POLITICAL COMMUNICATION
https://doi.org/10.1080/10584609.2026.2679476

Routledge
Taylor & Francis Group

The Prevalence and Consequences of Women's Algorithmic Underrepresentation in Cross-National Political Google Searches

Tobias Rohrbach[a], Mykola Makhortykh[b], and Maryna Sydorova[b]

[a]Department of Communication and Media Research, University of Fribourg, Fribourg, Switzerland; [b]Institute of Communication and Media Studies, University of Bern, Bern, Switzerland

ABSTRACT

Search engines like Google have become major information gatekeepers that use artificial intelligence (AI) to determine who and what voters find when searching for political information. Yet there is concern that biases in algorithmically curated representation of women politicians in online search results mirror and amplify structural inequalities and therefore further reduce women's electoral chances. We empirically assess the prevalence and consequences of algorithmic gender bias in political Google searches in a series of four studies. First, we conduct two multinational algorithm audits of political image searches (study 1: 56 countries and 6,363 images; study 2: 11 countries and 152,098 images) to assess women's quantitative representation in search results. We find that women politicians' algorithmic representation on average accurately mirrors the actual gender composition of legislative bodies while consistently remaining below absolute gender parity. Second, we show in two online experiments from three different samples (1,388 respondents in total) how more extreme cases of algorithmic underrepresentation can influence human perceptions of the political reality and actively reinforce a white and masculinized view of politics. These misperceptions act as a casual conduit through which algorithmic biases can not only diminish the perceived chances of winning an election for minoritized candidates but also result in voters feeling that their voices matter less. Together, the results highlight the risk that, under certain conditions, algorithmic systems like search engines can negatively impact electoral processes.

KEYWORDS

Algorithms; gender bias; representation; search engines; artificial intelligence

Though political landscapes around the world are slowly changing, the face of contemporary political decision-making is still disproportionately male and, in the Global North, disproportionately white. In 2024, global descriptive representation for women in parliaments is 26.7% (Inter-parliamentary Union, 2024). In the U.S. House of Representatives, 29% of members are women and 26% are nonwhite,[1] making the 118th Congress the most ethnically diverse and gender-balanced to date in U.S. history (Pew Research Center, 2023a, 2023b). Yet according to census data, the shares of women and nonwhite persons in the general U.S. population are 50.5% and 41.1% respectively (U.S. Census Bureau, 2023).

CONTACT Tobias Rohrbach tobiasrohrbach@hotmail.com Depatment of Communication and Media Research, University of Fribourg, Bd. de Pérolles 90, 1700 Fribourg, Switzerland

Supplemental data for this article can be accessed online at https://doi.org/10.1080/10584609.2026.2679476

Decades of scientific exploration into the mechanisms of the underrepresentation of minoritized groups indicate that political structures, gender perceptions, and the media environment are multicausal drivers of gender and race inequalities in politics (Bast et al., 2022; Bos et al., 2022; Dolan & Hansen, 2018; Heldman et al., 2005; Kanthak & Woon, 2015; Rohrbach et al., 2023; Wängnerud, 2009). In this article, we turn to artificial intelligence (AI) driven search engines like Google to examine the possibility of them becoming a novel driver of political exclusion of underrepresented groups. Specifically, we propose and test a framework of algorithmic representation to delineate the role of search engines in constructing masculinized and white views of politics as well as delineate the consequences of such views for political perceptions.

Search engines have become major information gatekeepers that use artificial intelligence (AI)[2] to determine who and what voters find when searching for political information (Trielli & Diakopoulos, 2022; Urman et al., 2022; White & Horvitz, 2015). How information is represented on search engines has been shown to influence perceptions of political campaigns (Epstein & Robertson, 2015) and individual vote choices (Diakopoulos et al., 2018). Selection and ranking of information by search engines – and, crucially, biases of such algorithmic curation,[3] including the systematic under- and misrepresentation of gender and racial groups (Makhortykh et al., 2021; Noble, 2018; Pradel, 2021; Urman & Makhortykh, 2022) – in turn shape perceptions of political realities. Specifically, search engine outputs that underrepresent women or nonwhite politicians may reinforce inequalities in politics by reifying the collective stereotypical representation of politicians as white and male (Bateson, 2020; Corbett et al., 2022; Stokes-Brown & Dolan, 2010; Vlasceanu & Amodio, 2022).

In two cross-national algorithmic audits (studies 1 and 2) of Google image searches with political queries, we provide mixed evidence of women's algorithmic underrepresentation. On the one hand, women are underrepresented in Google searches in absolute terms, as they consistently appear less frequently in search results than men, with their representation remaining well below the societal 50% mark for gender parity. On the other hand, however, the results also show that the quantitative representation of women in Google search output accurately reflects their presence in lower and upper chambers of most legislative bodies, suggesting that search engines do not introduce additional gender bias relative to existing structural inequalities faced by women. We then examine the possible implications of search engines reiterating this unequal representation, using a more extreme form of algorithmic underrepresentation in two online experiments (studies 3 and 4). Our findings show that more extreme forms of underrepresentation of women or nonwhite politicians in Google search outputs leads voters to underestimate these groups' descriptive representation by roughly 10% points. Crucially, mediation analyses suggest that this perceptual bias regarding descriptive representation results in undesirable political perceptions concerning the viability of politicians from minoritized groups.

The results advance scientific understanding of how AI gatekeepers reflect structural inequalities in politics – and how, under certain conditions, they can amplify them by exacerbating biases in political perceptions and decision- making. Our framework of algorithmic representation adds a crucial component to understanding the enduring structural disadvantages of women and people of color in an increasingly AI-driven political landscape. It contributes to current public and interdisciplinary scientific research concerning algorithmic fairness and injustice (Birhane, 2021; Kalluri, 2020; Weinberg, 2022)

and also provides a reference point for future studies on the topic, capturing the performance of search engines in terms of their representation of political actors at specific point in time. Such insight is integral for raising societal awareness of the discriminatory tendencies inherent in AI-driven systems within increasingly digital political spaces (Friesen et al., 2021). Moreover, it provides an empirical basis for developing new regulation for preventing risks associated with the growing adoption of AI and is thus relevant for a broad range of stakeholders, including policymakers and industry representatives, but also civil society and human right advocacy groups.

Search Engines and Women's Algorithmic and Descriptive Representation

Collective descriptive representation is the quality of political institutions to represent their citizens on the basis of shared political and sociodemographic characteristics (Atkeson & Carrillo, 2007). Increases in descriptive representation of minoritized groups have been connected to policy outcomes that benefit their substantive interests, thus decreasing structural inequalities (Hessami & da Fonseca, 2020; Wängnerud, 2009). The inclusion of women in governing bodies also has a symbolic quality (Stokes-Brown & Dolan, 2010). Whether governments adequately approximate the demographic composition of their population "sends important signals about who should (and should not) participate in politics and the degree to which certain groups and interests will receive a fair hearing" (Stauffer, 2021, p. 1226). Descriptive representation gives rise to a range of political perceptions, such as impressions of who is qualified to be a political leader and more symbolic ways of feeling (un)fairly represented (Atkeson & Carrillo, 2007). At its very core, collective descriptive representation therefore acts as a visibility mechanism that highlights political perspectives with which citizens can identify.

We contend that politicians' representation on search engine outputs – that is, their algorithmic representation – is connected to their descriptive representation in two ways. The first proposition is that the prevalence of gender biases on search engines mirrors structural inequalities in politics. Most people never encounter politicians in person but only through intermediaries, such as media coverage and, increasingly, in AI-curated digital spaces. Research on the gendered mediation thesis has long argued that the transformation of a political reality into a mediated reality happens through various journalistic filters (Gidengil & Everitt, 2003; Rohrbach et al., 2023; Thomas et al., 2020). These filters take the form of gendered organizational structures of newsrooms, journalists' sourcing and writing practices and their professional norms (Riedl et al., 2022). This filtered mediation presents audiences with coverage that quantitatively and qualitatively upholds a masculinized (and white) view of politics. These views matter as they mark the core of gendered political socialization (Bos et al., 2022) and ultimately shape women's political ambition to emerge as candidates (Fox & Lawless, 2011; Kanthak & Woon, 2015).

In "post-broadcast democracies," the human gatekeeping influence of journalists is increasingly substituted for AI gatekeeping (Stier et al., 2022). Studies show that search engines like Google have become key sources of political information (Pradel, 2021; Trevisan et al., 2018; Urman et al., 2022). Moreover, in some cases, people place more trust in the information obtained through their internet searches than in traditional news media (Edelman Trust Institute, 2024). The widespread use and perceived trustworthiness of search engines are cause for concern, as research has documented that such algorithmic

systems can be biased in political contexts (Epstein & Robertson, 2015; Trielli & Diakopoulos, 2022). A number of studies have specifically documented the "representational harm" (Fabris et al., 2020, p. 5) arising from search engines' tendency to underrepresent women (Pradel, 2021; Vlasceanu & Amodio, 2022) as well as people of color (Makhortykh et al., 2021; Noble, 2018).

