A Deeper Investigation of the Importance of Wikipedia Links to Search Engine Results

NICHOLAS VINCENT, Northwestern University, USA
BRENT HECHT, Northwestern University, USA

A growing body of work has highlighted the important role that Wikipedia's volunteer-created content plays in helping search engines achieve their core goal of addressing the information needs of hundreds of millions of people. In this paper, we report the results of an investigation into the incidence of Wikipedia links in search engine results pages (SERPs). Our results extend prior work by considering three U.S. search engines, simulating both mobile and desktop devices, and using a spatial analysis approach designed to study modern SERPs that are no longer just "ten blue links". We find that Wikipedia links are extremely common in important search contexts, appearing in 67-84% of desktop SERPs for *common* and *trending* queries, but less often for *medical* queries. Furthermore, we observe that Wikipedia links often appear in "Knowledge Panel" SERP elements and are in positions visible to users without scrolling, although Wikipedia appears less often and in less prominent positions on mobile devices. Our findings reinforce the complementary notions that (1) Wikipedia content and research has major impact outside of the Wikipedia domain and (2) powerful technologies like search engines are highly reliant on free content created by volunteers.

4

CCS Concepts: • **Human-centered computing → Empirical studies in collaborative and social computing**

KEYWORDS
Wikipedia; search engines; user-generated content; data leverage

ACM Reference Format:
Nicholas Vincent and Brent Hecht. 2021. A Deeper Investigation of the Importance of Wikipedia Links to Search Engine Results. In *PACM on Human Computer Interaction*, Vol. 5, No. CSCW1, Article 4, April 2021. ACM, New York, NY, USA. 18 pages. https://doi.org/10.1145/3449078

1 INTRODUCTION

Previous work has highlighted critical interdependencies between Wikipedia and Google Search [11,20,28,39]. For many query categories, Wikipedia links appear more than any other website on Google's search engine results pages (SERPs) [39]. This suggests that Wikipedia—a resource made by volunteers—plays an essential role in helping Google Search achieve its core function of addressing information needs. It also raises the stakes of Wikipedia-related research findings, with Wikipedia's collaboration processes, contribution patterns, and content outcomes being imperative to the success of Google Search.

In this paper, we replicate and extend earlier work that identified the importance of Wikipedia to the success of search engines but was limited in that it focused only on desktop results for Google and treated those results as a ranked list, i.e. as "ten blue links" [3]. Like this earlier work, we collect the first SERP for a variety of important queries and ask, "How often do Wikipedia links appear and where are these links appearing within SERPs?" However, our work includes three critical extensions that provide us with a

This was work was funded by NSF grants 1815507 and 1707296.
Authors email addresses: nickvincent@u.northwestern.edu, bhecht@northwestern.edu

https://doi.org/10.1145/3449078

deeper view into the role Wikipedia links play in helping search engines achieve their primary goals. First, we consider three popular search engines – Google, Bing, and DuckDuckGo – rather than just Google. Second, in light of extensive search engine use on mobile devices [43], we simulate queries from both desktop and mobile devices. Third, we built software to allow us to use a spatial approach to search auditing that is more compatible with modern SERPs, which have moved away from the traditional "10 blue links"; SERPs now have "knowledge panels" in a second, right-hand column, "featured answers", social media "carousels", and other elements that influence user experience. Our spatial approach to SERP analysis, which considers the location of each link within a SERP, also allows us to visually compare the representation of our SERP data to screenshots of SERPs. This is important because search engines are proprietary, opaque, and subject to frequent changes (e.g. Google has previously redesigned their SERPs [31]), and it is critical to validate that data being used for quantitative search auditing analysis accurately reflects how SERPs appear to users. We are sharing our data and code for replication purposes and to support additional search audits taking this approach.[1]

After running our audit, we found that the results identified in prior work regarding Google desktop search largely extend to other search engines and, with a smaller effect size, to mobile devices. For desktop SERPs, Wikipedia appeared in 81-84% of SERPs for our *common* queries, 67-72% for our *trending* queries, and 16-54% for our *medical* queries. Results are similar, though somewhat lower, for mobile devices. Furthermore, using our spatial analysis approach, we observe that for many of these queries, Wikipedia appears in the prominent area of SERPs visible without scrolling, through both Knowledge Panel elements and traditional blue links. We did, however, observe that Wikipedia appears less prominently for mobile SERPs, though they are still quite prominent and, in several situations, appear more often than they do in the desktop Google SERP case.

Our results have implications for a variety of constituencies. Past work suggested that Wikipedia was very important to desktop Google users. On one hand, this means search users stand to benefit from the high-quality knowledge produced by the Wikipedia community. On the other hand, this means search results are impacted by Wikipedia's weaknesses such as content biases. Our results suggest that this relationship extends beyond just desktop users of Google, although the presence of Wikipedia is smaller in mobile SERPs.

Finally, our results also reinforce the importance of volunteer-created data to the success of intelligent technologies. As we discuss below, by measuring how often and where Wikipedia links appear in SERPs, we can better understand to what degree and how the volunteer labor underlying Wikipedia fuels search engines, some of the most lucrative, important and widely used intelligent technologies. These findings suggest a need to rethink the power dynamics and economic relationships between people who generate content and the intelligent technologies that rely on this content.

[1] See https://github.com/nickmvincent/se-scraper for archival repository of code used for the original dataset collection. See https://github.com/nickmvincent/LinkCoordMin for more recently updated SERP collection software.

2 RELATED WORK

This paper builds on research that has studied search engines and SERPs, as well as research that has specifically studied the role of Wikipedia articles in search results.

In McMahon et al.'s 2017 experimental study [20], a browser extension hid Wikipedia links from Google users and the authors observed a large drop in SERP click-through rate, possibly the most important search success metric [6]. This study, and follow-up work from Vincent et al. [39], which measured the prevalence of user-generated content in SERPs, were motivated by a call from the Wikimedia Foundation to study the re-use of Wikipedia content [35]. A key concern was raised by McMahon et al. regarding how Google SERPs provided peer-produced content with attribution. McMahon et al. also discussed the "paradox of reuse": if SERPs do not provide searchers with links to Wikipedia (or other peer production platforms), they may deprive peer production communities of new members. We return to questions around attribution and reuse in SERPs below in our Discussion.

These projects are part of a broader body of search auditing literature, which has used auditing techniques (i.e. collecting the outputs of search engines) to study aspects of search such as personalization and the special components, such as "Knowledge Panels", that appear in SERPs [7,15,29]. Notably, in a study that focused on political SERPs, Robertson et al. [29] identified a large number of Wikipedia articles in the SERP data they collected, in line with the results from Vincent et al.'s user-generated content-focused study. Related work has focused specifically on the complexities of SERP elements like the Knowledge Panel [19,30]. Lurie and Mustafaraj found that Wikipedia played a critical role in populating the "Knowledge Panel" shown in news-related Google SERPs: queries about news sources with Wikipedia articles frequently had Knowledge Panels in the corresponding SERPs [19]. Rothschild et al. studied how Wikipedia links in Google Knowledge Panel components influenced user perception of online news [30].

