



Original Manuscript

“Foreign beauties want to meet you”: The sexualization of women in Google’s organic and sponsored text search results

new media & society

1–22

© The Author(s) 2022



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/14614448221099536

journals.sagepub.com/home/nms



Aleksandra Urman 

University of Zurich, Switzerland

Mykola Makhortykh 

University of Bern, Switzerland

Abstract

Search engines serve as information gatekeepers on a multitude of topics dealing with different aspects of society. However, the ways search engines filter and rank information are prone to biases related to gender, ethnicity, and race. In this article, we conduct a systematic algorithm audit to examine how one specific form of bias, namely, sexualization, is manifested in Google’s text search results about different national and gender groups. We find evidence of the sexualization of women, particularly those from the Global South and East, in search outputs in both organic and sponsored search results. Our findings contribute to research on the sexualization of people in different forms of media, bias in web search, and algorithm auditing as well as have important implications for the ongoing debates about the responsibility of transnational tech companies for preventing systems they design from amplifying discrimination.

Keywords

Algorithm auditing, bias, critical Internet studies, gender, Google, race, sexualization, web search

Corresponding author:

Aleksandra Urman, Social Computing Group, Department of Informatics, University of Zurich, 8006 Zurich, Switzerland.

Email: urman@ifi.uzh.ch

Introduction

Search engines are major information gatekeepers that determine how information is filtered and what sources are prioritized in response to user queries. However, similar to other complex algorithmic systems, they are prone to biases that can arise for different reasons such as the quality of the data on which an algorithm is trained or societal/individual prejudices affecting design decisions (Bozdag, 2013). Bias in web search outputs can be broadly defined as the distortion of the information landscape (Grimmelmann, 2010; Mowshowitz and Kawaguchi, 2002), for instance, through a systematic skewness of results toward specific perspectives (Kay et al., 2015) or the imbalances in the prevalence of specific sources in search outputs (Goldman, 2008). One specific form of bias in search outputs is social bias—a systematic misrepresentation of individuals or groups that can amplify their discrimination and/or marginalization (Otterbacher et al., 2017). In this article, we focus on this form of bias.

Such systematic misrepresentations are concerning because search results can influence public opinion and perceptions of the social reality (Epstein and Robertson, 2015; Fisher et al., 2015; Kulshrestha et al., 2019). This is aggravated by the users' tendency to trust the output of search engines (Pan et al., 2007). In 2021, 56% of the population globally said they trust the news they find on search engines, putting search engines ahead of all other news sources, including traditional media in terms of user trust in news sources (Edelman Trust Barometer, 2021). In practice, however, web search results are not impartial and often reiterate existing racial and gender biases (Noble, 2018). Consequently, the reiteration of social biases via search engines has the potential for amplifying societal discrimination and exacerbating prejudices toward specific groups which in some cases might even provoke violence against them (Wright and Tokunaga, 2016).

Previously, Noble (2018) has found that women—especially women of color—tend to be sexualized and sexually objectified in both organic and sponsored¹ web search results. Building on this important work, we systematically assess whether there is evidence of sexualization of different groups in web search results. While Noble's (2018) work was largely focused on the US context—that is, on different groups that are part of the US population such as Black women, Asian women, and Black boys—we aim to look at the problem from a broader global perspective, incorporating the representation of different *nationalities* into the analysis.

Specifically, we use algorithm impact auditing (Mittelstadt, 2016) to examine the differences in the ways people of different nationalities/ethnicities/religions are represented in Google's organic and sponsored text search results. We focus on Google since it dominates the search market with over 90% of the market share (Statcounter, 2020). Our analysis is centered on several central research questions:

RQ1. Is there evidence of sexualization of different groups in sponsored search results?

RQ2. Are there discrepancies in the level of the sexualization of different groups (women vs men, as well as different national/ethnic/religious groups, varies²) in organic text search results?

RQ3. How do the findings about sponsored and organic results relate to each other?

Related work

Sexualized representations of people in different media and their effects

“Sexualization” refers to the portrayal of people (usually women) in a way that is focused on their appearance and sexual appeal (Ward, 2016). Exposure to such portrayals can affect people’s beliefs and behaviors, for instance, by leading to greater body dissatisfaction, especially among (young) women (e.g. Johnson et al., 2007; Mulgrew et al., 2014; Reichart Smith, 2016). People portrayed in a sexualized manner are also attributed less personhood by others (Loughnan et al., 2013; Ward, 2016). Furthermore, exposure to sexualized portrayals of women can increase sexist attitudes and behaviors (Hitlan et al., 2009) and tolerance of sexual violence (Dill et al., 2008; Romero-Sanchez et al., 2017).³

Given the tangible negative effects of sexualized representations, it is crucial to study their prevalence in popular forms of digital and analogue media (Ward, 2016). There is an abundance of evidence that sexualized portrayals of people are widespread across different types of popular cultural products—from magazines and music videos to reality TV and video games—as well as in advertisements (Burgess et al., 2007; Flynn et al., 2015; Frisby and Aubrey, 2012; Ward, 2016).

Importantly, research shows that women are sexualized more often than men (Ward, 2016). This gender-based disparity often interacts with the one concerning different national and ethnic groups since women from certain groups (e.g. Black or East Asian ones) tend to be sexualized more than others. This phenomenon can be referred to as “sexoticism” (Mukkamala and Suyemoto, 2018; Schaper et al., 2020) and has been observed by Noble (2018) in Google’s text search results. In this article, we build on Noble’s (2018) work and make several contributions. First, we contribute to the work on sexualized portrayals of people in different media by looking at search engines that are highly trusted and influential in terms of opinion formation (Fisher et al., 2015). Second, we examine the disparities in the sexualization of people based on their nationality and ethnicity, expanding Noble’s (2018) work by incorporating national groups into the analyses. Finally, we contribute to the field of algorithm impact auditing.

Auditing web search results

Methodologically, algorithm impact auditing studies fall into three categories: those that rely on manually generated data (i.e. collected from individual users or generated by the researchers through querying search engines) (Robertson et al., 2018a, 2018b; Steiner et al., 2020), those that rely on virtual agents simulating users’ browsing behavior to generate and collect the data (Haim et al., 2017; Makhortykh et al., 2020; Trielli and Diakopoulos, 2019; Unkel and Haim, 2019; Urman et al., 2021), and those that combine these two approaches (Hannak et al., 2013; Puschmann, 2019). We opt for the second approach as it allows isolating most external factors that might influence search results and is easily scalable.

Algorithm auditing is frequently used to examine biases in search outputs. In web search, biases can arise from how search results are interpreted by the users (usage biases) and the way the results are filtered and ranked by the search engine (retrieval biases). In this article, we focus on the latter.

Retrieval biases stem from the ways search engines select and rank results. As noted earlier, in this study we focus on the *social bias*. Web search engines can perpetuate such bias both via personalized recommendations and default filtering and ranking mechanisms. For instance, Google's search results were shown to perpetuate common racial and gender stereotypes in the United States (Noble, 2018). Another example of social bias found on Google is a systemic reiteration of stereotypes in the representation of women in image results related to different professional occupations (Kay et al., 2015).