Applying the filtering logic of gendered mediation to search engines, our first expectation is that political searches on Google will result in an *absolute* underrepresentation of women in image search results (H1). In this paper, we use the term absolute underrepresentation to describe search results where women make up less than 50% of depicted individuals, comparing their algorithmic to the demographic benchmark of (near-)gender parity. Furthermore, if the masculinized political and social reality translates into an algorithmic underrepresentation of women on Google, then the extent of this underrepresentation should vary across societies with different political benchmarks regarding women's inclusion in politics. In addition to an absolute underrepresentation, we expect search engine algorithms to mirror gendered role distributions by underrepresenting women *relative* to the political benchmark of their descriptive representation (H2).[4] The two forms of algorithmic underrepresentation differ in their normative implications. Women's absolute underrepresentation for political queries suggests a societal issue; that is, search engines contribute to the persisting construction of the political sphere as a masculinized sphere in which women are wrongfully cast as a minority. In contrast, women's relative underrepresentation is a much stricter measure of gender bias that also incorporates the idea of algorithmic accuracy; that is, whether or not search engines adequately portray the existing political reality.

H1: Women are algorithmically underrepresented in (gender-neutral) political Google image search outputs.

H2: Women's algorithmic underrepresentation is positively correlated with their descriptive representation across countries but does not exceed it.

The second proposition is that exposure algorithmic representation has the potential to shift perceptions of collective descriptive representation, which reinforces existing inequalities in the political realm. Recent work on strategic discrimination (Bateson, 2020), also referred to as pragmatic bias (Corbett et al., 2022), describes the tendency of voters to withhold electoral support for minority candidates not on the basis of their own direct prejudice but out of a strategic calculus of these candidates' perceived ability to win the election (Ashworth et al., 2024). While direct race or gender bias in the electorate may be declining (Juenke & Shah, 2016; Rohrbach, 2024; Schwarz & Coppock, 2022), studies on strategic discrimination have consistently found that participants perceive women and Black candidates as less electable and, consequently, are less likely to vote for them (Bateson, 2020; Corbett et al., 2022; Green et al., 2022).

We posit that algorithmic representation operates within the mechanism of strategic discrimination in two steps. In a first step, the algorithmic representation of minoritized groups in search engine outputs influences individual perceptions of descriptive representation of those groups. Such misperceptions matter, as "public opinion on descriptive representation is likely to explain political behavior at both the individual and systemic

levels" (Dolan & Sanbonmatsu, 2009, p. 410). We thus expect that search engine output serves as a political heuristic feeding the gendered perception gap, which is the "systematic overestimation by men and underestimation by women of their electability" (Ashworth et al., 2024, p. 290). As most voters have no clear idea of the actual demographic composition of political institutions (Stauffer, 2021), search engine users will approximate the algorithmically curated reality with the political reality (H3):

> **H3:** Algorithmic underrepresentation of women and nonwhite politicians lowers estimations of their descriptive representation.

In a second step, these algorithmically driven misperceptions regarding the actual inclusion of minoritized groups affect electability assessments (H4a). Bateson (2020) identifies the perceived ability of a candidate to raise funds, to generate media coverage and to win over voters as determinants of electability. Voter perceptions of these dimensions remain vague and difficult to empirically capture. We therefore broaden this approach and introduce the representation in the output of Google searches as a key variable in voters' strategic calculus. Because search engines are so regularly and widely used as well as trusted, they present a particularly plausible source for individuals to form their "second-order preferences" – that is, the perception of the ideological and demographic candidate profile the abstract median voter would electorally support (Green et al., 2022, p. 888). Furthermore, information searches online are closer to individuals' daily life than, for instance, concerns about politicians' fundraising capabilities.

In addition to lowering perceived electability of candidates belonging to minoritized groups, there is reason to believe that underestimations of descriptive representation also diminish voters' external efficacy, understood as the perception that government is responsive to their concerns (H4b; Atkeson & Carrillo, 2007). Specifically, Stauffer (2021) convincingly showed for the case of women politicians that it is voters' subjective perception of women's inclusion rather than their objective reality that shapes political evaluations (see also Dolan & Sanbonmatsu, 2009). We replicate this argument and extend it to nonwhite politicians:

> **H4:** Underestimated descriptive representations in turn lower voters' evaluations of (a) the electability of women and nonwhite politicians and (b) voters' external efficacy.

Measuring the Extent of Women's Algorithmic Representation

To investigate women's algorithmic representation (H1, H2), we conducted two algorithm audits of Google image searches.[5] Algorithm auditing has become the dominant research method for "diagnosing problematic behavior in algorithmic systems" (Bandy, 2021, p. 1). Such audits often involve the creation of virtual agents used to simulate the user interaction with an algorithmic system under controlled conditions and to generate system outputs (Urman et al., 2022).

The audits focus on Google image searches for several reasons. First, analyses of Google trends have linked voters' information search behavior to political events and election outcomes, thus highlighting its relevance of the search engine for political

decision-making (Trevisan et al., 2018; Urman & Makhortykh, 2023). Second, longitudinal analyses indicate that the average citizen is spending less time reading and more time viewing images, indicating a rising preference for visual information (American Academy of Arts & Sciences, 2019). Politicians in turn respond to voters' visual preference by increasingly relying on visual communication for their political image building (Bast et al., 2022; Boomgaarden et al., 2016; Carpinella & Bauer, 2021). Third, previous research suggests that visual information environments are particularly likely to "influence users' views of the social world" (Otterbacher et al., 2017, p. 6630). Early audits of the representation of professional occupations found that Google images searches reinforce gender stereotypical associations and portray minority gender in ways suggesting lower professionalism (Kay et al., 2015, April 18 - 23) or entail a negativity backlash (Otterbacher et al., 2017). A recent large-scale study on both Google text and image searches for professional occupations suggests that images were particularly likely to prime and amplify gender bias and thus "come at a critical social cost" (Guilbeault et al., 2024, p. 6).

Study 1: Auditing Gender Bias in Google Image Search

Methods

For the first study, we manually simulated user activity in 56 different countries with a local IP address (accessed through a VPN service called Le VPN; https://www.le-vpn.com/) in the first week of August 2023. The simulation consisted of deploying a virtual agent – that is, a computer script – that was automated to conduct Google image search queries in each geographical location. For consistency, we used the ".com" version of Google due to the possibility of not having a language- specific version of the search engine for all cases. We conducted a Google image search query for the country's lower or single chamber of parliament with the following pattern: [name of legislative body][person] (see Table B1 in the Supplemental Materials). For bicameral systems, the query was repeated for the upper chamber as well, resulting in a total of 84 queried legislative §bodies. In line with Vlasceanu and Amodio (2022), we used the country's dominant language to conduct the query and added the term "person" to obtain results related to people rather than buildings. We chose the abstract term "person" because it carries no or only very weak gendered connotations in most languages. In contrast, more specific terms, such as citizen, politician, or member of parliament would require choosing a grammatical gender in most languages which would likely influence search results (for a discussion see Vlasceanu & Amodio, 2022). For the United States, the queries thus read: "house of representatives person" and "senate person." For each conducted query, we collected the first 75 images and extracted the number of persons and their gender by means of computer vision using the commercial Amazon Rekognition platform (https://aws.amazon.com/rekognition/R).[6] This methodology relies on the Rekognition's ability to detect and analyze faces of humans within an image to extract specific attributes. Specifically, the DetectFaces method was utilized to identify and localize faces within an image, providing bounding box coordinates and confidence scores for each detected face. For each detected face, the Gender attribute, as returned by the DetectFaces response, was used to classify individuals as male or female based on their physical appearance. This yielded a data set of 6,363 images depicting 58,343 persons. We then merged the data at the country-level with women's actual descriptive representation in

these legislative bodies Inter-parliamentary Union (2024). We used the clean browser and cleared its history after each query to prevent the possible impact of the browser history on search outputs.

Our main measure of algorithmic representation is the share of women of all depicted persons in each image (e.g., 50% in an image with two women and two men). We additionally repeat all analyses with the absolute number of depicted women and men as well as a dummy variable predicting the presence of at least one man or woman in each image.

Results

We first test our expectations that political queries on search engines will result women's *absolute* underrepresentation in Google image outputs compared to men (H1). We run generalized linear mixed effects models with by-country random intercepts to predict the share, presence, and number of women in output images. In line with our expectation, the results show a consistent algorithmic underrepresentation of women on all three measures and in search queries for both lower and upper chambers (see Table 1). While far from the demographic benchmark of gender-parity, the extent of women's algorithmic underrepresentation (lower: 29.2%, *CI*[26.8%-31.7%]; upper: 28.1%, *CI*[26.3%-31.4%]) lies a few percentage points above the political benchmark of women's average global descriptive representation (26.7%).

Next we turn to our hypothesis that women's algorithmic representation tracks with the gendered distribution into the political roles, as measured by women's descriptive representation (H2). For this, we assess the correlation between women's algorithmic and actual descriptive representation. We find significant positive correlations in both chambers (lower: $r(56) = 0.38$, $p = .003$; upper: $r(28) = 0.57$, $p = .002$). These associations indicate that the proportion of women in Google images is higher (lower) in countries and chambers with more (fewer) elected women (see Figure 1A).[7]

Contrary to expectation, the evidence indicates a mixed pattern of *relative* algorithmic representation, including both algorithmically driven under- and overrepresentation. Figure 1B illustrates the differences between women's actual and algorithmic representation for all queries. We find no significant difference for the majority of countries and chambers

Table 1. Overview of women's algorithmic representation in Google image search outputs.

