

Location, Location, Location: The Impact of Geolocation on Web Search Personalization

Chloe Kliman-Silver
Brown University
chloe@cs.brown.edu

Aniko Hannak
Northeastern University
ancsaaa@ccs.neu.edu

David Lazer
Northeastern University
d.lazer@neu.edu

Christo Wilson
Northeastern University
cbw@ccs.neu.edu

Alan Mislove
Northeastern University
amislove@ccs.neu.edu

ABSTRACT

To cope with the immense amount of content on the web, search engines often use complex algorithms to personalize search results for individual users. However, personalization of search results has led to worries about the *Filter Bubble Effect*, where the personalization algorithm decides that some useful information is irrelevant to the user, and thus prevents them from locating it.

In this paper, we propose a novel methodology to explore the impact of location-based personalization on Google Search results. Assessing the relationship between location and personalization is crucial, since users' geolocation can be used as a proxy for other demographic traits, like race, income, educational attainment, and political affiliation. In other words, *does location-based personalization trap users in geolocal Filter Bubbles?*

Using our methodology, we collected 30 days of search results from Google Search in response to 240 different queries. By comparing search results gathered from 59 GPS coordinates around the US at three different granularities (county, state, and national), we are able to observe that differences in search results due to personalization grow as physical distance increases. However these differences are highly dependent on what a user searches for: queries for local establishments receive 4-5 different results per page, while more general terms exhibit essentially no personalization.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval; H.3.5 [Information Systems]: Online Services—*web-based services*

Keywords

Search; Personalization; Geolocation; Internet Filter Bubble

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

IMC'15, October 28–30, 2015, Tokyo, Japan.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3848-6/15/10 ...\$15.00.

DOI: <http://dx.doi.org/10.1145/2815675.2815714>.

1. INTRODUCTION

Search engines are the primary gateway to information in the developed world. Thus, it is no surprise that Google has been the most visited site on the Internet for several years now [1]; it receives more than 48,000 queries every second [2]. The importance of content and ordering in search results is exemplified by Europe's recent Right to be Forgotten ruling [21], as well as the thriving Search Engine Optimization (SEO) industry [23].

To cope with the immense amount of content on the web, search engines use complex algorithms to *personalize* search results for individual users [9]. In many cases, personalized search results are useful: if two people on opposite ends of the US search for "coffee shop" they should probably be shown search results for local cafés.

However, personalization of search results has also led to worries about the *Filter Bubble Effect*, where the algorithm decides that some useful information is irrelevant to the user, and thus prevents them from locating it [19]. This issue is particularly concerning in the context of political and news-related information: personalization based on a user's political preferences may trap them in an "echo-chamber" where their pre-existing beliefs are constantly reinforced.

Motivated by concerns about Filter Bubbles, our prior work set out to explore which factors triggered personalization in Google Search [11]. We found that Google infers users' geolocation based on their IP address, and that *location-based personalization* caused more differences in search results than any other single feature. However, while these initial findings are intriguing, many questions remain, such as: does location-based personalization impact all types of queries (e.g., politics vs. news) equally? At what distance do users begin to see changes in search results due to location? Answering these questions is crucial, since users' geolocation can be used as a proxy for other demographic traits, like race, income-level, educational attainment, and political affiliation. In other words, *does location-based personalization trap users in geolocal Filter Bubbles?*

In this paper, we propose a novel methodology to explore the impact of location on Google Search results. We use the JavaScript `Geolocation` API [12] to present arbitrary GPS coordinates to the mobile version of Google Search. Google personalizes the search results based on the location we specified, giving us the ability to collect search results from any location around the globe. Although we focus on

Progressive Tax
Impose A Flat Tax
End Medicaid
Affordable Health And Care Act
Fluoridate Water
Stem Cell Research
Andrew Wakefield Vindicated
Autism Caused By Vaccines
US Government Loses AAA Bond Rate
Is Global Warming Real
Man Made Global Warming Hoax
Nuclear Power Plants
Offshore Drilling
Genetically Modified Organisms
Late Term Abortion
Barack Obama Birth Certificate
Impeach Barack Obama
Gay Marriage

Table 1: Example *controversial* search terms.