We build on the existing scholarship on social bias in web search results and, based on previous research, expect to find evidence of sexualization of women, especially women of certain national and ethnic groups, in search results.

Method

Search queries

As a first step in creating the set of search queries, we compiled a list of national and ethnic attributes. We based it on the most comprehensive list of nationalities we could find—namely, the one provided by the UK government for company filings (Government). We expanded it by adding several regional categories (e.g. “Eastern European”) and several ethnic categories as described in the Racial and Ethnic Categories for the NIH diversity program (NIH). In addition, we included terms referring to the followers of several major religions (e.g. “Christian” and “Muslim”). The full list consists of 242 terms referring to specific national/ethnic/regional/religious groups.⁴ Despite our best efforts, certain groups might not be included and future research will benefit from the further expansion of the list.

As a second step, we combined each term from the list with each of the following seven terms describing gender/age groups: “people,” “women,” “men,” “females,” “males,” “girls,” and “boys.” It resulted in 1694 queries used for data collection; examples include “Afghan boys,” “Norwegian men,” and “Muslim women.”

Data collection

To collect the data, we used *rvest* package (Wickham and RStudio, 2020). All the queries were executed from R console using the following pattern: “https://www.google.com/search?q=” + search query. We used R console to minimize potential personalization effects that arise from the usage of specific browsers (Makhortykh et al., 2020).⁵ For each query, we retrieved full HTML of the first page of Google text search results as it would be seen by regular users. From there, we extracted top 10 organic search results, the accompanying sponsored content (ads), and other results such as image panels or “People also ask” section when present. We focused on the first page of search output because users typically pay more attention to top results (Pan et al., 2007; Urman and Makhortykh, 2021).

For the data collection, we used two separate machines with identical OS (Ubuntu Linux 18.04) and R (4.0.2) versions. To mitigate the effects of researchers' own IP addresses, we used a commercial VPN provider ExpressVPN. This also allowed us to

run the analysis across two locations simultaneously to make sure that any derived observations are not specific to a given country. The VPN was connected to a server in Washington DC (the United States) on one machine and to a server in Dublin (Ireland) on another one. The resulting data were separated into two data sets based on the server location and all the analyses on these data sets were performed separately.

The choice of the specific locations is motivated by several reasons. First, collecting English-language results allows us to contextualize our findings on social bias against other studies using English queries (e.g. Kay et al., 2015; Noble, 2018; Otterbacher et al., 2017). Our decision to include specifically Ireland and the United States is guided by the 2020 Global Gender Gap report of the World Economic Forum (2020). Because search engine companies often attribute bias in web search results to existing societal prejudices (Noble, 2018), we decided to examine whether there are differences in the observations, with respect to gender, between the countries that rank the best (Ireland) and the worst (the United States) by gender parity (World Economic Forum, 2020) among primarily English-speaking countries. Finally, we selected the servers located in the capitals of these countries for comparability purposes.

To make sure that the observations are persistent, we conducted a longitudinal experiment. The data were collected in 3-day intervals for 9 days in total over a period from 21 July 2020 to 5 August 2020. We opted for the interval-based rather than continuous data collection since, as we learned during a testing period, after several days of continuous querying Google would flag suspicious activity and temporarily (e.g. until a CAPTCHA is entered) block the requests.

On each day of the data collection, the two machines would start the scripts simultaneously and use identical query lists. However, between the daily iterations, the order of execution of the queries was reshuffled to compensate for the effects that previously executed queries can have on the subsequent results (Hannak et al., 2013).

Analysis

RQ1: sexualization in sponsored content. First, we extracted the sponsored content for each of the queries for each day of data collection. The extraction was based on the HTML page structure—organic and sponsored results on Google are accompanied by different HTML tags, making it straightforward to automatically distinguish between the two. We aggregated the number of advertisements for each query across all days and calculated the average number of ads per day of data collection (i.e. total number of ads per query divided by 9). Then, we selected the queries—referred to as “ad queries” below—for which at least one ad per day appeared on average. We checked the distribution of these ad queries by age–gender group—“people,” “men,” “women,” “boys,” “girls,” “males,” and “females.” This allowed us to see whether the prevalence of sponsored content is skewed in terms of gender.

We then qualitatively examined the list of top 20 ad queries (based on the average number of ads) and the corresponding ads for each age–gender group to identify which groups these queries refer to. By doing so, we checked whether the prevalence of sponsored content systematically differs based on one of the outlined group-related characteristics (gender/region/nationality, etc).

To make sure that the observed skewness does not stem from advertisers simply targeting the most popular queries, we compared the relative popularity of top 10 ad queries in each sample with the queries that were not accompanied by sponsored content. For that, we utilized Google Trends data from the United States and Ireland for the period of data collection. Google Trends only returns metrics of relative popularity (e.g. one query vs another) and allows comparing up to five queries at the same time. Because comparison of one of the top ad queries with four non-ad-queries would be inconclusive, we repeated the analysis 7 times for each of the top ad queries comparing it with 28 non-ad-queries in total. The non-ad-queries were chosen at random. After each comparison round, we recorded the rank of the top ad query relative to the four non-ad queries (with 1 = the ad query was more popular than the four non-ad-queries and 5 = the ad query was the least popular). Then, we computed the average rank of the ad query versus the non ad-queries across 7 runs. Since the top ad queries attracted by far more ads than the non-ad-queries, if the advertisers indeed simply targeted the most popular search queries, the top ad queries should consistently emerge as the most popular.

Finally, we extracted the top domains in the sponsored content, checked corresponding websites, and read through the ads accompanying links to these domains to infer their main topics to establish whether the content of ads includes sexualized portrayals of different groups.

RQ2: prevalence of sexualization in organic content across different population groups. To examine whether women or men are sexualized in search results more, we relied on word embeddings. Word embedding-based models represent each word from a corpus they are trained on as a multidimensional vector with the geometry of vectors capturing the semantic relationships between words (Garg et al., 2018) and have been shown to reflect biases of the underlying corpus (Bolkubasi et al., 2016). Although such bias encoding is problematic for the practical application of word embeddings, it can be useful for analytical purposes—specifically, to quantify biases in a collection of texts (e.g. Bolkubasi et al., 2016; Garg et al., 2018). Following these studies, we used word embeddings to examine the discrepancies in how different groups are portrayed in the collected organic search results with a specific focus on their sexualization.

To conduct the analysis, we first trained two—one on the US-based and one on the Ireland-based corpus of results—word2vec (Mikolov et al., 2013) models using word2vec R package (Wijffels, 2020). We followed preprocessing steps that are commonly used in studies employing a similar methodology (Kroon et al., 2021), namely, converting all words to lowercase, removing numbers, and splitting the texts into sentences. The US corpus-based vocabulary contained 27,822 unique terms (and corresponding embedding vectors), while the Irish one contained 27,577 terms. Our corpora thus were smaller than ones typically used for training word embeddings models. Nonetheless, we argue that the trained embeddings are suitable for our analytical task, as we ensure that they capture the concepts of interest (gender and nationalities/regional divisions) sufficiently (see Online Appendix A for the details on this).

