through collaborative technologies, which makes local news even more relevant to CSCW research.

This work was also motivated by social computing research that has aimed to support journalism by characterizing how technologies impact the distribution of news media. For example, Morgan et al. [60] explored how social media platforms affect news sharing practices, and some algorithm audits [53,70] have explored how search engine results treat news content. A recent study by Fischer et al. [28] focused specifically on local news, and showed that Google often excludes local news sources from search results. Overall, social computing research in this area has addressed local news production and distribution, but has not addressed geotargeted advertising, a key factor driving the economic crisis in local journalism [57].

2.2 Targeting Errors in Online Advertising

Our study joins a growing body of literature that asks: *how effective is targeted advertising*? While tracking-based targeted advertising has received critical scholarly attention on the basis of invasive data collection [17,83,90], discrimination [3,22,78], and overall effectiveness [24,13,27,35,41,58], most closely related to our work are studies of targeting errors on advertising platforms, which can occur in various forms. Studies have shown that all targeting attributes are prone to errors, including demographic information like vocation, age, and sex [80,82], as well as profile information related to user interests and preferences [7,66]. Here, we focus on errors related to location, since location-based targeting was one of the key features of online advertising platforms that threatened the business model of local journalism [1].

Our work is motivated most directly by a 2012 study from Jones et al. [48], which concluded that geotargeted advertisements on Google "targeted correctly in just half the cases." Jones et al. focused on the accuracy of medical recruitment ads in the United Kingdom, and found that reported postcodes from users often did not align with the target postcode in Google AdWords. More recently, informal surveys have also suggested geotargeting errors may be common. Search

Engine Land, a publication focused on digital marketing, conducted a survey [14] in 2015, finding that Google Analytics produced inaccurate locations for 45% of all respondents, with an average error of 145 miles among desktop users in the United States. Another informal survey in 2016 compared HTML5 geolocation to Google Analytics' geolocation, concluding that "relying on Google Analytics below regional levels is risky" [37,38].

Our work builds on the Jones et al. study and the informal surveys in a number of ways. Most basically, given its critical potential implications for local news, the Jones et al. study is an exemplary candidate for the growing efforts toward more replication work in HCI (e.g., see "repliCHI" [85]). This paper seeks to provide that replication in the context of a new, larger market and – importantly –geopositioning and advertising technology that incorporates over seven years of improvements (e.g. [8,32,54,56]). Our results show substantial differences from those in Jones et al., thus reinforcing the importance of replication. At the same time, our work is not limited to replication: it also represents an extension of prior work, expanding on Jones et al. and the informal surveys along several dimensions. For instance, our survey focuses on residential relationships, but also collects additional information about relationships people have to a given location (e.g. familiarity due to work or other reasons). We also follow a suggestion from Jones et al. for decreasing sampling bias by creating a separate campaign for each individual target location. This helped us avoid skewed sampling that can arise when aggregating all locations into a single campaign.

2.3 Why Geotargeting Errors May Occur

The literature on positioning sheds light on the causes for potential geotargeting errors; they may result from positioning system errors (e.g. IP address positioning [36,67,74]), human mobility [88], intentional user obfuscation [12], and/or policy-related challenges [26,33,64]. For example, while Google's privacy policy [94] mentions using IP addresses to infer user locations, this approach is known to have limited accuracy beyond country-level positioning, producing more errors at the state and city levels [34,40,67,74]. Some users also obfuscate their location intentionally, in order to protect their privacy [55,62,71,75]. Furthermore, recent changes in public policy (i.e. GDPR, CCPA) have made it more difficult to collect some geographic information from users [33], and software changes in Apple iOS and Google Android now require more active consent from users to collect location information. Based on early evidence, these policy changes and software changes have made accurate geotargeted advertising more challenging for platforms [26,49]. At the same time, Aly et al. [4] demonstrate that location data is a "unique and sensitive commodity for location-based services and advertising." Considering the value of location data alongside the growing challenges in collecting it, our study extends previous efforts to investigate geotargeting errors in online advertising.

3 METHODS

3.1 Survey Design and Deployment

3.1.1 Display Ads on Google. We studied geotargeting errors in the context of display ads (which show up alongside other content on websites across the internet), due to the strong interdependence between news websites and display advertising. Rather than directly selling "native" advertisements to appear on their website, local news organizations now tend to give their advertising slots to third parties, such as Google, which serve as intermediaries connecting advertisers to audiences. If an advertiser pays $1.00 when a user clicks a Google display ad on a news website, the news website receives $0.68 and Google receives $0.32 [95]. According to a 2016 analysis [13], 91% of news websites rely on this kind of third-party display advertising, compared to just 12% of all websites. Google has also shared results from internal studies [69] showing a strong interdependence between news publishers and third-party display ads.

Since Google and Facebook dominate the market for online display advertising [29], they were the two primary options for our study. Facebook does offer third-party display advertising through a “audience network” of other apps and websites, however, most of their revenue comes from display ads within their own platform [46]. Google’s display advertising spans millions of websites in the "Google Display Network," and captured $7.95 billion in annual revenue as of 2019 [29]. For this study, we decided to focus on geotargeting errors in the Google Display Network, though Facebook advertising remains a promising site for future work.

3.1.2 Home Targeting via ZIP Codes. Geotargeted advertising can be viewed as a family of different positioning problems, each corresponding to a different type of relationship between a user and a location. For example, advertisers may want to reach people who *live in* a target location (as in Aly et al. [4]), people who are *frequently in* a target location (such as what Google’s targeting offers [96]), or people who are *interested in* a target location (another Google targeting option [96]). Our study focuses on the *lives in* use case, following the “home targeted ads” scenario utilized by Aly et al. [4] which involves “a business that wants to deliver ads to people whose home is in a certain geospatial region.”

Advertisers wanting to deliver home targeted ads at a spatial scale smaller than the city can use either ZIP code targeting or pin and radius targeting.1 We chose to analyze ZIP code targeting because it is a common scale for marketing and communication efforts, with firms such as Harte Hanks, Fair Isaac Corporation (FICO), Claritas, and Nielsen using ZIP codes as a primary means of customer segmentation.2 Common types of advertising campaigns that use ZIP code segmentation include businesses licensed only in certain locations (as in the roofing example from Aly et al. [4]), local events, and businesses seeking to build regular local customers/members (e.g. grocery stores, restaurants, coffee shops, and fitness centers). Advertising platforms cater to this widespread use of ZIP codes: Google Ads (formerly AdWords) has offered ZIP code targeting since 2012 [42], Facebook since 2011 [16], and Yahoo ads since 2008 [59].

3.1.3 U.S. ZIP Code Sampling. Following many studies of geographic positioning (e.g. [11,50,86,89]), we focus on a specific geographic area, in our case the United States. Later in the paper, we note how future work may extend our approach to address our research questions in other countries.

Since data collection relied on purchased advertisements, we could not collect samples from all ZIP codes in the United States. As such, we developed a sampling scheme to estimate the overall error rate, summarized in Table 1. The initial dataset came from the U.S. Census Bureau, which linked ZIP code Tabulation Areas (ZCTAs) to ZIP codes. We then linked each ZIP code to a Nielsen Designated Marketing Area (DMA) using data from [77]. This allowed us to account for the various relationships that ZIP codes have to administrative boundaries (e.g. some ZIP codes overlap county borders).