Google Search in the US, our methodology is general, and could easily be applied to other search engines like Bing.

Using our methodology, we collected 30 days of search results from Google Search in response to 240 different queries. By selecting 75 GPS coordinates around the US at three granularities (county, state, and national), we are able to examine the relationship between distance and location-based personalization, as well as the impact of location-based personalization on different types of queries. We make the following observations:

- As expected, the differences between search results grows as physical distance between the locations of the users increases.
- However, the impact of location-based personalization changes depending on the query type. Queries for *politicians’* names (e.g., “Joe Biden”) and *controversial* topics (“abortion”) see minor changes, while queries for *local* terms (“airport”) are highly personalized.
- Surprisingly, only 20-30% of differences are due to Maps embedded in search results. The remainder are caused by changes in “normal” search results.
- Also surprisingly, the search results for *local* terms are extremely noisy, i.e., two users making the same query from the same location at the same time often receive substantially different search results.

Outline. The rest of the paper is organized as follows: in Section 2, we give an overview of our data collection methodology, and then present analysis and findings in Section 3. We discuss related work in Section 4 and conclude in Section 5.

2. METHODOLOGY

Our goal is to explore the relationship between geolocation and personalization on Google Search. Thus, we require the ability to send identical queries to Google Search, at the same moment in time, from different locations. In this section, we explain our methodology for accomplishing these goals. First, we introduce the locations and search terms used in our study. Next, we explain our technique for

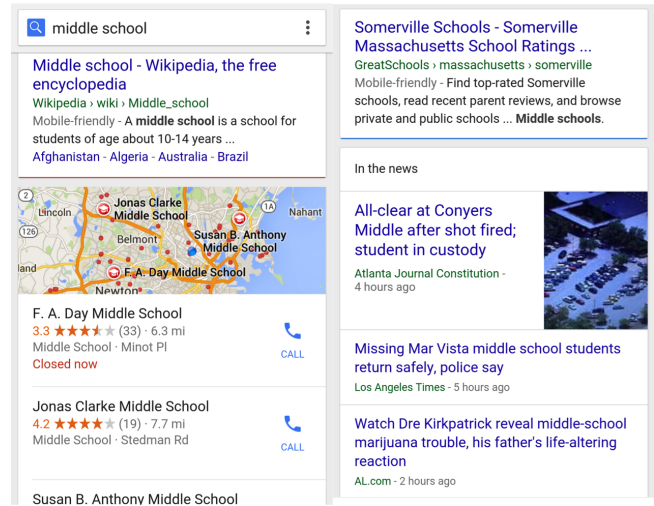


Figure 1: Example search results from the mobile version of Google Search.

querying Google Search from arbitrary locations, and how we parsed Google Search results. Finally, we discuss how we quantify differences between pages of search results.

2.1 Locations and Search Terms

Locations. First, we must choose the locations in which to execute queries. We decided to focus our study on Ohio, since it is known to be a “battleground” state in US politics. This property is important, since we want to examine whether demographics like political affiliation correlate with location-based personalization.

Overall, we picked 66 locations for our study spread across three *granularities*. For *nation*-level, we chose the centroids of 22 random states in the United States. For *state*-level, we chose the centroids of 22 random counties within Ohio. On average, these counties 100 miles apart. Finally, for *county*-level, we chose the centroids of 15 voting districts in Cuyahoga County, which is the most populous county in Ohio. On average, these voting districts are 1 mile apart. By examining locations in different granularities, we will be able to observe changes in search results across small, medium, and large-scale distances. This also gives us the ability to compare search results served in places with different demographics characteristics.

Search Terms. Next, we must select search terms for our study. We built a corpus of 240 queries that fall into three categories: 33 *local* queries, 87 *controversial* queries, and 120 names of *politicians*. *Local* queries correspond with physical establishments, restaurants, and public services such as “bank”, “hospital”, and “KFC”. We chose these terms because we expect them to produce search results that are heavily personalized based on location, i.e., we treat them as an upper-bound on location-based personalization.