Search auditing work has also focused on the topic of tech monopolies: an audit by Jeffries and Yin found that much of the content in Google SERPs was Google's own products (e.g. modern SERP components likes Maps and answer boxes) [12]. This work provides motivation for search auditors to consider modern SERP components. Additionally, it has motivated development on a "Simple Search" tool that returns to a "ten blue links" SERP [36].

The study that we most directly replicate and extend is Vincent et al.'s 2019 study, which used a ranking-based analysis to examine the role user-generated content played in search results for the desktop version of Google search. It found that the Wikipedia domain appeared more than any other domain [39]. The results from Vincent et al.'s study provided motivation for us to focus specifically on Wikipedia while considering more search engines, mobile devices, and a spatial analysis approach.

It is also valuable to note that in addition to the academic literature on Wikipedia and search, a number of studies conducted by search engine optimization (SEO) companies have also measured how frequently Wikipedia appeared in SERPs [4,5,11,28]. Estimated incidence rates range from 34% to 99%, depending on the choice of queries (a point to which we will return below).

Our analysis was additionally motivated by prior work that studied how users interact with SERPs. Prior work has shown that visual attention and clicks are directed towards higher-ranking links, especially the first two links. Papoutsaki et al. replicated three seminal search behavior studies using a webcam eye tracker [22]. In addition to replicating the concentration of attention and clinks on top-ranked links, they found in some cases that ads in right-hand column of a SERP can attract about as much visual attention as the 3^{rd} and 4^{th} links.

In designing our spatial analysis, we were influenced by the above findings relating to the role Wikipedia plays in populating Google's Knowledge Panel, the robust results regarding the importance of top ranked links, and the potential value of right-hand column. Below, we describe how we used these findings to define and interpret spatial incidence rates.

3 METHODS

As mentioned above, modern SERPs include SERP components that are much more complex than "ten blue links". These multifaceted SERPs pose challenges to ranking-based approaches used in prior work, such as Vincent et al.'s study which used Cascading Style Sheets (CSS) styles to differentiate different result types (e.g. blue link vs. Knowledge Panel vs. News Carousel) and create a ranked list of these results [39]. The critical challenge with "ranking list" approaches is that they struggle to account for (1) the fact that SERP elements come in a variety of sizes and (2) the fact that SERP elements appear in a separate right-hand column. Furthermore, to differentiate results with CSS, researchers must manually identify which styles are associated with different result types, and these styles are liable to change.

In contrast to this prior work, we consider every link (i.e. each HTML <a> element) in each SERP and look at the coordinate for each link, in pixels, relative to the top left corner of the SERP, an approach specifically designed to handle modern SERPs. There are several other high-level benefits to this approach. It is easier to analyze multiple search engines (the approach requires substantially fewer hard-coded rules based on CSS styles), and we can visually validate that the representation of SERP data we analyze reflects the appearance of modern, component-heavy SERPs.

Below, we detail our data collection, validation, and analysis pipeline, and expand on the benefits of this spatial approach to SERP analysis.

3.1 Search Engines and Queries

In this work, we focus on three U.S. based search engines: Google, Bing and DuckDuckGo. Google and Bing are the two most popular search engines in the United States, while DuckDuckGo is the most popular privacy-focused search engine [46]. Market share estimates from different analytics services vary. In January 2020, analytics company StatCounter estimated Google served 81% of U.S. desktop queries, Bing served 12%, and DuckDuckGo served 1.5%, and for mobile devices Google served 95%, Bing served 1.5%, and DuckDuckGo served 1.2% [47]. Data from marketing research company ComScore suggests Bing has a larger market share; it reports that Google served about 62% of desktop search queries while Bing served about 25% [48].

Query selection is a critical and challenging aspect of search auditing work. Query datasets are highly proprietary, and in the past sharing queries has harmed users' privacy [2]. For these reasons, companies do not make query data available to sample from. What constitutes an "important" search query may change with current events and other factors. The results of a search audit are highly contextual based on queries used. To address these challenges, we followed best practices in prior work [7,39] closely and aimed to use queries that we believe to very important to a large fraction of search engine users.

Specifically, deferring to prior work [39], we consider queries from three important query categories: *common* queries, *trending* queries, and *medical* queries. *Common* queries are made very frequently, and thus are important by virtue of query volume. *Trending queries* are those that received a spike in interest, typically relating to news and current events. Thus, SERPs returned for *trending* queries can influence how search engine users learn about what is happening in the world. Finally, *medical* queries concern medical issues, and have the potential to impact users' health decisions. Notably, past work looking at Google SERPs found a lower prevalence of Wikipedia links for *medical* queries than other search categories [39]. In all three cases, the queries we used in our study are queries that we can plausibly assume are made very frequently, i.e., the corresponding SERPs were seen by a large group of search engine users.

For each query category, we needed to seek public sources of query data because query datasets are highly proprietary. Although search engine operators do not share their raw data about query volume, search engine optimization (SEO) companies collect data to estimate query volumes. Thus, to create a list of *common* queries, we took the top 100 queries by volume in October 2019, as estimated and made public by SEO company ahrefs.com [34] (we use their list that filters out "not safe for work" queries). To create a list of *trending* queries, we took all 282 queries from Google's public set of top trending queries from 2018 (we prepared our queries in 2019 and collected data in February of 2020) [44]. Unlike the ahrefs data, Google's trending query data contains no information about query volume – it provides the top trending queries for pre-selected topics (e.g., politics, music, etc.). While five of these query categories were in Spanish, the rest were English language queries (we used all queries regardless of language). Finally, to obtain important *medical* queries, we used a list of the top 50 medical queries made public by Bing for previous information retrieval research [33]. In total, we consider 432 different important queries.

Although we designed our query sets to be similar to those used in prior work, our query categories also map to different types of search intent. In the search literature, queries are often classified as "navigational" (a user wishes to navigate to some website, e.g. Facebook, via a SERP), "informational" (a user wishes to learn information about the query), and "transactional" (a user wishes to make some transaction) [10]. Past work suggests that users can have large differences in search behavior for different query intents, spending more time examining results for informational queries [22]. While there is not a clear objective mapping between individual queries and user intent categories (i.e. a single query could be associated with different user intents), considering approximate associations between our query categories and user intents is useful in identifying potential implications of our results for search users.

The *common* category appears to correspond to navigational queries, i.e. most queries refer to a major websites and companies. The top five queries are "facebook", "youtube", "amazon", "gmail", and "google". On the other hand, the *trending* category appears to primarily correspond to informational queries. Example trending queries include "World Cup" (a global sports event), "thank u, next" (a popular music album released in 2018), "How to vote", "Alexandria Ocasio-Cortez" (a U.S. politician), and "What is Fortnite". Finally, the *medical* queries we use also seem to be *informational*, e.g. "indigestion", "how to lose weight", "common cold symptoms", "acid reflux", and "can't sleep". The full list of queries is available with the code package[1].