[1] Some platforms previously allowed advertisers to draw custom boundaries around target areas, but removed the feature in 2019 after charges of discrimination from the U.S. Department of Housing, for more details see [87]

[2] For an example ZIP code-based advertising tool, see https://claritas360.claritas.com/mybestsegments/

Sample phase	Number of ZIP codes
All U.S. ZIP codes (tabulation areas)	32,907 (100%)
ZIP codes with news desert classification	29,293 (89%)
ZIP codes with population over 51	28,697 (87%)
Sampled ZIP codes	400 (1%)
ZIP codes yielding survey participants	99 (<1%)

Table 1: the steps used to sample ZIP codes.

Next, to facilitate analysis of news deserts, we removed ZIP codes that straddle news desert counties and other counties (news deserts are classified at the county level [1]). This step allowed us to explore whether "news deserts" may also be "advertising deserts" that lack both local news and local advertising, in which case our results would present a clear opportunity for the journalism industry and for future research. The step removed 11% of the set. Lastly, we removed the 2% of ZIP codes with the smallest populations: 51 people or fewer (e.g. ZIP code 99656 in Georgetown, Alaska). None of our pilot advertisements reached respondents in such ZIP codes. The final set before random sampling comprised 28,697 ZIP codes, 9,085 (32%) of which were in news deserts.

Following the suggestion by Jones et al., we aimed to collect a balanced random sample by collecting one response per ZIP code. This was to ensure sampling from more rural ZIP codes, which was a significant challenge in the Jones et al. study. We aimed to collect a total of 400 samples, and used a sampling method in the Pandas python library to select 400 random ZIP codes from the set detailed in the previous paragraph. While densely-populated urban areas have more ZIP codes than rural areas with sparse populations, the ZIP codes were non-contiguous and spaced far apart, so we presumed no substantial spatial autocorrelation. We also confirmed that the geographic distribution of survey participants aligned with the distribution of the initial random sample. The sample of 400 ZIP codes included ZIP codes from 50 states and 295 counties, and 23% of these ZIP codes were in "large central metro" urban areas according to urban-rural classifications from the National Center for Health Statistics (NCHS) [61]. Among all respondents who partially completed the survey (see 3.1.6), the represented ZIP codes (N=99) followed a similar distribution, coming from 39 different states, 99 counties, and 24% "large central metro" urban areas.

3.1.4 Survey Design. We designed a survey to collect rich data for analyzing positioning errors, following similar positioning studies that utilized surveys to evaluate the accuracy of user data and user location [7,48,50,66,80,82]. The survey was implemented in Qualtrics and approved by our Institutional Review Board, including a consent page that detailed the research and provided contact information for the first author. Similar to surveys in prior related work [12,51,62], a pilot test showed that many users were apparently hesitant to share their location, so we revised the survey accordingly. The revised survey followed a more nuanced design that collected geographic information without requiring sensitive data from users. For example, we removed a question that recorded exact coordinates of the respondent's location (via HTML5 and GPS), a step which deterred many users in pilot tests, likely due to privacy concerns [44,75].

The results in this paper are from the final survey format, which first displayed an attention check in the form of a multiple choice question, asking what kind of device was being used (verified using the "Device Type" reported by Qualtrics). Then, the survey displayed the aforementioned consent page before collecting the following information from respondents (the full instrument is included as supplementary material):

- Whether they recognized the target ZIP code (and their relationship to it, if recognized)
- The ZIP code of their current residence
- Whether they were currently in the target county and target state

The question about relationship to the target ZIP code asked whether the respondent worked in, traveled to, or had some other relationship with the ZIP code, how frequently they visited (daily, weekly, monthly, or annually), and whether they were currently in the ZIP code. The question about being currently in the target county and state explicitly asked about the appropriate county and state (e.g. "are you currently in California?" and "are you currently in Santa Clara county, California?").

3.1.5 Survey Deployment via Google Display Ads. To deploy the survey through the Google Display Network, we created a separate display campaign for each of the target ZIP codes, as suggested by Jones et al. [48]. Each campaign's targeting was set to "people in or regularly in" the target ZIP code (not the default "people in, or who show interest in" the target location). This setup linked each target ZIP code to one campaign, such that Google only distributed ads in that campaign to a single ZIP code (exemplified in Figure 1). For this reason, the target ZIP code is equivalent to Google's reported location. We also ensured this was the case by manually verifying that Google did not report impressions in "Other Locations" for each campaign that yielded a sample. In other words, Google's location reports suggest that all impressions came from within the target ZIP code for each respective campaign.

At first, each campaign was set to use standard delivery settings: the "maximize clicks" bidding strategy (default and recommended by Google), an all-day schedule, showing on all devices, and using no content or audience targeting beyond the ZIP code. We also blocked our ads from appearing in mobile app games (on iOS and Android), which produced extremely low conversion rates in pilot tests. Furthermore, our budget for total campaign spending was $1600, based on the target sample size of 400 and an average cost per sample less than $4.00 in pilot campaigns. However, it was extremely challenging to collect a large sample size given this budget, leading us to adjust campaign settings.

In the early phases of deployment, we paid for many clicks that did not yield survey participation, so we adjusted some campaign settings in attempt to lower the cost per sample (summarized in Table 2). We placed Google conversion trackers on each page of the survey and changed the campaign bidding strategy to "maximize conversions," which corresponded to survey participation (i.e. answering the attention check). We also implemented an ad schedule, after finding that many fruitless clicks occurred between 12am and 2am each day. According to several sources [92,97] in the online advertising business, this pattern reflects competing advertisers attempting to deplete our campaign budgets at the beginning of each day, thus preventing our campaigns from bidding at key times later in the day. Our schedule ran all ads from 7am to 10pm U.S. Central Time. Finally, we implemented a block list from the digital marketing agency WebMechanix that included 162 websites [98] known to provide extremely low conversion rates, as well as a placement list comprising only apps and websites where our ads had produced conversions. The placement list was most effective in lowering the cost per sample, though it was still quite expensive, as detailed in the results.

Setting	Intention
Block iOS and Android games	Avoid apps prone to accidental clicks
“Maximize conversions” bidding	Increase survey participation rate
Schedule from 7am to 10pm	Prevent competing bidders from diminishing budget [92,97]
Website block list	Avoid websites known to produce low conversions [98]
Website placement list	Target websites that already produced conversions

Table 2: Settings used during survey deployment to increase survey participation.

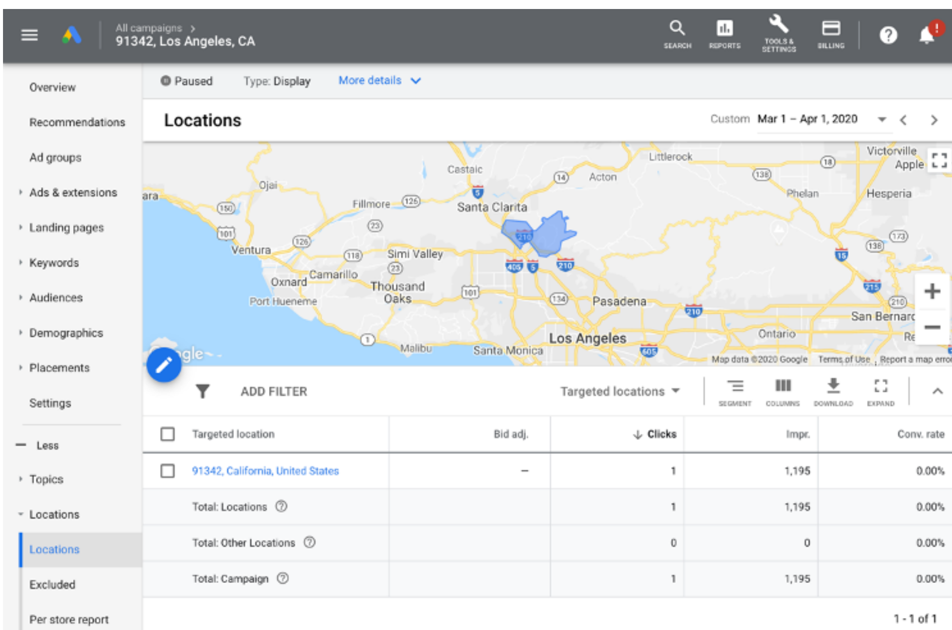

Figure 1: Screenshot from the Google Ads interface, showing a campaign targeted to a ZIP code in California. We manually verified that Google did not report impressions in “Other Locations” for each campaign yielding a sample.

