

Gender Stereotype Reinforcement: Measuring the Gender Bias Conveyed by Ranking Algorithms

Alessandro Fabris, Alberto Purpura, Gianmaria Silvello, Gian Antonio Susto

Department of Information Engineering, University of Padua, Italy.

Abstract

Search Engines (SE) have been shown to perpetuate well-known gender stereotypes identified in psychology literature and to influence users accordingly. Similar biases were found encoded in Word Embeddings (WEs) learned from large online corpora. In this context, we propose the *Gender Stereotype Reinforcement* (GSR) measure, which quantifies the tendency of a SE to support gender stereotypes, leveraging gender-related information encoded in WEs.

Through the critical lens of construct validity, we validate the proposed measure on synthetic and real collections. Subsequently, we use GSR to compare widely-used Information Retrieval ranking algorithms, including lexical, semantic, and neural models. We check if and how ranking algorithms based on WEs inherit the biases of the underlying embeddings. We also consider the most common debiasing approaches for WEs proposed in the literature and test their impact in terms of GSR and common performance measures. To the best of our knowledge, GSR is the first specifically tailored measure for IR, capable of quantifying representational harms.

Keywords: Fairness, Gender Stereotypes, Information Retrieval, Search Engines, Word Embeddings

1. Introduction

In a world with zettabytes of data, SEs become the gatekeepers of information. The continuous growth of internet-based content, the maturity of Information Retrieval (IR - the scientific field underlying SEs), along with seamless user experience, contribute to their widespread use. Since the early 2000s, SEs have been utilized by over 90% of internet users¹ and, in the past decade, they have consistently been reported as the most trusted source for general news and information². Constant availability of information has shaped expectations and cognitive processes of SE users [77]. Such factors concur to the importance and relevance of SEs in acquiring knowledge and culture, including perceptions about gender and stereotypes [46].

Stereotypes can be modelled as associative networks of concepts [73]. They may arise from co-occurrence of features [48], such as membership to a group and display of certain traits and roles, which become linked in a Bayesian fashion based on culture and direct observation [40]. Women and men are particularly salient categories, recognizable since an early age, and available for stereotypical association with traits, behaviors and events [50]. In turn, even when outspokenly rejected, gender stereotypes influence the lives of women and men both descriptively and prescriptively, shaping the qualities, priorities and needs that members of each gender are expected to possess [22]. During their lives, individuals are frequently exposed to information about gender, through direct experience and indirect information coming from social interactions and cultural representations [20], often portrayed by the media.

Email addresses: fabrisal@dei.unipd.it (Alessandro Fabris), purpuraa@dei.unipd.it (Alberto Purpura), silvello@dei.unipd.it (Gianmaria Silvello), gianantonio.susto@dei.unipd.it (Gian Antonio Susto)

¹<https://www.pewresearch.org/internet/2012/03/09/search-engine-use-2012/>

²https://www.edelman.com/sites/g/files/aatuss191/files/2019-03/2019_Edelman_Trust_Barometer_Global_Report.pdf

Cultivation theory [30], historically focused on television, posits that increasing exposure to a medium and its contents leads to a progressive alignment to the beliefs, culture and reality depicted in the televised world. Within this framework, the way women are depicted in primetime television has been studied; recent analysis highlights persistent representational stereotypes related to physical appearance and warmth [76], confirmed by public opinion [20, 79]. According to cultivation theory, heavy viewers are likely to be influenced in their perception of the real world, due to the availability heuristic [74]: in judging frequency and normality (e.g. of women being affectionate), they resort to the examples that come to their mind, the media being a potential source of information to recall. The availability heuristic has been proposed and verified as a general shortcut in human cognitive processes [81], and recently studied as a bias that arises while exploring result pages from SEs [61].

Inevitably, SEs influence users, helping them to link topics, concepts and people as they read, browse and acquire knowledge. They therefore can play an important role in countering or reinforcing stereotypes. For instance, search results on Google images were found to reflect current gender differences in occupation, with a tendency to slight exaggeration [46]; at the time of the study, searching images of a job with a female-to-male ratio of 1:4 in the employed population, such as software engineer, would yield pictures depicting women in less than 20% of the results. Moreover, manipulation of female-male representation in search results about a job, artificially increasing the presence of one gender in images, significantly impacted people’s perception about gender ratios in that occupation [46]. A study on Bing photos found a greater frequency of women in depictions of warm traits (e.g. sensitive), while men are more common in searches about competence traits (e.g. intelligent) [62]. These results highlight the importance of measuring and countering bias in SEs, as recently pointed out by critical race and gender studies scholarship [58].

Gender stereotypes held by people are commonly measured in two ways: directly, on the basis of individual agreeing with statements about gender and specific traits [20]; indirectly, via Implicit Association Tests (IAT) between mental representations of objects [35] or assessment of attitude through priming [23]. Indirect tests are appealing as they allow an unobtrusive assessment of attitudes towards groups (determined e.g. by gender and ethnicity) and can measure association of categories, such as **women**, with words from a specific domain, such as **family**, even when subconscious. Large text corpora sourced from the web, such as Wikipedia, have been found to echo some of the above biases: as an example, Wikipedia entries related to women are more likely to mention marriage- and sex-related contents and events [34]. Interestingly, Word Embeddings can be used to detect gender-related biases in the corpus they have been trained on [18, 29].

Word Embeddings (WEs) are vectorial representations of words computed automatically using different supervised and unsupervised machine learning approaches [45, 54]. Most frequently, they are learnt from large text corpora available online (such as Wikipedia, Google News and Common Crawl, capturing semantic relationships of words based on their usage. Recent work [7] shows that WEs retain the stereotypical associations from their training corpora, encoding a full spectrum of biases from the IAT, including gender-related ones about career and family, science and arts. Additional problematic depictions of men and women have been identified in these WEs, including sexist analogies (such as **woman** – **man** \simeq **midwife** – **doctor** \simeq **whore** – **coward** [6]) and representation of jobs skewed with respect to gender, in ways that reflect current gender gaps in the US workforce [17, 29, 70]. For this reason, WEs have been proposed as an unobtrusive measurement tool of the average bias of the many contributors to these corpora and, generalizing, from the society they live in [29] or the language they speak [18]. Based on co-occurrence with intrinsically gendered terms within the text corpora (such as **woman** and **man**), a *genderedness* score can be derived for each word in the embedding space. Among words with a high score, some are duly gendered (**hers**, **his**), while others reflect an accidental status quo aligned with stereotype (**hygienist**, **electrician** - stereotypically female and male, respectively). This bias, undesirable when WEs are part of a socio-technical system, is an interesting property we can leverage to measure gender stereotypes in SEs.

In this work, we propose the *Gender Stereotype Reinforcement* (GSR) measure that is specifically tailored to quantify the tendency of a SE’s ranked list to support gender stereotypes. GSR exploits gender bias encoded in WEs to detect and quantify the extent to which a SE responds to stereotypically gendered queries with documents containing stereotypical language of same polarity.

Firstly, we validate the word-level genderedness score against well-studied gender stereotypes and subsequently, we employ a basic compositional model to quantify whether retrieved documents are connected

to queries along stereotypical lines. We operationalize GSR based on this model, and verify its ability to capture *direct* and *indirect gender stereotype* reinforcement on synthetic and real collections.

Secondly, we audit IR ranking algorithms from different families: (i) lexical models, including BM25 [71], Query Likelihood Model (QLM - [87]) and `tf-idf` [72]; (ii) semantic models, such as those using Word2Vec with additional compositionality (`w2v_add`) and with self-information (`w2v_si`) [82]; (iii) neural architectures, including Deep Relevance Matching Model (DRMM - [36]), and Match Pyramid (MP - [63]). We measure each system’s performance and GSR on the Text REtrieval Conference (TREC)³ Robust04 [38] curated and widely-used news-based collection. We also perform qualitative analysis of queries with the highest genderedness score and we find these queries to mirror gender stereotypes studied in psychology literature. Moreover, we analyze these ranking models to investigate whether semantic and neural models inherit problematic gender associations from underlying WEs, while verifying the neutrality of lexical models in this regard. We also investigate the tradeoff between performance and fairness for these families of models.

Thirdly, we assess the impact of debiasing WEs [6] in the context of IR, confirming recent findings that gender-related information is redundantly encoded along multiple directions [33]. We conclude that the genderedness score, estimated by the proposed GSR measure, is a good proxy for gender bias [29]; and, that treating WEs to neutralize it is not sufficient to enforce a real lack of gender bias in word representations and downstream tasks.

Finally, we discuss the construct validity and reliability of our measurement model [42, 52]. We decouple GSR as a *construct* (the unobservable theoretical abstraction we aim to characterize), from its *operationalization* (how we measure it), and elucidate the underlying assumptions and properties it should capture. In such context, we argue that some clustering of language along a gendered dimension captured by WEs is inevitable due to domain-specificity of language. The reliability of GSR is evaluated by testing its stability when computed based on WEs learned from different corpora and learning architectures, ensuring that the measurement is robust and not overly dependent on choices of training set and architecture.

Contributions of this work include:

1. GSR measure tailored for SEs and its evaluation within the *construct validity* framework;
2. audit, in terms of GSR, of several widely-known and used ranking algorithms;
3. estimation of the impact of different WE debiasing approaches, both on ranking effectiveness and countering gender bias.

Outline. The rest of this paper is organized as follows. Section 2 describes related works from different domains as IR, Natural Language processing (NLP), algorithmic auditing, social psychology and validity theory. The GSR measure is described in Section 3, preceded by a detailed definition of the abstract construct we aim to quantify, and followed by a toy example that favors a discussion of its key properties. Experiments on real and synthetic IR collections are reported in Section 4, while Section 5 summarizes our conclusions and outlines future works.

2. Related work

2.1. WEs and neural models in IR

Word2Vec [54] was the first widely used WE model. Word2Vec can learn similar representations for terms used in similar contexts in the training data, i.e. corpora of millions of documents in natural language. In addition, as the embedded word representations learned with Word2Vec reflect the usage distribution of respective terms, they have been employed as a proxy for the semantic similarity of terms in many NLP applications. The popularity of Word2Vec also paved the way to other machine learning approaches to obtain embedded word representations such as GloVe [67] or FastText [45]. WE models were soon adopted in the IR domain, promoting the exploration of deep learning approaches for document retrieval [55].

³<http://trec.nist.gov/>

Lexical approaches such as `tf-idf` [72], `QLM` [87] and `BM25` [71] were the first and most popular techniques adopted for document retrieval. Nevertheless, these retrieval models do not take into account terms which are not contained in the user query nor their semantics. For this reason, the potential offered by embedded word representations – i.e. the possibility to represent the meaning of a term and compare it to others in a measurable way – was soon put to use by newly proposed retrieval models.

The simplest WE-based document retrieval approach in our experiments was named `w2v-add` [82]. In this case, we compute a query and a document representation averaging the WEs of the terms they contain, and then rank documents according to their cosine similarity to the query vector. This approach however reduces the query/document representation problem to the core. For example, it does not take into account the relative importance of each term. `w2v-si` solves this problem: queries and documents representations are obtained computing a weighted average of the word vectors of their terms and then rank the documents as in the previous case. Each term weight corresponds to its self-information (si) which is a term specificity measure similar to `IDF` [14].

Among the first most successful deep learning models for IR, there is `DRMM` [36]. `DRMM` uses embedded representation of words to compute the similarity between every pair of terms in a user query and each document in a ranked list.