The mappings of our query categories to user intents described above allows us to discuss our results in light of these user intents. For instance, what does it mean to have a high incidence rate of Wikipedia links for navigational queries? Many people searching for "facebook" (the most searched query, with an estimated monthly volume of 232 million searches) are likely not seeking information about Facebook. Nevertheless, as we will see in our results and discuss further below, desktop users are likely to be exposed to Wikipedia content as part of their navigational query, whereas mobile users are not.

3.2 Data Collection

We collected data programmatically using software we built that extends the *se-scraper* JavaScript library. The *se-scraper* library uses *puppeteer* to run a headless Chrome browser that can make search queries, save SERP results as HTML, and record screenshots of each page.[2] The screenshot functionality is important to our ability to validate that the SERPs collected by a headless browser resemble real SERPs. For each query, we collect only the first SERP.

Critically, the package we built stores the coordinates of each clickable hyperlink (specifically, HTML <a> elements with nonzero width and height) within a SERP. The package extracts the top coordinate (how many pixels from the top edge of the SERP is the top border of the link) and left coordinate (how many pixels from the left edge of the SERP is the left border of the link) of each link using JavaScript's *Element.getBoundingClientRect* function.

Our software can emulate mobile devices by providing a mobile user-agent string and using *puppeteer*'s "Devices" API. To collect mobile SERPs, we simulate an iPhone 10 running the default Safari browser[3]. To simulate desktop queries, we use the *se-scraper*'s default user-agent, which corresponds to Google Chrome running on a Windows 10 machine.

Past work has suggested that while geography plays a major role in the personalization of SERPs [15], this personalization does not heavily impact the incidence rate of user-generated content like Wikipedia for different locations with the United States [39]. As such, we made queries from a single urban location in the United States, which made it possible to include mobile devices and multiple search engines while limiting the time cost associated with data collection.

3.3 Data Validation

[2] *se-scraper*: https://github.com/NikolaiT/se-scraper,
puppeteer: https://github.com/puppeteer/puppeteer/
[3] We also ran additional experiments to compare iPhone results to Android results and saw very few differences. See "results_notebook_android_v_ios.html" at https://github.com/nickmvincent/LinkCoordMin

Given the proprietary and dynamic nature of search engines, SERP data can change its format unexpectedly and it is therefore critical to manually validate that the representation of SERP data being analyzed (e.g., links and their coordinates) resembles a SERP as consumed by users. Search engines frequently change the appearance of SERPs or add entirely new features, e.g. special components, as shown in Figure 1, that are very different than "ten blue links" [3]. For instance, Google recently substantially redesigned their SERPs, moving link URLs above "blue links" [31]. SERP analysis software designed to analyze SERPs that works well one day could easily fail the next when a major design change—like Google's recent revamp—is rolled out.

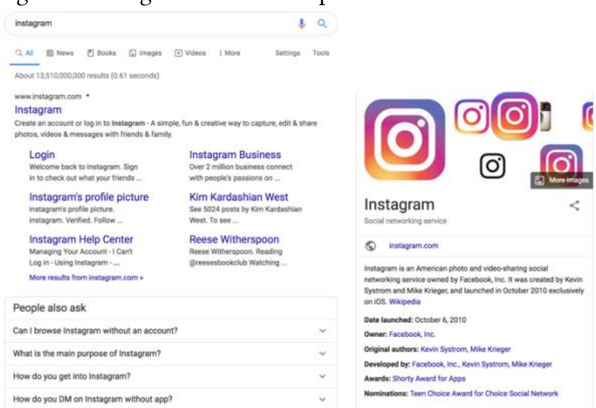

Figure 1. A Google SERP for the *common* query "Instagram" with special components, including a right column "Knowledge Panel" with a Wikipedia link.

Our software is designed to allow an auditor to validate that the representation of SERP data matches the appearance of SERPs. As described above, our software extracted all link elements and their coordinates. We created a visual representation of links at their coordinates and compared this representation to a screenshot of the SERP produced by our data collection software. The purpose of this step is to ensure that our spatial representation of SERP links reflects how a person would view the page, i.e., we ensure that our extracted link coordinates match the actual appearance of the SERP. We performed this validation for 5 random queries per configuration (device, search engine, and query category), for a total of 90 samples. This process is focused on ensuring that the extracted link data (i.e. links and their associated coordinates in the page) matches the screenshots captured during data collection.

Throughout our research process, we also engaged in manual verification checks that entailed comparing our programmatically generated SERP screenshots to SERPs generated by a separate human-operated web browser. Our goal here is to not ensure an exact match (the personalization of SERPs means that two SERPs made for the same query might not match exactly [7]), but rather to ensure that there are no major structural issues with data collection in the context of constantly changing complex SERP design.

This process was effective in identifying and fixing data collection errors in light of the challenges posed by SERP data. For instance, using our visual validation and sanity check process, we discovered several data collection nuances and inconsistencies, e.g., SERPs that partially loaded, were blocked by a "location access" prompt, or failed to load entirely.

3.4 SERP Analysis

Our analysis is based on the coordinate-based representation of SERP links described above. Specifically, our data represents hyperlinks as a triplet consisting of domain, left coordinate, and top coordinate. For instance, an English Wikipedia link with left coordinate at 200 pixels and top coordinate at 300 pixels would be represented as the triplet (*en.wikipedia, 200, 300*).

Based on our representation of SERPs, we define a series of Wikipedia *incidence rates* that measure how often Wikipedia links appear. To begin, we measure the *full-page* incidence rate, or the fraction of SERPs in which Wikipedia links appear for a combination of device, search engines, and query category. We leverage our spatial analysis to better understand where Wikipedia links appear on a SERP, an important consideration for evaluating how salient these links are.

More specifically, we define four spatial incidence rates: *above-the-fold*, *left-hand*, *right-hand*, and *left-hand-above-the-fold*. The above-the-fold incidence rate is meant to capture how often links appear "above the fold", the portion of the SERP that is visible without scrolling. We know from prior work that top-ranked links receive much more visual attention and clicks [22,42], but we also know that modern SERPs do not appear as a ranked list of ten, equally sized blue links – the rank 1 element could be a large SERP element that takes up most of the viewport (e.g. a map element). The above-the-fold incidence rate is a useful heuristic that measures how often links appear in prominent positions, while accounting for the impact of large SERP elements.

The left-hand incidence rate is a proxy for how often Wikipedia links appear in the left column (where traditional "blue links" and other SERP components appear), while the right-hand incidence rate is a proxy for how often Wikipedia links appear in Knowledge Panel components. Finally, we additionally define a left-hand-above-the-fold incidence rate to understand how often Wikipedia links appear above the fold without appearing in the Knowledge Panel. In other words, the left-hand-above-the-fold incidence rate helps us easily answer the question: "Is Wikipedia primarily appearing above the fold because it appears in Knowledge Panel components?"

To create an operational definition of left-hand and right-hand incidence rates, we manually identified a vertical dividing line between the left column and the right column containing the Knowledge Panel and related components. To obtain this line, we examined panel elements for all three search engines and found that a vertical line at 780 pixels was able to cleanly divide the left column and right column for all three search engines. We do not consider left-hand or right-hand incidence rates for our mobile SERPs, as mobile results are presented in a single column.