To examine whether certain groups are sexualized and objectified in organic search results, we measured the strength of the association between different appearance-related adjectives used in previous research (Garg et al., 2018) and group-related terms relying

on normalized association score (NAS)-, mean average cosine similarity (MAC)-, and relative norm distance (RND)-based scores commonly used for such analysis (Caliskan et al., 2017; Kroon et al., 2021; Manzini et al., 2019) and implemented with sweater R package (Chan, 2022). The details on this are presented in Online Appendix B.

RQ3: examining the relation between sexualization in organic and sponsored content. To assess whether there is a potential connection between the organic and sponsored content, we first examined the domains of the organic search results displayed for ad queries—as defined in the “Examining sponsored content” subsection—and for queries that were not accompanied by sponsored content. We qualitatively checked the 20 most frequently displayed domains for each group of queries to infer whether there are differences. As a next step, we compared the proportions of sponsored domains⁶ among organic domains retrieved for ad queries and those not accompanied by sponsored content. We performed a statistical analysis of the results using a test of equality of proportions with the null hypothesis that the proportion of sponsored domains is the same in organic results displayed for the ad queries and the non-ad queries.

Results

RQ1: sexualized sponsored content and its prevalence in relation to queries referring to different population groups

We find that for both US and Ireland locations, the vast majority (over 80%) of ads are displayed for women-related queries (Table 3). This highlights a major gender-based disparity, with women-related queries being treated as the ones which are “better selling” among Google’s advertisers and thus more sexualized and commodified.

Through qualitative analysis, we established that over 90% of the advertisements link to (racialized) dating and/or so-called “mail-order bride” (see Lloyd, 1999) websites, thus confirming our expectation about the presence of (racialized) sexualization of women in sponsored results. The advertisements that constitute the other 10% of the content address a multitude of topics ranging from different health issues (e.g. promoting health treatments aimed at the nationals of a specific country) to the websites of non-governmental organizations (NGOs) and charities working in certain countries (usually, in Africa and South America).

In Table 1, we list the top 10 domains with the highest number of ads linking to them in each country-sample. All of the domains refer to websites that market themselves as international “dating” sites. As evident from their domain names, these sites explicitly promote “dating” women from the Global South and Global East (Dados and Connell, 2012). Such sites often provide surrogate dating—some offer translation services or even have employees masking themselves as potential “brides” who communicate with the users (Liu, 2015). Such “mail-order bride” and surrogate dating content directly relates to the phenomenon of sexoticism (Schaper et al., 2020) as it reinforces neocolonial sexist stereotypes as in the case of “international dating” websites (Liu, 2015). It primarily targets Western European and North American men who are willing to “buy” a bride from one of the developing countries, commodifying women, relationships, and sex (Liu, 2015). Such forms of

Table 1. Top 10 domains by the total number of ads linking to them in each subsample.

Domains (US)	Total number of ads linking to the domain	Domains (Ireland)	Total number of ads linking to the domain
www.afrointroductions.com	463	www.afrointroductions.com	404
www.latinamericancupid.com	460	www.asiandating.com	349
www.asiandating.com	347	www.internationalcupid.com	307
www.latinamericancupid.com	300	www.globalladies.com	117
www.internationalcupid.com	223	www.filipinocupid.com	82
www.russiancupid.com	155	www.dating-rating.com	80
www.caribbeancupid.com	134	www.russiancupid.com	76
www.globalladies.com	97	www.muslima.com	62
www.thedatingcritic.com	89	www.czech-ladies.com	61
www.muslima.com	84	www.chinalovecupid.com	57

Table 2. Top 10 queries by the average number of accompanying ads in the US and Ireland locations.

Query, US	Average number of ads	Average position by popularity in Google Trends	Query, Ireland	Average number of ads	Average position by popularity in Google Trends
Ukrainian females	6.8	4.00	Ukrainian females	7	4.29
Russian females	5.2	4.86	Russian women	6.7	2.00
Ukrainian women	5.2	3.00	Asian females	6.5	3.86
Russian women	4.7	2.14	Russian females	6.4	3.57
Asian females	4.5	4.29	Ukrainian women	6.3	3.29
Caribbean females	4.5	2.57	Filipino females	5.8	4.14
Chinese females	4.2	3.71	Chinese females	5.1	3.29
Eastern European females	4.1	4.14	Eastern European women	4.8	1.71
Filipino females	4	3.14	Asian women	4.8	2.29
Eastern European women	3.7	1.71	Filipino women	4.7	2.86

The average position of each of these queries relative to four randomly selected queries that were not accompanied by ads in terms of relative popularity according to Google Trends is also reported (1 = the query was more popular than all non-ad queries; 5 = the query was less popular than all non-ad queries).

commodification promote the sexualization of women, especially those from the Global East and Global South, and amplify global sexual exploitation. Furthermore, by utilizing the stereotype of neocolonial benevolence in the form of the “white man’s burden” (Laforteza, 2007), this content also reiterates racial discrimination.

The explicit focus on “mail-order bride” websites in the sponsored content is accompanied by the regional disparities in the number of ads accompanying different queries (Table 2). The queries with the largest number of ads refer to women from the developing countries in Asia, Eastern Europe, and South and Central America. However, there are some differences across the two subsets, for instance, the US subset has a higher prevalence of queries referring to women from South and Central America. We suggest this can be explained by the geographical proximity of the United States to South and Central American countries.

Based on the Google Trends data, we established that the selection of queries that the advertisers targeted is not dependent on their popularity with the users: all queries with the highest number of ads with the exception of “Eastern European women” were ranked consistently lower in Google Trends than the randomly selected queries that did not attract ads at all (Table 2). Thus, the advertisers do not simply choose queries that are the most popular, but rather target references to the women from the Global South and Global East utilizing neocolonial and sexotic stereotypes (Laforteza, 2007).

Men-related ad queries are way less prevalent (Table 3). Among the few men-related ad queries, there are still regional disparities. Similar to the women-related ad queries, they predominantly refer to men from the Global South and East. In addition, even when referring to men, the ads often link to the websites that offer “meeting” women from the respective countries or, in some cases, both—women and men.

In the US subset, only a single men-related query (i.e. “Ukrainian males,” average n of ads=2.3) is in the list of top 50 ad queries. The advertisements accompanying it, though, still link to the websites about Ukrainian women. One example ad reads, “#1 Ukraine Dating Site—7 Ukrainian Women to Every Man.” In the Irish subset, there are four men-related queries among the top 50 ad queries: “Ukrainian males” (average n of ads=3.4), “Asian males” ($n=2.5$), “Irish males” ($n=2.1$), and “Thai males” ($n=2.1$). Whereas the ads for “Ukrainian males” again lead to the websites offering Ukrainian “mail-order brides” (for instance, “Ukrainian Wives—100% Privacy Protected”), those for “Asian males” and “Thai males” offer meeting both men and women (e.g. “Meet Thailand Men—1178 New Single Thai Females”). Finally, those for “Irish men” focus on finding male dating partners (e.g. “Date Irish Men—Search Date Irish Men”).