Another paradigmatic approach in the Neural IR field is `MatchPyramid` (`MP` - [63]). This approach, originally proposed as a document classification model, was also successfully applied to the ranked task.

For our study, we select these two approaches, being popular in IR and easy to use. Moreover, they allow us to evaluate the impact of diverse WEs in different Neural IR architectures.

2.2. Gender stereotype in WEs and SEs

A convincing body of research shows that WEs learnt on large corpora of text available online encode cultural aspects, some of which undesirable. Among them, worth noting at the core of this work are gender-related biases which comprise: sexist analogies [6], stereotypical association of gender with science and arts [7], representation of occupations correlated to differences in female and male employment [17, 29, 70], gender roles in career and within the family [7]. Communion (also called warmth) and agency are two further dimensions consistently associated with gender [20], analyzed in Section 3.2.1; in line with this stereotype, we find warm traits (e.g. “emotional”) to have female polarity in the embedding space, while agentic traits (e.g. “aggressive”) are more commonly associated to men (Section 3.2.1). A wealth of studies in psychology and labor economics literature confirms the presence of the above-mentioned biases in society [10, 16, 20, 25, 39, 59, 60], making their presence in WEs particularly interesting.

Several of these biases, found in SEs, potentially reinforce gender stereotype through powerful and pervasive search tools available to the public. Kay et al. [46] show that gender bias in image search results is exaggerated: the gender distribution for Google image results about jobs are correlated to and amplify differences in female and male employment. Bing images associate agentic traits to men and warm traits to women [62]. `Monster` and `CareerBuilder` were audited, displaying group unfairness against female candidates in 1/3 of the job titles surveyed [11]. This does not imply that these SEs are likely to have the same biased WEs as part of their algorithmic machinery. Rather, finding that known gender biases in SEs are also encoded in vectorial representations of words suggests that WEs can be used as a tool to measure gender bias in SEs.

In this respect, Bolukbasi et al. [6] find that gender-related information for each word in the embedding space is mostly confined within a single dimension:

1. They propose ten word pairs to define gender: `she-he`, `her-his`, `woman-man`, `Mary-John`, `herself-himself`, `daughter-son`, `mother-father`, `gal-guy`, `girl-boy`, `female-male`.
2. For each pair they compute the difference between the two word vectors, obtaining ten candidate vectors (dimensions) to encode gender.
3. They stack the ten vectors into a single matrix, on which they perform a principal component analysis, finding 60% of the variance explained by first principal component, subsequently treated as the gender

subspace w_g . We dub *genderedness* score of a word w , its scalar projection along the gender subspace

$$g(w) = \frac{w \cdot w_g}{|w||w_g|}, \quad (1)$$

and use it as a building block to operationalize GSR (Section 3.2.1).

The sign and magnitude of $g(w)$ determine the polarity and strength of gender-association for word w - e.g. $g(\text{sister}) = 0.31$, $g(\text{brother}) = -0.22$. After identifying a gender subspace (or direction w_g), Bolukbasi et al. [6] remove gender-related information from most words via orthogonal projection. Only intrinsically gendered word pairs (such as **she**, **he**) retain a non-zero component in the gender subspace. Prost et al. [70] propose a *strong* variant of this approach where the procedure applies to the whole vocabulary. This family of debiasing techniques seem limited and imperfect [33], with gender information redundantly encoded along multiple dimensions, and thus hard to eradicate. Confirmation of this statement is given in the context of SEs and gender stereotype in Section 4.2.3, where we assess the impact of regular and strong debiasing with respect to performance and GSR of IR models based on WEs.

2.3. Fairness and diversity in IR

Fairness in information retrieval and recommendation is an area of increasing interest for academia and industry, with entire tracks⁴, workshops⁵ and corporate teams devoted to such a complex topic. Efforts in the field are aimed at emphasizing the social context which SEs and recommender systems inhabit and influence.

We borrow from Ekstrand et al. [21] in sketching a taxonomy of fairness in search along two dimensions: the people benefiting from our efforts and the type of harm we are trying to prevent. Based on their position in the information pipeline, fairness can benefit:

1. consumers of documents (e.g. SE users), when focused on user privacy [27], content diversity [28] or targeted advertisement, which may imply unequal opportunity for different segments of the population [9];
2. providers of documents, who deserve an equal chance to be read, viewed and clicked [75, 83, 84, 86];
3. information subjects (mentioned in documents or queries), who may be present in contexts where they would rather not appear or, conversely, neglected or censored out against their will [2, 46, 58, 62].

Harms can be:

1. distributional, when related to a resource of interest, such as education opportunities, jobs, access to credit, possibility of parole [86], exposure [75], or, more generally, attention of consumers using a SE over time [4], from which the above-mentioned resources depend when decision-making is not fully automated;
2. representational, likely to take place when individuals and groups are unable to self-determine their image, which may end up being stereotyped, inadequate or offensive [1, 58, 62]. In the context of SEs, representational harms typically refer to information subjects, and our work is no exception: GSR by SEs is firstly a type of problematic representation of women and men who happen to be information subjects in search and browsing experiences of SE users.

A subfield of research, often referred to as *fair ranking*, addresses distributional harms for providers [4, 84] and consumers [9]. This work is typically aimed at minimizing disparities in the outcomes of similar individuals (*individual fairness*) or groups determined by a protected attribute such as gender, ethnicity, religion (*group fairness*). These approaches perform re-ranking of results which have been retrieved and ranked by a supposedly biased algorithm. An alternative paradigm aims at directly modifying the retrieval

⁴<https://fair-trec.github.io/>

⁵<http://bias.disim.univaq.it/>

algorithm. Gerritse [31]⁶ studies the impact of debiasing WEs [6] in algorithms of query reformulation based on Word2Vec embeddings. This work is the closest to our evaluation of the effects of debiasing in Section 4.2.3, where we perform a complementary analysis on different IR algorithms which are purely based on WEs.

Metrics and approaches from fair ranking can also be employed to measure and favor a diversified topical coverage [28], where political leaning or sentiment take on the role of a protected attribute which should have reasonable diversification across search results. This flavor of fairness overlaps with diversity and novelty research from the IR community [8, 12, 85].

In some areas, such as political search in social media, it is interesting to evaluate how diversity and bias in search results can be influenced by (implicit) bias in queries. For instance, Kulshrestha et al. [47] find that Twitter’s response to queries about US political candidates tends to give better ranking to tweets from sources with the same political leaning as the candidate. Although different in methods and objective, our work is conceptually similar as we are interested in evaluating how a construct measured on queries (stereotypical genderedness) relates to the same construct measured on search results.

It should be noted that the taxonomy we presented is far from complete. A thorough categorization of ongoing efforts to audit and improve the fairness and trustworthiness of SEs would be as complex as the underlying socio-technical systems. Further considerations may include a temporal dimension and a spectrum to quantify division within a community, as in the case of research on echo chambers and filter bubbles [26]. User interfaces also play a key role: responsible augmentation of search results may be important to convey information about fact checking [88] and controversy [90]; query auto-completion can lead to problematic results [58], while panels which summarize results for users within the SE may reduce click-through rates for the websites from which information is extracted.

2.4. Construct validity and reliability

Construct validity, in its modern connotation, is a unified view on the desired properties for a measure aimed at quantifying a given construct that enables an overall judgement about adequacy and appropriateness based on empirical evidence and theoretical rationales [51]. Embedded in this definition is a clear distinction between, on the one hand, the unobservable theoretical attribute we are trying to evaluate (the *construct*, e.g. “teacher quality”), with its context and underlying theme and, on the other, the way the construct becomes operational through a measurement model (the *operationalization*).

We follow Jacobs and Wallach [42], who describe seven components of construct validity, which we summarize below:

1. *Face validity*. How plausible does the measurement model look compared to the construct? Answers to this question are highly subjective and little more than a preliminary step.
2. *Content validity*. Is there a coherent understanding of the theoretical construct? Is the selected operationalization in accordance with it?
3. *Convergent validity*. Does our measurement agree with other measurements of the same construct?
4. *Discriminant validity*. What else is the measurement capturing? Are there other constructs which are justifiably or unexpectedly correlated with the proposed measurement?
5. *Predictive validity*. Are any other properties likely to be influenced by our construct? Is our operationalization of the construct related to those properties as expected?
6. *Hypothesis validity*. Are the construct and its operationalization meaningful and useful, so that they can be used to test hypotheses and raise new questions?
7. *Consequential validity*. Should our measure be used? In which context can it be employed and what would the be consequences?

⁶We refer to the extension of this work discussed in ECIR 2020 workshop on Algorithmic Bias in Search and Recommendation (<http://bias.disim.univaq.it/>) whose proceedings are currently unavailable.

Section 3.1 describes in detail Gender Stereotype Reinforcement (GSR) as a construct, referring to supporting literature from social psychology, which deals with the common understanding of GSR and its *content validity* as a construct. The reliability of our operationalization with reference to the construct is addressed through discussion (Section 3.2) and experiments (Section 4.2.5). Considering key properties of GSR, Section 3.5 discusses its *discriminant validity*, tied to domain-specificity of language, along with its *convergent validity* in a wider context of fairness metrics. *Consequential validity* and *hypothesis validity* are linked with current limitations and future work, discussed in Section 5. In the absence of a user study *predictive validity* cannot be properly discussed. Within the context of gender stereotypes in SE, the only user study we are aware of centers on image retrieval [46], while our proposed measure deals with textual data. Due to its subjective nature, we do not specifically address *face validity*.

We also discuss GSR *reliability*, a more familiar concept to computer scientists. It depends on stability of measured quantity, precision of measurement tool, and process noise, to determine how robust repeatable and reliable a measure is; Section 4.2.4 is devoted to this aspect.

3. Proposed approach

We articulate our approach, untangling the definition of a *construct*, i.e. the phenomenon we want to study, from its subsequent *operationalization*, which details how the phenomenon can be measured from observed data [42, 52].

3.1. Construct

Our aim is to quantify to what extent a SE can reinforce gender stereotypes in users. We call this construct *Gender Stereotype Reinforcement* (GSR), resorting to supporting concepts from psychology literature before giving a formal definition. This incremental process is important to establish the *content validity* of GSR as a construct.

Definition 3.1. Stereotype

Stereotypes are beliefs about groups of individuals with a common trait, widely held by a population of interest. Their appearance is likely influenced by the strength of an observational link, i.e. how often one position along a dimension (such as gender) co-occurs with another (such as warmth) [48].

Stereotypical associations picked up by individuals can be attributed to culture and socialization [40]. Bayesian principles are thought to be at play in the acquisition of culture, which is often screened and mediated by search technology, whose trustworthiness is generally taken for granted [37]. In other words, our cognition is receptive to repeated co-occurrence of topics and entities. It may therefore end up forming links between them, also thanks to the media and technology we interact with on a daily basis.

Definition 3.2. Gender Stereotype

A gender stereotype is a generalised view or preconception about attributes or characteristics, or the roles that are or ought to be possessed by, or performed by, women and men. ⁷

Stereotypes about gender have been studied in a variety of contexts, including school [16], workplace [5], parenthood [15] and search for romantic partners [65], with respect to several aspects such as depiction, perception (of self and others) and outcomes. Common themes have been identified through decades of scholarship, including *agency* and propensity to science, *communion* and importance of appearance [22].

As a well-researched example, historical meta-analysis over seven decades confirms agency and communion as consistently and increasingly salient in U.S. opinion polls about gender differences [20]. Agency, perceived as predominantly male, refers to drive for achievement, while communion is related to caring for others and is increasingly associated to women.