For *politicians*, we selected 11 members of the Cuyahoga County Board, 53 random members of the Ohio House and Senate, all 18 members of the US Senate and House from Ohio, 36 random members of the US House and Senate not from Ohio, Joe Biden, and Barack Obama. For national figures like Barack Obama, we do not expect to see differences in search results due to location; however, it is not clear how

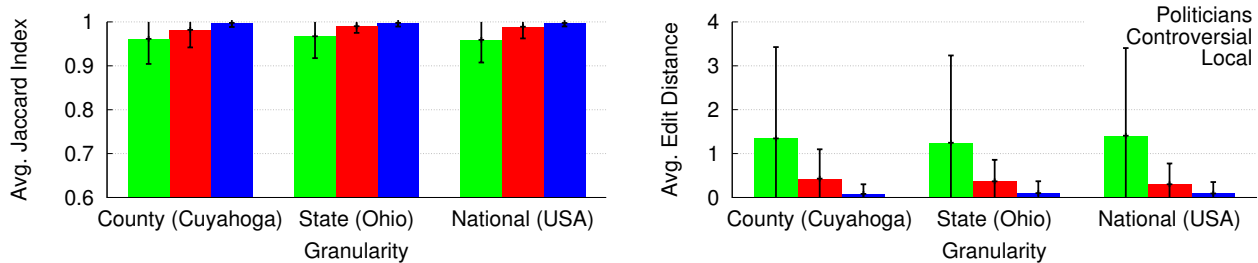


Figure 2: Average noise levels across different query types and granularities. Error bars show standard deviations.

Google Search handles queries for state- and county-level officials inside and outside their home territories.

Finally, our *controversial* terms are news or politics-related issues like those shown in Table 1. We chose these terms because it would be concerning if Google Search personalized search results for them based on location. To avoid possible external bias, we picked search terms that, to the best of our knowledge, were not associated with specific news-worthy events at the time of our experiments. Although we cannot completely rule out the possibility that exogenous events impacted the search results, we note that such an event would impact each treatment equally, and thus would likely not impact our findings.

2.2 Data Collection and Parsing

Our methodology for gathering data from Google Search is based on the techniques presented in our prior work et al. [10,11], with one key difference. As in prior work, we use PhantomJS [20] to gather data, since it is a full implementation of a WebKit browser. We wrote a PhantomJS script that takes a search term and a latitude/longitude pair as input, loads the mobile version of Google Search, executes the query, and saves the first page of search results.

Unlike prior work [11], we targeted the mobile version of Google Search because it uses the JavaScript `Geolocation` API [12] to query the user’s precise location. By overriding the `Geolocation` API in our PhantomJS script, we can feed the coordinates specified on the command line to Google Search, thus giving us the ability to run queries that appear to Google as if they are coming from any location of our choosing. We distributed our query load over 44 machines in a single /24 subnet to avoid being rate-limited by Google. Finally, all of our experimental treatments were repeated for 5 consecutive days to check for consistency over time.

Validation. To make sure that Google Search personalizes search results based on the provided GPS coordinates rather than IP address, we conducted a validation experiment. We issued identical controversial queries with the same exact GPS coordinate from 50 different Planet Lab machines across the US, and observe that 94% of the search results received by the machines are identical. This confirms that Google Search personalizes search results largely based on the provided GPS coordinates rather than the IP address. Furthermore, Google Search reports the user’s precise location at the bottom of search results, which enabled us to manually verify that Google was personalizing search results correctly based on our spoofed GPS coordinates.

Browser State. To control for personalization effects due to the state of the browser, all of our treatments were

configured and behaved identically. The script presented the User-Agent for Safari 8 on iOS, and all other browser attributes were the same across treatments, so each treatment should present an identical browser fingerprint. Furthermore, we cleared all cookies after each query, which mitigates personalization effects due to search history, and prevents Google from “remembering” a treatments prior location. Lastly, we note that prior work has shown that Google Search does not personalize search results based on the user’s choice of browser or OS [11].