We do not observe any discrepancies in the prevalence of sponsored content with regard to religious groups—in both subsamples, there is no sponsored content associated with them.

RQ2: sexualization of different genders and population groups in organic search results

Gender-based discrepancies in sexualization. We report the results of the NAS-based analysis of the association between different sex-, appearance-, and dating-related terms in Figure 1 (the US sample) and Figure 2 (the Irish sample). In both Figures,

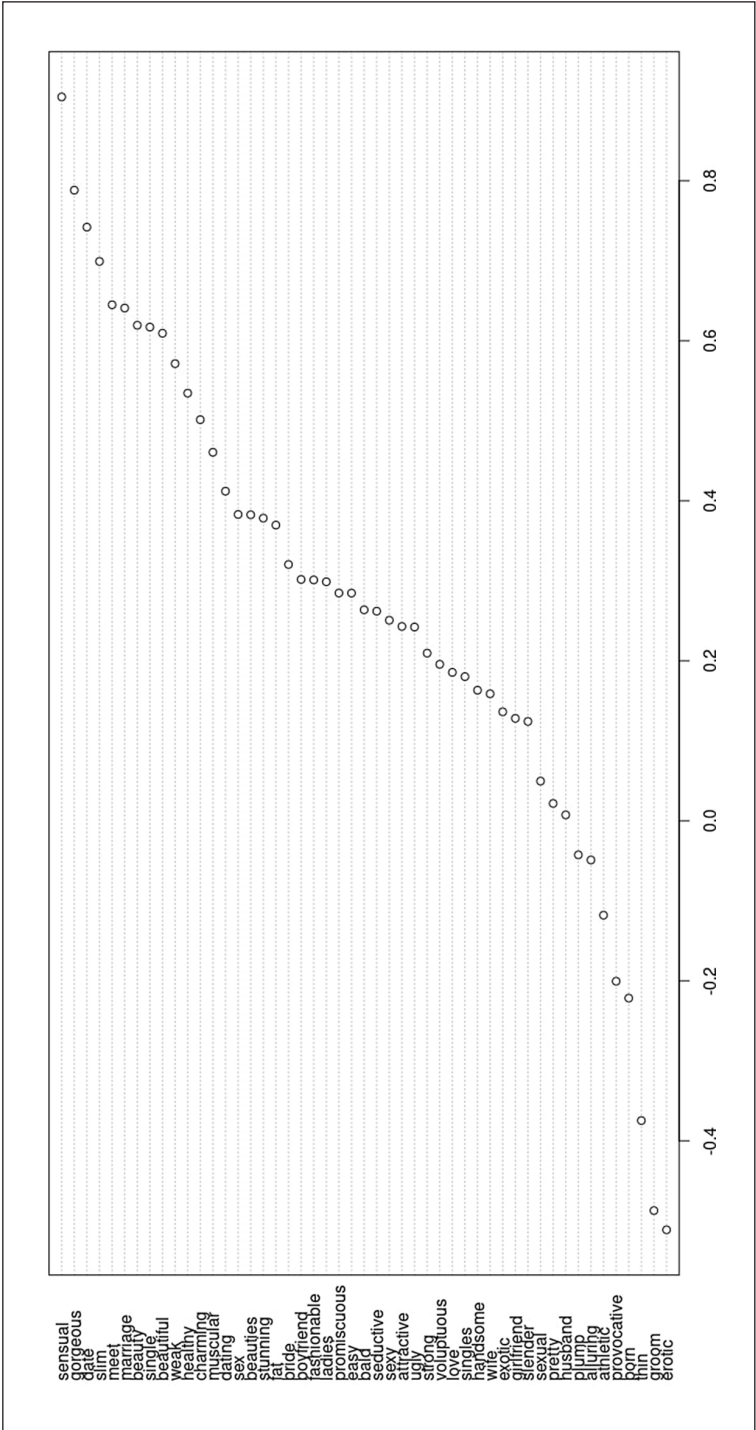


Figure 1. Results of NAS association analysis, US sample. Words with scores > 0 are associated with women; words with scores < 0 are associated with men. The further away the score is from 0, the stronger is the association.