⁷<https://www.ohchr.org/en/issues/women/wrgs/pages/genderstereotypes.aspx>

Definition 3.3. Direct gender stereotype

Association of a stereotypically gendered concept with people of the respective gender.

This applies to any sentence where preconceptions about one gender are directly associated to a member of that gender, mentioned through either a noun (**man**), adjective (**his**), pronoun (**he**) or name (**John**).

Example: *She is affectionate.*

Definition 3.4. Indirect gender stereotype

The link of a stereotypically gendered concept with another stereotypically gendered concept, commonly associated to the same gender.

This definition is based on a view of culture, social constructs and stereotypes as networks of concepts [32, 66] and implicit associations [3, 23, 35]. Co-occurrence of stereotypical characteristics and traits, commonly associated to one gender, may reinforce a link in a network of stereotypes about women and men. To exemplify, we argue that beliefs about stereotypically female (male) jobs are likely to fall on women (men). Research from social cognition and political science highlights that networks of stereotypes associated with protected attributes, such as gender and ethnicity, can play a role in a person’s perception, without them being aware of it [3]. This may happen to a person, even if they sincerely dislike said stereotype [23].

Example: *The nurse is affectionate*

Characterization of the GSR construct.

Given the above terminology, we characterize GSR in the context of IR as the SE’s tendency to respond to stereotypically gendered queries with documents containing stereotypical language with the same polarity. We defer a thorough definition, complete with mathematical formalization, to Definition 3.8.

In societal systems, GSR is measured by the agreement of human constituents with gender stereotype descriptors [69, 80]. In operationalizing this construct, we aim to quantify the impact of SEs on the perception of gender: more specifically, its alignment to existing direct and indirect stereotypes encoded in culture and language. Search results may end up reinforcing gender stereotypes if, when responding to potentially stereotypical queries, their language is skewed along gendered lines with matching polarity.

Intuitively, the influence that people around us may exert can be regarded as the societal counterpart of documents and their language in the context of SEs. An example is a SE which, responding to a query about nursing, displays documents with a strong representation of women (*direct* stereotype), or emphasis on attributes related to communion (*indirect* stereotype).

3.2. Operationalization

After defining our construct, we show how it can be made operational. This entails illustrating our assumptions and their interplay with the building blocks of our measurement model [42]. We begin by defining the basic concepts in the context of search.

Definition 3.5. Ranked list

The response of a SE to a query, i.e. a ranked list \mathcal{L} of documents, decreasingly ordered by (estimated) relevance with respect to a given user query.

Definition 3.6. Search history

A set of (query, ranked list) pairs representing the interactions of one or more users with a SE.

Stereotype formation may be conceptualized as an acquisition of culture and associations of a particular kind, taking place through *repeated* interaction. The response to a single query, though anecdotally interesting, is less informative than a set of responses to different queries. Hence, we refer to a *search history*, a somewhat overloaded expression, which potentially encompasses every past user interaction with a SE, including the pages they visited along with very detailed logs of click behavior, browsing and permanence.

Our usage of the expression is different in two ways. (1) It applies to any subset of user interactions with a SE, including for instance only recent ones. We do not require a complete list of queries issued and results shown. (2) The level of granularity and depth of logging entailed by our definition is minimal. This work is

aimed at auditing and modeling SEs rather than users. For this reason we do not require click logs, which are user-dependent and thus accidental with respect to our analysis. More in general, Definition 3.6 adapts to data coming from multiple users in a bundled and anonymized fashion, as well as data collected and curated by a practitioner. These differences are important to correctly assess the applicability, practicality and ethics of our operationalization.

To summarize the following sections, we assume that a strong correlation between *genderedness* of queries and of ranked document list in a search history reinforces gender stereotype. In the following, we gradually introduce related quantities; the adopted notation is summarized in Table 1.

\mathcal{Q}	set of queries in search history
\mathcal{D}	set of available documents
\mathcal{L}	ranked list of documents
$N = \mathcal{Q} $	number of queries in search history
w	a word
$q \in \mathcal{Q}$	a specific query
$d \in \mathcal{D}$	a specific document
$g(w)$	genderedness of word w
$g(q)$	genderedness of query q
$g_q(d)$	genderedness of document d retrieved for q
$g_q(\mathcal{L})$	genderedness of ranked list \mathcal{L} retrieved for q
r_k	rank of document d_k in list \mathcal{L}
μ_q	average genderedness of queries from \mathcal{Q}
$\sigma_{g(q)}^2$	variance in genderedness of queries
$\mu_{q,\mathcal{L}}$	average genderedness of ranked lists of documents
$m_s(\mathcal{Q}, \mathcal{D})$	GSR for system s on collection $(\mathcal{Q}, \mathcal{D})$.

Table 1: Notation for proposed measure.

3.2.1. Measuring gender stereotype

Stereotypes about gender are plentiful and pervasive, likely due to the fact that the underlying categories (especially the classical female-male dichotomy) are available to our cognition from an early age on a daily basis. Preferential association of a concept or topic to men or women is measured by surveying a population of individuals. The study of gender-based associations thus depends on resources, time available and research agendas.

Increasing evidence from the field of NLP shows that, among the powerful results and interesting properties of WEs, their geometry captures well-known stereotypes related to gender [6, 7, 19, 29, 64, 70]. Techniques have been proposed to isolate a word’s genderedness along a single direction [6]. Based on this approach, each word is associated to a “gender score” consisting of a signed scalar value. In a convention employed hereafter, a strongly positive (negative) score will be a proxy for a strong association to female (male) gender. The upper part of Figure 1 depicts, as a simplified example, the projection of the word **beauty**, which is strongly positive and thus associated to female gender.⁸

To validate genderedness, encoded by Equation 1, as a score of perceived masculinity/femininity, we test it against known gender stereotypes. Two commonly studied constructs in psychology literature are *agency* and *communion* [20, 39], alternatively dichotomized as *competence* and *warmth* [25]. Agency, stereotypically associated to men, is related to ability and drive to pursue one’s goals, while displaying leadership and assertiveness. Communion, prevalent in female stereotypes, relates to a person’s orientation towards others and their well-being, suggesting propensity for caring, nurturing, compassion and emotion.

⁸For obvious reasons, a figure can only represent 2 out of the 300 dimensions in which the **w2v** embedding is encoded.

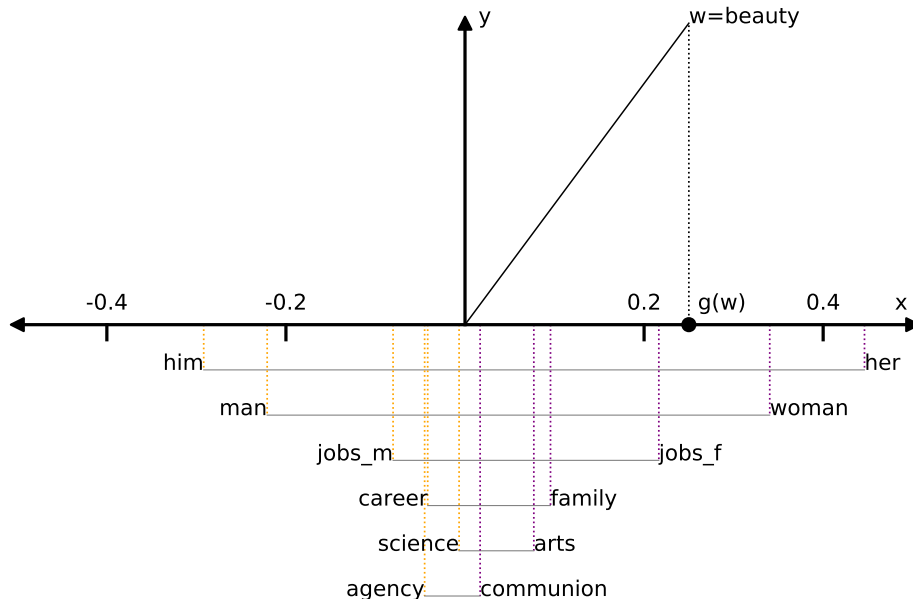


Figure 1: Computation of genderedness for $w = \text{beauty}$. Gender direction is on x axis [6], while y axis represents subspace orthogonal to gender. w displays a significant component along the gender direction, hence we consider it stereotypically female in the embedding space and the underlying text corpus. See, in Appendix A, Table A.4 for *agency vs communion* ($p = 2.2 \times 10^{-2}$), Table A.5 for *science vs arts* ($p = 1.8 \times 10^{-3}$), Table A.6 for *career vs family* ($p = 5.5 \times 10^{-4}$), Table A.7 for *jobs_m vs jobs_f* ($p = 0$). P-values are computed with four one-tailed permutation tests on the genderedness of the words in each table.

Attitude towards mathematics and sciences have been measured implicitly [16, 60] and explicitly [16]. Studies provide evidence of cognitive link between math and male gender from an early age. This association is often studied in opposition to arts (and language) which are found to be predominantly associated with female gender [59, 60].

Career orientation, in opposition to family, is another dimension related to gender [59]. Career can also be broken down into sector. Some professions have a very high male representation, while other work is overwhelmingly carried out by women [10].

In considering research on gender stereotypes four opposing associations emerged, which are described above. We compute their genderedness as follows: for *agency vs communion* we summarize the genderedness of either construct with the average genderedness of adjectives in Table A.4, taken from [20]. With the same averaging procedure, we follow [59] for terms related to *science vs arts* (Table A.5) and [60] for *career vs family* (Table A.6). Finally, we sample the 20 most gendered single-word jobs from [10], shown in Table A.7, and perform the same computation, dubbing this comparison *jobs_m vs jobs_f*.

Results are summarized in the lower part of Figure 1, where we also report the projections of **woman**, **man**, **her**, **his** for comparison. All four stereotypes are confirmed, with male clusters’ projections (orange) falling to the left of their female counterparts (purple). According to one-tailed permutation tests, the dichotomy *agency vs communion* is the least gendered, significant at $p = 2.2 \times 10^{-2}$. Interestingly, the strongest association with gender is *jobs_m vs jobs_f* ($p = 0$), stemming from census data and representing occupations with extreme skew in gender distribution.

We conclude that projection along the gender subspace (although potentially noisy for single terms) is, on average, a suitable proxy for stereotypical association with gender.

3.2.2. Modeling stereotype in query-document pairs

Semantic memory is a specific aspect of human memory which holds general knowledge about concepts. It is regarded as a widely distributed neural network [66]. Associative network structures, often referred

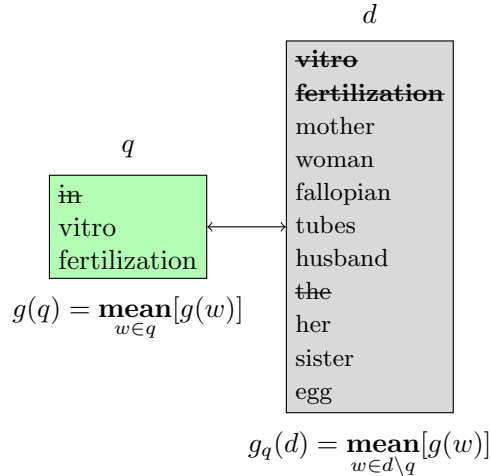


Figure 2: Concepts from query associated with concepts from document along gender dimension. Before computing the average genderedness of query and documents, stop words are removed (struck through font), and query terms which explicitly appear in a document are neglected (bold and struck through).

to as schemas, are commonly used in neuroscience as models representing complex constructs that guide behavior [32]. This suggests that any acquisition of knowledge and culture resides in part in the formation of rich networks of concepts. The acquisition and articulation of stereotypes are not conceptually different: a recent line of work employs network analysis to study stereotypical associations as clusters and subclusters of concepts [73].