	Study 1		Study 2	
Outcomes	Lower	Upper	Lower	Upper
Share women	29.2	28.1	22.5	27.0
	(26.8–31.7)	(26.3–31.4)	(18.2–26.9)	(22.3–31.8)
Presence women	44.5	46.1	41.2	39.2
	(40.8–48.2)	(41.2–50.7)	(36.1–46.4)	(33.0–45.4)
Presence men	61.4	67.0	70.3	67.9
	(57.7–65.2)	(63.6–70.4)	(64.1–76.5)	(59.4–76.3)
Number women	2.9	2.3	3.1	1.2
	(2.4–3.4)	(1.9–2.8)	(2.2–4.0)	(0.8–1.5)
Number men	7.1	5.3	8.2	3.1
	(5.8–8.2)	(4.1–6.5)	(5.7–10.6)	(2.1–4.0)
$N_{images}/k_{countries}$	4243/56	2120/28	92,213/11	59,885/9

Notes. Numbers are average predictions for the fixed effect intercepts from separate models per representation measure. All models are clustered around the country level; models in study 2 are additionally clustered around the location of the VPN ($k_{location} = 11$) and the virtual agents ($k_{agent} = 220$).

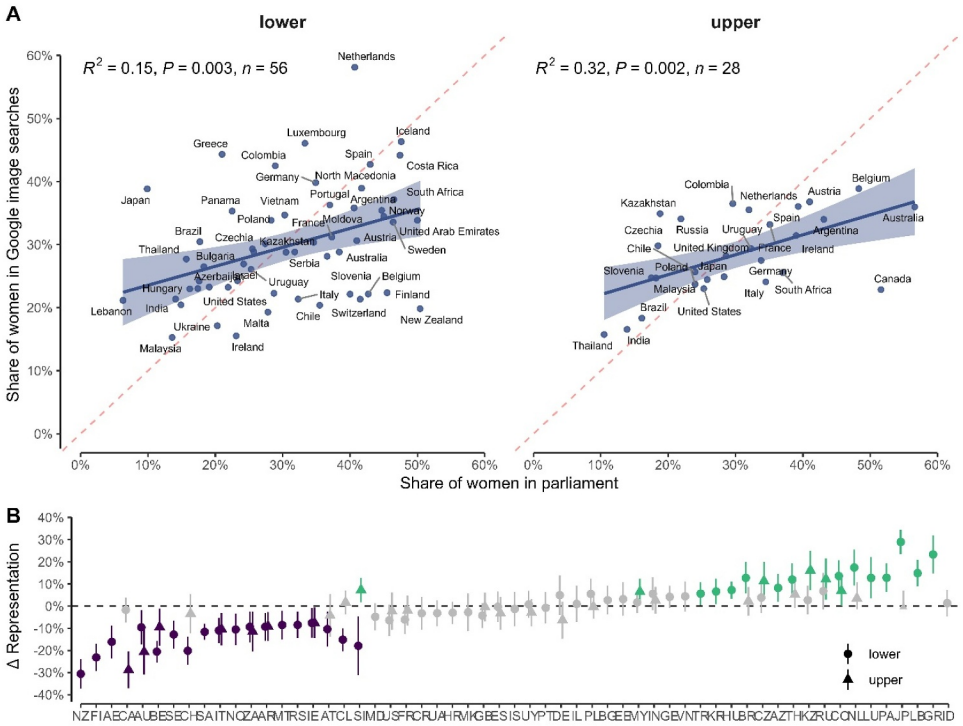

Figure 1. Associations between women's algorithmic and descriptive representation in parliamentary chambers across 56 countries. *Notes*. Panel A depicts women's algorithmic (y-axis) and descriptive representation (x-axis) in study 1. Panel B depicts the mean differences in representation per country and chamber.

(e.g., the U.S. House of Representatives), indicating that search engines accurately mirror women's inclusion in most legislative institutions. Google's search algorithm introduces additional bias in women's underrepresentation in 20 cases. For instance, there is a relative underrepresentation of women in the output for the U.S. Senate query by −5.8% percentage points (*CI*[−11.6%- −0.07%]). However, women are algorithmically overrepresented relative to their descriptive representation in 21 cases.

Study 2: Internal Replication and Audit Robustness Check

Methods

Study 2 constitutes an internal replication of the first study with the goal to probe the robustness of the previous audit (a) for the time of data collection and (b) the noise due to randomization of search engine outputs (Bandy, 2021; Urman et al., 2022). Data collection for study 2 included queries for 20 legislative bodies (nine bicameral and two unicameral countries) and took place in March 2024, almost a year after the first study. We deployed 20 virtual agents – rather than just one – which simultaneously conducted the same queries to account for random noise in each Google search. To account for geographical influences on Google searches, we

paralleled this procedure in each of the 11 countries by modeling the location of the virtual agents through the set of IP addresses provided by Google Compute Engine,[8] which resulted in a total of 4,400 queries. For each query we collected and coded up to the first 50 images in the Google search output. This resulted in a data set of 152,098 images depicting 1,324,560 persons.

Results

For the analysis we repeat the generalized linear mixed effects models from the previous study with random intercepts for the country of the query, the location in which the agent performed the query, and the agent id. We first again test for the presence of women's absolute underrepresentation in Google image searches (H1). Confirming results from the previous study, women are underrepresented on all measures compared to the absolute benchmark of gender-parity (see Table 1). For example, women account for 22.5% (*CI*[18.2%-26.9%]) of depicted persons in searches for lower parliamentary chambers and 27.0% (*CI*[22.3%-21.8%]) in those for upper chambers. We then assess whether the distribution of women's representation on Google image search outputs mirrors their actual inclusion in governments across countries (i.e., relative underrepresentation; H2). Like in study 1, we find a positive association of women's representation on Google search output and their actual presence in legislative bodies. Specifically, we find that a percentage point increase in women's descriptive representation indeed results in an increase in their algorithmic representation by half a percentage point ($b = 0.50$, *CI*[0.46:0.53], $p < .001$; see Table C2.2 in the Supplemental Materials).[9]

Experimental Effects of Algorithmic Representation

So far, the observational studies have established that women are algorithmically underrepresented in Google image searches in absolute terms and that this underrepresentation tends to mirror women's descriptive (under)representation in relative terms. To investigate the consequences of exposure to algorithmic underrepresentation (H3, H4a, H4b), we conducted two online experiments with exposure to more extreme forms of underrepresentation in Google search output as treatments. In both studies, we vary the composition of the image search output between-subjects to be either balanced (i.e., equal representation) or biased (i.e., underrepresentation of minoritized candidates). Whereas study 3 varies the algorithmic representation across gender lines, study 4 additionally accounts for race and intersectional biases.

Study 3: Consequences of Algorithmic Gender Bias

Methods

We recruited a US ($n = 304$) and a UK ($n = 303$) sample through Prolific. Both samples are gender balanced (see section B2 in the Supplemental Materials for a breakdown of sample characteristics as well as a priori sample size calculations).

We constructed the stimuli to resemble the first 11 images of actual Google searches as closely as possible (for a similar approach see Vlasceanu & Amodio, 2022). Both stimuli show screenshots of a Google image search output, keeping the same layout and search bar.

The stimuli only differed in the gender composition of depicted persons (1:1 vs. 1:9). For the manipulation of depicted persons, we used nine images of real persons that we collected for the Swiss, and Australian queries in study 1. This decision was made to ensure the perceived realness of the images while reducing familiarity bias due to participants' recognition of politicians. Image selection for each treatment prioritized similarities in content and context, aside from gender composition. We adapted the stimuli to the samples in two ways to bolster external validity. First, we varied the depicted query to be either a query for the House of Representatives (US sample) or the House of Commons (UK sample). Second, we incorporated an image of the vacant House of Representatives (or House of Commons) and another of the institution's official flag. To prevent participant skipping, a 20-second mandatory exposure was enforced (allused stimuli are included in section B3 in the Supplemental Materials). Figure 2 shows an example of the gender bias condition that was used for the UK sample.

We measured three outcomes. First, we ask participants to provide a pre- and post-treatment estimate of descriptive representation of women politicians using a slider scale from 0 to 100. While the pre-treatment measure represents participants' baseline estimation of women's presence in legislative bodies, the post-treatment measure captures estimation influenced by exposure to algorithmic (under)representation (for a similar approach see Vlasceanu & Amodio, 2022). Second, participants rated the perceived electability of women compared to that of men using a single slider scale ranging from −5 (women less electable) to +5 (women more electable to win; see Bateson, 2020). Finally, we assessed participants' external efficacy beliefs by asking their agreement to two statements on a 5-point scale: (1) "Government officials care what people like me think" and (2) "People like me have a say in what the government does" (see Stauffer, 2021). We include the following measures as statistical controls: participant gender, age, search engine use ("If you think about a - regular day, how many times do you use search engines (like Google or Bing) on a - given day?"), ideology (two items from 1 left/liberal to 10 right/conservative), and party vote ("If there were elections today, which party would you vote for?").