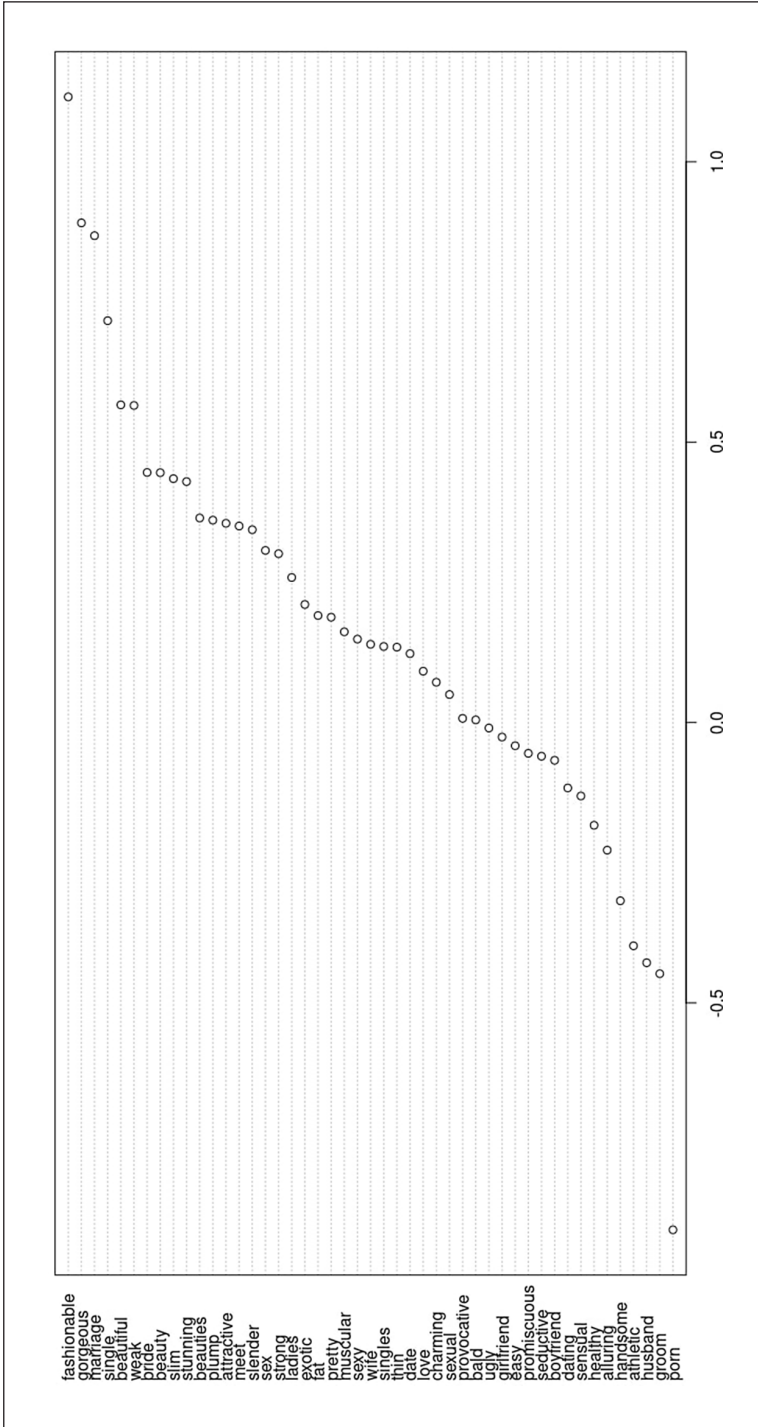


Figure 2. Results of NAS association analysis, Ireland sample.

Table 3. Prevalence of queries corresponding to different groups among ad queries.

Group term	US	Ireland
people	0%	0%
women	21.3%	26.2%
men	4.3%	2.3%
females	49.4%	47.7%
males	12.2%	10%
girls	12.2%	13.8%
boys	0.6%	0%

words with scores > 0 are associated with women and words with scores < 0 are associated with men. The further away the score is from 0, the stronger is the association. Most sex-/dating-/appearance-related terms in both samples are associated with women rather than men. The fact that certain terms (e.g. athletic or erotic or porn) are strongly associated with men suggests the presence of some level of the sexualization of men. However, since most terms, including such as “sexy” or “gorgeous,” are associated with women, there is a gender-based disparity in the prevalence of sexualization, with women sexualized more than men.

Sexualization of different population groups

In Table 4, we list 20 terms with the highest numbers of associations with sexualized terms and high sexualization ratios (see “Method” section for details). In both samples, the group terms with the highest sexualization ratios predominantly refer to groups from the Global South and Global East, albeit there are also references to the Global North, especially in the Irish subsample. Specific religious (e.g. Muslim) and ethnic groups (e.g. Black) tend to have comparatively few strong associations with sexualization terms, suggesting that in our samples sexualization prevalence goes along national but not ethnic or religious lines.

Is sexualization both gendered and racialized? In Table 5, we list first the number of population group terms associated with women and men in each sample and then the shares of such terms with high and very high sexualization ratios (see “Method” section and Online Appendix B). In both samples, the number of men-associated population group terms is higher than the number of women-associated terms. At the same time, the share of population group terms with high or very high sexualization ratios is higher among women-associated terms in both samples, with this gender-based discrepancy being particularly high in the US sample. We suggest this indicates that the nationality terms-based sexualization disparities described in the previous section are gendered.

RQ3: connection between sponsored and organic content

We examined the potential connection between organic and sponsored content by checking for the exact matches and topical overlaps in sponsored and organic content. Top 20

Table 4. Twenty population group terms with the highest number of strong associations (assoc.) with sexualization terms and high sexualization ratios.

US		Ireland					
Group term	Sexualized assoc.	All assoc.	Sexualization ratio	Group term	Sexualized assoc.	All assoc.	Sexualization ratio
peuvian	17	42	0.40	luxembourger	15	29	0.52
ukrainian	17	56	0.30	dutch	15	49	0.31
moroccan	17	64	0.27	scandinavian	14	42	0.33
serbian	16	53	0.30	slovakian	13	49	0.27
colombian	16	55	0.29	moroccan	13	58	0.22
slovakian	16	56	0.29	vietnamese	12	35	0.34
nicaraguan	14	38	0.37	swedish	12	37	0.32
scandinavian	14	39	0.36	norwegian	12	40	0.30
dutch	13	54	0.24	latina	12	43	0.28
filipino	12	41	0.29	colombian	11	43	0.26
danish	12	42	0.29	serbian	10	37	0.27
albanian	12	56	0.21	peruvian	10	42	0.24
scottish	11	31	0.35	singaporean	9	34	0.26
brazilian	11	36	0.31	ukrainian	9	36	0.25
ethiopian	10	22	0.45	polish	8	25	0.32
panamanian	10	28	0.36	italian	8	29	0.28
italian	10	29	0.34	taiwanese	8	34	0.24
portuguese	10	29	0.34	armenian	8	36	0.22
latina	10	31	0.32	estonian	7	27	0.26
cuban	10	37	0.27	cuban	7	29	0.24

Table 5. Number of population group terms associated with women and men in both samples; shares of women and men-associated terms with high (>11.1%) and very high (>22.2%) sexualization ratios.

US sample		Ireland sample					
Terms associated. . . with women (NAS >0.2)	with men (NAS <-0.2)	with men (NAS <-0.2)	with women (NAS >0.2)	with men (NAS <-0.2)	with men (NAS <-0.2)	with women (NAS >0.2)	with men (NAS <-0.2)
31	47	47	38	66	66	66	66
with sexualization ratio. . .	>22.2%	>11.1%	>22.2%	>11.1%	>22.2%	>11.1%	>22.2%
>11.1%	80.7%	38.7%	53.2%	0%	73.6%	36.8%	68.2%
							24.2%

domains from the organic search results retrieved for the ad queries and for queries without sponsored content are listed in Table 6 (the US sample) and Table A1 in Online Appendix C (the Irish sample). In both samples, there is evidence of search concentration in the results, with YouTube and Wikipedia accounting for over 7% of all links in the organic content. Besides these sources, each sample contains a high share of links to media such as the *Guardian* or BBC as well as data sources and scientific databases such as JSTOR. However, mail-order bride and racialized dating websites are also present. They are more prominent in the results for the ad queries, although they also appear in the results for queries that were not accompanied by ads (see Table 6 and Table A1).

We obtained contradictory results on the difference in the proportion of sponsored domains in organic results for the ad queries and non-ad ones. In the US sample, there was a significant ($p < .01$) difference with sponsored domains accounting for 2.9% of all organic content links for ad queries and for 2.3% of the rest. In the Irish sample, sponsored domains accounted for very low shares of organic content—0.06% for ad queries and 0.05% for the rest—with no significant difference in the proportions between the two samples ($p = .92$).

Limitations

First, we looked for search outputs for only two locations—United States and Ireland—both of which are Western democracies from the Global North. Future research can benefit from diversifying the set of locations for which search outputs are collected and using a wider range of IPs to reduce the likelihood that search result consistency is something other than a form of caching on Google’s side. Similarly, it will be beneficial to include queries in different languages to measure how important is the role of the language in reiterating misrepresentations in web search. Another limitation of our study is that we took a simplistic approach to gender, treating it as a binary category. Although this is common for studies that examine gender biases online (Chen et al., 2018; Kay et al., 2015; May et al., 2019) due to the nature of the data and the complexity of the construct of gender, we acknowledge that this is a fundamental limitation of the studies exploring gender-related differences, including ours. Furthermore, our focus on text results is limiting as it does not account for the potential perpetuation of bias in the other

Table 6. Top 20 domains retrieved in organic search results for ad queries and those not accompanied by sponsored content with corresponding share of times each domain was linked to, the US sample.