We are interested in modeling the potential association of concepts, with a tendency to cluster along a gendered dimension. Search technologies play an important role in helping users to build links between concepts. When issuing a query, SE users are likely receptive to the formation of new links between concepts from their query and information found in ranked lists [46]. If a document d retrieved for a query q (e.g. **nurse**) contains terms mostly aligned with the genderedness of q (e.g. **care**, **woman**, **Mary**) it may end up reinforcing gender stereotype through an association of such concepts.

In order to assess the stereotypical gender agreement between q and d , we compute their average genderedness $g(q)$ and $g_q(d)$ as schematized in Figure 2. Both queries and documents are represented as bag-of-words (stop words are removed). $g(q)$ is subsequently computed as the average genderedness of remaining query terms. For $g_q(d)$, query terms are removed (bold in Figure 2) before performing the same averaging procedure. Our goal is to model the alignment of query-document concepts along stereotypically gendered lines. When computing $g_q(d)$, we therefore neglect document terms which also appear in the query, to remove the spurious bias due to redundant self-linking. For this reason, $g_q(d)$ depends on the query, as illustrated by subscript q .

3.2.3. Computing the genderedness of a ranked list

As generally known, users seldom dabble into result pages beyond the first one, and the likelihood of view decreases with document rank [24, 43, 44]. Performance metrics in IR have taken this aspect into account, assigning more importance to top-ranked rather than low-ranked documents [43, 56]. A widely-adopted evaluation measure, based on this user model, is the Discounted Cumulative Gain [43], which weighs documents according to a coefficient that decreases with rank in a logarithmic fashion. This weighing scheme is applied to measure the effectiveness of a ranked list based on the relevance and position of the documents within it. Our approach is identical, except for our focus on genderedness rather than relevance.

Figure 3 shows a ranked list \mathcal{L} of documents d_i , retrieved for a query q . A vector of weights \bar{w} is computed with a rank-based logarithmic discount and normalized. The genderedness of a ranked list $g_q(\mathcal{L})$ is calculated as the weighted average of the genderedness of documents in \mathcal{L} with weight vector \bar{w} .

	q	rank	weight	gend.
{	d_7	1	$\frac{1}{\log_2(1+1)}$	$g_q(d_7)$
	d_5	2	$\frac{1}{\log_2(2+1)}$	$g_q(d_5)$

	d_2	k	$\frac{1}{\log_2(k+1)}$	$g_q(d_2)$

$$g_q(\mathcal{L}) = \bar{w} \cdot \bar{g}_q \cdot \frac{1}{W}$$

Figure 3: Genderedness of ranked list is computed as a weighted average of the genderedness of each document retrieved, with weight computed according to a rank-based logarithmic discount. Note that in the calculation of $g_q(\mathcal{L})$, without any loss of generality, we opt for a base-2 logarithm. W is a normalizing constant, i.e. the sum of elements in \bar{w}

Definition 3.7. Genderedness of a ranked list

Let $\bar{w} = [w_1, \dots, w_K]$ be a vector of weights such that K is the length of the ranked list, $W = \sum_{k=1}^K w_k$ and $w_k = 1/\log_2(r_k + 1), k \in [1, K]$. Then, the genderedness of \mathcal{L} is defined as

$$g_q(\mathcal{L}) = \frac{1}{W} \sum_{k=1}^K w_k \cdot g_q(d_k), \tag{2}$$

with r_k being the rank of document d_k , and $g_q(d_k)$ its genderedness.

As a toy example, which will be expanded and further discussed in Section 3.4, suppose we have the following setting with a single-term query and two retrieved documents:

$q = \text{electrician}$

$d_1 = \text{The man is an electrician.}$

$d_2 = \text{The woman is an electrician.}$

$\mathcal{L} = [d_1, d_2]$.

Then, according to Definition 3.7, the genderedness of ranked list \mathcal{L} is computed as follows:

$$g_q(\mathcal{L}) = \frac{1}{W} \left[\frac{1}{\log(2)} g_q(d_1) + \frac{1}{\log(3)} g_q(d_2) \right] = -3.8 \times 10^{-3},$$

where d_2 is less important in this weighted average, being the last document in \mathcal{L} . Its genderedness $g_q(d_2)$ is thus discounted accordingly, and the negative value of $g_q(d_1)$, albeit smaller in modulo than that of $g_q(d_2)$, ends up prevailing.

3.2.4. From ranked list to search history

Multiple search results constitute a search history which may reinforce gender stereotypes. If the language of documents in ranked lists (more specifically their genderedness) consistently agrees with that of user’s queries, it is reasonable to assume that the search history supports concept clustering along a gender-stereotypical dimension.

More precisely, given a set of queries \mathcal{Q} and a set of ranked lists (one per query) returned by a system s , we compute a linear fit between query genderedness $g(q)$ and ranked list genderedness $g_q(\mathcal{L})$, considering it a summary of the GSR carried out by s on \mathcal{Q} .

3.3. Measurement

Below is a summary of the steps to measure GSR:

- Genderedness of a word w is measured as its projection along the gender direction (Equation 1).
- Genderedness $g(q)$ of a query is defined as average genderedness of the terms in the query, after removing stop words.
- Genderedness $g_q(d)$ of a document d relevant for a query q , is computed as average genderedness of its terms, neglecting stop words and query terms.
- Genderedness $g_q(\mathcal{L})$ of a ranked list \mathcal{L} is computed as a weighted average of documents’ genderedness. An inverse logarithmic function of rank determines the weight of each document.
- Given a set of queries \mathcal{Q} , a set of ranked lists (one per query) retrieved by a system s from a collection \mathcal{D} , and the linear fit between query genderedness and ranked list genderedness, the GSR is the slope $m_s(\mathcal{Q}, \mathcal{D})$ of the linear fit.

Hence, GSR is formally defined as follows:

Definition 3.8. Gender Stereotype Reinforcement (GSR)

Let \mathcal{Q} with cardinality N be a set of queries, $g(q_i)$, $i \in [1, N]$ the genderedness of $q_i \in \mathcal{Q}$; let \mathcal{D} be a corpus of documents and \mathcal{L}_i the ranked list provided by a system s for the query q_i over \mathcal{D} . Then, GSR of s on collection $(\mathcal{Q}, \mathcal{D})$ is defined as:

$$m_s(\mathcal{Q}, \mathcal{D}) = \frac{1}{\sigma_{g(q)}^2} \frac{1}{N} \sum_{i=1}^N (g(q_i) - \mu_q)(g_{q_i}(\mathcal{L}_i) - \mu_{q,\mathcal{L}}). \quad (3)$$

GSR weighs the extent to which a SE responds to stereotypically gendered queries with documents containing stereotypical language with the same polarity. In the above equation, $\mu_q, \sigma_{g(q)}^2$ are query genderedness mean and variance and $\mu_{q,\mathcal{L}}$ is the average genderedness of ranked lists of documents.

We chose slope instead of correlation since the latter quantifies the predictability of the genderedness of a ranked list, given that of the query. The former also captures the extent to which highly “female” and “male” queries are answered with completely different language along the gender dimension.

3.4. Toy example

We build a toy document collection to show how GSR captures gender stereotypes.

Setup: from [10] we sample the single-word jobs with the widest gender gap (Table A.7).

- \mathcal{Q} is the set of (single-word) queries of occupations considered hereafter.
 High female representation: hygienist, secretary, hairdresser, dietician, paralegal, receptionist, phlebotomist, maid, nurse, typist.
 High male representation: stonemason, roofer, electrician, plumber, carpenter, firefighter, millwright, welder, machinist, driver.
- \mathcal{D} is the set of all documents deriving from permutations of “The ⟨person⟩ is a ⟨job⟩”, with ⟨person⟩ $\in \{\text{man, woman}\}$ and ⟨job⟩ from all occupation entries in Table A.7.
- \mathcal{N} is a neutral retrieval system returning, for each query, both documents (female and male) in which the query term appears (Figure 4, center, $m_{\mathcal{N}}(\mathcal{Q}, \mathcal{D}) = 0$).




Figure 4: GSR on toy dataset $(\mathcal{Q}, \mathcal{D})$ for different retrieval systems: stereotypical (S), neutral (N), counter-stereotypical (CS). GSR is the slope of linear fit, taking values $m_S(\mathcal{Q}, \mathcal{D}) = 1.61$, $m_N(\mathcal{Q}, \mathcal{D}) = 0$, $m_{CS}(\mathcal{Q}, \mathcal{D}) = -1.61$.

- **S** is a retrieval system returning, for each query, only the stereotypical document (Figure 4, left, $m_S(\mathcal{Q}, \mathcal{D}) = 1.61$). For instance, given a query about a job with a high male representation, **S** would only provide documents mentioning men.
- **CS** is a retrieval system returning, for each query, only the counter-stereotypical document (Figure 4, right, $m_{CS}(\mathcal{Q}, \mathcal{D}) = -1.61$). Contrary to **S**, given a query about a job with a high male representation, **CS** would only provide documents mentioning women.

Discussion: Figure 4 shows the behavior of three synthetic search engines measured by GSR. Each dot represents a query q , with its genderedness $g(q)$ on the x axis and the genderedness of documents retrieved for q ($g_q(\mathcal{L})$) on the y axis. GSR is the slope of the linear fit. **CS** and **S** are quite extreme, as they only return documents that challenge gender gap in occupations or fully reinforce it, whereas **N** is neutral. GSR successfully captures this aspect with zero slope for **N** and significantly non-zero slopes for **S** and **CS**, equal in magnitude and opposed in sign. The magnitude of GSR for **S** and **CS** is very large compared, for instance, against GSR values of real IR algorithms on a news collection, in the order of magnitude of 10^{-2} (Table 2). This depends on (1) the collection $(\mathcal{Q}, \mathcal{D})$ itself, especially conceived for *direct gender stereotype*, (2) the systems **S** and **CS** which are extreme as they respond to job-related queries with documents mentioning women or men in accordance with, or in opposition to, stereotypes related to gender gaps in the occupations mentioned within a query. Overall, this experiment shows that GSR is suited to capture direct gender stereotypes. This is further confirmed by experiments on a shared IR collection (Section 4.2.5).

3.5. Key properties

The toy example presented above defines a controlled setting where we can test the *convergent validity* of our construct with metrics of algorithmic fairness and diversity in IR [12, 28]. The IR task can be framed as a binary classification problem – i.e. classifying documents as relevant or non-relevant – with a binary protected attribute encoding whether a document is stereotypical or not. A document is deemed stereotypical for a query if it displays genderedness of same polarity. We assume a search history (in this context: solution to classification problem) to be reasonable if, for every query q_i , the documents included in the ranked list \mathcal{L}_i contain the query term.

Hence, in our controlled setting, for each query, such as **driver**, a maximum of two documents can be retrieved, namely **The woman is a driver** (counter-stereotypical) and **The man is a driver** (stereotypical), i.e. the ones which contain the term **driver** from the complete permutation described above. This is a sensible assumption and makes enumeration feasible.

We enumerate every reasonable solution, and for each compute GSR along with the percentage of stereotypical documents among the retrieved ones, equivalent to *statistical parity fairness* from Gao and Shah




Figure 5: Agreement between GSR and percentage of stereotypical documents among retrieved ones, equivalent to *statistical parity fairness* [28]. Pearson’s $r = 0.92$, $p < 1e-40$.