Creating an operational definition for above-the-fold incidence is more challenging, as devices and settings (e.g., browser zoom, browser resizing, etc.) affect what content is visible without scrolling. To address this uncertainty, we considered multiple device viewport sizes to allow us to obtain lower bound, middle ground, and upper bound estimates of above-the-fold incidence.

For mobile devices, we considered a range of viewport heights corresponding to different devices [45]. Our lower bound was 667 pixels, the height of smaller iPhone 6/7/8 devices, and close to the height of Galaxy S7 devices. Our middle ground viewport estimate was 736 (corresponding to iPhone 6/7/8 plus and Galaxy S8/9 devices). Our upper bound viewport estimate was 812 pixels (corresponding to the large iPhone X and Google Pixel 3 devices).

For desktop devices, for our middle ground estimate, we consider the common viewport height of 768 pixels [47]. Our lower bound scenario corresponds to a window at 110% zoom (i.e. zoomed in so that less content shows above the fold), our middle ground corresponds to 100% zoom (the default), and our upper bound corresponds to 90% zoom (i.e. zoomed out so that more content shows).

Overall, this range-based approach means that while our above-the-fold incidence rates make some assumptions about browsing configuration, we can observe how robust our results are against variations in these assumptions. Of course, these lower and upper bounds could be chosen based on factors other than a desktop user's zoom level (e.g., a user might resize their window without changing their zoom level, use external monitors, etc.). Our approach here is select a reasonable starting point, and future work might attempt to incorporate more such factors.

One important consideration for above-the-fold incidence is how it relates to ranking-based measurements like top-k incidence rate. For instance, based on findings that most clicks go to the top three ranked items in a SERP, past work has looked at the "top-three incidence rate" [39]. Unfortunately, there is no exact conversion between above-the-fold incidence and top-k incidence rates because SERP components in a post- "ten blue links" era have highly variable lengths. Instead, above-the-fold incidence provides a measurement approach that accounts for these variable lengths. Future work might consider more specific spatial incidence rates that take into account empirical data about click patterns for important search contexts (for instance, consider only sections of the SERP that receive on average a large number of clicks).

4 RESULTS

4.1 Full-page incidence rates

We begin by reporting how often Wikipedia appeared in our (first page) SERPs. Figure 2 shows full-page incidence rates across all combinations of devices, search engines, and query categories. As described above, we considered two devices, three search engines, and 432 total queries, so in total we collected 2592 SERPs (2 devices x 3 search engines x 432 queries). Below, we report all our results broken down by our three query categories.

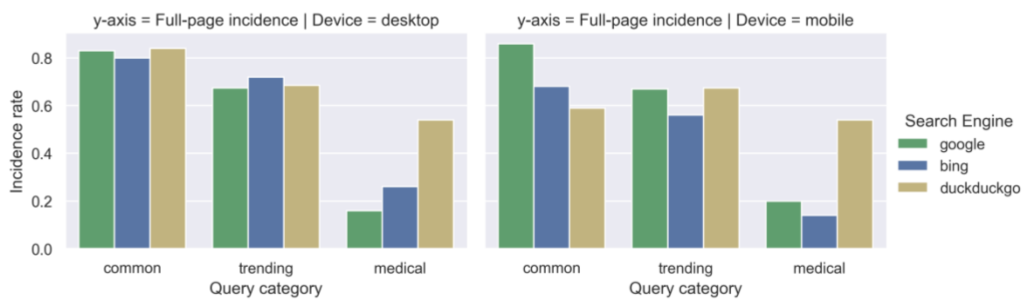

Figure 2: Full-page incidence rates for Wikipedia links. Left plot shows desktop results and right plot shows mobile results. Results are faceted by query category (along the x-axis) and search engine (color).

Looking at desktop full-page incidence across categories rates (the left half of Figure 2), we see that Wikipedia links were present in many *common* and *trending* SERPs but appear much less frequently in *medical* SERPs. Specifically, across search engines, Wikipedia appears in 80-84% of *common* SERPs and 67-72% of *trending SERPs*. However, for *medical* desktop queries, Wikipedia only appears in 16% of Google SERPs, 24% of Bing SERPs, and 54% of DuckDuckGo SERPs. Looking at the Google results, the high incidence for *common* and *trending* queries and low incidence for *medical* queries replicated Vincent et al.'s prior work which used similar query categories [39]. Furthermore, for *common* and *trending* queries, Wikipedia's large incidence rate extends across search engines.

The results in Figure 2 mean for the query categories we studied, the only major difference across search engines in Wikipedia's full-page incidence rate for desktop SERPs was for *medical* queries: DuckDuckGo shows many more Wikipedia links for these queries.

Comparing our desktop full-page incidence rates to mobile incidence rates (the right half of Figure 2), we see similar results. For Google, the largest difference in mobile and desktop full-page incidence rates is 0.04 (for the *medical* category). Bing's mobile vs. desktop differences are slightly larger (0.12-0.16). Finally, DuckDuckGo shows the largest difference for mobile results: for *common* queries, the incidence rate is 0.25 lower.

4.2 Spatial incidence rates

While the full-page incidence rates presented above provide new insight into how Wikipedia helps to serve search queries across different search engines and devices, these results do not provide insight into where Wikipedia is appearing. For this, we turn to our spatial measurements: above-the-fold, left-hand, right-hand, and left-hand-above-the-fold incidence rates. These measurements give us crucial insight into how Wikipedia links appear in SERPs in a post- "ten blue links" world.

The first spatial measurements we look at are left-hand and right-hand incidence rates. These are shown in Figure 3. Overall, right-hand incidence rates are higher than left-hand incidences rates and thus closer to full-page rates. For instance, Wikipedia's right-hand incidence rate for *trending* queries (shown in the right half of Figure 3) ranges across search engines from 77-83%, very close to the full-page range of 80-84%. This suggests Knowledge Panel-style elements are a critical source of Wikipedia links in SERPs – but not the only source, as left-hand incidence rates are still substantial, e.g. 61-66% for *trending* queries.

Next, we look at our above-the-fold incidence rates. The middle ground above-the-fold rates are shown in Figure 4, while the lower bounds and upper bounds (corresponding to smaller and larger viewports) are included in Table 1, which provides a summary of all our desktop incidence rates. In general, the lower bounds and upper bounds were relatively close to the middle ground estimate. A primary take-away is that for many cases, particularly for desktop SERPs, the *above-the-fold incidence rate* is only slightly lower than the *full-page incidence rate*. This means that not only is Wikipedia appearing frequently, but it is also appearing frequently in the most prominent area of our SERPs.

Table 1. Desktop Wikipedia incidence rates for each search engine and query category.