3.1.6 Survey Response Data. Despite challenges related to participation and suspicious click activity in recruiting participants via Google Display ads, our campaigns yielded participation from 111 respondents before reaching our maximum budget of $1600. This sample size is comparable to that of Kariryaa et al. [50], which ran a similar survey over Twitter and collected data from 132 respondents. While we expected and planned for 400 responses, 97% of clicks we paid for did not yield any survey participation, thus restricting our sample size. Following Google’s terminology [100], we refer to the 97% of clicks as suspicious click activity, and provide more details about them in section 4.4.

Out of 190 respondents who viewed the survey according to Qualtrics, 111 passed the attention check and consented to participating in the study, and 91 completed the full survey. Respondents who consented to participation but did not complete the survey (N=20) were included in results when possible (also similar to Kariryaa et al.). For example, some respondents stated they did not live in the target ZIP code but exited before providing their ZIP code. We included their responses when computing error frequency but not error severity, since we could not compute the distance of the error.

3.2 Outcome Metrics

To explore errors in geotargeted Google Display Ads, we focused on two outcome measures: frequency of errors and severity of errors. These two outcome measures mapped to our two research questions.

For percentage metrics, we follow suggestions and standard practice for statistical reporting in HCI research [23], reporting each point estimate with a 95% confidence interval (CI). All point estimates and confidence intervals were calculated through Efron's bootstrapping method [25], which "consists of generating many alternative datasets from the experimental data by randomly drawing observations with replacement," [23] then estimating sampling error based on the variability across these datasets. We calculated each confidence interval with 10,000 resampling iterations (in line with recommendations for 95% CIs [52]) using open-source software by Beecher et al. [9]. We also used this software for difference-of-proportions tests to validate our results and check for robustness in section 4.3. Research has shown the bootstrapping method is versatile to many kinds of distributions, and provides accuracy with a sample size of 20 or more [52]. It is becoming more common to use bootstrapping when calculating confidence intervals from survey data [73], including in CSCW literature [55].

3.2.1 RQ1: How frequent are geotargeting errors in the Google Display Network? Considering our focus on home targeted advertisements, we first tabulated users who lived in the target location. Our main metric of interest was the non-resident rate, which is simply the percentage of respondents who did not live in the target ZIP code. Note that we focused on residency based on the home targeting use case from Aly et al. [4], even though Google does not provide targeting for "people who live in" a target location.

Because of this platform limitation, we also considered "people in or frequently in" the target location, which is a targeting option Google provides when selling geotargeted advertisements. We evaluated the accuracy of this option using the non-visitor rate, which accounts for other relationships to the target location that respondents shared in the survey. In particular, respondents who lived in the target ZIP code or visited the target ZIP code monthly, weekly, or daily were considered "people frequently in" the target location, and respondents who were currently in the target ZIP code were considered "people in" the target location.

3.2.2 RQ2: How severe are geotargeting errors in the Google Display Network? To explore the severity of targeting errors, we measured the distance between the target ZIP code and the self-reported ZIP code where each respondent lived. First, we measured the distance between ZIP code centroids, using coordinates from the GeoNames database [84]. The database uses a variety of data sources, is actively maintained by the National Geospatial-Intelligence Agency, and is commonly used in HCI research (e.g. [47,79]). As an additional metric, we also calculated neighbor contiguity between the target and reported ZIP code. We used queen polygon contiguity, which includes neighbors that share borders and vertices [5], and calculated it manually due to the sample size. Queen contiguity was 1 if the reported ZIP code was a neighbor of the target ZIP code, 2 if the reported ZIP code was a neighbor of a neighbor of the target ZIP code, and so on. As a final way of measuring the severity of geotargeting errors, we calculated the non-resident rate for the county and state that included the target ZIP code (a scale-based error measurement versus a distance-based one). In cases where the reported ZIP code overlapped multiple counties including the target county, respondents were considered residents of the target county.

4 RESULTS

4.1 RQ1: How frequent are geotargeting errors?

4.1.1 Non-Resident Rate. Our first metric of interest was the non-resident rate, which we found to be 41% (CI: 32% - 50%, N=46 out of 111). In other words, when targeting display ads to a random ZIP code in the United States, approximately two in five respondents did not live in the target ZIP code. This may present issues for local advertising campaigns aiming to reach residents of specific ZIP codes.

In terms of replication, our results provide a similar characterization to the 2012 study by Jones et al., with the confidence interval for our estimate overlapping the 50% estimate from Jones et al. Notably, our study was conducted in the United States, which has fewer residents per ZIP code tabulation area (less than 11,000) than the United Kingdom has per postal code (over 500,000). This may help explain the apparent continuity in error rates despite years of advancements in geographic positioning.

4.1.2 Non-Visitor Rate. In our results, 35% (CI: 26% - 44%, N=39 out of 111) of respondents were neither residents nor visitors of the target ZIP code. Of respondents who lived in a different ZIP code but recognized the target ZIP code (N=11), 4 visited daily, 3 visited weekly, 1 visited monthly, and 3 visited annually, and only one respondent stated they were currently in the target ZIP code. Based on this data, 35% of respondents were not “people in or frequently in” the target ZIP code, the audience which Google’s advertising platform claims to reach.

4.2 RQ2: How severe are geotargeting errors?

While the results established that a substantial number of respondents did not live in or visit the target ZIP code, analysis of error severity showed that most of these errors were minor, to the point of insignificance for many advertising goals. Table 3 shows that the non-resident rate decreases at less local geographic scales: 11% of respondents did not live in the target county, 7% did not live in the target DMA, and 2% did not live in the target state. In terms of neighboring ZIP codes, we found that among respondents who did not live in the target ZIP code and reported their actual ZIP code (N=37), 67% lived just one or two ZIP codes away. This is reflected in Figure 2, which shows the distribution of errors in terms of queen contiguity steps.

Most errors were also minor with respect to distance. Of the 37 respondents who did not live in the target ZIP code and reported the ZIP code of their current residence, the median error distance between the reported ZIP code and target ZIP code centroids was 17 kilometers (interquartile range: 9km – 55km). In the worst case, one respondent lived in a ZIP code 223 kilometers away from the target ZIP code. The distribution function of all errors is shown in Figure 3, which further indicates that major errors were fairly uncommon: 60% of errors were less than 20 kilometers, and 81% of errors were less than 100 kilometers.

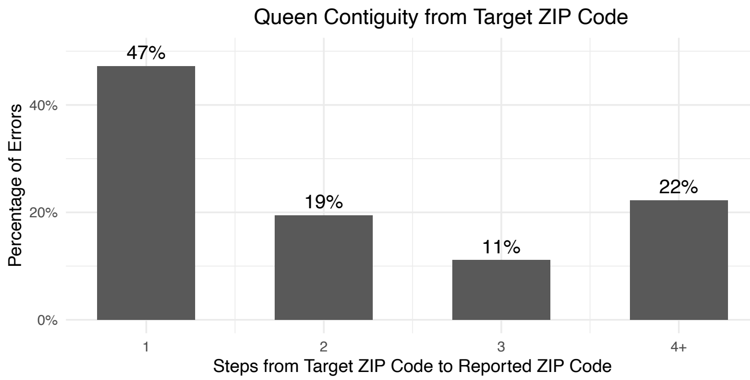

Figure 2: Queen contiguity steps for non-residents of the target ZIP code who reported their actual ZIP Code (i.e. errors only, N=37). Calculated manually, with "1" indicating the reported ZIP was a neighbor of the target ZIP, "2" indicating the reported ZIP was a neighbor of a neighbor of the target ZIP, etc. In 67% of all errors, respondents only lived one or two ZIP codes away from the target ZIP.