[28], already employed in the context of fair ranking to enforce equal exposure of SE users to different topics. Results from Figure 5 show a strong agreement between these quantities, which stems from the very definition of slope coefficient:

$$m_{x,y} = \frac{1}{\sigma_x^2 N} \sum_{i=1}^N (x - \mu_x)(y - \mu_y) = \frac{1}{\sigma_x^2 N} \left[\sum_S (x - \mu_x)(y - \mu_y) + \sum_{CS} (x - \mu_x)(y - \mu_y) \right], \quad (4)$$

where we explicitly partitioned retrieved documents into stereotypical (S) and counter-stereotypical (CS). This partition is equivalent to topical group assignment for which statistical parity fairness enforces equal user exposure [28]. Specializing Equation 4 for GSR, we get

$$m_s(\mathcal{Q}, \mathcal{D}) = \frac{1}{\sigma_g^2(q)NW} \left[\sum_{(q_i, d_k) \in S} \frac{(g(q_i) - \mu_q)(g_{q_i}(d_k) - \mu_{q,\mathcal{L}})}{\log_2(r_k + 1)} + \sum_{(q_i, d_k) \in CS} \frac{(g(q_i) - \mu_q)(g_{q_i}(d_k) - \mu_{q,\mathcal{L}})}{\log_2(r_k + 1)} \right], \quad (5)$$

where we have used Equations 2 and 3. Documents d_k which are stereotypical for a query q_i (first summation in Equation 5) bring a positive contribution to the slope coefficient $m_s(\mathcal{Q}, \mathcal{D})$, while counter-stereotypical documents bring a negative one. Equal exposure (50% stereotypical documents) does not entail neutrality (GSR=0) and vice versa, however the two measurements are clearly correlated (Pearson’s $r = 0.92$, significant at $p < 1e-40$). Indeed GSR is a measure of *weighted* statistical parity, between stereotypical and counter-stereotypical documents, with weight proportional to genderedness of query times genderedness of document. Although not central in this toy example, document position in ranked list r_k is a further weighing factor through logarithmic discount.

A discussion focused on *discriminant validity* (Section 2.4) of GSR is due. In any sensible text corpus, words from the same domain (e.g. medicine) and *a fortiori* subdomain (e.g. gynecology) are likely to co-occur and their word vectors will end up close to one another, duly capturing their semantic proximity. At the same time, in order to satisfy an information need (e.g. query “in vitro fertilization”), it will be necessary to employ the specific language of relevant fields to which the query pertains (such as gynecology). For

this reason, some query-document agreement should be expected in language and, more specifically in genderedness. Thus, in non-trivial settings, any reasonable SE is expected to have positive GSR. Figure 2 is a real example of this aspect, depicting a query and a relevant document (both represented as bags-of-words, the latter subsampled for brevity) taken from Robust04 collection [38]. The document surely contains domain-specific language, however it is also centered around female entities, echoing old gender roles in the framing of involuntary childlessness [57].

In other words, if a kernel of truth is present in some stereotypes [68], positive GSR captures a kernel of relevance, and its value is fundamentally influenced by documents available to respond to a query. Bearing such aspect in mind, it is fundamental to provide a baseline GSR for relevant documents. System N is an example of such baseline in the trivial setting of Section 3.4. Upward deviation of GSR from this baseline (as computed differentially or through a ratio) is regarded as the SE’s contribution towards reinforcement of gender stereotype. The baseline GSR for relevant documents captures a mixture of *historical bias* [78] and inevitable domain-specificity of language. For this reason, when measuring GSR, it is important to have a list of relevant documents as a baseline, which is to be externally validated and reasonably regarded as a ground truth.

4. Evaluation

Methodology

GSR captures both problematic query-document associations by a SE, and domain-specificity of language embedded in a collection of documents and queries. Hence, a perfect SE, that retrieves all and only the relevant documents for each query, is expected to have positive GSR. This stems from the fact that a document is more likely to be relevant for a query if it contains specific language from the query domain - and words from the same domain tend to cluster together in the WE space, and consequently in the gender subspace.

We can use GSR of the perfect SE as a baseline against which to compare real SEs. In other words, a SE can be said to counter gender stereotypes, even if it displays positive GSR, as long as its GSR is smaller than that of the perfect SE. Conversely, a SE which reinforces gender stereotypes will have a larger GSR. For this reason, we perform tests on shared test collections based on the Cranfield paradigm [13], for which relevance judgements have been provided by qualified human assessors.

Our source code and data are publicly available for reproducibility purposes.⁹

4.1. Synthetic dataset

In this section we introduce a synthetic example, similar in setting to the toy example presented in Section 3.4, exemplifying *indirect* gender stereotype and GSR’s ability to capture it.

Hypothesis: GSR can measure *indirect* gender stereotypes (Definition 3.4) stemming from clustering of concepts and language segregation along the gender direction, such as the association of stereotypically gendered occupations and traits. For example, GSR should highlight situations where SEs respond to queries about jobs with strong male representation with documents that focus on traits related to agency.

Setup: We build a synthetic dataset of queries \mathcal{Q} and documents \mathcal{D} where a SE might promote gender stereotypes, or counter them. We simulate three SEs, designed to be stereotypical, counter-stereotypical or neutral and check whether GSR captures this aspect. Below we summarize the dataset (\mathcal{Q}, \mathcal{D}) and the simulated SEs working on the dataset.

- \mathcal{Q} is the set of (single-word) queries consisting of occupations with large gap in gender representation from table A.7.

⁹<https://github.com/alessandro-fabris/gsr>




Figure 6: GSR on synthetic dataset for different retrieval systems: stereotypical (S), neutral (N), counter-stereotypical (CS). GSR is the slope of linear fit, taking values $m_S(Q, \mathcal{D}) = 0.16$, $m_N(Q, \mathcal{D}) = 0$, $m_{CS}(Q, \mathcal{D}) = -0.16$.

- \mathcal{D} is the set of all documents deriving from permutations of “The ⟨job⟩ is ⟨adjective⟩”, with ⟨job⟩ from Table A.7 and ⟨adjective⟩ from Table A.4. The adjectives considered are commonly used to assess gender stereotypes held by a population, and are descriptive of *communion* (commonly considered a female trait) and *agency* (often associated with males).
- N is a neutral retrieval system (search engine) returning, for each query, all documents in which the query term appears.
- S is a retrieval system returning only the stereotypical document in which the query term appears. Predominantly female (male) jobs are therefore associated to communion (agency) adjectives - e.g. *The plumber is hardworking*.
- CS is a retrieval system returning only the counter-stereotypical document. Predominantly female (male) jobs are associated to agency (communion) adjectives.

Discussion: Results are summarized in Figure 6. System S reinforces indirect stereotypes, since it links occupation and personality roles along gender-stereotypical lines, strengthening gender clusters. The proposed measure successfully captures this aspect, along with the neutral nature of N and the counter-stereotypical nature of CS. GSR can detect indirect gender stereotypes captured by underlying WEs.

4.2. Real document collection

We demonstrate our approach on a widely-used TREC evaluation collection: TREC 2004 Robust Track [38], dubbed hereafter Robust04. This collection consists of about 528K news documents and 249 queries. The domain, news, is one where the relevance of web SEs is well established in mediating user access [41].

4.2.1. Preliminary qualitative analysis

Objective: We evaluate qualitatively whether Robust04 is an interesting collection for GSR analysis, i.e. whether the most gendered queries according to Word2Vec ($w2v$), contributing the most to GSR in Equation 3, contain recognizable gender stereotypes. We have selected this candidate collection for three main reasons: (1) the large number of queries ($N = 249$) for which relevance judgements from human annotators are available; (2) the domain, news, where the importance of SEs in mediating user access is wide and well-established [41], and (3) its relevance within the IR community.

Setup: We restrict our analysis to topic titles, inspecting the 10 most “female” and “male” queries according to $w2v$. These are the queries q_i whose title has the highest and lowest *genderedness* $g(q_i)$, computed as the average projection of query terms onto the gender direction of $w2v$ [6]. The most gendered queries are depicted in Figures 7a, 7b.



Figure 7: Most gendered queries from Robust04 under $w2v$. The text is printed with color-coded gradient where strongly male words are orange, strongly female words are purple, neutral words are white. Terms’ projection along gender direction can be read below each word. Stop words are removed from queries.

Discussion: Among the most “female” queries, few are *intrinsically gendered*, such as *women in parliaments* (topic 321) and *women clergy* (topic 445). Some more queries are *biologically gendered* (such as *postmenopausal estrogen britain* and *osteoporosis* – topics 356 and 403 respectively), describing topics biologically associated to women. The remaining queries can be described as *culturally gendered*. Some are associated to disorders with apparently higher incidence on the female population (*agoraphobia* and *anorexia nervosa bulimia* – topics 348 and 369). The final three (*quilts income*, *child labor*, *in vitro fertilization* – topics 418, 440, 368) seem to capture unnecessary or even harmful stereotypes related to communion [20] and gender roles [57]. Topic 440 (*child labor*) highlights a limit of GSR and the underlying word representations. In the presence of a polysemous word, its embedding encodes a mixed representation of the different uses of said word. In this example, the embedding for *labor* has been influenced by the meaning related to giving birth, while the query has a different intent. Contextualized approaches may be useful to mitigate this issue.

All “male” queries seem to be *culturally gendered*, containing terms loosely related to agency (such as *dismantling*, *heroic*, *evasion*), and occupation (*retirement*, *term*), contrary to communion (*dangerous*, *arsenal*, *crime*, *traps*), and more frequently associated to men (*cigar*, *rap*).

Overall, most of these queries associate gender (as encoded by $w2v$) with undesirable and harmful concepts. For this reason, Robust04 is a reasonable collection to study GSR.

4.2.2. GSR on Robust04

Hypothesis: SEs based on WEs, such as $w2v$ and FastText (ftt - [53]), have higher GSR than purely lexical ones, i.e. they are more prone to support gender stereotypes due to the problematic gender information embedded in word representations.

Setup: We evaluate GSR for three families of well-known IR systems, which could serve as a basis for a SE:

- **Lexical:** these algorithms are based on matching query terms to document terms, without any information about semantics. In this group we include three models inspired by different key paradigms: the widely-used probabilistic model BM25 [71], a popular language model [87] (i.e. Language Modelling with Bayesian smoothing and a Dirichlet Prior) called QLM, and the classic vector space model *tf-idf* [72].
- **Semantic:** IR systems based on WEs have been proposed [82], with the idea of exploiting the latent

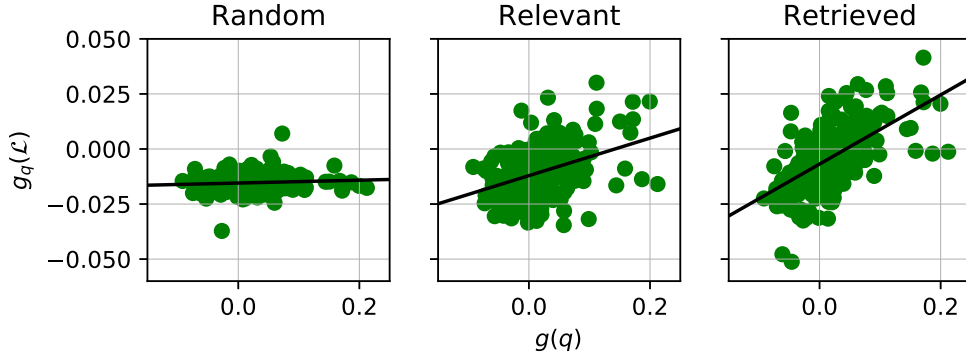


Figure 8: GSR for relevant, random and retrieved docs. The x axis represents genderedness of queries $g(q)$, while y axis represents the genderedness of ranked document list $g_q(\mathcal{L})$. GSR is the slope of the linear fit of the scatter plot.

relationship between words encoded by the embeddings. We test `w2v_add`, `w2v_si` [82] and `ftt_add` (`w2v_add`'s counterpart based on `ftt` WEs) for this family.