Search Engine	Query Category	Full page	Left hand	Right hand	Above-the-fold incidence (lower bound - upper bound)	Left-hand above-the-fold incidence (lower bound - upper bound)
google	common	0.83	0.7	0.78	0.80 (0.80 - 0.80)	0.05 (0.05 - 0.06)
google	medical	0.16	0.16	0	0.08 (0.08 - 0.12)	0.08 (0.08 - 0.12)
google	trending	0.67	0.66	0.46	0.54 (0.49 - 0.56)	0.33 (0.28 - 0.37)
bing	common	0.8	0.35	0.77	0.76 (0.76 - 0.76)	0.01 (0.01 - 0.03)
bing	medical	0.26	0.26	0.04	0.08 (0.08 - 0.08)	0.08 (0.08 - 0.08)
bing	trending	0.72	0.61	0.6	0.60 (0.59 - 0.62)	0.22 (0.18 - 0.26)
duckduckgo	common	0.84	0.49	0.83	0.83 (0.83 - 0.83)	0.14 (0.08 - 0.19)
duckduckgo	medical	0.54	0.52	0.44	0.44 (0.44 - 0.44)	0.18 (0.18 - 0.20)
duckduckgo	trending	0.68	0.66	0.61	0.64 (0.63 - 0.64)	0.45 (0.40 - 0.48)

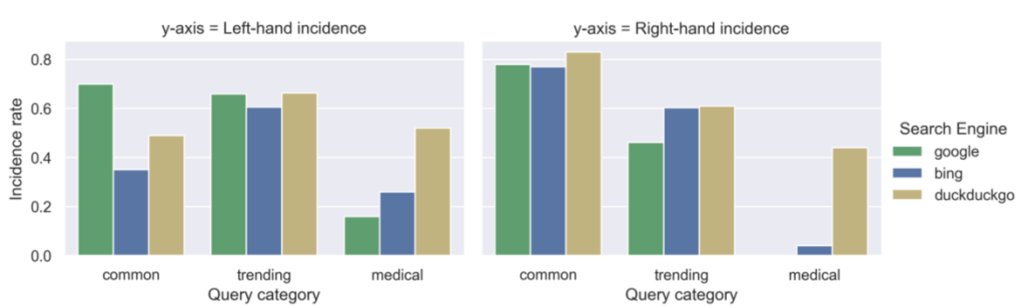

Figure 3: Left-hand incidence rates (left plot) and right-hand incidence rates (right plot) for Wikipedia links.

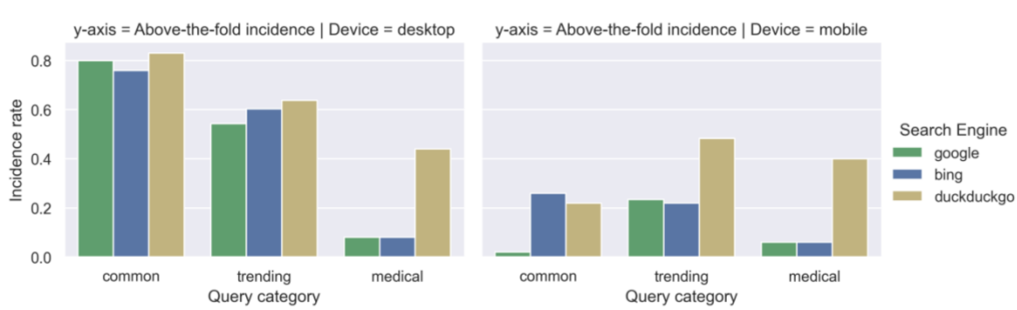

Figure 4: "Middle ground" above-the-fold incidence rates for Wikipedia links.

Table 2. Mobile Wikipedia incidence rates for each search engine and query category.

Search Engine	Query Category	Full-page incidence	Above-the-fold incidence (lower bound - upper bound)
google	common	0.86	0.02 (0.02 - 0.04)
google	medical	0.2	0.06 (0.06 - 0.08)
google	trending	0.67	0.23 (0.21 - 0.26)
bing	common	0.68	0.26 (0.22 - 0.30)
bing	medical	0.14	0.06 (0.06 - 0.06)
bing	trending	0.56	0.22 (0.21 - 0.24)
duckduckgo	common	0.59	0.22 (0.17 - 0.24)
duckduckgo	medical	0.54	0.40 (0.40 - 0.40)
duckduckgo	trending	0.67	0.47 (0.47 - 0.52)

One exception to this trend is that Google and Bing have much lower above-the-fold incidence rates for mobile devices (shown in the right half of Figure 4) than desktop devices, even when considering upper bound scenarios, i.e., even for mobile devices with large screens (lower and upper bounds for mobile above-the-fold incidence rates are shown in Table 2). However, these incidence rates are still very high (over 20% for some cases). For context, a 20% incidence rate is likely highly desirable for many websites. Given the growing use of search engines from mobile devices, differences between desktop and mobile SERPs are important to consider in understanding the importance of Wikipedia to SERPs. As we will discuss below, the mobile use case is made even more complicated by features like "Siri Suggestions", which loads website suggestions as a user types search queries in Safari (before the query is sent to Google or other search engines). This may be one way in which Wikipedia-addressable information needs are siphoned away from search engines before they get a chance to accept mobile user queries.

To fully understand the implications of Wikipedia's above-the-fold incidence rates, it is valuable to additionally consider left-hand-above-the-fold incidence rates. As described above, this incidence rate allows us to identify if Wikipedia is only appearing above the fold because it appears in Knowledge Panels. Looking at these incidence rates, shown in Table 1, we see they are substantially lower than above-the-fold rates. For *common* and *medical* queries, no left-hand-above-the-fold incidence rate is above 22%, and even for trending queries the largest rate is 45%. This result suggests that Wikipedia is primarily appearing above the fold because of the Knowledge Panel, especially for *common* queries. Only for *trending* queries does Wikipedia appear frequently above the fold as a "blue link" in the left column of our SERPs. This also provides a potential explanation for the lower above-the-fold incidence rates on mobile devices, which do not have a right column to highlight Knowledge Panel elements.

5 DISCUSSION

Our results above suggest that the important role Wikipedia has been observed to play in serving search queries extends beyond just Google, beyond just the "10 blue links", and beyond just desktop search, although the effect is smaller on mobile devices. Indeed, Wikipedia articles appear very frequently across search engines, and they often appear "above-the-fold", driven in part by Knowledge Panel SERP components.. Below, we discuss the implications of Wikipedia's prominent role in search results. We additionally discuss how this prominence relates to search intents. Finally, we revisit the limitations and caveats attached to a targeted study like the one we present here.

5.1 Wikipedia's Volunteer-created Content has Impact Outside Wikipedia

Our results reinforce an idea with important implications for Wikipedia editors and those who research Wikipedia: Wikipedia content has a huge impact well beyond the wikipedia.org website. More specifically, our results highlight that the properties of Wikipedia content will also define the effectiveness of web search, at least in many important domains. This means that the large body of research that has sought to understand or improve Wikipedia likely has broader implications that has been widely understood thus far.