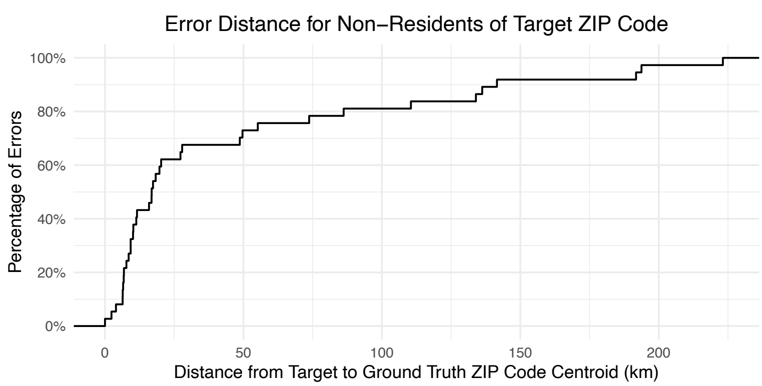

Figure 3: Error distance for non-residents of the target ZIP code who reported their actual ZIP code (i.e. errors only, N=37). Most respondents who did not live in the target ZIP code lived in a ZIP code less than 25km away (measured using centroid coordinates).

Geographic Scale	Non-Resident Rate
Target ZIP Code	41% (CI: 32% - 50%)
Neighboring ZIP	17% (CI: 11% - 24%)
Target County	11% (CI: 5% - 17%)
Target DMA	7% (CI: 3% - 13%)
Target State	2% (CI: 0% - 5%)

Table 3: Non-resident rate for five geographic scales, suggesting that geotargeting errors are quite common in terms of ZIP codes, but far less common in larger regions.

4.3 Robustness Checks

We also checked if our results varied across important categories, for example, if errors were more common in news deserts. To do this, we evaluated the difference of proportions between categorical groups in the data, using the bootstrapping software by Beecher et al. [9] to estimate true difference in proportions. The software calculates a significance value (i.e. p-value) by generating B samples from the full distribution, measuring the test statistic t_i for each sample, then calculating the proportion of samples that are more extreme than the observed overall value (t_{obs}):

$$\hat{p} = \frac{\sum_{i=1}^{B} I(t_i \geq t_{obs}) + 1}{B + 1}$$

We did not find sufficient evidence to reject the null hypotheses that error rates were equal between news deserts and non-news deserts, both in terms of the non-resident rate and the non-visitor rate. We also did not find evidence of variation in terms of other important categories: operating system (iOS vs. Android), time of response (before or after median response), relative location in the United States (east vs. west-coast ZIP codes), or population level (above or below population of median ZIP code).

4.4 Additional Findings in Google Display Ads

An unexpected but potentially interesting result from our study was that many clicks did not result in survey participation. In a related study that deployed a location survey via Twitter's online advertising [50], 35% of all clicks resulted in participation. But over the course of our study, only 2-7% of all clicks resulted in participation (including partial participation), with variation as we adjusted ad campaign settings (see Figure 4). Here, we provide additional details about these suspicious clicks, since they hint at a potential barrier for advertisers to effectively reach local audiences, as well as a potential opportunity for local news organizations.

Even after implementing conversion tracking, block lists, and ad schedules (as summarized in Table 2), 93% of clicks did not result in any survey participation. This is unexpected given that our advertisements explicitly led to a survey, and Google takes measures to ensure that clicks represent "real users with genuine interest." We expected most users who intentionally clicked on an advertisement for a research survey would at least answer the attention check.

Figure 4 shows the steep drop-off between clicks and participation. Before adjusting the campaigns, only 59% of clicks led to viewing the survey landing page (according to a Google conversion tracker), and only 2% of clicks led to any participation (according to Qualtrics). Qualtrics records partially complete responses[3] of any kind, beginning with the multiple-choice attention check on the first page of the survey. Following the adjustments detailed in section 3.1.5 and Table 2, 79% of clicks resulted in viewing the survey according to Google, and 7% of clicks led to some degree of survey participation according to Qualtrics. Both were modest improvements. However, while some potential participants likely exited the survey after viewing the first question, the percent of clicks leading to participation remained much lower than expected based on related work (7% in our data compared to 35% in the related work on Twitter [50]).

5 Discussion

Our results show that people reached by Google's geotargeted display ads often live in or near the target location, although there may be other barriers to reaching this audience. The findings prompt a discussion of implications for advertisers and local news organizations, as well as some important limitations of the study.

[3] Qualtrics temporarily records empty responses as "0% complete," but deletes them after inactivity, so they are not included in the results.

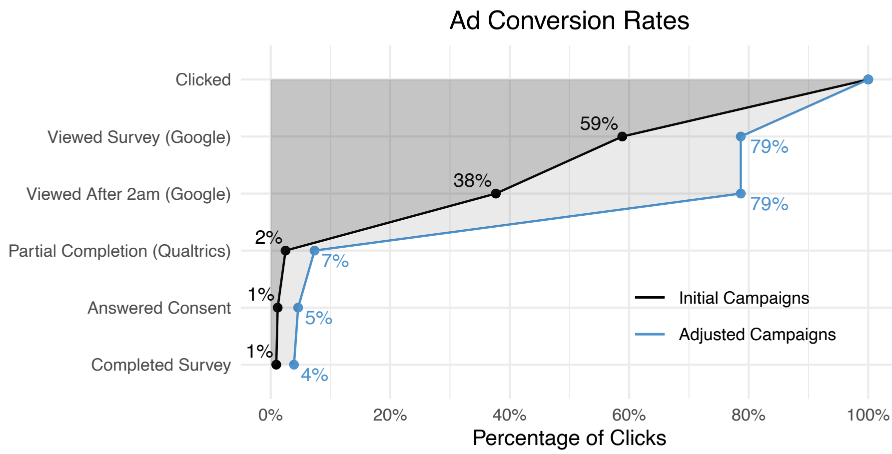

Figure 4: Conversion rates at key pages of the survey. Many clicks did not result in viewing the landing page, though adjusted campaigns (with ad schedules and placement lists) increased participation rates.

5.2 Google Display Ads for Local Advertising

Geotargeted Google display ads will likely serve the goals of many local advertisers: only 17% of our respondents lived outside the target county, and just 2% lived outside the target state. But while the people reached were fairly local, we encountered several barriers to reaching them that advertisers and news websites should weigh when considering advertising platforms. Even after following Google's recommendations and implementing best practices from the digital advertising industry, the vast majority of clicks we paid for were not from "real users with genuine interest," as 21% of clicks did not view the survey landing page according to Google's tracking, and 93% did not register any kind of participation according to Qualtrics.

We posit three potential explanations for the low conversion/participation rates: user behavior, campaign deficiencies, and platform deficiencies. As with any survey, we did not expect 100% of people who clicked on the ad to participate. We did, however, expect the rate to be much higher than 2-7% (Figure 4), based on related work. For example, a similar study that deployed a survey on Twitter [50] yielded participation from 35% of all clicks.