Controlling for Noise. Unfortunately, not all differences in search results are due to personalization; some may due to noise. As in our prior work [10,11], we take the following precautions to minimize noise:

1. All queries for term t are run in lock-step, to avoid changes in search results due to time.
2. We statically mapped the DNS entry for the Google Search server, ensuring that all our queries were sent to the same datacenter.
3. Google Search personalizes search results based on the user’s prior searches during the last 10 minutes [11]. To avoid this confound, we wait 11 minutes between subsequent queries.

However, even with these precautions, there may still be noise in search results (e.g., due to A/B testing). Thus, for each search term and location, we send two identical queries at the same time. By comparing each result with its corresponding *control*, we can measure the extent of the underlying noise. When comparing search results from two locations, any differences we see above the noise threshold can then be attributed to location-based personalization.

Parsing. As shown in Figure 1, Google Search on mobile renders search results as “cards”. Some cards present a single result (e.g., “Somerville Schools”), while others present a meta-result (e.g., locations from Google Maps or a list of “In the News” articles). In this study, we parse pages of search results by extracting the first link from each card, except for Maps and News cards where we extract all links. Thus, we observe 12–22 search results per page.

2.3 Measuring Personalization

As in our prior work [11], we use two metrics to compare pages of search results. First, we use *Jaccard Index* to examine the overlap: a Jaccard Index of 0 represents no overlap between the pages, while 1 indicates they contain the same search results (although not necessarily in the same order). Second, we use *edit distance* to measure reordering of search

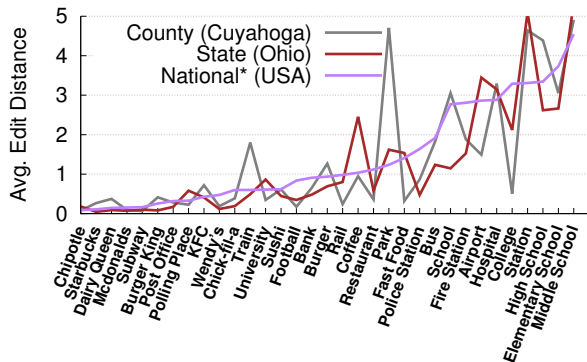


Figure 3: Noise levels for local queries across three granularities.

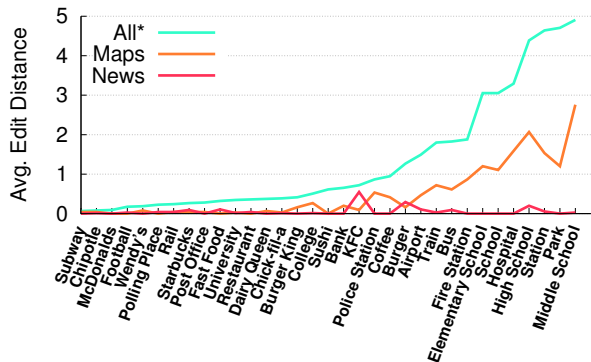


Figure 4: Amount of noise caused by different types of search results for local queries.

results. Edit distance calculates the number of additions, deletions, and swaps necessary to make two lists identical.

3. ANALYSIS AND FINDINGS

Using the methodology described in Section 2, we collected 30 days of data from Google Search. We executed the 120 local and controversial queries once per day for five straight days in the county, state, and national locations (so, 15 days total). We then repeated this process with the 120 politicians. Using this dataset, we analyze the impact of location-based personalization on Google Search results.

3.1 Noise

To start, we examine whether there is noise in our search results. To calculate noise, we compare the search results received by treatments and their controls, i.e., two browsers that are running the same queries at the same time from the same locations.

Unlike prior work [11], we find that Google Search results are noisy. Figure 2 shows the average Jaccard Index and edit distance for all treatment/control pairs broken down by granularity and query types (values are averaged over all queries of the given type over 5 days). We make three observations. First, we see that local queries are much noisier than controversial and politician queries, in terms of result composition (shown by Jaccard) and reordering (shown by edit distance). Second, not only do local queries have more differences on average, but we also see that they have more variance (indicated by the standard deviation error bars). Third, we observe that noise is independent of location, i.e., the level of noise is uniform across all three granularities.