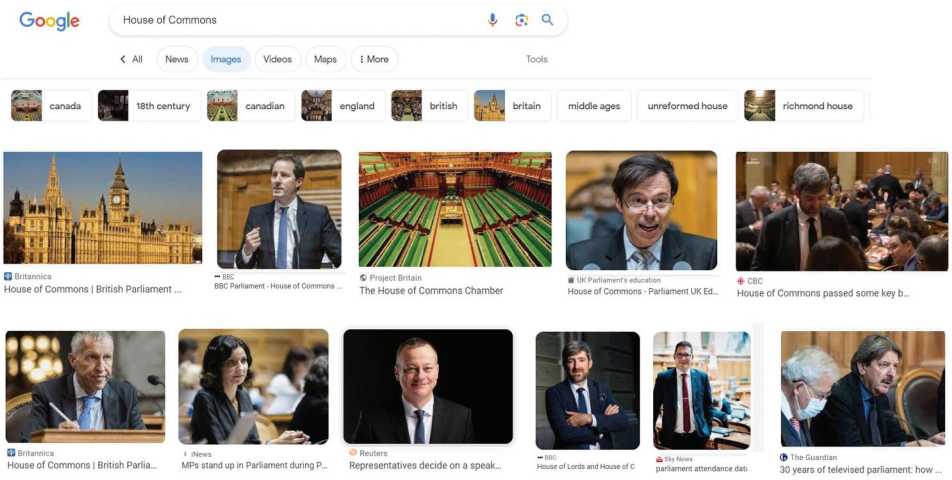

Figure 2. Example of experimental stimulus showing women's underrepresentation in Google image search output used in study 3.

Results

We first assess our hypothesis (H3) that exposure to algorithmic underrepresentation, in particular under more extreme forms of absolute underrepresentation, influences perceptions of political realities by driving people to underestimate the actual descriptive representation of minoritized groups in legislative bodies (i.e., shaping the perception of women's presence in actual governments). In both samples, we predict participants' estimations of women's descriptive representation in a generalized linear mixed effects model with output condition (gender-balanced vs. underrepresentation) and time (pre- vs. post-treatment) as fixed effects and with by-participant random intercepts. Our hypothesis was supported by a significant condition by time interaction in both samples (U.S.: $b = 10.99$, $SE = 0.72$, $t = 15.34$, $p < .001$; UK: $b = 11.58$, $SE = 0.57$, $t = 20.37$, $p < .001$). This interaction is illustrated in Figure 3A, showing that participants have reasonable baseline perceptions of women's inclusion in politics prior to treatment. Whereas exposure to gender-balanced search engine output led to slight overestimations of women's representation relative to their actual presence in legislative bodies, this more extreme form of algorithmic underrepresentation resulted in substantial underestimations of women's inclusion in politics. Figure 3B shows the distribution of the net perceptual bias, calculated by subtracting women's actual descriptive representation from participants' posttreatment estimations. Participants in the algorithmic underrepresentation on average misjudged women's actual representation in the lower chamber by 10.9% points ($SD = 12.1$), compared to the equal representation condition ($M = 0.09$, $SD = 11.3$), $t(588) = 11.4$, $p < .001$, Hedges' $g = 0.93$.

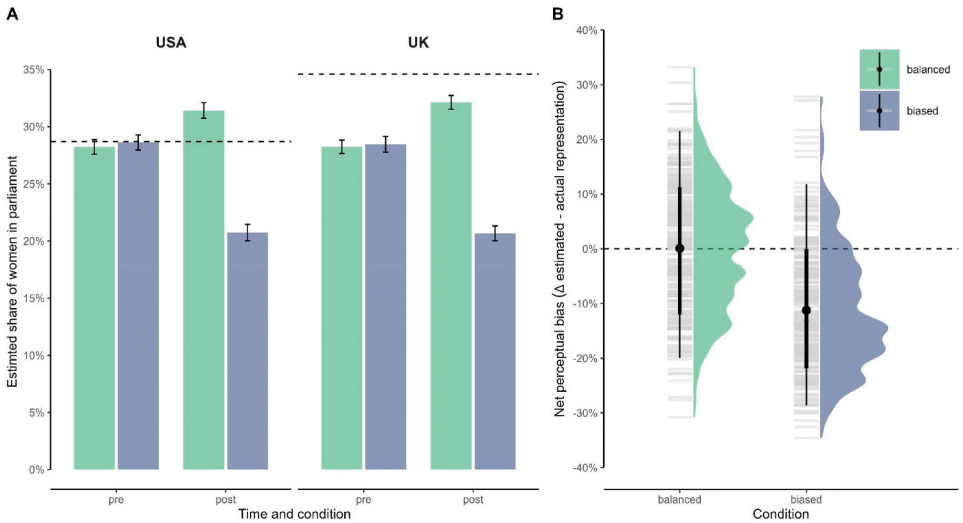

Figure 3. Experimental effects of exposure to algorithmic representation on women's estimated descriptive representation in study 3. *Notes*. Panel A shows the predicted pre- and posttreatment estimations of women's descriptive representation in both experimental conditions and samples. The dashed lines denote women's actual descriptive representation. Panel B depicts the distribution of the net perceptual bias for both conditions.

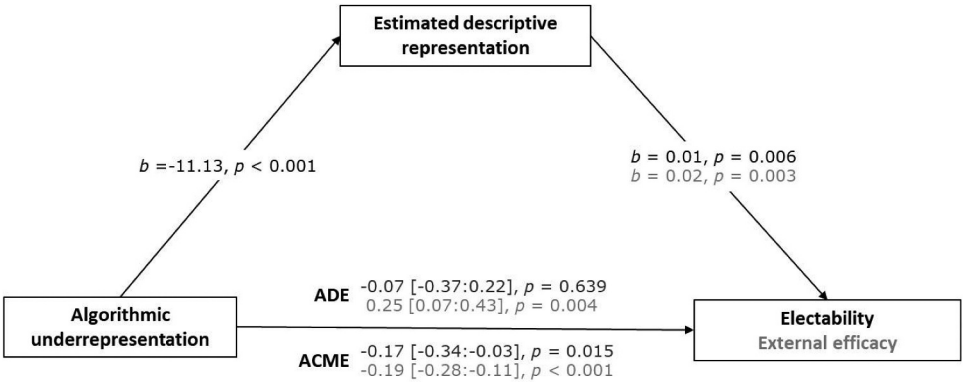

Figure 4. Path diagram of the indirect effect of algorithmic representation in study 3. *Notes*. The horizontal path shows the average direct effect (*ADE*) and the average causal mediation effect (*ACME*) along with their 95% confidence intervals based on 10,000 bootstrapping samples.

Next, we test the strategic bias hypothesis, which states that the under estimations of the descriptive representation minoritized groups – resulting from the algorithmic underrepresentation treatments as documented above – result in lower electability perceptions of candidates belonging to these groups (H4a) and in lower external efficacy beliefs (H4b). For that purpose, we estimate the indirect effect of algorithmic underrepresentation (vs. balanced) on participants' electability assessments of women politicians as well as their own external efficacy through the (mis)estimation of women's descriptive representation (see Figure 4). We find significant negative average causal mediation effects (*ACME*) for both outcomes (electability: −0.17, *CI*[−0.34:-0.03]; external efficacy: −0.19, *CI*[−0.28:-0.11]). Together, the evidence shows that exposure to biases in search engine output drastically lowers voters' baseline perceptions of the role of women in political realities which in turn present a key source of women's electability assessments and voters' efficacy beliefs.

Study 4: Consequences of Algorithmic Gender, Race, and Intersectional Biases

Methods

The previous study focused on the specific case of women's representation in politics. The final study broadens the focus of the consequences of biased algorithmic representation to nonwhite and intersectional politicians (i.e. nonwhite women; Bateson, 2020; Stokes-Brown & Dolan, 2010). Study 4 uses the same experimental "no bias" and "gender bias" conditions of the previous study but adds two more conditions: a "race bias" condition in which nonwhite politicians are underrepresented in Google outputs and an intersectional "double bias" condition in which both women and nonwhite politicians are underrepresented. The design of the stimuli for study 4 identical to that of study 3 except that some white persons were replaced with nonwhite persons (see Figure 2). We again used Prolific to recruit a gender-balanced sample of 781 US participants.

Study 4 uses the same three measures but with a few improvements. First, we drop the pre-treatment measure of estimated descriptive representation to avoid potential anchoring effects.[10] Moreover, participants are now asked to estimate three shares – that of women, nonwhite (Black, Asian, or with Hispanic origins), and female nonwhite politicians. Second,

we employ a more realistic 5-point question to measure perceived electability that is tailored to the study context (Bateson, 2020; Corbett et al., 2022): "Do you think it will be harder or easier for a [woman/member with Black, Asian, or Hispanic origins] member to win re-elections for the House of Representatives, compared to a [man/white member]?" In addition to this assessment of general electability, we now include a ranking exercise to capture perceptions of the specific electability of candidates. This exercise presented participants with a list of eight randomly ordered names that they were told represent candidates for primary elections of the party closest to them. The names were selected from a pretested list to convey a white or nonwhite male or female identity (two names per category; see Crabtree et al., 2023). Participants then had to guess candidates' chances of winning the elections by dragging each name to a box with the labels high (=1), average (=0), or little (=−1) chances of winning the election. We then averaged the rankings of different name combinations to derive electability measures for (a) women, (b) nonwhite, and (c) nonwhite women candidates. Third, we use the same measures of external efficacy as in the previous study.