Top domains (ad queries)	Share of links to each domain in the organic results (ad queries)	Top domains (other queries)	Share of links (other queries)
en.wikipedia.org	0.110	en.wikipedia.org	0.126
www.youtube.com	0.109	www.youtube.com	0.076
www.shutterstock.com	0.028	www.pinterest.com	0.032
www.pinterest.com	0.021	www.everyculture.com	0.025
www.facebook.com	0.017	www.shutterstock.com	0.023
mylatinabride.com	0.016	www.britannica.com	0.020
theculturetrip.com	0.014	www.facebook.com	0.016
www.cnn.com	0.013	www.bbc.com	0.015
www.gettyimages.com	0.008	www.gettyimages.com	0.011
www.tripadvisor.com	0.008	theculturetrip.com	0.011
www.jstor.org	0.008	www.quora.com	0.011
www.theguardian.com	0.008	www.theguardian.com	0.009
russiansbrides.com	0.007	www.amazon.com	0.009
www.npr.org	0.007	www.ncbi.nlm.nih.gov	0.008
www.quora.com	0.007	www.hr.org	0.008
www.everyculture.com	0.007	www.jstor.org	0.007
www.mylatinlife.com	0.007	www.rosebrides.com	0.007
www.nytimes.com	0.007	www.nytimes.com	0.007
www.bbc.com	0.006	www.tripadvisor.com	0.006
www.amazon.com	0.006	www.babynamewizard.com	0.005

Links to mail-order bride and racialized dating websites are bolded.

elements of the search result page. Our preliminary analysis showed that there are no significant differences between population groups (women and men; queries that attracted sponsored content and those that did not) in how often these other elements (e.g. image panels) are displayed. Although examining the content of these other page elements is beyond the focus of this study, we believe such examination is an important direction for future research. Finally, since we focused on algorithm impact auditing, not functionality auditing, we cannot draw conclusions about the sources of observed misrepresentations in search outputs, although we believe examining this would be a fruitful direction for future work and would help shed light on whether Google's results merely reflect existing stereotypes or further distort them.

Discussion

Our analysis reveals gender-based discrepancies in the levels of sexualization in both organic and sponsored search results. Women are sexualized more than men; thus, we observe a gender-based discrepancy similar to that found in other media (Ward, 2016). We also find discrepancies in the levels of the sexualization of different national groups,

with women from the Global East and South being sexualized more than women from the Global North, especially in sponsored search results. Such a discrepancy leads to the reinforcement of sexotic (Schaper et al., 2020) stereotypes and the reiteration of the “colonial fantasies” (Yeegenoglu, 1998), which treat women outside the Global North as objects expected to be conquered as exemplified by search results such as “‘Asian Girls Are Tighter’: Dispelling the Myth of Vagina Size.”

Regardless of the source of observed misrepresentations, they can be harmful to individuals and societies. Even if these biases do not introduce, but merely reflect existing stereotypes, they can further reinforce them as exposure to sexualized representations of people can increase sexist attitudes and has numerous other adverse effects (Ward, 2016). For instance, the persistent prioritization of “mail-order bride” websites in sponsored and organic results can foster the stigmatization of women from certain countries as mail-order brides, including those coming to the Global North as immigrants or refugees. It is particularly concerning in the case of groups that are already stigmatized as sexualized objects in the public imagination in the Global North, such as in the case of Eastern European (Maydell, 2017; Shpeer and Howe, 2020; Tankosic, 2020) and (South) East Asian women (Ricordeau, 2017; Suphsert, 2021), especially as queries related to them attract a disproportionate share of mail-order bride sponsored content.

By accepting such sponsored content, Google can contribute to the reinforcement of existing stigmas through prioritizing sexualized and stereotypical representations of women from Global East and South. Our analysis does not definitively prove the presence of a direct connection between sponsored and organic content but lends preliminary evidence for it. For instance, we find that racialized dating websites–related content is more frequently featured on the first page of organic search results for queries that are accompanied by sponsored content in the US sample—but not in the Irish one. One explanation behind the discrepancy could be the difference in the user numbers: it is possible that in Ireland search volume and sponsored content clicking is not high enough to affect the organic results in a major way. Although our results are not definitive, we believe they warrant further investigation since a connection between sponsored and organic content might have important implications. For example, it would open an opportunity for advertisers with big budgets to drive certain types of content higher in the organic search results, thus manipulating search outputs.

Our observations also highlight some implications for the Google web search functionality and commercial model. While Google itself is not responsible for the existence of websites that sexualize certain groups, it reiterates existing biases by prioritizing such websites in its search outputs for the broadly formulated queries that do not per se refer to sexualization. Because of the breadth of the queries we used, it is hardly possible that sexualized outputs are the only ones available in the database of websites indexed by Google. However, sexualized outputs, including “mail-order bride” websites, still resurface in top organic search results. The cooperation between Google and businesses behind such websites is questionable by itself because it can be viewed as a form of exploitation that relies on the reinforcement of biases toward national and social groups for commercial purposes. While it is probably not realistic to expect Google to break out from such cooperation, it can at least keep such websites and associated ads away from outputs related to more general queries—for example, those queries where a user is

searching for information about “Ukrainian women” generally, not for “Ukrainian women to marry.” Otherwise, such advertising profit-driven results can “lay groundwork for implicit bias,” as noted by Noble (2018).

Instead of prioritizing results reinforcing sexist biases, the top organic search results related to more general queries can promote more neutral reference sources and outlets exposing the stigmatization of vulnerable groups. The differences between the results for “Black girls” described by Noble (2018), and the results retrieved for the same query in our data set, demonstrate that it is possible for Google to debias its search results. While the removal of sexualized results from the first page in relation to Black women is laudable, we argue that similar debiasing should occur for other general queries related to people, especially those at high risk of being stigmatized such as women, minorities, or immigrant populations.

Finally, our observations are important for the ongoing debate about the algorithmic fairness and its applicability to the Global South and East (Mohamed et al., 2020; Sambasivan et al., 2021), in particular in the context of the responsibility of transnational tech companies for preventing systems they design from amplifying discrimination. Similar to other applications of artificial intelligence (AI) (Mohamed et al., 2020), web search is embedded in existing power relationships often resulting in the reiteration of unfair treatment of certain groups that can facilitate their exploitation (Sambasivan et al., 2021). Our findings exemplify this showing that Google’s search algorithms currently tend to render some people (predominantly women from the Global South and East) as commodities, while others (men from the Global North) as customers and consumers. Thus, our observations highlight the importance of integrating the cultural and socio-economic dimensions in the algorithmic fairness debate more prominently.