Search Terms. Given the high standard deviations for local queries, we pose the question: *do certain search terms exhibit more noise than others?* To answer this, we calculate the Jaccard Index and edit distance for each search term separately. Figure 3 shows the local queries along the x-axis, with the average edit distance for each query along the y-axis. The three lines correspond to search results gathered at different granularities; for clarity, we sort the x-axis from smallest to largest based on the national locations.

Figure 3 reveals a divide between the queries: brand names like “Starbucks” tend to be less noisy than generic terms like “school”. We observe similar trends for Jaccard Index. We examine this observation further next, when we look at the impact of different types of search results.

Search Result Types. To isolate the source of noise, we analyze the types of search results returned by Google Search. As described in Section 2.2, Google Search returns “typical” results, as well as Maps and News results. We suspect that Maps and News results may be more heavily impacted by location-based personalization, so we calculate the amount of noise that can be attributed to search results of these types separately. Intuitively, we simply calculate Jaccard and edit distance between pages after filtering out all search results that are not of type t .

Figure 4 shows the amount of noise contributed by Maps and News results for each query, along with the overall noise. Figure 4 focuses on the edit distance for local queries at county granularity, but we see similar trends at other granularities, and for Jaccard values. We observe that Maps results are responsible for around 25% of noise (calculated as the total number of search result changes due to Maps, di-

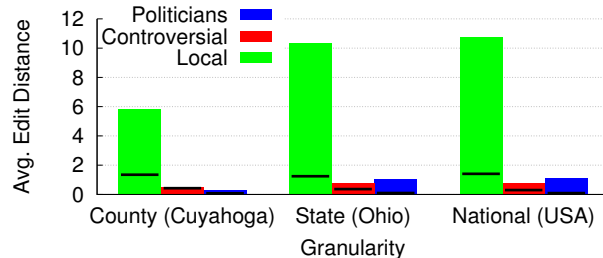
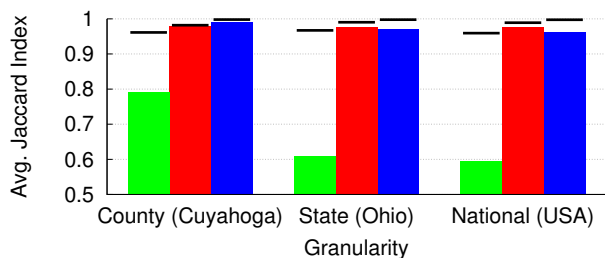


Figure 5: Average personalization across different query types and granularities. Black bars shows average noise levels from Figure 2.

vided by the overall number of changes), while News results cause almost zero noise. After some manual investigation we found that most differences due to Maps arise from one page having Maps results and the other having none. However, we also found cases where both queries yield Maps that highlight a different set of locations. Surprisingly, searches for specific brands typically do not yield Maps results, hence the low noise levels for those search terms.

Although we do not show the findings here due to space constraints, we observe the reverse effect for *controversial* queries: 6-17% of noise in such queries is due to News, while close to 0 is due to Maps. However, as Figure 2 shows, the level of noise in *controversial* queries is low overall.

3.2 Personalization

Now that we have quantified the noise in our dataset, we focus on answering the following two questions. First, *do certain types of queries trigger more personalization than others?* Second, *how does personalization change as the distance between two locations grows?*

Figure 5 shows the average Jaccard Index and edit distance values for each query category at each granularity. Values are averaged across all queries of the given types across 5 days. Recall that in the previous section, we were comparing treatments to their controls in order to measure noise; in this section, we are comparing all pairs of treatments to see if search results vary by location. For the sake of comparison, the average noise levels seen in Figure 2 are shown as horizontal black lines in Figure 5.

The first takeaway from Figure 5 is that *local* queries are much more personalized than *controversial* and *politicians* queries. The Jaccard index shows that 18-34% of the search results vary based on location for *local* queries, while the edit distance shows that 6-10 URLs are presented in a different order (after subtracting the effect of noise). *Controversial* and *politician* queries also exhibit small differences in Figure 5, but the Jaccard and edit distance values are very close to the noise-levels, making it difficult to claim that these changes are due to personalization.