Another potential explanation involves deficiencies in our campaigns. Many dynamic and unpredictable factors affect online advertising, including the COVID-19 pandemic (e.g. [99]), which began during our study. Furthermore, many advertisers target and bid on precise audience attributes (i.e. age, gender, and interests), but the only targeting our campaigns used was location-based. Also, our account did not accumulate enough conversions to utilize Google's more advanced bidding strategies such as pay-per-conversion or target cost per conversion, which may have improved participation rate. At the same time, our observations were consistent across hundreds of campaigns and locations, and in many ways our strategies and expenditures emulate the experience of small businesses. Thus, we would expect small businesses that attempt to use geotargeted display advertising would encounter similar challenges in reaching local audiences.

Platform deficiencies provide the last potential explanation for the low conversion rates in our data. Google's documents describe a variety of measures in place to mitigate fraudulent clicks, but some patterns we observed were difficult to explain through user behavior or campaign deficiencies. These patterns include the high volume of fruitless clicks between 12am and 2am, and the substantial drop-off between clicking the ad and viewing the landing page. Together, these patterns suggest we paid for many fraudulent and/or accidental clicks. Whether fraudulent or accidental, these suspicious clicks present a barrier for advertisers attempting to reach local audiences. This led us to revisit an important motivation for this study related to cost-effective advertising for local journalism.

5.3 Toward Cost-Effective Advertising on Local News Websites

The externalities of online advertising, particularly with regard to the business of local journalism, provided key motivation for this study. Considering that many news websites now rely on third-party display ads for revenue [12], we looked specifically at Google's third-party display advertising, asking whether it provides accurate local advertising that warrants its disruption to local newspaper advertising. Our results suggest that most people reached through Google's geotargeted display ads do indeed live in or near the target location – in contrast to findings in prior work – and this likely will meet many advertisers' needs.

Still, the 41% non-resident rate in target ZIP codes may provide an opportunity for local news websites to provision cost-effective local advertisements without relying on platforms. That is, "native" display ads on news websites could potentially reach residents more effectively than Google's display ads, especially if local news websites could leverage accurate, fine-grained location information from subscribers (e.g. from first-party information such as billing addresses). Additionally, suspicious clicks in our data suggest there may be yet another opportunity for local news sites. During our campaigns, we paid for thousands of clicks (97% of total) that did not result in survey participation, and were likely accidental or fraudulent. Local news websites handle much less traffic than the Google Display Network, and thus could be less prone to suspicious clicks. At the same time, local newspapers' small scale is a key limitation that could deter advertisers who seek to reach larger audiences. Local news websites may also lack the customer support infrastructure, demographic targeting options, and other technical features which have helped Google dominate the display advertising market. Nonetheless, our findings suggest that in some cases, native advertising on local news sites could be preferable to third-party advertising, reinforcing the importance of ongoing research in this area.

5.4 Limitations

Our study suffers from several limitations that are important to highlight. First, for a number of reasons, we collected fewer samples than intended before maximizing our budget, and samples were limited to the United States and mostly to mobile devices. In future studies, researchers might explore geotargeting accuracy with more samples, on additional device types, and/or in different countries. Second, our approximation of "people in or frequently in" a location does not account for people who may have frequently visited a ZIP code, but did not recognize it. Future work may explore ways to collect more precise location information, such as additional incentives for respondents. Maps and visual aids may also improve location information from users. At the same time, some limitations are inherent to surveys, as noted in prior research (e.g. [7,48,50,66,80,82]) including the analogous study by Kariryaa et al. [50]. Namely, surveys are always prone to response bias and potentially inaccurate information from participants.

Finally, while we spent a modest budget and invested significant time into optimizing our campaigns, they still may not emulate the campaigns used by real-world advertisers, especially large ones. On one hand, small businesses likely face similar budget constraints to those we faced in our study, so in some ways we emulated the experiences and goals of local advertisers. At the same time, we likely spent a much smaller budget compared to some corporate advertisers using geotargeted display ads. We also changed many default settings and took several measures to reduce suspicious clicks, which may have altered the population surveyed. In future work, researchers may benefit from partnering with local advertisers and designing field experiments, as Blake et al. [10] did with eBay.

6 Conclusion

This work has shown that, despite existing evidence to the contrary, people reached by Google display ads tend to live in or near the target location. Across random ZIP codes in the United States, a survey deployed via Google's display network found that 41% of respondents lived outside the target ZIP code, but only 11% lived outside the target county, and only 2% lived outside the target state. In other words, errors were fairly common but not very severe, such that Google display advertising will likely serve as an effective positioning system for most advertising purposes. However, the results showed other potential deficiencies, including suspicious clicks which in some cases may have suggested fraudulent behavior aimed at depleting our campaign budgets. Future work should continue exploring targeting errors and potentially fraudulent activity on advertising platforms, especially given the urgent need to stabilize the business of local journalism.

ACKNOWLEDGMENTS

This work is supported by National Science Foundation Grants IIS-1815507 and IIS-1707296. We are grateful to the anonymous reviewers for their guidance in the revision process, as well as technical and methodological help from Nicholas Vincent, Hanlin Li, and Allen Yilun Lin.

REFERENCES

[1] Penelope Muse Abernathy. 2018. *The Expanding News Desert.* University of North Carolina Press Chapel Hill.

[2] Tanja Aitamurto, Mike Ananny, Chris W. Anderson, Larry Birnbaum, Nicholas Diakopoulos, Matilda Hanson, Jessica Hullman, and Nick Ritchie. 2019. HCI for Accurate, Impartial and Transparent Journalism: Challenges and Solutions. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (CHI EA '19), Association for Computing Machinery, Glasgow, Scotland Uk, 1–8. DOI:https://doi.org/10.1145/3290607.3299007

[3] Muhammad Ali, Giridhari Venkatadri, Filipe Nunes Ribeiro, George Arvanitakis, Fabrício Benevenuto, Krishna P Gummadi, Patrick Loiseau, Alan Mislove, Sorelle A Friedler, and Christo Wilson. 2018. Potential for Discrimination in Online Targeted Advertising. In *Proceedings of Machine Learning Research*, 1–15.

[4] Heba Aly, John Krumm, Gireeja Ranade, and Eric Horvitz. 2018. On the Value of Spatiotemporal Information: Principles and Scenarios. ACM, 179–188. DOI:https://doi.org/10.1145/3274895.3274905

[5] Luc Anselin, Ibnu Syabri, and Youngihn Kho. 2010. GeoDa: An Introduction to Spatial Data Analysis. In *Handbook of Applied Spatial Analysis: Software Tools, Methods and Applications*, Manfred M. Fischer and Arthur Getis (eds.). Springer, Berlin, Heidelberg, 73–89. DOI:https://doi.org/10.1007/978-3-642-03647-7_5

[6] Afdhel Aziz. Facebook Ad Boycott Campaign 'Stop Hate For Profit' Gathers Momentum And Scale: Inside The Movement For Change. *Forbes.* Retrieved September 23, 2020 from https://www.forbes.com/sites/afdhelaziz/2020/06/24/facebook-ad-boycott-campaign-stop-hate-for-profit-gathers-momentum-and-scale-inside-the-movement-for-change/

[7] Muhammad Ahmad Bashir, Umar Farooq, Maryam Shahid, Muhammad Fareed Zaffar, and Christo Wilson. 2019. Quantity vs. Quality: Evaluating User Interest Profiles Using Ad Preference Managers. In *Proceedings 2019 Network and Distributed System Security Symposium.* DOI:https://doi.org/10.14722/ndss.2019.23392

[8] Christine Bauer, Christine Strauss, C Bauer, and C Strauss. 2016. Location-based advertising on mobile devices A literature review and analysis. *Management Review Quarterly* 66, (2016), 159–194. DOI:https://doi.org/10.1007/s11301-015-0118-z