- Neural: WEs can be fed as input to neural networks, which in turn learn to match the signals of user queries with that of relevant documents. We consider Deep Relevance Matching Model (DRMM - [36]), and Match Pyramid (MP - [63]). The embeddings used as input to these systems are `w2v` trained on Google News [53].

If a query has K relevant documents, according to the assessors' judgments, we compute GSR for each system on the top- K documents retrieved by it. This makes GSR of the perfect SE (retrieving all and only the K relevant documents) directly comparable to that of the systems at hand.

Results: As a preliminary illustration, Figure 8 depicts GSR for three systems.

- On the left, a search engine retrieving random documents.
- In the middle, a perfect search engine which retrieves all and only relevant documents, also ranking them perfectly according to relevance judgements.
- On the right, a search engine based on `w2v_add` [82].

As mentioned in Section 3.5, while discussing *discriminant validity*, positive GSR can be associated with relevance. This is shown by the positive GSR of the perfect SE in the middle pane, compared against the near-zero GSR of random retrieval. This is not surprising; it is due to a combination of language specificity and *historical bias* [78] potentially present in news coverage, as discussed in Section 3.5. For this reason, hereafter we will report for comparison the GSR of the perfect SE.

Figure 9 depicts GSR for the search results of eight different retrieval systems. Each panel contains a scatter plot of 249 different points (one for each query in Robust04), along with their linear fit (solid) and the linear fit of the perfect SE for comparison (dashed). Panels (4)-(6), depicting `w2v_add`, `w2v_si`, `ftt_add`, confirm that semantic SEs have higher GSR than lexical ones, namely `QLM`, `tf-idf` and `BM25`, depicted in panels (1)-(3). This fact is easily to explain: SEs based on gender-biased WEs inherit the bias and tend to reinforce it.

Interestingly, neural systems based on the same word representation (MP, DRMM), shown in panels (7) and (8) respectively, seem to dampen this effect, thanks to successful tuning of weights during training, which reduces the importance of the (biased) gender direction in `w2v`.

As expected, semantic models based on biased WEs are likely to reinforce gender stereotypes, even when based on an IDF-inspired weighting scheme (as in the case of `w2v_si`), aimed at assigning greater importance to terms that bear more information. On the other hand, lexical models have low GSR, comparable to that of the ideal SE.




Figure 9: GSR for different systems on Robust04 according to gender direction of $w2v$. The x axis represents the genderedness of queries $g(q)$, while the y axis represents the genderedness of ranked document lists $g_q(\mathcal{L})$. GSR is the slope of the linear fit through the scatter plot (solid). The dashed line is the linear fit of the perfect SE, reported for comparison.

4.2.3. Debiasing moderately reduces GSR

Objective: The gender direction along which we measure GSR can be removed from the embeddings, by means of orthogonal projection [6]. We evaluate the impact of this operation on performance and GSR of semantic and neural SEs based on WEs.

Setup: For each system which relies on WEs, we repeat the previous retrieval task with three different versions of $w2v$ embeddings. *Regular* embeddings are the original version trained on Google News. *Debiased* embeddings are obtained by eliminating the gender direction from neutral words while maintaining it for gendered word such as *woman* [6]. *Strong debiased* embeddings take this procedure a step further, eliminating the gender component from each word [70]. The same debiasing procedures are applied to FastText embeddings fed to `ftt_add`.

Discussion: GSR values are reported in Table 2, under the header **GSR ($w2v$)**. Debiasing is effective in reducing GSR for systems where it is particularly high, namely the semantic ones, purely based on WEs (`w2v_add`, `w2v_si`, `ftt_add`). However, even for these systems the reduction is quite weak, ranging between 10%-25%, as gender information leaks along different directions, orthogonal to the one that is eliminated through debiasing. This aspect has been previously studied [33], to conclude that “the gender-direction provides a way to measure the gender-association of a word but does not determine it”. Our results confirm that this is true for a measure based on the gender direction such as GSR. Furthermore, *strong debiasing* brings no major advantage compared to simple debiasing.

The impact on performance is very limited, as shown in Table 3, reporting Mean Average Precision (MAP), precision for the top-10 ranked documents (P@10) and Normalized Discounted Cumulative Gain for the top-100 ranked documents (nDCG@100). We focus on systems based on WEs, leaving aside lexical ones, since our interest is to evaluate the impact of debiasing on classical performance measures. Our results, which are in line with prior art [49], show that debiasing (both regular and strong) produce negligible changes to average performance.

How does such small impact on performance coexist with a significant impact on GSR (shown in Table 2)? Figure 10 answers this question, depicting the Kendall τ distance for ranked lists of documents (top-100) retrieved by `w2v_add` before and after debiasing (on the y axis), against the genderedness of the respective query in absolute value (on the x axis). From the plot an expected property of debiasing WEs emerges in

System	GSR (<i>w2v</i>)		GSR (<i>ftt</i>)	
	absolute	relative	absolute	relative
perfect	8.5×10^{-2}	0%	9.3×10^{-2}	0%
w2v_add				
regular	16×10^{-2}	84%	15×10^{-2}	62%
debiased	14×10^{-2}	62%	14×10^{-2}	53%
strong debiased	14×10^{-2}	62%	14×10^{-2}	53%
w2v_si				
regular	15×10^{-2}	77%	15×10^{-2}	64%
debiased	13×10^{-2}	55%	14×10^{-2}	53%
strong debiased	13×10^{-2}	54%	14×10^{-2}	53%
ftt_add				
regular	12×10^{-2}	46%	15×10^{-2}	58%
debiased	11×10^{-2}	35%	13×10^{-2}	44%
strong debiased	11×10^{-2}	35%	13×10^{-2}	43%
MP (<i>w2v</i>)				
regular	9.0×10^{-2}	6%	9.6×10^{-2}	4%
debiased	9.0×10^{-2}	6%	9.7×10^{-2}	5%
strong debiased	9.0×10^{-2}	6%	9.7×10^{-2}	5%
DRMM (<i>w2v</i>)				
regular	11×10^{-2}	28%	12×10^{-2}	24%
debiased	11×10^{-2}	25%	11×10^{-2}	21%
strong debiased	11×10^{-2}	26%	11×10^{-2}	21%
<i>lexical</i>				
QLM	8.9×10^{-2}	4%	10×10^{-2}	12%
tf-idf	9.5×10^{-2}	11%	10×10^{-2}	14%
BM25	9.4×10^{-2}	11%	10×10^{-2}	14%

Table 2: GSR measured according to *w2v* and *ftt* with 2 significant figures. Raw GSR values are shown (dubbed absolute), along with relative values, obtained from the former, as a percentage of the GSR value for the perfect search engine. Agreement between *w2v* and *ftt*: Spearman’s $\rho = 0.96$, $p < 1e-10$.

the context of IR algorithms: the most impacted queries are the ones with high genderedness, which are also the most important ones for GSR. If most queries have a low gender score, then the impact on aggregated performance will be insignificant.

4.2.4. Reliability

Objective: Word representations learnt with different techniques and corpora such as *w2v* (Google News) and *GloVe* (Wikipedia) have already been shown to exhibit similar bias along the gender direction [6, 7, 29]. To test the reliability of GSR, we check how dependent it is on a specific WE implementation. We do so by computing GSR based on FastText (*ftt*) WEs, trained on Common Crawl and check its agreement with *w2v*-based GSR.

Setup: The procedure to isolate a gender direction and projecting word vectors from *w2v* onto it [6] is perfectly applicable to different WEs. We compute GSR according to *ftt* embeddings, and compare it against results from previous sections obtained with *w2v*.

Discussion: As a preliminary check, we compute the correlation between query genderedness measured by *w2v* and *ftt*. Figure 11 is a scatter plot of the genderedness of 249 queries from Robust04 under *w2v* and *ftt*, which shows a strong correlation between the two (Pearson’s $r = 0.78$, $p < 1e-40$). This preliminary check confirms that *w2v* and *ftt* are likely to encode stereotypically gendered concepts in similar ways.

Table 2 shows the values of *ftt*-based GSR, in columns 3 and 4. SEs can be ranked according to GSR

System	MAP	nDCG@100	P@10
w2v_add			
regular	0.067	0.170	0.174
debiased	0.068 (+2%)	0.171 (+0%)	0.176 (+1%)
str. deb.	0.068 (+1%)	0.171 (+0%)	0.175 (+0%)
w2v_si			
regular	0.093	0.213	0.216
debiased	0.094 (+1%) [‡]	0.213 (+0%)	0.217 (+0%)
str. deb.	0.094 (+1%) [†]	0.213 (+0%)	0.217 (+0%)
ftt_add			
regular	0.056	0.144	0.150
debiased	0.056 (+0%)	0.144 (+0%)	0.148 (-1%)
str. deb.	0.056 (+0%)	0.144 (+0%)	0.147 (-2%)
MP (w2v)			
regular	0.151	0.283	0.287
debiased	0.148 (-2%)	0.279 (-1%)	0.285 (-1%)
str. deb.	0.148 (-2%)	0.279 (-1%)	0.285 (-1%)
DRMM (w2v)			
regular	0.260	0.423	0.456
debiased	0.259 (-1%) [‡]	0.422 (+0%)	0.454 (+0%)
str. deb.	0.259 (-0%) [†]	0.421 (-1%) [†]	0.457 (+0%)

Table 3: Impact of regular debiasing [6] and strong debiasing [70] on performance of models based on WEs. A Student’s t test is computed between regular and debiased versions of the same algorithm, with significance at $p = 0.05$ and $p = 0.01$ denoted by \dagger, \ddagger respectively.




Figure 10: Impact of regular debiasing on **w2v_add**: absolute value of query genderedness, on the x axis, and difference between top-100 documents retrieved by **w2v_add** before and after debiasing, on the y axis, as measured by Kendall τ distance. Pearson’s $r = 0.81$, $p < 1e - 50$.




Figure 11: Genderedness of Robust04 queries, according to **w2v** (x axis) and **ftt** (y axis). Correlation: Pearson’s $r = 0.78$, $p < 1e - 40$.

scores computed with **ftt** and **w2v**. We regard the correlation of these scores as a measure of the reliability of GSR across different WEs. In other words, we would like the ranking determined by **w2v**-based GSR and **ftt**-based GSR to agree as much as possible, as measured by Spearman’s rank coefficient (ρ). The values in Table 2 (either absolute or relative) yield Spearman’s $\rho = 0.96$, with a p -value $p < 1e - 10$. We conclude that GSR and the underlying gender direction is fairly reliable across **w2v** and **ftt** embeddings, despite the different text corpora from which they were learnt (Google News and Common Crawl respectively).

We anticipated that SEs based on `w2v` (namely `w2v_add`, `w2v_si`) would have a higher score when GSR is measured according to the same `w2v` WEs, than when GSR is `ftt`-based. Similarly, `ftt_add` has higher GSR when computed according to `ftt` than according to `w2v`. This is not surprising, given each SE based on word representations inherits the peculiar biases of its underlying WEs, which are the same biases that GSR captures. Despite this, `w2v`-based and `ftt`-based GSR show a solid overall agreement.