Results

We first again assess the hypothesized negative effect of exposure to a more extreme form of algorithmic underrepresentation on the estimated representation of minoritized groups (H3). We run separate linear regression models for the three estimation measures and insert the experimental conditions as dummy-coded predictors (along with control variables; see section C4 in the Supplemental Materials for the full models). Figure 5 illustrates the results of these models. It shows, on the one hand, that algorithmic representation with the balanced gender and race composition of depicted politicians results in slight overestimations of the descriptive representation of all minoritized groups. On the other hand, algorithmic underrepresentation of women and/or nonwhite politicians results in

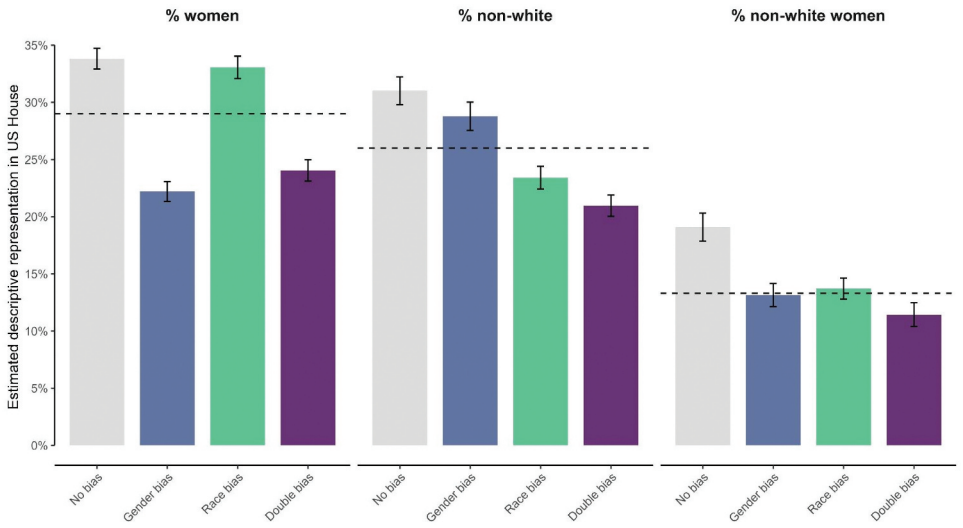

Figure 5. Experimental effects of exposure to algorithmic representation on women's estimated descriptive representation in study 4. *Notes*. Bars show the predicted post- treatment estimations of descriptive representation across experimental conditions. The dashed lines denote the minoritized groups' actual descriptive representation.

a consistent underestimation of their descriptive representation. As in the previous study, exposure to gender-biased Google output decreases the perception of women's presence in the U.S. House of Representatives ($b = -11.53$, $SE = 1.32$, $p < .001$). We find the same pattern for participants in the double bias condition where both women and nonwhite politicians are underrepresented ($b = -9.49$, $SE = 1.32$, $p < .001$). Conversely, the perceived descriptive representation of nonwhite politicians is negatively impacted by the algorithmic underrepresentation of nonwhite politicians (race bias: $b = -10.62$, $SE = 1.52$, $p < .001$; double bias: $b = -12.47$, $SE = 1.53$, $p < .001$). Finally, gender, race, and the intersectional bias in search engine representation all significantly lower the perceived presence of nonwhite women in the House of Representatives, compared to the balanced control condition. However, only the combined underrepresentation of women and nonwhite politicians results in a significant underestimation of intersectional politicians relative to their actual descriptive representation.

Finally, we look for evidence in our data for the causal strategic discrimination mechanism, which states that these algorithmically induced misestimations of descriptive representation result in lower ratings of electability for candidates from minoritized groups (H4a) and participants' own external efficacy beliefs (H4b). We conduct the same mediation analysis as in the study 3, except that we align the treatment and mediator with the minoritized group, resulting in a three-by-three grid (see Figure 6). For example, we test the treatment effect of the race bias (vs. no bias) condition on participants' assessments of nonwhite politicians' general and

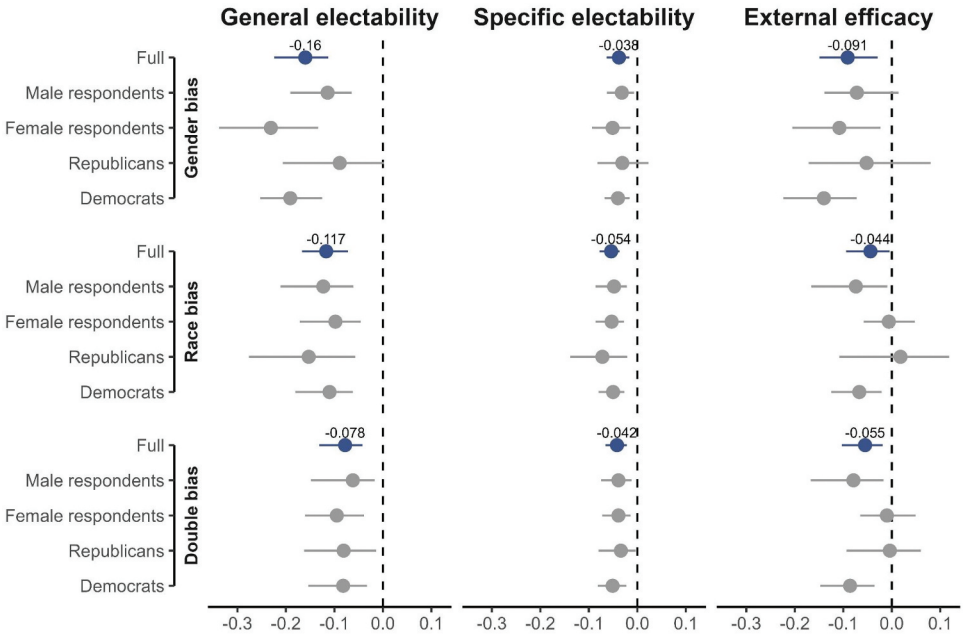

Figure 6. Summary of indirect effects of algorithmic representation from a series of simple mediation analyses using the full sample (blue) and subgroups (gray). *Notes*. Estimates represent average causal mediation effects (*ACME*) for different algorithmic underrepresentation treatments (rows) and outcomes (columns).

specific electability as well as participants' external efficacy. We find consistent evidence of the same causal pattern across all nine models, indicated by significant negative average causal mediation effects (ACME; see blue rows in Figure 6). In other words, the algorithmic underrepresentation of a group diminishes their perceived presence in politics which in turn translates into lower general and specific electability assessments as well as diminished beliefs in participants' external efficacy.

As a last step, we explore whether the *ACME*s vary across different subgroups (see gray rows in Figure 6), namely only female, male, Republican, or Democratic respondents. Though the magnitude of effects varies between subgroups, their direction remains unchanged for the general and specific electability measures. However, we find that the hypothesized depressing effect of algorithmic underrepresentation does not extend to Republican participants in any of the bias conditions. No matter what minoritized group is underrepresented in Google search output, this never translates into reduced feelings of having their voices heard for Republicans, though it similarly affects their misestimations of minoritized groups' descriptive representation. Furthermore, male respondents are similarly unaffected by the algorithmic underrepresentation of women candidates. A somewhat curious finding emerges in the form of statistically insignificant ACMEs for female respondents in the double bias (and race bias) conditions; their external efficacy beliefs remain unchanged after seeing nonwhite women candidates being underrepresented.

General Discussion and Conclusion

What is the role of search engines in upholding views of politics as a masculine and white domain? In this article, we proposed a framework of algorithmic representation of minoritized groups that theorizes this question as a chain of four intertwined types of processes, which we then empirically tested across four empirical studies. Overall, our findings stress the importance of a nuanced treatment of the concept of the algorithmic underrepresentation and distinguishing between the different forms of representation, especially when considering what is normatively desired performance of complex AI-powered systems like search engines. On the one hand, we find that Google image searches algorithmically underrepresent women in their output compared to the societal benchmark of absolute gender parity. On the other hand, women's algorithmic underrepresentation in Google image searches generally tracks in relative terms with their collective descriptive representation in 84 legislative bodies across 56 countries, thus largely reflecting women's relative underrepresentation within these bodies. This finding highlights that Google image search, while showing fewer women than men politicians in absolute terms, does not disproportionately underrepresent women relative to their actual inclusion in legislative bodies This disparity between different forms of underrepresentation raises the question of whether search engines are expected to correct absolute underrepresentation, even if it reflects actual descriptive (under)representation, or reiterate the actual representation of political reality even if it results in reiteration of the underrepresentation of specific groups (and, potentially, become another mechanism for impeding potential changes).

An important consideration regarding the potential (normative) implications of the algorithmic underrepresentation arises from our findings regarding the experimental component of the study. Specifically, we found – using a more extreme version of underrepresentation-based stimuli – that exposure to gender, race, and intersectional bias in

search engine outputs consistently affects perceptions of the political reality by decreasing the estimated presence of politicians belonging to these groups among individual search engine users. We also find that these misperceptions act as a casual conduit through which algorithmic biases not only diminish the perceived chances of winning an election for minority candidates but also result in voters feeling that their voices matter less.