A particularly significant implication is that the biases of Wikipedia content will impact search results. A large literature has sought to understand the biases of Wikipedia content (e.g. [8,13,21,26,41]). Although we did not specifically analyze content biases in our data, the fact that we observed large incidence rates for Wikipedia in general suggest that these biases are likely reflected in web search results. It is likely the case that for queries about content that is underrepresented on Wikipedia (e.g., articles about women), there will be less high-quality Wikipedia content available to answer such queries. In other words, gaps in Wikipedia's coverage or quality could lead to gaps in search engines' ability to address user information needs. This is a critical area for future work, which might explicitly how specific biases (e.g., gender, language, content about rural areas, etc.) extend to search technologies and if there are opportunities for search engines to help address biases and gaps.

There are many active efforts aimed at filling in Wikipedia's knowledge gaps (e.g. [14,27]). Because these gaps may be hurting search engines, our results suggest that search engines may have a substantial incentive to help address these gaps. At minimum, search engines should continue to attribute Wikipedia content and avoid cutting off Wikipedia from direct traffic. Looking forward, identifying opportunities to leverage the shared interests of search engines and the Wikipedia community could be fruitful for continuing to improve Wikipedia's coverage. These opportunities could include organizing events (e.g., an "edit-a-thon"), making design changes to show search users how to become Wikipedia editors, or other forms of engagement and support. For instance, if a user searches for information about a topic that is not well covered by Wikipedia, search engines could display a SERP element that indicates that the user can add information they find from their search to Wikipedia. Of course, any intervention will need to happen in dialogue with the Wikipedia community, as there are many ways in which design changes or other interventions could work against their intended outcome (e.g. if a design change leads to increased vandalism of Wikipedia, which might be particularly likely for underrepresented topics). These are issues that would be best handled through open discussion and careful consideration of existing Wikipedia practices.

5.2 For-profit Intelligent Technologies and Volunteer Content

Another important implication of the results above is that they highlight the essential role that volunteer-created content plays in fueling highly profitable intelligent technologies. Viewed through this lens, our results suggest that Wikipedia is providing critical content that helps search engines to succeed at their missions. In line with prior work [20,39]), it appears search engines would provide worse services in a world without Wikipedia. Alternatively, perhaps search engines would pay people to create a similar resource, potentially reducing the profitably of the industry. While users who primarily search from mobile devices might be less affected in a world without Wikipedia, a huge number of search results would be worse, and even mobile searchers would likely see worse search results (recall that although incidence rates were lower on mobile, they are still substantial). Relatedly, the fact that Wikipedia appears frequently in DuckDuckGo's search results suggests Wikipedia may be particularly valuable as a search result in a more privacy-focused world (DuckDuckGo does not collect any personal information).

Viewed through the lens that creating content like Wikipedia articles is a form of volunteer labor, Wikipedia can be understood as providing a "subsidy" of free labor to search engines [20,39]. However, it is also important to note that search engines provide a critical stream of traffic (both readers and potential editors) to Wikipedia [20]. This means that Wikipedia itself certainly sees benefits from its relationships with search engines and similar technologies.

That said, our results support the argument that the donations made by search engine companies to Wikipedia (e.g., [23]) likely represent only a tiny fraction of the value the Wikipedia community's labor has created for these companies. Additionally, given the Wikipedia community's important role in helping search engines serve their users, there is also an argument that the Wikipedia community should have agency in how its content is used, e.g. in discussions about how prominent Wikipedia links should be in knowledge panels [35]. An immediate concrete implication of this argument is that search engines should continue to attribute Wikipedia prominently, and when using Wikipedia content in SERP elements they should make hyperlinks easy to find. While we cannot be completely sure that search engines are always attributing Wikipedia when they present Wikipedia content, our high incidence rate results suggest that search engines are attributing Wikipedia frequently. Looking to the future, designers of search engines might seek input from communities like Wikipedia when designing new SERP elements.

More generally, our results also provide an important data point for the growing discussion about the power dynamics between tech companies and the public, and the role of data in these dynamics. Scholars [1,25,38–40] and the media [9,24] have recently expressed interest in identifying and highlighting the dependencies that large tech companies have on data coming from the public with the goal of changing the power dynamics between the public and these companies. This might mean creating new economic relationships (e.g. getting paid for data in some manner [24]) or using these dependencies for "data leverage" in order to change tech company behavior [37]. Our results highlight an underappreciated dependency, and a large one at that: peer-produced data. Moreover, the dependency in this study is of a different type than is often considered in these discussions. "Data leverage" has typically been considered in terms of personal data (e.g., data about users) [9,24] and interaction data (e.g., "trace" data that is generated as people browse websites and apps) [17,25,38]. However, our results reinforce the importance of considering data from user-generated content communities as well.

The Wikipedia community itself is unlikely to engage in direct action against search engines. The manner in which search engines use Wikipedia content is legal, and, assuming Wikipedia is attributed (as it was in our results), in line with the stated goals of Wikipedia. Therefore, it does not seem likely that Wikipedia itself would engage in extreme forms of protest against search engines (e.g., blocking search engine web crawlers or intentionally vandalizing content in protest). Instead, our results suggest that a primary role of Wikipedia in discussions around data leverage is as a key example case in which the value of the public's data contributions to intelligent technologies, computing, and related fields is measurable. Ideally, we would also measure the value of the public's contributions to private datasets (like implicit feedback used to rank search results), but there are strong incentives in place that make this unlikely (e.g. many valuable datasets are proprietary and contain personally identifiable information [2]).

The general approach of studying peer production and commons datasets to measure the value of data contributions can apply to other platforms (e.g., OpenStreetMap, citizen science platforms), and future work along these lines will be highly valuable to inform discussion around data leverage and power dynamics between users and tech companies. Additionally, political action by Wikipedia is not completely out of the question. Wikipedia engaged in a "blackout" by temporarily shutting down English Wikipedia to protest the Stop Online Piracy Act [49], and so there remains a possibility that extreme situations (e.g. if new search engines features create major harms to Wikipedia) could lead to similar actions.

It is also important to note that the results in this paper likely barely scratch the surface in terms of the ways that search engine companies rely on Wikipedia. In addition to user-facing links, search engines also use Wikipedia as foundational elements of their "knowledge graphs" and all capabilities that are dependent on their knowledge graphs [32]. Wikipedia is also used in machine translation, another service often baked into search engine results (e.g. [18]) along with other applications (e.g. [16]).

5.3 Wikipedia in SERPs and User Intent

As mentioned above, it is useful to consider our results in light of the user intent that may be associated with our query categories. While our *common* queries can be seen as *navigational* (e.g., "facebook", "youtube",), Wikipedia links are appearing for a huge number of these queries. This means even for users making these queries with navigational intent, they are still being exposed to relevant Wikipedia content. For instance, a user searching "facebook" with the intent to visit Facebook will be still be briefly exposed to Wikipedia text displayed in the Knowledge Box. This type of navigational exposure to Wikipedia content may be an interesting aspect of search engine-Wikipedia relationships for future study. In particular, how does Wikipedia's prominence affect the experience of people using search navigationally? Do users trying to login to Facebook end up reading about Facebook on Wikipedia?