Acknowledgements

We thank Prof. Dr. Anikó Hannák and Stefania Ionescu for their feedback and suggestions on the original draft of this manuscript.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Aleksandra Urman  <https://orcid.org/0000-0003-3332-9294>

Mykola Makhortykh  <https://orcid.org/0000-0001-7143-5317>

Supplemental material

Supplemental material for this article is available online.

Notes

1. Paid advertisements displayed alongside “regular” (organic) search results; typically, they look similar to the organic results but are accompanied by an “Ad” tag.

2. For the sake of brevity, we refer to these as “population groups” below.
3. For other adverse effects of sexualized portrayals, see Ward (2016).
4. See <https://figshare.com/s/b6519f8bc1555b4418cb>
5. While we cannot state with certainty that the retrieved results corresponded to a “default” non-personalized selection, we have evidence supporting this. During pre-testing, we established that the results obtained through R console are almost identical (with minor variation that might be attributed to randomization effects) to those obtained in clean “Private Browsing” sessions on our computers via Firefox and Chrome browsers, and in turn dissimilar to those we received within “regular” browsing from our default browsers when logged in to our Google accounts.
6. Those that appeared at least 10 times in all collected sponsored results; we used a benchmark of 10 to reduce sparsity. There were 43 such domains in the US sample and 51 domains in the Irish sample.

References

- Bolukbasi T, Chang KW, Zou J, et al. (2016) Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In: *Proceedings of the 30th international conference on neural information processing systems (NIPS'16)*, pp. 4356–4364. Available at: <https://proceedings.neurips.cc/paper/2016/file/a486cd07e4ac3d270571622f4f316ec5-Paper.pdf>
- Bozdag E (2013) Bias in algorithmic filtering and personalization. *Ethics and Information Technology* 15(3): 209–227.
- Burgess MCR, Stermer SP and Burgess SR (2007) Sex, lies, and video games: the portrayal of male and female characters on video game covers. *Sex Roles* 57(5): 419–433.
- Caliskan A, Bryson JJ and Narayanan A (2017) Semantics derived automatically from language corpora contain human-like biases. *Science* 356(6334): 183–186.
- Chan C (2022) *sweater: Speedy Word Embedding Association Test and Extras Using R*. Available at: <https://CRAN.R-project.org/package=sweate>
- Chen L, Ma R, Hannak A, et al. (2018) Investigating the impact of gender on rank in resume search engines. In: *Proceedings of the 2018 CHI conference on human factors in computing systems*. Available at: <https://cbw.sh/static/pdf/chen-chi18.pdf>
- Dados N and Connell R (2012) The global south. *Contexts* 11(1): 12–13.
- Dill KE, Brown BP and Collins MA (2008) Effects of exposure to sex-stereotyped video game characters on tolerance of sexual harassment. *Journal of Experimental Social Psychology* 44(5): 1402–1408.
- Edelman Trust Barometer (2021) Edelman. Available at: <https://www.edelman.com/trust/2021-trust-barometer>
- Epstein R and Robertson RE (2015) The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences* 112(33): E4512–E4521.
- Fisher M, Goddu MK and Keil FC (2015) Searching for explanations: how the Internet inflates estimates of internal knowledge. *Journal of Experimental Psychology: General* 144(3): 674–687.
- Flynn MA, Park SY, Morin DT, et al. (2015) Anything but real: body idealization and objectification of MTV docusoap characters. *Sex Roles* 72(5): 173–182.
- Frisby CM and Aubrey JS (2012) Race and genre in the use of sexual objectification in female artists’ music videos. *Howard Journal of Communications* 23(1): 66–87.
- Garg N, Schiebinger L, Jurafsky D, et al. (2018) Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115(16): E3635–E3644.

- Goldman E (2008) Search engine bias and the demise of search engine utopianism. *Yale Journal of Law and Technology* 8: 188–200.
- Government UK (2017) List of nationalities CSV file—GOV.UK. Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/664133/CH_Nationality_List_20171130_v1.csv/preview
- Grimmelmann J (2010) Some skepticism about search neutrality. *Faculty Scholarship*. Available at: https://digitalcommons.law.umaryland.edu/fac_pubs/1417
- Haim M, Arendt F and Scherr S (2017) Abyss or shelter? On the relevance of web search engines' search results when people google for suicide. *Health Communication* 32(2): 253–258.
- Hannak A, Sapiezynski P, Khaki AM, et al. (2013) Measuring personalization of web search. In: *Proceedings of the 22nd international conference on World Wide Web*. Available at: <https://arxiv.org/pdf/1706.05011.pdf>
- Hitlan RT, Pryor JB, Hesson-McInnis MS, et al. (2009) Antecedents of gender harassment: an analysis of person and situation factors. *Sex Roles* 61(11): 794–807.
- Johnson PJ, McCreary DR and Mills JS (2007) Effects of exposure to objectified male and female media images on men's psychological well-being. *Psychology of Men & Masculinity* 8(2): 95–102.
- Kay M, Matuszek C and Munson SA (2015) Unequal representation and gender stereotypes in image search results for occupations. In: *Proceedings of the 33rd annual ACM conference on human factors in computing systems (CHI'15)*. Available at: <https://www.csee.umbc.edu/~cmat/Pubs/KayMatuszekMunsonCHI2015GenderImageSearch.pdf>
- Kroon AC, Trilling D and Raats T (2021) Guilty by association: using word embeddings to measure ethnic stereotypes in news coverage. *Journalism & Mass Communication Quarterly* 98(2): 451–477.
- Kulshrestha J, Eslami M, Messias J, et al. (2019) Search bias quantification: investigating political bias in social media and web search. *Information Retrieval Journal* 22(1): 188–227.
- Laforteza E (2007) White geopolitics of neo-colonial benevolence: the Australia Philippine "Partnership." *The Australian Critical Race and Whiteness Journal* 3(1): 1–17. Available at: <https://researchers.mq.edu.au/en/publications/white-geopolitics-of-neo-colonial-benevolence-the-australia-phili>
- Liu M (2015) Surrogate dating and the translation of gendered meanings across borders: the case of China's e-mail-order brides. *Signs: Journal of Women in Culture and Society* 41(1): 29–53.
- Lloyd KA (1999) Wives for sale: the modern international mail-order bride industry comments. *Northwestern Journal of International Law & Business* 20(2): 341–368.
- Loughnan S, Pina A, Vasquez EA, et al. (2013) Sexual objectification increases rape victim blame and decreases perceived suffering. *Psychology of Women Quarterly* 37(4): 455–461.
- Makhortyk M, Urman A and Ulloa R (2020) How search engines disseminate information about COVID-19 and why they should do better. *Harvard Kennedy School Misinformation Review*. Available at: <https://misinforeview.hks.harvard.edu/article/how-search-engines-disseminate-information-about-covid-19-and-why-they-should-do-better/>
- Manzini T, Lim YC, Tsvetkov Y, et al. (2019) Black is to criminal as Caucasian is to police: detecting and removing multiclass bias in word embeddings. Available at: <http://arxiv.org/abs/1904.04047>
- May A, Wachs J and Hannak A (2019) Gender differences in participation and reward on stack overflow. *Empirical Software Engineering* 24(4): 1997–2019.
- Maydell E (2017) We need engineers, not Russian brides': cultural stereotypes in New Zealand print media and what the engineers have to say. *Communication Research and Practice* 3(4): 335–357.