A first implication of this study is the introduction of search engines as an important piece in the puzzle of persistent underrepresentation of minoritized groups in contemporary politics. The findings illustrate how widely used digital tools shape – and are shaped by – existing inequalities; they raise concerns about the future impacts of these tools, especially considering the rapid adoption of AI-driven systems in different societal sectors. This combination of structural and algorithmic explanations of political bias goes with and beyond more established candidate- and voter-centered approaches (see also Dolan & Hansen, 2018). On the supply-side, Juenke and Shah (2016) showed that voters can rarely "choose amongst a menu of racially or politically diverse candidates" (p. 84). Because the share of women in search engine output seems connected to political rather than societal benchmarks, search engines may compound the struggle for visibility experienced by minoritized candidates (Thomsen & King, 2020). Googling politics and not finding any role models who share your gender or skin color might stymy nascent political ambition in potential future candidates (Bos et al., 2022; Fox & Lawless, 2011; Kanthak & Woon, 2015), thus also affecting their visibility in search engine output. On the demand-side, recent studies indicate that overt forms of voter bias are disappearing (Rohrbach, 2024; Rohrbach et al., 2023; Schwarz & Coppock, 2022). By conceptually broadening the notion of strategic discrimination (Bateson, 2020; Green et al., 2022), we show that search engine outputs can act as the basis for voters' assessments of politicians' electability and thus present a more subtle and indirect form of voter bias. Empirically, our mediation analyses consistently support the notion of strategic discrimination as a causal driver of race- and gender-based prejudice in contemporary politics.

Second, we have theoretically and empirically outlined the circular logic of search engine algorithms in politics (see Savaget et al., 2019). Algorithms power AI models which are trained on data reflecting existing structural inequalities; in turn, algorithmic performance replicates the skewed social and political realities in that data, amplifying existing inequalities through their filtering behavior in virtual environments (Fabris et al., 2020; Makhortykh et al., 2021). All findings from our experimental conditions attest to this vicious cycle in which search engines underrepresent candidates from minoritized groups and thereby diminish their perceived electoral chances. However, this pattern reverses into a virtuous cycle in the conditions where we artificially reversed the logic of algorithmic representation in search engine output. Accordingly, algorithmic overrepresentation of minoritized groups boosts voters' estimations of their political presence in government. This finding extends existing work on public perceptions of political underrepresentation (Dolan & Hansen, 2018; Dolan & Sanbonmatsu, 2009; Stauffer, 2021) by positioning search engines as sources of public impression formation. In line with previous research (Stauffer, 2021), the evidence linking perceived inclusion of minority groups to increased feelings of external efficacy is especially noteworthy. In times of fatiguing democratic support around the world (Diamond, 2015; Gavras et al., 2022), algorithmic representation could play a vital role in fostering public support of political institutions. An exception to this pattern are Republican participants' efficacy beliefs, which remain unaffected by exposure to algorithmic underrepresentation of minoritized groups. While adding further evidence to the notion that Republican are just not as concerned about matters of equal

representation (see Badas & Stauffer, 2023), this finding also warrants future research into the differential impact of algorithmic systems across partisan lines. Third, our empirical evidence contributes to ongoing public debates and cross-disciplinary research on algorithmic fairness (Kalluri, 2020; Weinberg, 2022). Our findings are integral for increasing societal awareness about the discriminatory potential of AI-driven systems in the context of political participation (Savaget et al., 2019; Stier et al., 2022). By delineating the consequences of exposure to algorithmic representation, we showcase that the concept of "bias in, bias out" extends beyond a catchphrase; the functioning of proprietary AI-driven systems, such the ones used by Google, should be more actively incorporated into broader discussions on algorithmic governance and injustice (Birhane, 2021). Accordingly, our findings provide an empirical basis for developing new regulation for preventing risks associated with the growing adoption of AI and is thus relevant for a broad range of stakeholders, including policymakers and industry, but also civil society and human right advocacy groups.

Moreover, both our study and other related works (e.g. Vlasceanu & Amodio, 2022) demonstrate that algorithmic representation in search engines can result in the amplification of prejudices, thus worsening the status quo. Under these circumstances, we align with critical studies (e.g. Zajko, 2021) which suggest that we should not expect to easily fix the societal inequalities by making algorithmic systems blind to them, in particular for the cost of "more conventional kinds of political action" (Zajko, 2021, p. 1054). However, we still argue that addressing the underrepresentation of vulnerable groups by an algorithmic system is an important step in countering discrimination and preventing the self-reinforcing loop of biased treatment. At the same time, we need to note that the practical implementation of countering underrepresentation on an algorithmic level remains a challenging task, both practically and normatively. The principle of absolute equality of representation may be difficult to implement, especially considering that it can be applied to different criteria (e.g. ethnicity, race, gender, or religion). It also may be particularly important to apply to certain cases – for instance, political representation which is essential for the functionality of the public sphere – as applying it for all the cases indiscriminately and independently of the specific context may potentially create additional challenges.

Consequently, we acknowledge that mitigating the negative impacts of algorithmic underrepresentation is a non-trivial task. Considering that underrepresentation is a complex phenomenon which amalgamates different forms of algorithmic and social bias, it has to be addressed from different angles. From the technological point of view, inspiration to design search engines with regard to balanced representation of minoritized groups could be drawn from past work on the debiasing of algorithmic systems (e.g. Dash, 2023). In the case of search engines, such "debiasing" strategies could target transparent (and negotiable) representation benchmarks that are anchored in real-world structures. More broadly, this involves greater transparency in algorithmic outputs and enhanced ranking methods, such as evaluating image content instead of only the surrounding text (as what is currently the case for most search engines).From the societal point of view, it is crucial to help users of search engines to develop critical media literacies which can mitigate the effects of skewed algorithmic representation.

This study comes with several limitations. The algorithm audits center on the single case of Google image search algorithm. Insightful and widely used as this specific case may be, it remains unclear how our framework of algorithmic representation transfers to other AI-curated digital spaces. Future research should extend the focus to social media platforms like TikTok that heavily rely on AI to structure content resulting in the amplification of undesirable political

messages, including hate and disinformation (Weimann & Masri, 2023). Another limitation arises from the focus on images. Our visual analysis is based on a low information context in which voters rely on a predominantly heuristic model of political impression formation (Boomgaarden et al., 2016). How algorithmic representation affects political attitudes in contexts where people engage in more deliberative modes of information-seeking, using repeated queries for a range of text, news, and image searches, warrants further investigation. Though the manual coding suggests good reliability, the use of computer vision for automated gender detection in images is limited to a binary classification that disregards the complexity of gender identity (Scheuerman et al., 2019). For ethical and baseline-related reasons regarding the use of AI, our auditing studies also ignore the question of representation of race. Critical analysis and in-depth qualitative work could help establish a benchmark for how search engines (mis)represent nonwhite politicians. Much of our analysis captures dynamics on the level of political institutions and groups of candidates. A promising approach would be extending our framework to the level of individual candidates to study algorithmic representation in the context of politicians with concrete political histories and, for instance, partisan identities (Pradel, 2021). To further enhance external validity, future research could also combine survey or tracking methodology to collect users' actual queries – rather than imposing somewhat artificial ones like in our study – prior to setting up the design of algorithm audits. Finally, our experimental evidence documents how algorithmic underrepresentation decreases perceived electability specific contexts from US and UK politics and relying on rather extreme scenarios of algorithmic underrepresentation (using ratios of underrepresentation of 1:9). More work is needed to assess the contextual boundaries of these algorithmically induced electability penalties and to delineate its downstream electoral consequences, ideally using observational data from actual elections or longitudinal designs. Whereas the political landscape is only slowly becoming more inclusive toward women and nonwhite politicians, the algorithmic landscape has been evolving rapidly in all directions. This article captured how these two landscapes are interlocked to sustain white and masculine views of politics at present; tracking how these dynamics unfold in the future remains a major concern for the legitimacy of democratic structures and processes.

Notes

1. We use the term "non-white" in equivalence to the U.S. Census Bureau's category of "White alone (not Latino or Hispanic)" to refer to individuals who do not report as being of exclusively white Caucasian origin.
2. While the functionality of search engines is not by default enabled by AI and historically relied on simpler information retrieval algorithms, AI applications have been long used to improve the functionality of modern search engines, in particular Google. Hence, we refer to search engines in this study as AI-driven systems.
3. We understand bias in relation to computer systems as system performance that is systematically skewed toward particular individuals or group (see Friedman & Nissenbaum, 1996)
4. For the prevalence of algorithmic underrepresentation, we focus exclusively on gender and disregard the question of race for two reasons. First, measuring the share of nonwhite politicians in Google output would require labeling the race of depicted persons on a purely visual basis. In addition to strong ethical concerns, computational labeling approaches rely on face recognition algorithms which are themselves subject to inherent race biases (Scheuerman et al., 2019). Second, governments do not typically

provide information about the racial or ethnic background of their legislative branches, making it impossible to establish a benchmark for comparison.