Relatedly, given that both *trending* queries and *medical* queries are predominantly *informational* queries, the substantially lower incidence rates for *medical* queries compared to *trending* queries is striking. It seems Google and Bing in particular may have made design choices to highlight non-Wikipedia sources of medical information. Examining SERPs for our medical queries (which we note, were taken from a medical query dataset released by Bing itself), we see both Google and Bing present a special knowledge box SERP element for these queries. This suggests that search engine designers have implemented a feature to treat these medical queries differently than other queries and doing so impacts Wikipedia's incidence rate. Of course, it could be the case that differences in results are being driven by implicit feedback from search users, and not by platform designers.

5.4 Limitations and Future Work

It is important to emphasize the limitations of our targeted search audit. First, the choice of queries heavily impacts results such as incidence rates. For instance, it easy to construct a query set for which Wikipedia will never appear or for which Wikipedia will always appear (for instance, if we append the word "Wikipedia" to each of our queries, we can easily achieve a 100% Wikipedia incidence rate). We addressed this limitation by selecting queries that are made frequently and can be have large impacts on users (e.g. trending queries inform users about events in the world and medical queries may influence health decisions). Nonetheless, any incidence rate is conditional on query choice and investigating other query contexts is an important area of future work.

Although we considered (some) variation in device, search engine, and query category, there are other factors that impact SERPs. Most notably, search engines may have huge differences across languages and countries. This is a particularly ripe area for further study of the relationship between Wikipedia and search engines. Wikipedia covers a large number of languages, and future work could identify opportunities and pitfalls for the relationship between non-English Wikipedia and search engines. We note that although our trending queries did include some Spanish queries and corresponding Spanish es.wikipedia.org results, the Wikipedia links in our study were primarily English Wikipedia links.

Finally, future work should consider systems that modify the search process, e.g., "Siri Suggestions" or other "AI assistants". For instance, we observed "Siri Suggestions" that send users directly to Wikipedia pages without visiting SERPs several times during our confidence checks from a mobile device. These AI assistants may bypass search engines entirely, or substantially reduce the number of results a user receives. For instance, users who interact with AI assistants using a voice assistant may only get a "top 1" result, instead of 10 blue links or spatially situated SERPs. Assuming that many voice queries will tend to be more informational than navigational, it could be the case that Wikipedia's "incidence rate" for voice search and Siri Suggestions will also be very high. Future auditing work might specifically investigate voice search and AI assistants. In cases that involve a single, or few results, auditing these systems may actually be substantially simpler than auditing complex SERPs, although automating audits may be challenging.

5 CONCLUSION

In this work, we reported on a targeted investigation of how often—and where—Wikipedia links appear within SERsS for three major search engines and across desktop and mobile. By considering the precise location of links within the SERPs, we shed light on how often Wikipedia content appears "above the fold" and in "Knowledge Panel" elements that exist outside the typical ranked list of search results. We find evidence that Wikipedia's volunteer created content is important to all three search engines we studied, although the magnitude of this importance is heavily context dependent, with results varying across devices and types of queries.

6 ACKNOWLEDGMENTS

This was work was funded by NSF grants 1815507 and 1707296. Thanks to Nick Hagar, Daniel Trielli, Jack Bandy, Sohyeon Hwang, and Chau Tong for helpful feedback regarding the data collection software. Thanks to Isaac Johnson, Hanlin Li and colleagues in the Community Data Science Collective for discussions around the ideas in the paper.

REFERENCES

[1] Imanol Arrieta Ibarra, Leonard Goff, Diego Jiménez Hernández, Jaron Lanier, and E Weyl. 2018. Should We Treat Data as Labor? Moving Beyond "Free." American Economic Association Papers & Proceedings 1, 1 (2018).
[2] Michael Barbaro and Tom Zeller Jr. 2006. A Face Is Exposed for AOL Searcher No. 4417749. N.Y. Times (August 2006). Retrieved from https://www.nytimes.com/2006/08/09/technology/09aol.html
[3] Danqi Chen, Weizhu Chen, Haixun Wang, Zheng Chen, and Qiang Yang. 2012. Beyond Ten Blue Links: Enabling User Click Modeling in Federated Web Search. In Proceedings of the Fifth ACM International Conference on Web Search and Data Mining (WSDM '12), ACM, New York, NY, USA, 463–472. DOI:https://doi.org/10.1145/2124295.2124351
[4] Danny Goodwin. 2012. Wikipedia Appears on Page 1 of Google for 99% of Searches [Study] - Search Engine Watch. Retrieved from https://www.searchenginewatch.com/2012/02/13/wikipedia-appears-on-page-1-of-google-for-99-of-searches-study
[5] Danny Goodwin. 2012. Bing, Not Google, Favors Wikipedia More Often in Search Results [Study] - Search Engine Watch. Retrieved from https://www.searchenginewatch.com/2012/03/19/bing-not-google-favors-wikipedia-more-often-in-search-results-study
[6] Artem Grotov and Maarten de Rijke. 2016. Online learning to rank for information retrieval: Sigir 2016 tutorial. In Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, 1215–1218.
[7] Aniko Hannak, Piotr Sapiezynski, Arash Molavi Kakhki, Balachander Krishnamurthy, David Lazer, Alan Mislove, and Christo Wilson. 2013. Measuring personalization of web search. In Proceedings of the 22nd international conference on World Wide Web, ACM, 527–538.
[8] Benjamin Mako Hill and Aaron Shaw. 2013. The Wikipedia gender gap revisited: characterizing survey response bias with propensity score estimation. PloS one 8, 6 (2013), e65782.
[9] Chris Hughes. 2018. The wealth of our collective data should belong to all of us. The Guardian. Retrieved from https://www.theguardian.com/commentisfree/2018/apr/27/chris-hughes-facebook-google-data-tax-regulation
[10] Bernard J Jansen, Danielle L Booth, and Amanda Spink. 2007. Determining the user intent of web search engine queries. In Proceedings of the 16th international conference on World Wide Web, ACM, 1149–1150.
[11] Greg Jarboe. 2020. YouTube's Organic Visibility Tops Wikipedia in Google SERPs. Search Engine Journal (January 2020). Retrieved from https://www.searchenginejournal.com/youtube-organic-visibility-google-serps/341419
[12] Adrianne Jeffries and Leon Yin. 2020. Google's Top Search Result? Surprise! It's Google. The Markup. Retrieved from https://themarkup.org/google-the-giant/2020/07/28/how-we-analyzed-google-search-results-web-assay-parsing-tool
[13] Isaac L Johnson, Yilun Lin, Toby Jia-Jun Li, Andrew Hall, Aaron Halfaker, Johannes Schöning, and Brent Hecht. 2016. Not at home on the range: Peer production and the urban/rural divide. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, ACM, 13–25.
[14] Isaac Johnson, Florian Lemmerich, Diego Sáez-Trumper, Robert West, Markus Strohmaier, and Leila Zia. 2020. Global gender differences in Wikipedia readership. arXiv preprint arXiv:2007.10403 (2020).
[15] Chloe Kliman-Silver, Aniko Hannak, David Lazer, Christo Wilson, and Alan Mislove. 2015. Location, location, location: The impact of geolocation on web search personalization. In Proceedings of the 2015 ACM Conference on Internet Measurement Conference, ACM, 121–127.
[16] Maya Kosoff. 2018. YouTube Slaps a Feel-Good Band-Aid on Its Fake-News Problem. Vanity Fair (March 2018). Retrieved from https://www.vanityfair.com/news/2018/03/youtube-wikipedia-conspiracy-theory-video-problem
[17] Jaron Lanier and E Glen Weyl. 2018. A Blueprint for a Better Digital Society. Harvard Business Review (2018).
[18] Quoc V Le and Mike Schuster. 2016. A neural network for machine translation, at production scale. Retrieved from https://ai.googleblog.com/2016/09/a-neural-network-for-machine.html
[19] Emma Lurie and Eni Mustafaraj. 2018. Investigating the Effects of Google's Search Engine Result Page in Evaluating the Credibility of Online News Sources. In Proceedings of the 10th ACM Conference on Web Science, 107–116.
[20] Connor McMahon, Isaac L Johnson, and Brent Hecht. 2017. The Substantial Interdependence of Wikipedia and Google: A Case Study on the Relationship Between Peer Production Communities and Information Technologies. In ICWSM, 142–151.
[21] Daniel Oberhaus. 2017. Nearly All of Wikipedia Is Written By Just 1 Percent of Its Editors - Motherboard. Retrieved from https://motherboard.vice.com/en_us/article/7x47bb/wikipedia-editors-elite-diversity-foundation
[22] Alexandra Papoutsaki, James Laskey, and Jeff Huang. 2017. Searchgazer: Webcam eye tracking for remote studies of web search. In Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval, 17–26.