The second takeaway from Figure 5 is that personalization increases with distance. The change is especially high between the *county*- and *state*-levels, with 2 additional search results changed and 4 reordered. As expected, this indicates that differences due to location-based personalization grow with geographic distance.

Search Terms. Our next step is to examine how personalization varies across search terms. As before, we focus on *local* queries since they are most impacted by personalization. Figure 6 shows the edit distances for each *local* search term at each granularity (with the *x*-axis sorted by the *national*-level values). The significant increase in personalization between *county*- and *state*-level search results is again apparent in this figure.

Overall, we see that location-based personalization varies dramatically by query. The number of search results that change is between 5 and 17, where 17 is essentially all search results on the page. We also notice that (similar to our observations about noise) general terms such as “school” or “post office” exhibit higher personalization than brand names.

The analogous plots for *politicians* and *controversial* queries show similar trends as Figure 6, but with much lower overall personalization. However, there are a few exceptional




Figure 6: Personalization of each search term for *local* queries.

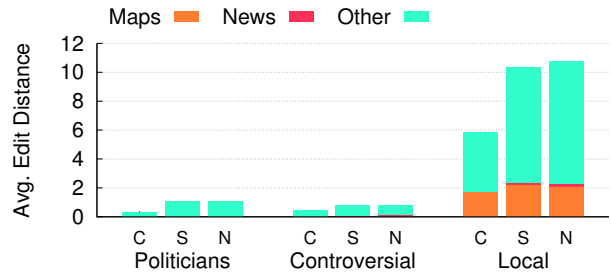


Figure 7: Amount of personalization caused by different types of search results.

search terms. In the case of *politicians*, these exceptions are common names such as “Bill Johnson” or “Tim Ryan”, so it is likely that the differences stem from ambiguity. In the case of *controversial* terms, the most personalized queries are “health”, “republican party”, and “politics”.

Search Result Types. It is not terribly surprising that Google personalizes Maps and News results based on location. However, we find that personalization of Maps and News results only explains a small portion of the differences we observe.

Figure 7 breaks down the overall edit distance values into components corresponding to News, Maps, and all other search results, for each granularity and query type. For *controversial* queries, 6-18% of the edit distance can be attributed to News results, and interestingly, this fraction increases from *county* to *nation* granularity. A different composition is seen for *local* queries: 18-27% of differences are caused by Maps results. The takeaway is that, surprisingly, the vast majority of changes due to location-based personalization impact “typical” results.

Consistency Over Time. Thus far, all of our plots have presented values averaged over 5 days. To determine whether personalization is consistent over time, we plot Figure 8. In this figure, we choose one location in each granularity to serve as the *baseline*. The red line plots the average edit distance when comparing the baseline to its control (i.e., the red line shows the noise floor); each black line is a comparison between the baseline and another location at that granularity. We focus on *local* queries since they are most heavily personalized.

Figure 8 shows that the amount of personalization is stable over time. *Politicians* and *controversial* terms show the

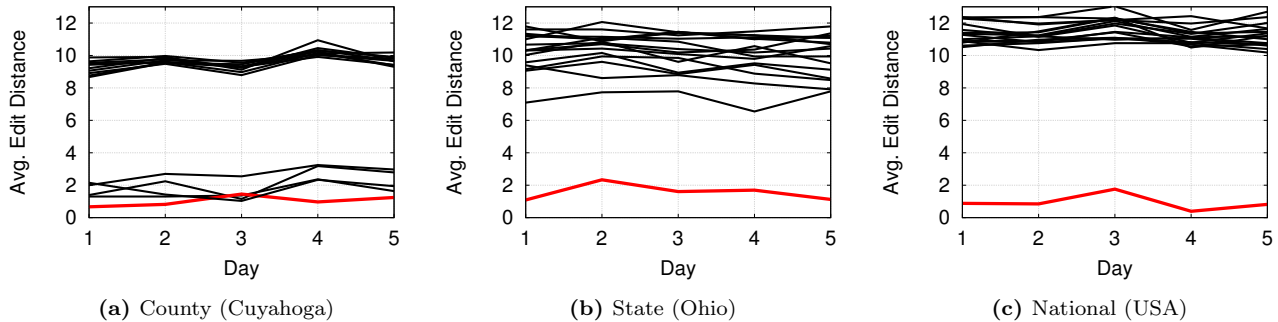


Figure 8: Personalization of 25 locations, each compared to a baseline location, for *local* queries. The red line compares two treatments at the baseline location (i.e., the experimental control), and thus shows the noise floor.

same trend but with lower personalization overall (findings not shown). As expected, we see a wide gulf between the baseline and other locations at *state* and *nation* granularity, since search results are extremely different at these long distances. However, interestingly, we see that some locations “cluster” at the *county*-level, indicating that some locations receive similar search results to the baseline.