5. All studies were pre-registered and received ethical clearance from the institutional review board. Note that the results reported in the manuscript deviated from the pre-registration in several ways. Specifically, the manuscript uses modified operationalizations across studies and omits additional hypotheses concerning alternative outcomes to assess the impact of algorithmic exposure (study 3). All deviations are documented in section A in the Supplemental Materials.
6. Amazon Rekognition predicts the gender of a depicted person as a male vs. female binary based on physical appearance. An algorithmic system itself, Amazon Rekognition's gender prediction has been shown to work best for (white) cis-gender women and men with true positive rates of 95% and 99% respectively (Scheuerman et al., 2019). Commercial gender detection models notoriously perform worse for women and nonwhite people but retain accuracy rates of 80% and better (Schwemmer et al., 2020). We compared the automatically annotated gender with a (binary) manual classification conducted by the authors ($n = 300$) and achieved satisfying reliability (Krippendorff's $\alpha = 0.92$).
7. These bivariate associations hold even after controlling for country- or query-level predictors (see section C1 in the Supplemental Materials for robustness checks of this finding).
8. We find substantial random variation in search engine behavior. Yet the pattern of women's consistent algorithmic underrepresentation holds across VPN locations and individual agent queries. We report more detailed analyses in section C2 of the Supplemental Materials.
9. Replicating the bivariate analysis from study 1, we find also find a positive correlation between women's algorithmic and descriptive representation across 20 legislative bodies, $r(20) = 0.33$, $p = .14$. Note that this bivariate association is not statistically significant, which might be a result of the low number of legislative bodies.
10. This change sought to preempt the concern that the pre-treatment estimation of descriptive representation serves as an initial reference point for the post-treatment estimation (see Mussweiler et al., 2000).

Acknowledgments

We would like to thank Amanda Friesen, members of the Australian Centre for Health Engagement, Evidence and Values as well as panelists of the 2024 annual conference of the International Communication Association and the Australian Political Studies Association for their insightful feedback on earlier version of this paper.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

Notes on contributors

Tobias Rohrbach is an assistant professor of digital journalism at the Department of Communication and Media Research (University of Fribourg). His research focuses on gender-related aspects across fields and disciplines, including strategic, political, and health communication and journalism research. His methodological expertise includes mixed methods designs, combining qualitative exploratory approaches with quantitative and computational analysis.

Mykola Makhortykh is an Alfred Landecker lecturer at the Institute of Communication and Media Studies (University of Bern), where he studies politics- and history-centred information behavior in online environments and how it is affected by the AI-driven systems, such as search engines and recommender systems. To achieve this goal, he combines traditional social science methods with

novel computational approaches. His other research interests include trauma and memory studies, armed conflict reporting, disinformation and computational propaganda research, cybersecurity and critical security studies, and bias in information retrieval systems.

Maryna Sydorova is a data engineer and a scientific programmer at the Institute of Communication and Media Science at the University of Bern. Before working at the Institute, Maryna worked as a freelance data scientist and cloud architect with a particular emphasis on AI. Among other things, she works as a full-stack developer responsible for the implementation of a cloud-based cross-platform algorithm audit infrastructure.

Data Availability Statement

All data and replication materials are available at https://doi.org/10.7910/DVN/VLIGJB.

Open scholarship

This article has earned the Center for Open Science badges for Open Data, Open Materials and Preregistered. The data and materials are openly accessible at https://doi.org/10.7910/DVN/VLIGJB.

References

American Academy of Arts & Sciences. (2019). *Time spent reading*. https://www.amacad.org/humanities-indicators/public-life/time-spent-reading

Ashworth, S., Berry, C. R., & Bueno de Mesquita, E. (2024). Modeling theories of women's underrepresentation in elections. *American Journal of Political Science*, *68*(1), 289–303. https://doi.org/10.1111/ajps.12785

Atkeson, L. R., & Carrillo, N. (2007). More is better: The influence of collective female descriptive representation on external efficacy. *Politics & Gender*, *3*(1), 79–101. https://doi.org/10.1017/S1743923X0707002X

Badas, A., & Stauffer, K. E. (2023). Descriptive representation, judicial nominations, and perceptions of presidential accomplishment. *Representation*, *59*(2), 249–270. https://doi.org/10.1080/00344893.2022.2037699

Bandy, J. (2021). Problematic machine behavior: A systematic literature review of algorithm audits. *Proceedings of the ACM on Human-Computer Interaction*, *5*(1), 1–34. https://doi.org/10.1145/3449148

Bast, J., Oschatz, C., & Renner, A.-M. (2022). Successfully overcoming the "double bind"? A mixed-method analysis of the self-presentation of female right-wing populists on Instagram and the impact on voter attitudes. *Political Communication*, *39*(3), 358–382. https://doi.org/10.1080/10584609.2021.2007190

Bateson, R. (2020). Strategic discrimination. *Perspectives on Politics*, *18*(4), 1068–1087. https://doi.org/10.1017/S153759272000242X

Birhane, A. (2021). Algorithmic injustice: A relational ethics approach. *Patterns*, *2*(2), 1–9. https://doi.org/10.1016/j.patter.2021.100205

Boomgaarden, H. G., Boukes, M., & Iorgoveanu, A. (2016). Image versus text: How newspaper reports affect evaluations of political candidates. *International Journal of Communication*, *10*, 2529–2555.

Bos, A. L., Greenlee, J. S., Holman, M. R., Oxley, Z. M., & Lay, C. J. (2022). This one's for the boys: How gendered political socialization limits girls' political ambition and interest. *The American Political Science Review*, *116*(2), 484–501. https://doi.org/10.1017/S0003055421001027

Carpinella, C., & Bauer, N. M. (2021). A visual analysis of gender stereotypes in campaign advertising. *Politics, Groups, and Identities*, *9*(2), 369–386. https://doi.org/10.1080/21565503.2019.1637353

Center, P. R. (2023a). 118th congress has a record number of women. Technical report. Technical report. Pew Research Center. https://www.pewresearch.org/short-reads/2023/01/03/118th-congress-has-a-record-number-of-women/?gad_source=1&gad_campaignid=23853897313&gbraid=0AAAAA-ddO9HpD65ohyS1X6WcN2lscLF8A&gclid=Cj0KCQjwoMXQBhDcARIsAH-eEtsbkB9BMOCFideVN1VMVdS0jdS0726sNnRFwBYwKUI1HeE05D8y_N8aAkibEALw_wcB

Center, P. R. (2023b). 118th U.S. congress continues to grow in racial, ethnic diversity. Technical report. Pew Research Center. https://www.pewresearch.org/short-reads/2023/01/09/u-s-congress-continues-to-grow-in-racial-ethnic-diversity/

Corbett, C., Voelkel, J. G., Cooper, M., & Willer, R. (2022). Pragmatic bias impedes women's access to political leadership. *Proceedings of the National Academy of Sciences*, *119*(6), 1–11. https://doi.org/10.1073/pnas.2112616119

Crabtree, C., Kim, J. Y., Gaddis, S. M., Holbein, J. B., Guage, C., & Marx, W. W. (2023). Validated names for experimental studies on race and ethnicity. *Scientific Data*, *10*(1), 130. https://doi.org/10.1038/s41597-023-01947-0

Dash, S. (2023). Fairness in image search: A study of occupational stereotyping in image retrieval and its debiasing, 1–24. https://doi.org/10.48550/arXiv.2305.03881

Diakopoulos, N., Trielli, D., Stark, J., & Mussenden, S. (2018). I vote for-how search informs our choice of candidate. In M. Moore & D. Tambini (Eds.), *Digital dominance: The power of Google, Amazon, Facebook, and Apple* (Vol. 22, pp. 320–341). Oxford University Press.

Diamond, L. (2015). Facing up to the democratic recession. *Journal of Democracy*, *26*(1), 141–155. https://doi.org/10.1353/jod.2015.0009

Dolan, K. A., & Hansen, M. (2018). Blaming women or blaming the system? Public perceptions of women's underrepresentation in elected office. *Political Research Quarterly*, *71*(3), 668–680. https://doi.org/10.1177/1065912918755972

Dolan, K. A., & Sanbonmatsu, K. (2009). Gender stereotypes and attitudes toward gender balance in government. *American Politics Research*, *37*(3), 409–428. https://doi.org/10.1177/1532673X08322109

Edelman Trust Institute. (2024). *2024 Edelman Trust Barometer*. https://www.edelman.com/trust/2024/trust-barometer

Epstein, R., & Robertson, R. E. (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), 4512–4521. https://doi.org/10.1073/pnas.1419828112

Fabris, A., Purpura, A., Silvello, G., & Susto, G. A. (2020). Gender stereotype reinforcement: Measuring the gender bias conveyed by ranking algorithms. *Information Processing & Management*, *57*(6), 1–24. https://doi.org/10.1016/j.ipm.2020.102377

Fox, R. L., & Lawless, J. L. (2011). Gendered perceptions and political candidacies: A central barrier to women's equality in electoral politics. *American Journal of Political Science*, *55*(1), 59–73. https://doi.org/10.1111/j.1540-5907.2010.00484.x

Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. *ACM Transactions on Information Systems*, *14*(3), 330–347. https://doi.org/10.1145/230538.230561

Friesen, A., Burge, R., & Britzman, K. (2021). Digital segregation: Gender, occupation, and access to politics. *Social Science Computer Review*, *39*(1), 38–55. https://doi.org/10.1177/0894439319851164

Gavras, K., Schoen, H., & Wuttke, A. (2022). Have Europeans grown tired of democracy? New evidence from eighteen consolidated democracies, 1981-2018. *British Journal of Political Science*, *52*(1), 416–428. https://doi.org/10.1017/S0007123420000149