- Mikolov T, Sutskever I, Chen K, et al. (2013) Distributed representations of words and phrases and their compositionality. Available at: <http://arxiv.org/abs/1310.4546>
- Mittelstadt B (2016) Automation, algorithms, and politics| auditing for transparency in content personalization systems. *International Journal of Communication* 10: 12.
- Mohamed S, Png MT and Isaac W (2020) Decolonial AI: decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology* 33(4): 659–684.
- Mowshowitz A and Kawaguchi A (2002) Assessing bias in search engines. *Information Processing & Management* 38(1): 141–156.
- Mukkamala S and Suyemoto KL (2018) Racialized sexism/sexualized racism: a multimethod study of intersectional experiences of discrimination for Asian American women. *Asian American Journal of Psychology* 9(1): 32–46.
- Mulgrew KE, Johnson LM, Lane BR, et al. (2014) The effect of aesthetic versus process images on men's body satisfaction. *Psychology of Men & Masculinity* 15(4): 452–459.
- NIH (n.d.) NOT-OD-15-089: racial ethnic categories definitions for NIH and diversity programs for other reporting purposes. Available at: <https://grants.nih.gov/grants/guide/notice-files/not-od-15-089.html>
- Noble SU (2018) *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press.
- Otterbacher J, Bates J and Clough P (2017) Competent men and warm women: gender stereotypes and backlash in image search results. In: *Proceedings of the 2017 CHI conference on human factors in computing systems (CHI'17)*, Denver, CO, 6–11 May.
- Pan B, Hembrooke H, Joachims T, et al. (2007) In Google we trust: users' decisions on rank, position, and relevance. *Journal of Computer Mediated Communication* 12(3): 801–823.
- Puschmann C (2019) Beyond the bubble: assessing the diversity of political search results. *Digital Journalism* 7(6): 824–843.
- Reichart Smith L (2016) What's the best exposure? Examining media representations of female athletes and the impact on collegiate athletes' self-objectification. *Communication & Sport* 4(3): 282–302.
- Ricordeau G (2017) The “mail-order bride” stigma: intermarried Filipino women and the Philippine public and political debates. In: Fresnoza-Flot A and Ricordeau G (eds) *International Marriages and Marital Citizenship: Southeast Asian Women on the Move*. London: Routledge, pp. 121–139.
- Robertson RE, Jiang S, Joseph K, et al. (2018a) Auditing Partisan audience bias within Google search. *Proceedings of the ACM on Human Computer Interaction* 2: 148:1–148:22.
- Robertson RE, Lazer D and Wilson C (2018b) Auditing the personalization and composition of politically-related search engine results pages. In: *Proceedings of the 2018 World Wide Web conference (WWW'18)*, Lyon, 23–27 April.
- Romero-Sanchez M, Toro-García V, Horvath MAH, et al. (2017) More than a magazine: exploring the links between lads' mags, rape myth acceptance, and rape proclivity. *Journal of Interpersonal Violence* 32(4): 515–534.
- Sambasivan N, Arnesen E, Hutchinson B, et al. (2021) Re-imagining algorithmic fairness in India and beyond. In: *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, Toronto, ON, Canada, 3–10 March.
- Schaper U, Beljan M, Eitler P, et al. (2020) Sexotic: the interplay between sexualization and exoticization. *Sexualities* 23(1–2): 114–126.
- Shpeer ME and Howe WT (2020) Exotic beauty, mail-order bride, secret agent: the stereotyped experiences of Russian women immigrating to the United States. *Russian Journal of Communication* 12(3): 306–322.

- Statcounter (2020) *Search Engine Market Share Worldwide*. Available at: <https://gs.statcounter.com/search-engine-market-share>
- Steiner M, Magin M, Stark B, et al. (2020) Seek and you shall find? A content analysis on the diversity of five search engines' results on political queries. *Information, Communication & Society* 25(2): 217–241.
- Suphsert L (2021) *The Experiences of Racial and Gender Stereotyping: The Case Study of Thai Women in Sweden*. Lund University. Available at: <http://lup.lub.lu.se/student-papers/record/9055626>
- Tankosic A (2020) Translingual identity: perpetual foreigner stereotype of the Eastern-European immigrants in Australia. *Australian Review of Applied Linguistics*. Epub ahead of print 26 November.
- Trielli D and Diakopoulos N (2019) Search as news curator: the role of Google in shaping attention to news information. In: *Proceedings of the 2019 CHI conference on human factors in computing systems (CHI'19)*, Glasgow, 4 May.
- Unkel J and Haim M (2019) Googling politics: parties, sources, and issue ownerships on Google in the 2017 German Federal Election Campaign. *Social Science Computer Review* 39: 844–861.
- Urman A and Makhortykh M (2021) You are how (and where) you search? *Comparative analysis of web search behaviour using web tracking data*. Available at: <http://arxiv.org/abs/2105.04961>
- Urman A, Makhortykh M and Ulloa R (2021) The matter of chance: auditing web search results related to the 2020 U.S. Presidential primary elections across six search engines. *Social Science Computer Review*. Epub ahead of print 8 April.
- Ward LM (2016) Media and sexualization: state of empirical research, 1995–2015. *The Journal of Sex Research* 53(4–5): 560–577.
- Wickham H and RStudio (2020) *rvest: Easily Harvest (Scrape) Web Pages*. Available at: <https://CRAN.R-project.org/package=rvest>
- Wijffels J (2020) *Distributed Representations of Words* [R package word2vec version 0.3.3]. Available at: <https://CRAN.R-project.org/package=word2vec>
- World Economic Forum (2020) Global gender gap report 2020. Available at: <https://www.weforum.org/reports/gender-gap-2020-report-100-years-pay-equality/>
- Wright PJ and Tokunaga RS (2016) Men's objectifying media consumption, objectification of women, and attitudes supportive of violence against women. *Archives of Sexual Behavior* 45(4): 955–964.
- Yegenoglu M (1998) *Colonial Fantasies: Towards a Feminist Reading of Orientalism*. Cambridge: Cambridge University Press.

Author biographies

Aleksandra Urman is a postdoctoral researcher at Social Computing Group, University of Zurich. Aleksandra's research interests include political communication on social media, algorithmic biases, and computational research methods.

Mykola Makhortykh is a postdoctoral researcher at the University of Bern, where he studies information behavior in online environments. Before moving to Bern, Mykola defended his PhD dissertation at the University of Amsterdam on the relationship between digital platforms and war remembrance in Eastern Europe and worked as a postdoctoral researcher in Data Science at the Amsterdam School of Communication Research, where he investigated the effects of algorithmic biases on digital news consumption.