Demographics. To investigate why certain locations cluster at the *county*-level, we examined many potential correlations between all pairs of *county*-level locations. This included correlations based on distance (i.e., do closer locations tend to cluster), as well as 25 demographic features like population density, poverty, educational attainment, ethnic composition, English fluency, income, *etc.* Unfortunately, we were unable to identify any correlations that explain the clustering of locations. Based on this analysis, it appears that Google Search does not use demographic features to implement location-based personalization.

4. RELATED WORK

Search Personalization. Many researchers have investigated strategies for personalizing search engines in order to increase the quality of results [8, 17, 18]. Dou et al. and Micarelli et al. survey several different personalization techniques [4, 14] to determine what features improve search results the most. Several studies have specifically focused on the importance of location in search personalization: [3, 26] use linguistic tools to infer geo-intention from search queries, while [25, 26] focuses on location relevance of webpage content to the given search query.

Auditing Algorithms. In contrast to studies that aim to develop new personalization algorithms, a recent line of work measures deployed personalization systems to understand their impact on users. Latanya Sweeney examined Google AdSense and uncovered that the system serves ads in a racially biased manner [22]. Our prior work [11] as well as Bobble [24] examine how Google Search personalizes search results, and find that geolocation is one of the features used by the algorithm. However, these studies only examine the impact of IP address geolocation, and only at course-grained locations (e.g., different states and countries). Other studies have examined the effects of algorithmic personalization on the Facebook News Feed [5, 6], e-commerce [10, 15, 16], and online ads [7, 13].

5. CONCLUDING DISCUSSION

In this paper, we present a detailed analysis of location-based personalization on Google Search. We develop a novel methodology that allows us to query Google from any location around the world. Using this technique we sent 3,600 distinct queries to Google Search over a span of 30 days from 59 locations across the US.

Our findings show that location does indeed have a large impact on search results, and that the differences increase as physical distance grows. However, we observe many nuances to Google’s implementation of location-based personalization. First, not all types of queries trigger the algorithm to the same degree: *politicians* are essentially unaffected by geography; *controversial* terms see small changes due to News; and *local* terms see large differences due to changes in Maps and normal results. Second, not all queries expected to trigger location-personalization do: for example, search results for brand names like “Starbucks” do not include Maps.

Finally, and most surprisingly, we also discover that Google Search returns search results that are very noisy, especially for *local* queries. This non-determinism is puzzling, since Google knows the precise location of the user (during our experiments), and thus should be able to quickly calculate the closest set of relevant locations.

Much work remains to be done. Our methodology can easily be extended to other countries and search engines. We also plan on further investigating the correlations between demographic features and search results. Additional content analysis on the search results may help us uncover the specific instances where personalization algorithms reinforce demographic biases.

The full list of query terms, as well as our source code and data, are all open-source and available at our website:

<http://personalization.ccs.neu.edu>

Acknowledgements

We thank the anonymous reviewers and our shepherd, Matthew Luckie, for their helpful comments. We also thank Arash Molavi Kakhki for developing the JavaScript reimplementation of the Geolocation API used in this project. This research was supported in part by NSF grants CNS-1054233, CNS-1319019, and CHS-1408345. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