[23] Ben Parr. 2010. Google Gives $2 Million to Wikipedia's Foundation. Retrieved from https://mashable.com/2010/02/16/google-wikipedia-donation
[24] Eduardo Porter. 2018. Your Data Is Crucial to a Robotic Age. Shouldn't You Be Paid for It? New York Times. Retrieved from https://www.nytimes.com/2018/03/06/business/economy/user-data-pay.html
[25] Eric A Posner and E Glen Weyl. 2018. Radical Markets: Uprooting Capitalism and Democracy for a Just Society. Princeton University Press.
[26] Joseph Reagle and Lauren Rhue. 2011. Gender bias in Wikipedia and Britannica. International Journal of Communication 5, (2011), 21.
[27] Miriam Redi, Martin Gerlach, Isaac Johnson, Jonathan Morgan, and Leila Zia. 2020. A Taxonomy of Knowledge Gaps for Wikimedia Projects (First Draft). arXiv preprint arXiv:2008.12314 (2020).
[28] Luke Richards. 2018. Why Wikipedia is still visible across Google's SERPs in 2018 - Search Engine Watch. Retrieved from https://www.searchenginewatch.com/2018/11/13/why-wikipedia-is-still-visible-across-googles-serps-in-2018
[29] Ronald E Robertson, David Lazer, and Christo Wilson. 2018. Auditing the Personalization and Composition of Politically-Related Search Engine Results Pages. In Proceedings of the 2018 World Wide Web Conference on World Wide Web, International World Wide Web Conferences Steering Committee, 955–965.
[30] Annabel Rothshild, Emma Lurie, and Eni Mustafaraj. 2019. How the Interplay of Google and Wikipedia Affects Perceptions of Online News Sources. In Computation+ Journalism Symposium.
[31] Jonathan Shieber. 2020. Google backtracks on search results design. TechCrunch (January 2020). Retrieved from https://techcrunch.com/2020/01/24/google-backtracks-on-search-results-design
[32] Amit Singhal. 2012. Introducing the knowledge graph: things, not strings. Official google blog 16, (2012). Retrieved from https://googleblog.blogspot.com/2012/05/introducing-knowledge-graph-things-not.html
[33] Luca Soldaini, Andrew Yates, Elad Yom-Tov, Ophir Frieder, and Nazli Goharian. 2016. Enhancing web search in the medical domain via query clarification. Information Retrieval Journal 19, 1–2 (2016), 149–173.
[34] Tim Soulo. 2019. Top Google searches (as of October 2019). Retrieved from https://ahrefs.com/blog/top-google-searches
[35] Dario Taraborelli. 2015. The Sum of All Human Knowledge in the Age of Machines: A New Research Agenda for Wikimedia. ICWSM-15 Workshop on Wikipedia, a Social Pedia: Research Challenges and Opportunities,.
[36] Maddy Varner and Sam Morris. 2021. Introducing Simple Search – The Markup. The Markup. Retrieved from https://themarkup.org/google-the-giant/2020/11/10/introducing-simple-search
[37] Nicholas Vincent and Brent Hecht. 2020. Can "Conscious Data Contribution" Help Users to Exert "Data Leverage" Against Technology Companies?
[38] Nicholas Vincent, Brent Hecht, and Shilad Sen. 2019. "Data Strikes": Evaluating the Effectiveness of New Forms of Collective Action Against Technology Platforms. In Proceedings of The Web Conference 2019.
[39] Nicholas Vincent, Isaac Johnson, Patrick Sheehan, and Brent Hecht. 2019. Measuring the Importance of User-Generated Content to Search Engines. In Proceedings of AAAI ICWSM 2019.
[40] Nicholas Vincent, Hanlin Li, Nicole Tilly, Stevie Chancellor, and Brent Hecht. 2021. Data Leverage: A Framework for Empowering the Public in its Relationship with Technology Companies. In ACM FAccT 2021 (formerly FAT*).
[41] Claudia Wagner, David Garcia, Mohsen Jadidi, and Markus Strohmaier. 2015. It's a Man's Wikipedia? Assessing Gender Inequality in an Online Encyclopedia. In ICWSM, 454–463.
[42] Ryen W White, Fernando Diaz, and Qi Guo. 2017. Search result prefetching on desktop and mobile. ACM Transactions on Information Systems (TOIS) 35, 3 (2017), 1–34.
[43] 2015. It's Official: Google Says More Searches Now On Mobile Than On Desktop - Search Engine Land. Retrieved from https://searchengineland.com/its-official-google-says-more-searches-now-on-mobile-than-on-desktop-220369
[44] 2018. Google Trends. Retrieved from https://trends.google.com/trends/hottrends
[45] 2018. Popular Screen Resolutions - Media Genesis. Retrieved from https://mediag.com/blog/popular-screen-resolutions-designing-for-all
[46] 2020. Web search engine - Wikipedia. Retrieved from https://en.wikipedia.org/wiki/Web_search_engine#Market_share
[47] 2020. StatCounter Global Stats - Browser, OS, Search Engine including Mobile Usage Share. Retrieved from https://gs.statcounter.com
[48] 2020. ComScore US Search Market Share. Retrieved from https://www.comscore.com/Insights/Rankings?country=US#tab_search_share
[49] Protests against SOPA and PIPA - Wikipedia. Retrieved from https://en.wikipedia.org/wiki/Protests_against_SOPA_and_PIPA

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