[9] Spencer Beecher, Don van der Drift, David Martin, Lindsay Vass, Sergey Goder, Benedict Lim, and Matt Langner. 2020. *bootstrapped Github Repository.* Facebook Incubator. Retrieved August 17, 2020 from https://github.com/facebookincubator/bootstrapped

[10] Thomas Blake, Chris Nosko, and Steven Tadelis. 2015. Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment. *Econometrica* 83, 1 (2015), 155–174. DOI:https://doi.org/10.3982/ecta12423

[11] Andrew L. Brooks and Coye Cheshire. 2012. Ad-itudes: twitter users & advertising. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work Companion* (CSCW '12), Association for Computing Machinery, Seattle, Washington, USA, 63–66. DOI:https://doi.org/10.1145/2141512.2141543

[12] A J Bernheim Brush, John Krumm, and James Scott. 2010. Exploring End User Preferences for Location Obfuscation, Location-Based Services, and the Value of Location. In *UbiComp*. Retrieved from http://research.microsoft.com/en-us/um/people/jckrumm/GPSData2009/

[13] Ceren Budak, Sharad Goel, Justin M. Rao, and Georgios Zervas. 2016. Understanding Emerging Threats to Online Advertising. In *ACM Conference on Economics and Computation*. DOI:https://doi.org/10.2139/ssrn.2505643

[14] Steve Cameron. 2015. *Does Google Really Know Where You Are? For Nearly Half Of You, The Answer Is No.* Retrieved from https://searchengineland.com/google-really-know-230001

[15] Keith Collins. 2017. *Google collects Android users' locations even when location services are disabled — Quartz*. Retrieved from https://qz.com/1131515/google-collects-android-users-locations-even-when-location-services-are-disabled/

[16] Josh Constine. *Facebook's New Zip Code Ad Targeting Could Boost Local Advertising Revenue*. Retrieved September 11, 2019 from https://www.adweek.com/digital/zip-code-ad-targeting/

[17] Nick Couldry and Ulises A Mejias. 2019. Data Colonialism: Rethinking Big Data's Relation to the Contemporary Subject. *Television and New Media* 20, 4 (2019), 336–349. DOI:https://doi.org/10.1177/1527476418796632

[18] Dharma Dailey and Kate Starbird. 2014. Journalists as Crowdsourcerers: Responding to Crisis by Reporting with a Crowd. *Comput Supported Coop Work* 23, 4 (December 2014). DOI:https://doi.org/10.1007/s10606-014-9208-z

[19] Dharma Dailey and Kate Starbird. 2017. Social Media Seamsters: Stitching Platforms and Audiences into Local Crisis Infrastructure. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (CSCW '17), Association for Computing Machinery, New York, NY, USA, 1277–1289. DOI:https://doi.org/10.1145/2998181.2998290

[20] Elizabeth Daly, Sheena Erete, Rosta Farzan, Gary Hsieh, Cliff Lampe, Claudia Lopez, Andres Monroy-Hernandez, Daniele Quercia, Raz Schwartz, and Amy Voida. 2015. Supporting Cities, Neighborhoods, and Local Communities with Information and Communication Technologies. In *Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing* (CSCW'15 Companion), Association for Computing Machinery, New York, NY, USA, 277–281. DOI:https://doi.org/10.1145/2685553.2685556

[21] Joshua P Darr, Matthew P Hitt, and Johanna L Dunaway. 2018. Newspaper Closures Polarize Voting Behavior. *Journal of Communication* 68, 6 (2018).

[22] Amit Datta, Anupam Datta, Jael Makagon, Deirdre K Mulligan, and Michael Carl Tschantz. 2016. Discrimination in Online Advertising: A Multidisciplinary Inquiry. *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* 81, (2016), 20–34.

[23] Pierre Dragicevic. 2016. Fair Statistical Communication in HCI. In *Modern Statistical Methods for HCI*, Judy Robertson and Maurits Kaptein (eds.). Springer International Publishing, Cham, 291–330. DOI:https://doi.org/10.1007/978-3-319-26633-6_13

[24] Gilad Edelman. Ad Tech Could Be the Next Internet Bubble. *Wired*. Retrieved October 15, 2020 from https://www.wired.com/story/ad-tech-could-be-the-next-internet-bubble/

[25] B. Efron and R. Tibshirani. 1986. Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy. *Statistical Science* 1, 1 (1986), 54–75. Retrieved August 20, 2020 from https://www.jstor.org/stable/2245500

[26] eMarketer Editors. 2019. *Location-Based Advertising Is Becoming More Costly*. Retrieved from https://www.emarketer.com/content/location-based-advertising-is-becoming-more-costly

[27] Ayman Farahat Yahoo and Michael Bailey. 2012. How Effective is Targeted Advertising? In *WWW 2012*. Retrieved from http://delivery.acm.org/10.1145/2190000/2187852/p111-farahat.pdf

[28] Sean Fischer, Kokil Jaidka, and Yphtach Lelkes. 2020. Auditing local news presence on Google News. *Nature Human Behaviour* (September 2020).

[29] Lauren Fisher. 2019. *Digital Display Advertising 2019 - eMarketer Trends, Forecasts & Statistics*. Retrieved August 29, 2019 from https://www.emarketer.com/content/digital-display-advertising-2019

[30] David Flatow, Mor Naaman, Ke Eddie Xie, Yana Volkovich, and Yaron Kanza. 2015. On the Accuracy of Hyper-local Geotagging of Social Media Content. In *WSDM 2015 - Proceedings of the 8th ACM International Conference on Web Search and Data Mining*, 127–136. DOI:https://doi.org/10.1145/2684822.2685296

[31] Pengjie Gao, Chang Lee, and Dermot Murphy. 2018. Financing Dies in Darkness? The Impact of Newspaper Closures on Public Finance. In *SSRN*.

[32] Song Gao and Sathya Prasad. 2016. Employing spatial analysis in indoor positioning and tracking using wi-fi access

points. In *Proceedings of the Eighth ACM SIGSPATIAL International Workshop on Indoor Spatial Awareness* (ISA '16), Association for Computing Machinery, Burlingame, California, 27–34. DOI:https://doi.org/10.1145/3005422.3005425

[33] Yola Georgiadou, Rolf de By, and Ourania Kounadi. 2019. Location Privacy in the Wake of the GDPR. *ISPRS International Journal of Geo-Information* 8, 3 (2019), 157. DOI:https://doi.org/10.3390/ijgi8030157

[34] Manaf Gharaibeh, Anant Shah, Bradley Huffaker, Han Zhang, Roya Ensafi, and Christos Papadopoulos. 2017. A Look at Router Geolocation in Public and Commercial Databases. 463–469. DOI:https://doi.org/10.1145/3131365.3131380

[35] Brett R. Gordon, Florian Zettelmeyer, Neha Bhargava, and Dan Chapsky. 2017. A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook. *SSRN 38*. DOI:https://doi.org/10.2139/ssrn.3033144

[36] Bamba Gueye, Steve Uhlig, and Serge Fdida. 2007. Investigating the Imprecision of IP Block-Based Geolocation. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Springer Berlin Heidelberg, Berlin, Heidelberg, 237–240. DOI:https://doi.org/10.1007/978-3-540-71617-4_26

[37] Stéphane Hamel. 2016. *Case Study: Accuracy and Precision of Google Analytics Geolocation.* Retrieved July 17, 2019 from https://radical-analytics.com/case-study-accuracy-precision-of-google-analytics-geolocation-4264510612c0

[38] Stéphane Hamel. *Benchmarking the Accuracy and Precision of Google Analytics Geolocation.* Retrieved July 17, 2019 from https://www.stephanehamel.net/playground/galocation

[39] Brent Hecht and Darren Gergle. 2010. On the "Localness" of User-Generated Content. In *CSCW*.