6. REFERENCES

- [1] Alexa Top 500 Global Sites. <http://www.alexa.com/topsites>.
- [2] J. Burn-Murdoch. US web statistics released for May 2012: which sites dominate, and where do we go for online news? *The Guardian*, 2012.
- [3] P. N. Bennett, F. Radlinski, R. W. White, and E. Yilmaz. Inferring and Using Location Metadata to Personalize Web Search. *SIGIR*, 2011.
- [4] Z. Dou, R. Song, and J.-R. Wen. A Large-scale Evaluation and Analysis of Personalized Search Strategies. *WWW*, 2007.
- [5] M. Eslami, A. Aleyasen, K. Karahalios, K. Hamilton, and C. Sandvig. FeedVis: A Path for Exploring News Feed Curation Algorithms. *CSCW*, 2015.
- [6] M. Eslami, A. Rickman, K. Vaccaro, A. Aleyasen, A. Vuong, K. Karahalios, K. Hamilton, and C. Sandvig. “I always assumed that I wasn’t really that close to [her]”: Reasoning about invisible algorithms in the news feed. *CHI*, 2015.
- [7] S. Guha, B. Cheng, and P. Francis. Challenges in Measuring Online Advertising Systems. *IMC*, 2010.
- [8] S. Gauch, J. Chaffee, and A. Pretschner. Ontology-based personalized search and browsing. *Web Intelligence and Agent Systems*, 1, 2003.
- [9] Google. Personalized Search Graduates from Google Labs. *News From Google Blog*, 2005. http://googlepress.blogspot.com/2005/11/personalized-search-graduates-from_10.html.
- [10] A. Hannak, G. Soeller, D. Lazer, A. Mislove, and C. Wilson. Measuring Price Discrimination and Steering on E-commerce Web Sites. *IMC*, 2014.
- [11] A. Hannak, P. Sapiezynski, A. M. Kakhki, B. Krishnamurthy, D. Lazer, A. Mislove, and C. Wilson. Measuring Personalization of Web Search. *WWW*, 2013.
- [12] HTML5 Geolocation API. <http://dev.w3.org/geo/api/spec-source.html>.
- [13] M. Lecuyer, G. Ducoffe, F. Lan, A. Papancea, T. Petsios, R. Spahn, A. Chaintreau, and R. Geambasu. XRay: Enhancing the Web’s Transparency with Differential Correlation. *USENIX Security*, 2014.
- [14] A. Micarelli, F. Gaspiretti, F. Sciarrone, and S. Gauch. Personalized Search on the World Wide Web. *The Adaptive Web*, Peter Brusilovsky, Alfred Kobsa, and Wolfgang Nejdl, eds., Springer-Verlag, 2007.
- [15] J. Mikians, L. Gyarmati, V. Erramilli, and N. Laoutaris. Detecting Price and Search Discrimination on the Internet. *HotNets*, 2012.
- [16] J. Mikians, L. Gyarmati, V. Erramilli, and N. Laoutaris. Crowd-assisted Search for Price Discrimination in E-Commerce: First results. *CoNEXT*, 2013.
- [17] M. G. Noll and C. Meinel. Web Search Personalization via Social Bookmarking and Tagging. *Proc. of The Semantic Web and 2nd Asian Conference on Asian Semantic Web Conference*, 2007.
- [18] A. Pretschner and S. Gauch. Ontology based personalized search. *ICTAI*, 1999.
- [19] E. Pariser. *The Filter Bubble: What the Internet is Hiding from You*. Penguin Press, 2011.
- [20] PhantomJS. 2015. <http://phantomjs.org>.
- [21] Right to be Forgotten ruling. http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet_data_protection_en.pdf.
- [22] L. Sweeney. Discrimination in Online Ad Delivery. *SSRN*, 2013.
- [23] D. Y. Wang, M. Der, M. Karmai, L. Saul, D. McCoy, S. Savage, and G. M. Voelker. Search + Seizure: The Effectiveness of Interventions on SEO Campaigns. *IMC*, 2014.
- [24] X. Xing, W. Meng, D. Doozan, N. Feamster, W. Lee, and A. C. Snoeren. Exposing Inconsistent Web Search Results with Bobble. *PAM*, 2014.
- [25] B. Yu and G. Cai. A query-aware document ranking method for geographic information retrieval. *GIR*, 2007.
- [26] X. Yi, H. Raghavan, and C. Leggetter. Discovering Users’ Specific Geo Intention in Web Search. *WWW*, 2009.