To sum up, GSR is stable across WEs that differ in architecture tweaks and choice of reasonably large training text corpus from the web domain.

4.2.5. GSR and direct stereotype

Hypothesis: In order to interpret results from GSR, we investigate its relationship with explicit mentions of female and male entities. For every document, we compute a binary measure of “intrinsic genderedness”. For the sake of simplicity, a document is considered:

- Intrinsically male, if it contains more male mentions than female ones.
- Intrinsically female, if it contains more female mentions than male ones.
- Neutral, otherwise.

We hypothesize that a high GSR will result in associating stereotypically gendered queries (such as the ones depicted in Figures 7a, 7b) to intrinsically gendered documents of the same polarity. In other words, GSR should capture direct gender stereotypes (Definition 3.3), taking large values for SEs which associate stereotypically female (male) concepts with female (male) entities. This hypothesis relates to the *content validity* of GSR, as we would expect our measure to capture this form of direct bias.

Setup: To assess intrinsic genderedness of documents, male and female names are sourced from `nlk`’s names corpus.¹⁰ Gendered nouns, adjectives and titles are obtained starting from definitional pairs and gender-specific words in Appendix C of Bolukbasi et al. [6], of which we only keep words that specifically refer to a person. Under this criterion, `aunt` is considered a female entity, whereas `pregnancy` is not. The resulting word list referring to gendered entities is reported in Appendix B.

We compare the “perfect” search engine (dubbed P) (retrieving all and only the relevant documents for each query) against one based on `w2v_add`, which has the highest GSR among the tested systems. As a comparison, we also include QLM and MP (low GSR) and `ftt_add` (medium GSR). Each system is compared against P as follows:

- For each query q we compute the number of intrinsically female and male documents among the ones retrieved by each SE (dubbed $f(q)$ and $m(q)$ respectively). We use their ratio as a summary of representation gap in search results ($gap(q) = \frac{m(q)}{f(q)}$). For a given list of search results \mathcal{L} , $gap(q)$ quantifies the extent to which documents in \mathcal{L} tend to mention more male entities than female ones.
- For each query q , we compute $gap(q)$ under P and `sys`, the system at hand (`w2v_add`, `ftt_add`, QLM and MP).
- Their difference, $\Delta_{gap}(q) = gap(q)^{sys} - gap(q)^P$ summarizes over- or under-exposure of user to documents with male entities, compared against the ground truth of system P.
- Based on the sign of $\Delta_{gap}(q)$, we determine whether `sys` is favoring male documents (if positive) or female documents (if negative).
- To test our hypothesis, we compute $sgn(\Delta_{gap}(q))$ for every query and compare it against the genderedness $g(q)$ of said query.

We expect low $g(q)$ (stereotypically male queries) to be associated with over-representation of male entities, high $g(q)$ with under-representation. Furthermore, this relationship should be strong for `w2v_add`, weaker for `ftt_add`, and absent from QLM and MP

¹⁰<https://www.nltk.org/book/ch02.html>




Figure 12: On the y axis, in orange (purple) percentage of queries where intrinsically male (female) documents are over-represented among retrieved ones. The complementary set, depicted in gray, is the percentage of queries for which neither gender is over-represented ($\Delta_{gap}(q) = 0$). The x axis is a quantization on query genderedness $g(q)$. Top panes depict semantic systems `w2v_add` (high GSR) and `ftt_add` (medium GSR), bottom panes show systems from the lexical and neural family (QLM and MP - low GSR). No query in Robust04 has $g(q) < -0.1$, the bin is therefore empty.

Discussion: Figure 12a confirms our expectation for `w2v_add`, with a clear trend along the x axis. We expect said trend to be less evident for `ftt_add`, given its lower GSR, which is confirmed by Figure 12b.

QLM and MP (low GSR) are represented in Figures 12c and 12d. The former seems to have a weak trend similar to that of `ftt_add`, disconfirmed however by the last bin, which contains the most gendered queries ($|g(q)| > 0.1$) but does not significantly favor “intrinsically female” documents. The latter displays no trend along the x axis. In sum, as anticipated, no consistent trend is visible for systems with low GSR.

We conclude that GSR captures this form of direct gender stereotype: SEs with high GSR associate stereotypically gendered queries with documents mentioning people of the same gender.

5. Conclusions and Future Work

We defined Gender Stereotype Reinforcement (GSR) in SEs, a construct describing the tendency of a SE to reinforce direct and indirect biases about gender, which we made operational employing WEs as a measurement tool. We validated our approach against well-studied gender stereotypes from psychology literature, and exploited the framework of construct validity [42] to critically evaluate our novel measure. We found that GSR captures gender stereotypes, while also being influenced by the relevance of documents retrieved for each query. This is due to the domain-specificity of language: queries and relevant documents are likely to share some specific vocabulary, whose words cluster in the embedding space and, subsequently, along the gender subspace. This aspect can be compensated when assessor judgements are available. In this regard, TREC collection Robust04 [38] has proven to be a suitable dataset to measure the extent to which different IR algorithms reinforce gender stereotype. This is due to availability of relevance judgements, large amount of queries, interesting content of some queries from a gender stereotype perspective.

Subsequently, we studied how lexical, semantic and neural IR models reinforce gender stereotypes. We found that semantic models, based on biased WEs, are most prone to reinforcement of gender stereotypes, while neural models based on the same word representations can mitigate this effect, exhibiting low GSR, comparable to that of lexical systems. The reliability of these conclusions was tested with two different sets of WEs (Word2Vec and FastText), identifying strong agreement between the two measurements.

Finally, we assessed the impact of debiasing WEs on downstream IR tasks. Regular debiasing [6] and strong debiasing [70] have a similar effect, reducing GSR to a significant yet moderate extent. We conclude that the gender direction encoded by WEs is a useful proxy for the gender-related biases contained in the large online corpora they have been trained on. However, debiasing techniques based on projecting WEs orthogonally to the same gender direction are superficial and insufficient, due to redundant encoding of stereotypical information. This also explains the minimum impact debiasing has on model performance.

In sum, GSR can measure associations of documents and queries along gendered lines, detecting and quantifying polarization in the language used to respond to stereotypically female and male queries. We showed that GSR captures the difference in the number of stereotypical and counter-stereotypical documents within a search history, drawing a parallel with existing statistical parity metrics [28].

A limitation of our measurement is the compositional model employed to assemble word scores into document scores, which does not account for syntactic structure, thus neglecting important information, such as negation. A second drawback is the noisy nature of the gender information encoded in WEs, which should discourage the deployment of GSR on small collections, unless supported by human supervision. These observations are crucial to discuss the *consequential validity* of the proposed measure. If GSR were to be integrated as part of the ranking function of a SE, it would likely favour documents which appear to be gender-neutral or counter-stereotypical for the queries issued by users. Indeed, it would be possible for providers of documents to target our measure, ensuring that their documents are not flagged as stereotypical for some queries of interest. Moreover, intrinsically gendered queries, such as **women in parliament**, require special care; low GSR may contradict user preferences. For these reasons, we consider our operationalization of GSR a preliminary attempt to measure gender stereotype reinforcement in SEs, with limited consequential validity in fully automated contexts. Future work should include an exploration of different compositional models, based, for instance, on dependency parsers, and novel approaches to compute a gender score for words and phrases, including *ad-hoc* training [89]. Finally, it will be interesting to measure GSR in cross-lingual scenarios; grammatical gender may pose an additional challenge in some languages, especially for the isolation of gender information along a single direction.

To the best of our knowledge, GSR is the first measure in the domain of IR capable of quantifying a specific type of representational harm, namely gender stereotypes. This opens the possibility to quantitatively study the interplay between distributional and representational harms, which makes GSR very promising in terms of *hypothesis validity* and its future uses. In the context of job search, it would be meaningful to study this interplay, due to high stakes, proven existence of biased tools [11], and availability of datasets [17]. As noticed by Chen et al. [11], search results in resume SEs, which happen to be biased with respect to gender, may lead to a dual harm: an immediate one, for the providers of CVs, competing to appear in the current search, and a long-term one, for the perception and future decisions of recruiters.

References

- [1] Abbasi, M., Friedler, S. A., Scheidegger, C., Venkatasubramanian, S., 2019. Fairness in representation: quantifying stereotyping as a representational harm. In: Proc. of SIAM 2019. pp. 801–809.
- [2] Bamman, D., O’Connor, B., Smith, N., 2012. Censorship and deletion practices in chinese social media. *First Monday* 17 (3).
- [3] Berinsky, A. J., Mendelberg, T., 2005. The indirect effects of discredited stereotypes in judgments of jewish leaders. *American Journal of Political Science* 49 (4), 845–864.
- [4] Biega, A. J., Gummadi, K. P., Weikum, G., 2018. Equity of attention: Amortizing individual fairness in rankings. In: Proc of 41st ACM SIGIR. SIGIR 18. Association for Computing Machinery, New York, p. 405414.
- [5] Bobbitt-Zeher, D., 2011. Gender discrimination at work: Connecting gender stereotypes, institutional policies, and gender composition of workplace. *Gender & Society* 25 (6), 764–786.
- [6] Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., Kalai, A. T., 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In: *Advances in neural information processing systems*. pp. 4349–4357.
- [7] Caliskan, A., Bryson, J. J., Narayanan, A., 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356 (6334), 183–186.
- [8] Carpineto, C., D’Amico, M., Romano, G., 2012. Evaluating subtopic retrieval methods: Clustering versus diversification of search results. *Information Processing & Management* 48 (2), 358–373.
- [9] Celis, E., Mehrotra, A., Vishnoi, N., 2019. Toward controlling discrimination in online ad auctions. In: Proc. of ICML 2019. pp. 4456–4465.
- [10] Census Bureau, 2019. Current population survey. Accessed = 2020-02-12. URL <https://www.bls.gov/cps/cpsa2019.pdf>
- [11] Chen, L., Ma, R., Hannák, A., Wilson, C., 2018. Investigating the impact of gender on rank in resume search engines. In: Proc. of the 2018 chi conference on human factors in computing systems. pp. 1–14.
- [12] Clarke, C. L., Kolla, M., Cormack, G. V., Vechtomova, O., Ashkan, A., Büttcher, S., MacKinnon, I., 2008. Novelty and diversity in information retrieval evaluation. In: Proc. of SIGIR 2008. pp. 659–666.
- [13] Cleverdon, C., 1997. *The Cranfield Tests on Index Language Devices*. Morgan Kaufmann Publishers Inc., p. 4759.
- [14] Cover, T., Thomas, J., 2012. *Elements of information theory*. John Wiley & Sons.
- [15] Cuddy, A. J., Fiske, S. T., Glick, P., 2004. When professionals become mothers, warmth doesn’t cut the ice. *Journal of Social issues* 60 (4), 701–718.
- [16] Cvencek, D., Meltzoff, A. N., Greenwald, A. G., 2011. Math–gender stereotypes in elementary school children. *Child development* 82 (3), 766–779.
- [17] De-Arteaga, M., Romanov, A., Wallach, H., Chayes, J., Borgs, C., Chouldechova, A., Geyik, S., Kenthapadi, K., Kalai, A. T., 2019. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In: Proc. of FAT* 2019. ACM, pp. 120–128.
- [18] DeFranza, D., Mishra, H., Mishra, A., 2020. How language shapes prejudice against women: An examination across 45 world languages. *Journal of Personality and Social Psychology*.
- [19] Dev, S., Phillips, J., 2019. Attenuating bias in word vectors. In: Chaudhuri, K., Sugiyama, M. (Eds.), *Proc. of Machine Learning Research*. Vol. 89. PMLR, pp. 879–887.
- [20] Eagly, A. H., Nater, C., Miller, D. I., Kaufmann, M., Sczesny, S., 2019. Gender stereotypes have changed: A cross-temporal meta-analysis of us public opinion polls from 1946 to 2018. *American psychologist* 75 (3), 301–315.
- [21] Ekstrand, M. D., Burke, R., Diaz, F., 2019. Fairness and discrimination in retrieval and recommendation. In: Proc. of SIGIR 2019. ACM, p. 14031404.
- [22] Ellemers, N., 2018. Gender stereotypes. *Annual Review of Psychology* 69 (1), 275–298.
- [23] Fazio, R. H., Jackson, J. R., Dunton, B. C., Williams, C. J., 1995. Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of personality and social psychology* 69 (6), 1013–1027.
- [24] Ferrante, M., Ferro, N., Maistro, M., 2014. Injecting user models and time into precision via markov chains. In: Proc. of SIGIR 2014. ACM, p. 597606.
- [25] Fiske, S. T., Cuddy, A. J., Glick, P., Xu, J., 2002. A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. In: *Journal of Personality and Social Psychology*. pp. 878–902.
- [26] Flaxman, S., Goel, S., Rao, J. M., 03 2016. Filter Bubbles, Echo Chambers, and Online News Consumption. *Public Opinion Quarterly* 80 (S1), 298–320.
- [27] Fu, Z., Ren, K., Shu, J., Sun, X., Huang, F., 2015. Enabling personalized search over encrypted outsourced data with efficiency improvement. *IEEE transactions on parallel and distributed systems* 27 (9), 2546–2559.
- [28] Gao, R., Shah, C., 2020. Toward creating a fairer ranking in search engine results. *Information Processing & Management* 57 (1), to appear.
- [29] Garg, N., Schiebinger, L., Jurafsky, D., Zou, J., 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115 (16), 3635–3644.
- [30] Gerbner, G., Gross, L., Morgan, M., Signorielli, N., 1986. *Living with television: The dynamics of the cultivation process. Perspectives on media effects 1986*, 17–40.
- [31] Gerritse, E., 2019. Impact of debiasing word embeddings on information retrieval. In: Proc. of FDIA 2019. CEUR-ws, p. 5459.
- [32] Ghosh, V. E., Gilboa, A., 2014. What is a memory schema? a historical perspective on current neuroscience literature. *Neuropsychologia* 53, 104–114.