[40] Kashmir Hill. 2016. How an Internet Mapping Glitch Turned a Random Kansas Farm Into a Digital Hell. *Fusion* (2016). Retrieved from https://splinternews.com/how-an-internet-mapping-glitch-turned-a-random-kansas-f-1793856052 http://fusion.net/story/287592/internet-mapping-glitch-kansas-farm/

[41] Paul Hoban and Neeraj Arora. 2018. Measuring Display Advertising Response Using Observational Data: The Impact of Selection Biases. *SSRN Electronic Journal* (2018). DOI:https://doi.org/10.2139/ssrn.3264871

[42] Richard Holden. *Inside AdWords: Get Local with AdWords.* Retrieved September 11, 2019 from https://adwords.googleblog.com/2012/04/get-local-with-adwords.html

[43] Amelie Hössjer and Kerstin Severinson Eklundh. Making Space for a New Medium: On the Use of Electronic Mail in a Newspaper Newsroom. In *Computer Supported Cooperative Work: Vol 18, No 1.* Retrieved October 14, 2020 from https://dl.acm.org/doi/abs/10.1007/s10606-008-9082-7

[44] Luke Hutton, Tristan Henderson, and Apu Kapadia. 2014. Short Paper: "Here I am, now pay me!": Privacy Concerns in Incentivised Location-sharing Systems. In *WiSec.* DOI:https://doi.org/10.1145/2627393.2627416

[45] Tim Hwang. 2020. *Subprime Attention Crisis: Advertising and the Time Bomb at the Heart of the Internet.* FSG Originals. Retrieved October 15, 2020 from https://bookshop.org/books/subprime-attention-crisis-advertising-and-the-time-bomb-at-the-heart-of-the-internet/9780374538651

[46] Internet Advertising Bureau. 2019. *Internet Advertising Revenue Report.* Retrieved from http://www.iab.net/media/file/IAB_PWC_1999Q2.pdf

[47] Alan Jackoway, Hanan Samet, and Jagan Sankaranarayanan. 2011. Identification of live news events using Twitter. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks* (LBSN '11), Association for Computing Machinery, Chicago, Illinois, 25–32. DOI:https://doi.org/10.1145/2063212.2063224

[48] Ray B Jones, Lesley Goldsmith, Christopher J Williams, and Maged N.Kamel Boulos. 2012. Accuracy of Geographically Targeted Internet Advertisements on Google Adwords for Recruitment in a Randomized Trial. In *Journal of Medical Internet Research.* DOI:https://doi.org/10.2196/jmir.1991

[49] Seb Joseph. 2020. *Apple's new privacy features have further rattled the location-based ad market -.* Retrieved from https://digiday.com/marketing/apples-new-privacy-features-rattle-location-based-ad-market/

[50] Ankit Kariryaa, Isaac Johnson, Johannes Schöning, and Brent Hecht. 2018. Defining and Predicting the Localness of Volunteered Geographic Information using Ground Truth Data. 1–12. DOI:https://doi.org/10.1145/3173574.3173839

[51] Patrick Gage Kelley, Michael Benisch, Lorrie Faith Cranor, and Norman Sadeh. 2011. When Are Users Comfortable Sharing Locations with Advertisers? In *CHI*, 2449. DOI:https://doi.org/10.1145/1978942.1979299

[52] Kris N. Kirby and Daniel Gerlanc. 2013. BootES: An R package for bootstrap confidence intervals on effect sizes. *Behav Res* 45, 4 (December 2013), 905–927. DOI:https://doi.org/10.3758/s13428-013-0330-5

[53] Juhi Kulshrestha, Motahhare Eslami, Johnnatan Messias, Muhammad Bilal Zafar, Saptarshi Ghosh, Krishna P Gummadi, and Karrie Karahalios. 2017. Quantifying search bias: Investigating sources of bias for political searches in social media. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW*, 417–432. DOI:https://doi.org/10.1145/2998181.2998321

[54] Sangmee Lee, Ki Joon Kim, and S. Shyam Sundar. 2015. Customization in location-based advertising: Effects of tailoring source, locational congruity, and product involvement on ad attitudes. *Computers in Human Behavior* 51, (October 2015). DOI:https://doi.org/10.1016/j.chb.2015.04.049

[55] Hanlin Li, Nicholas Vincent, Janice Tsai, Jofish Kaye, and Brent Hecht. 2019. How Do People Change Their Technology Use in Protest?: Understanding "Protest Users." In *CSCW*.

[56] Hui Liu, Santosh Pandey, and Jiang Zhu. 2017. Human mobility rule-based device location tracking. Retrieved June 1, 2020 from https://patents.google.com/patent/US9565584B2/en

[57] Lucas Manfredi. Google may have to search for answers on newspaper closings. Retrieved August 20, 2020 from https://www.foxbusiness.com/media/rep-raskin-puts-googles-feet-to-the-fire-on-newspaper-closings

[58] Veronica Marotta, Vibhanshu Abhishek, and Alessandro Acquisti. 2019. *Online Tracking and Publishers' Revenues: An Empirical Analysis*. Carlson School of Management, University of Minnesota. Retrieved from http://www.heinz.cmu.edu/˜acquisti/cv.htm.

[59] Matt McGee. *Yahoo Adds ZIP Code Ad Targeting*. Retrieved September 11, 2019 from https://searchengineland.com/yahoo-adds-zip-code-ad-targeting-15193

[60] Jonathan Scott Morgan, Cliff Lampe, and Muhammad Zubair Shafiq. 2013. Is news sharing on Twitter ideologically biased? In *Proceedings of the 2013 conference on Computer supported cooperative work* (CSCW '13), Association for Computing Machinery, New York, NY, USA, 887–896. DOI:https://doi.org/10.1145/2441776.2441877

[61] National Center for Health Statistics. 2019. Urban Rural Classification Scheme for Counties. Retrieved October 15, 2020 from https://www.cdc.gov/nchs/data_access/urban_rural.htm

[62] Xinru Page, Bart P. Knijnenburg, and Alfred Kobsa. 2013. What a tangled web we weave: lying backfires in location-sharing social media. In *Proceedings of the 2013 conference on Computer supported cooperative work* (CSCW '13), Association for Computing Machinery, San Antonio, Texas, USA, 273–284. DOI:https://doi.org/10.1145/2441776.2441808

[63] Sahil Patel. 2019. *Facebook Reaches Proposed Settlement in Video Measurement Lawsuit*. Retrieved from https://www.wsj.com/articles/facebook-reaches-proposed-settlement-in-video-measurement-lawsuit-11570482031

[64] Tim Peterson. 2019. GDPR was just a warmup. CCPA will arrive with a bang. *Digiday*. Retrieved from https://digiday.com/marketing/gdpr-just-warmup-ccpa-will-arrive-bang/

[65] Pew Research Center. 2019. *Trends and Facts on Newspapers – State of the News Media*. Retrieved July 11, 2019 from https://www.journalism.org/fact-sheet/newspapers/ http://www.journalism.org/fact-sheet/newspapers/

[66] Pew Research Center. 2019. *Facebook Algorithms and Personal Data*. Retrieved from https://www.pewinternet.org/2019/01/16/facebook-algorithms-and-personal-data/

[67] Ingmar Poese, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. *IP Geolocation Databases: Unreliable?* Retrieved from http://delivery.acm.org/10.1145/1980000/1971171/p53-poese.pdf