Gidengil, E., & Everitt, J. (2003). Talking tough: Gender and reported speech in campaign news coverage. *Political Communication*, *20*(3), 209–232. https://doi.org/10.1080/10584600390218869

Green, J., Schaffner, B., & Luks, S. (2022). Strategic discrimination in the 2020 Democratic primary. *Public Opinion Quarterly*, *86*(4), 886–898. https://doi.org/10.1093/poq/nfac051

Guilbeault, D., Delecourt, S., Hull, T., Desikan, B. S., Chu, M., & Nadler, E. (2024). Online images amplify gender bias. *Nature*, *626*(8001), 1049–1055. https://doi.org/10.1038/s41586-024-07068-x

Heldman, C., Carroll, S. J., & Olson, S. (2005). "She brought only a skirt": Print media coverage of Elizabeth Dole's bid for the Republican presidential nomination. *Political Communication*, *22*(3), 315–335. https://doi.org/10.1080/10584600591006564

Hessami, Z., & da Fonseca, M. L. (2020). Female political representation and substantive effects on policies: A literature review. *European Journal of Political Economy*, *63*(101896), 1–9. https://doi.org/10.1016/j.ejpoleco.2020.101896

Inter-parliamentary Union. (2024). *Global and regional averages of women in national parliaments* Report. https://data.ipu.org/women-averages?month=12&year=2022&op=Show+averages&form_build_id=form-jad8uVieHieZXN3-R86CKhsMeY6FaVjrj9t3Kao4-Bk&form_id=ipu__women_averages_filter_form

Juenke, E. G., & Shah, P. (2016). Demand and supply: Racial and ethnic minority candidates in white districts. *The Journal of Race, Ethnicity, and Politics*, *1*(1), 60–90. https://doi.org/10.1017/rep.2015.2

Kalluri, P. (2020). Don't ask if artificial intelligence is good or fair, ask how it shifts power. *Nature*, *583* (7815), 169–169. https://doi.org/10.1038/d41586-020-02003-2

Kanthak, K., & Woon, J. (2015). Women don't run? Election aversion and candidate entry. *American Journal of Political Science*, *59*(3), 595–612. https://doi.org/10.1111/ajps.12158

Kay, M., Matuszek, C., & Munson, S. A. (2015, April 18 - 23). 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*, Seoul Republic of Korea. (pp. 3819–3828). https://doi.org/10.1145/2702123.2702520

Makhortykh, M., Urman, A., & Ulloa, R. (2021). Detecting race and gender bias in visual representation of AI on web search engines. In L. Boratto, S. Faralli, M. Marras, & G. Stilo (Eds.), *Advances in bias and fairness in information retrieval* (pp. 36–50). Springer. https://doi.org/10.1007/978-3-030-78818-6_5

Mussweiler, T., Strack, F., & Pfeiffer, T. (2000). Overcoming the inevitable anchoring effect: Considering the opposite compensates for selective accessibility. *Personality & Social Psychology Bulletin*, *26*(9), 1142–1150. https://doi.org/10.1177/01461672002611010

Noble, S. U. (2018). *Algorithms of oppression*. New York University Press.

Otterbacher, J., Bates, J., & Clough, P. (2017). *Competent men and warm women: Gender stereotypes and backlash in image search results* (pp. 6620–6631). https://doi.org/10.1145/3025453.3025727

Pradel, F. (2021). Biased representation of politicians in Google and Wikipedia search? The joint effect of party identity, gender identity and elections. *Political Communication*, *38*(4), 447–478. https://doi.org/10.1080/10584609.2020.1793846

Riedl, A., Rohrbach, T., & Krakovsky, C. (2022). "I can't just pull a woman out of a hat": A mixed-methods study on journalistic drivers of women's representation in political news. *Journalism and Mass Communication Quarterly*, *Online first*, *101*(3), 1–24. https://doi.org/10.1177/10776990211073454

Rohrbach, T. (2024). Are women politicians kind and competent? Disentangling stereotype incongruity in candidate evaluations. *Political Behavior*, *47*(1), 411–434. https://doi.org/10.1007/s11109-024-09956-4

Rohrbach, T., Aaldering, L., & Van Der Pas, D. J. (2023). Gender differences and similarities in news media effects on political candidate evaluations: A meta-analysis. *Journal of Communication*, *73*(2), 101–112. https://doi.org/10.1093/joc/jqac042

Savaget, P., Chiarini, T., & Evans, S. (2019). Empowering political participation through artificial intelligence. *Science & Public Policy*, *46*(3), 369–380. https://doi.org/10.1093/scipol/scy064

Scheuerman, M. K., Paul, J. M., & Brubaker, J. R. (2019). How computers see gender: An evaluation of gender classification in commercial facial analysis services. In *Proceedings of the ACM on Human-Computer Interaction 3*(CSCW). https://doi.org/10.1145/3359246

Schwarz, S., & Coppock, A. (2022). What have we learned about gender from CandidateChoice experiments? A meta-analysis of sixty-seven factorial survey experiments. *Journal of Politics*, *84*(2), 655–668. https://doi.org/10.1086/716290

Schwemmer, C., Knight, C., Bello-Pardo, E. D., Oklobdzija, S., Schoonvelde, M., & Lockhart, J. W. (2020). Diagnosing gender bias in image recognition systems. *Socius*, *6*, 1–17. https://doi.org/10.1177/2378023120967171

Stauffer, K. E. (2021). Public perceptions of women's inclusion and feelings of political efficacy. *The American Political Science Review*, *115*(4), 1226–1241. https://doi.org/10.1017/S0003055421000678

Stier, S., Mangold, F., Scharkow, M., & Breuer, J. (2022). Post post-broadcast democracy? News exposure in the age of online intermediaries. *The American Political Science Review*, *116*(2), 768–774. https://doi.org/10.1017/S0003055421001222

Stokes-Brown, A. K., & Dolan, K. A. (2010). Race, gender, and symbolic representation: African American female candidates as mobilizing agents. *Journal of Elections, Public Opinion and Parties*, *20*(4), 473–494. https://doi.org/10.1080/17457289.2010.511806

Thomas, M., Harell, A., Rijkhoff, S. A. M., & Gosselin, T. (2020). Gendered news coverage and women as heads of government. *Political Communication*, *38*(4), 388–406. https://doi.org/10.1080/10584609.2020.1784326

Thomsen, D. M., & King, A. S. (2020). Women's representation and the gendered pipeline to power. *The American Political Science Review*, *114*(4), 989–1000. https://doi.org/10.1017/S0003055420000404

Trevisan, F., Hoskins, A., Oates, S., & Mahlouly, D. (2018). The Google voter: Search engines and elections in the new media ecology. *Information Communication & Society*, *21*(1), 111–128. https://doi.org/10.1080/1369118X.2016.1261171

Trielli, D., & Diakopoulos, N. (2022). Partisan search behavior and Google results in the 2018 U.S. midterm elections. *Information Communication & Society*, *25*(1), 145–161. https://doi.org/10.1080/1369118X.2020.1764605

Urman, A., & Makhortykh, M. (2022). "Foreign beauties want to meet you": The sexualization of women in Google's organic and sponsored text search results. *New Media and Society*, *26*(5), 1–22. https://doi.org/10.1177/14614448221099536

Urman, A., & Makhortykh, M. (2023). You are how (and where) you search? Comparative analysis of web search behavior using web tracking data. *Journal of Computational Social Science*, *6*(2), 741–756. https://doi.org/10.1007/s42001-023-00208-9

Urman, A., Makhortykh, M., & Ulloa, R. (2022). The matter of chance: Auditing web search results related to the 2020 U.S. presidential primary elections across six search engines. *Social Science Computer Review*, *40*(5), 1323–1339. https://doi.org/10.1177/08944393211006863

U.S. Census Bureau. (2023). *QuickFacts: United States* [Dataset]. https://www.census.gov/quickfacts/fact/table/US/PST045223

Vlasceanu, M., & Amodio, D. M. (2022). Propagation of societal gender inequality by internet search algorithms. *Proceedings of the National Academy of Sciences*, *119*(29), 1–8. https://doi.org/10.1073/pnas.2204529119

Wängnerud, L. (2009). Women in parliaments: Descriptive and substantive representation. *Annual Review of Political Science*, *12*(1), 51–69. https://doi.org/10.1146/annurev.polisci.11.053106.123839

Weimann, G., & Masri, N. (2023). Spreading hate on TikTok. *Studies in Conflict and Terrorism*, *46*(5), 752–765. https://doi.org/10.1080/1057610X.2020.1780027

Weinberg, L. (2022). Rethinking fairness: An interdisciplinary survey of critiques of hegemonic ML fairness approaches. *The Journal of Artificial Intelligence Research*, *74*, 75–109. https://doi.org/10.1613/jair.1.13196

White, R. W., & Horvitz, E. (2015). Belief dynamics and biases in web search. *ACM Transactions on Information Systems*, *33*(4), 1–46. https://doi.org/10.1145/2746229

Zajko, M. (2021). Conservative AI and social inequality: Conceptualizing alternatives to bias through social theory. *AI & Society*, *36*(3), 1047–1056. https://doi.org/10.1007/s00146-021-01153-9