[68] Reid Priedhorsky, Aron Culotta, and Sara Y. Del Valle. 2014. Inferring the origin locations of tweets with quantitative confidence. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing* (CSCW '14), Association for Computing Machinery, Baltimore, Maryland, USA, 1523–1536. DOI:https://doi.org/10.1145/2531602.2531607

[69] Deepak Ravichandran and Nitish Korula. 2019. *Effect of disabling third-party cookies on publisher revenue*. Retrieved from https://services.google.com/fh/files/misc/disabling_third-party_cookies_publisher_revenue.pdf

[70] Ronald E Robertson, Shan Jiang, Kenneth Joseph, Lisa Friedland, David Lazer, and Christo Wilson. 2018. Auditing Partisan Audience Bias within Google Search. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–22. DOI:https://doi.org/10.1145/3274417

[71] Shruti Sannon, Natalya N Bazarova, and Dan Cosley. 2018. Privacy lies: Understanding how, when, and why people lie to protect their privacy in multiple online contexts. In *Conference on Human Factors in Computing Systems - Proceedings*. DOI:https://doi.org/10.1145/3173574.3173626

[72] Lee Shaker. 2014. Dead Newspapers and Citizens' Civic Engagement. *Political Communication* 31, 1 (2014).

[73] Jun Shao. 2003. Impact of the Bootstrap on Sample Surveys. *Statist. Sci.* 18, 2 (May 2003), 191–198. DOI:https://doi.org/10.1214/ss/1063994974

[74] Yuval Shavitt and Noa Zilberman. 2011. A Geolocation Databases Study. *IEEE Journal on Selected Areas in Communications* 29, 10 (December 2011), 2044–2056. DOI:https://doi.org/10.1109/JSAC.2011.111214

[75] Wonsun Shin and Trisha Tsui Chuan Lin. 2016. Who avoids location-based advertising and why? Investigating the relationship between user perceptions and advertising avoidance. *Computers in Human Behavior* 63, (October 2016), 444–452. DOI:https://doi.org/10.1016/j.chb.2016.05.036

[76] Natasha Singer and Kate Conger. 2019. Google Is Fined $170 Million for Violating Children's Privacy on YouTube. *The New York Times*. Retrieved May 11, 2020 from https://www.nytimes.com/2019/09/04/technology/google-youtube-fine-ftc.html

[77] Gaurav Sood. 2016. *Geographic Information on Designated Media Markets*. Harvard Dataverse. DOI:https://doi.org/10.7910/DVN/IVXEHT

[78] Latanya Sweeney. 2013. Discrimination in Online Ad Delivery. *Queue* 11, 3 (2013).

DOI:https://doi.org/10.2139/ssrn.2208240

[79] Benjamin E. Teitler, Michael D. Lieberman, Daniele Panozzo, Jagan Sankaranarayanan, Hanan Samet, and Jon Sperling. 2008. NewsStand: a new view on news. In *Proceedings of the 16th ACM SIGSPATIAL international conference on Advances in geographic information systems - GIS '08*, 1.

[80] Michael Carl Tschantz, Serge Egelman, Jaeyoung Choi, Nicholas Weaver, and Gerald Friedland. 2018. The Accuracy of the Demographic Inferences Shown on Google's Ad Settings. In *WPES*, 33–41. DOI:https://doi.org/10.1145/3267323.3268962

[81] Heli Väätäjä and Paul Egglestone. 2012. Briefing news reporting with mobile assignments: perceptions, needs and challenges. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work* (CSCW '12), Association for Computing Machinery, Seattle, Washington, USA, 485–494. DOI:https://doi.org/10.1145/2145204.2145280

[82] Giridhari Venkatadri, Piotr Sapiezynski, Elissa M Redmiles, Alan Mislove, Oana Goga, Michelle Mazurek, and Krishna P Gummadi. 2019. Auditing Offline Data Brokers via Facebook's Advertising Platform. In *WWW '19*. DOI:https://doi.org/10.1145/3308558.3313666

[83] Sarah Myers West. 2019. Data Capitalism: Redefining the Logics of Surveillance and Privacy. *Business and Society* 58, 1 (2019), 20–41. DOI:https://doi.org/10.1177/0007650317718185

[84] Mark Wick and Bernard Vatant. 2012. The geonames geographical database.

[85] Max L. L. Wilson, Paul Resnick, David Coyle, and Ed H. Chi. 2013. RepliCHI: the workshop. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '13), Association for Computing Machinery, New York, NY, USA, 3159–3162. DOI:https://doi.org/10.1145/2468356.2479636

[86] Michael G Wing, Aaron Eklund, and Loren D Kellogg. 2005. Consumer-Grade Global Positioning System (GPS) Accuracy and Reliability. *Journal of Forestry* 103, 4 (2005), 169–173. DOI:https://doi.org/10.1093/jof/103.4.169

[87] Jeanine Worden, Kathleen Pennington, and Ayelet Weiss. 2019. *U.S. Department of Housing and Urban Development v. Facebook, Inc.* Retrieved from https://www.hud.gov/sites/dfiles/Main/documents/HUD_v_Facebook.pdf

[88] Fengli Xu, Guozhen Zhang, Zhilong Chen, Jiaxin Huang, Yong Li, Diyi Yang, Ben Y Zhao, and Fanchao Meng. 2018. Understanding Motivations behind Inaccurate Check-ins. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018). DOI:https://doi.org/10.1145/3274457

[89] Paul A Zandbergen. 2009. Accuracy of iPhone locations: A comparison of assisted GPS, WiFi and cellular positioning. In *Transactions in GIS*, John Wiley & Sons, Ltd (10.1111), 5–25. DOI:https://doi.org/10.1111/j.1467-9671.2009.01152.x

[90] Shoshana Zuboff. 2015. Big other: Surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology* 30, 1 (2015), 75–89. DOI:https://doi.org/10.1057/jit.2015.5

[91] 2017. Google's YouTube to undergo MRC audits for video viewability measurement. *Marketing Land*. Retrieved May 12, 2020 from https://marketingland.com/google-youtube-mrc-audit-video-viewability-207223

[92] 2019. What is Click Fraud? | Click Fraud Protection Software. *ClickGUARD™*. Retrieved June 2, 2020 from https://www.clickguard.com/what-is-click-fraud/

[93] Google DoubleClick suspended from Media Rating Council accreditation - Business Insider. Retrieved May 12, 2020 from https://www.businessinsider.com/google-doubleclick-suspended-from-media-rating-council-accreditation-2016-10

[94] Privacy Policy – Privacy & Terms – Google. Retrieved July 13, 2020 from https://policies.google.com/privacy

[95] AdSense revenue share - AdSense Help. Retrieved July 13, 2020 from https://support.google.com/adsense/answer/180195?hl=en

[96] About advanced location options - Google Ads Help. Retrieved May 31, 2020 from https://support.google.com/google-ads/answer/1722038?hl=en

[97] 4 Powerful Ways to Eliminate Click Fraud in Your Account. Retrieved June 2, 2020 from https://www.wordstream.com/blog/ws/2015/08/17/click-fraud

[98] 55,000 GDN Sites To Exclude From Your Placement Targeting. Retrieved June 2, 2020 from https://www.webmechanix.com/display-network-list

[99] The coronavirus pandemic may lead to more political ads online — Quartz. Retrieved May 12, 2020 from https://qz.com/1820187/the-coronavirus-pandemic-may-lead-to-more-political-ads-online/

[100] Troubleshooting invalid clicks - Google Ads Help. Retrieved from https://support.google.com/google-ads/troubleshooter/2557048?hl=en

Received June 2020; revised October 2020; accepted December 2020.