- [33] Gonen, H., Goldberg, Y., 2019. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. In: Proc. of NAACL 2019. ACL, pp. 609–614.
- [34] Graells-Garrido, E., Lalmas, M., Menczer, F., 2015. First women, second sex: Gender bias in wikipedia. In: Proc. of HT 2015. pp. 165–174.
- [35] Greenwald, A. G., McGhee, D. E., Schwartz, J. L., 1998. Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology* 74 (6), 1464–1480.
- [36] Guo, J., Fan, Y., Ai, Q., Croft, W. B., 2016. A deep relevance matching model for ad-hoc retrieval. ACM Press, pp. 55–64.
- [37] Halavais, A., 2008. Search Engine Society. Polity Press.
- [38] Harman, D., 1992. The darpa tipster project. *SIGIR Forum* 26 (2), 2628.
- [39] Hentschel, T., Heilman, M. E., Peus, C. V., 2019. The multiple dimensions of gender stereotypes: a current look at mens and womens characterizations of others and themselves. *Frontiers in psychology* 10 (11).
- [40] Hinton, P., 2017. Implicit stereotypes and the predictive brain: cognition and culture in biased person perception. *Palgrave Communications* 3 (1), 1–9.
- [41] Hong, S., Kim, N., 2018. Will the internet promote democracy? search engines, concentration of online news readership, and e-democracy. *Journal of Information Technology & Politics* 15 (4), 388–399.
- [42] Jacobs, A. Z., Wallach, H., 2019. Measurement and fairness.
- [43] Järvelin, K., Kekäläinen, J., 2002. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems* 20 (4), 422446.
- [44] Joachims, T., Granka, L., Pan, B., Hembrooke, H., Gay, G., 2017. Accurately interpreting clickthrough data as implicit feedback. *SIGIR Forum* 51 (1), 411.
- [45] Joulin, A., Grave, E., Bojanowski, P., Mikolov, T., 2017. Bag of tricks for efficient text classification. In: Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers. ACL, pp. 427–431.
- [46] Kay, M., Matuszek, C., Munson, S. A., 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. ACM, pp. 3819–3828.
- [47] Kulshrestha, J., Eslami, M., Messias, J., Zafar, M. B., Ghosh, S., Gummadi, K. P., Karahalios, K., 2017. Quantifying search bias: Investigating sources of bias for political searches in social media. In: Proc. of CSCW 2017. ACM, p. 417432.
- [48] Le Pelley, M. E., Reimers, S. J., Calvini, G., Spears, R., Beesley, T., Murphy, R. A., 2010. Stereotype formation: Biased by association. *Journal of Experimental Psychology: General* 139 (1), 138–161.
- [49] Marchesin, S., Purpura, A., Silvello, G., 2019. Focal elements of neural information retrieval models. an outlook through a reproducibility study. *Information Processing & Management*, to appear.
- [50] Martin, C. L., Ruble, D. N., 2010. Patterns of gender development. *Annual review of psychology* 61, 353–381.
- [51] Messick, S., 1995. Validity of psychological assessment: Validation of inferences from persons’ responses and performances as scientific inquiry into score meaning. *American psychologist* 50 (9), 741–749.
- [52] Messick, S., 1998. Test validity: A matter of consequence. *Social Indicators Research* 45 (1-3), 35–44.
- [53] Mikolov, T., Grave, E., Bojanowski, P., Puhersch, C., Joulin, A., 2018. Advances in pre-training distributed word representations. In: Proc. of LREC 2018.
- [54] Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., Dean, J., 2013. Distributed representations of words and phrases and their compositionality. In: Advances in neural information processing systems. pp. 3111–3119.
- [55] Mitra, B., Craswell, N., 2018. An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval* 13 (1), 1–126.
- [56] Moffat, A., Zobel, J., 2008. Rank-biased precision for measurement of retrieval effectiveness. *ACM Transactions on Information Systems* 27 (1), 1–27.
- [57] Mumtaz, Z., Shahid, U., Levay, A., 2013. Understanding the impact of gendered roles on the experiences of infertility amongst men and women in punjab. *Reproductive health* 10 (3), 1–10.
- [58] Noble, S. U., 2018. Algorithms of oppression: How search engines reinforce racism. NYU Press.
- [59] Nosek, B. A., Banaji, M. R., Greenwald, A. G., 2002. Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics: Theory, Research, and Practice* 6 (1), 101–115.
- [60] Nosek, B. A., Banaji, M. R., Greenwald, A. G., 2002. Math= male, me= female, therefore math≠ me. *Journal of personality and social psychology* 83 (1), 44–59.
- [61] Novin, A., Meyers, E., 2017. Making sense of conflicting science information: Exploring bias in the search engine result page. In: Proc. of CHIIR 2017. ACM, p. 175184.
- [62] Otterbacher, J., Bates, J., Clough, P., 2017. Competent men and warm women: Gender stereotypes and backlash in image search results. In: Proc. of CHI 2017. ACM, p. 66206631.
- [63] Pang, L., Lan, Y., Guo, J., Xu, J., Wan, S., Cheng, X., 2016. Text matching as image recognition. In: Proc. of AAAI 2016. AAAI Press, p. 27932799.
- [64] Papakyriakopoulos, O., Hegelich, S., Serrano, J. C. M., Marco, F., 2020. Bias in word embeddings. In: Proc. of FAT* 2020. ACM, p. 446457.
- [65] Park, L. E., Young, A. F., Eastwick, P. W., 2015. (psychological) distance makes the heart grow fonder: Effects of psychological distance and relative intelligence on mens attraction to women. *Personality and Social Psychology Bulletin* 41 (11), 1459–1473.
- [66] Patterson, K., Nestor, P. J., Rogers, T. T., 2007. Where do you know what you know? the representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience* 8 (12), 976–987.
- [67] Pennington, J., Socher, R., Manning, C., 2014. Glove: Global vectors for word representation. In: Proc. of EMNLP 2014.

- pp. 1532–1543.
- [68] Penton-Voak, I. S., Pound, N., Little, A. C., Perrett, D. I., 2006. Personality judgments from natural and composite facial images: More evidence for a kernel of truth in social perception. *Social cognition* 24 (5), 607–640.
 - [69] Plante, I., De la Sablonnière, R., Aronson, J. M., Théorêt, M., 2013. Gender stereotype endorsement and achievement-related outcomes: The role of competence beliefs and task values. *Contemporary Educational Psychology* 38 (3), 225–235.
 - [70] Prost, F., Thain, N., Bolukbasi, T., 2019. Debiasing embeddings for reduced gender bias in text classification. In: Proc. of the 1st ACL Workshop on Gender Bias for Natural Language Processing. ACL, pp. 69–75.
 - [71] Robertson, S. E., Zaragoza, U., 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Foundations and Trends in Information Retrieval (FnTIR)* 3 (4), 333–389.
 - [72] Salton, G., McGill, M. J., 1983. *Introduction to Modern Information Retrieval*. McGraw-Hill, New York, USA.
 - [73] Sayans-Jimnez, P., van Harreveld, F., Dalege, J., Rojas Tejada, A. J., 2019. Investigating stereotype structure with empirical network models. *European Journal of Social Psychology* 49 (3), 604–621.
 - [74] Shrum, L. J., 1995. Assessing the social influence of television: A social cognition perspective on cultivation effects. *Communication Research* 22 (4), 402–429.
 - [75] Singh, A., Joachims, T., 2018. Fairness of exposure in rankings. In: Proc. of KDD 2018. ACM, p. 22192228.
 - [76] Sink, A., Mastro, D., 2017. Depictions of gender on primetime television: A quantitative content analysis. *Mass Communication and Society* 20 (1), 3–22.
 - [77] Sparrow, B., Liu, J., Wegner, D. M., 2011. Google effects on memory: Cognitive consequences of having information at our fingertips. *science* 333 (6043), 776–778.
 - [78] Suresh, H., Gutttag, J. V., 2019. A framework for understanding unintended consequences of machine learning. arXiv preprint arXiv:1901.10002.
 - [79] Tantleff-Dunn, S., Barnes, R. D., Larose, J. G., 2011. It’s not just a “woman thing:” the current state of normative discontent. *Eating disorders* 19 (5), 392–402.
 - [80] Tresh, F., Steeden, B., de Moura, G. R., Leite, A. C., Swift, H. J., Player, A., 2019. Endorsing and reinforcing gender and age stereotypes: The negative effect on self-rated leadership potential for women and older workers. *Frontiers in psychology* 10.
 - [81] Tversky, A., Kahneman, D., 1973. Availability: A heuristic for judging frequency and probability. *Cognitive psychology* 5 (2), 207–232.
 - [82] Vuli, I., Moens, M.-F., 2015. Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings. In: Proc. of SIGIR 2015. ACM, p. 363372.
 - [83] Wu, Y., Zhang, L., Wu, X., 2018. On discrimination discovery and removal in ranked data using causal graph. In: Proc. of KDD 2018. pp. 2536–2544.
 - [84] Yang, K., Stoyanovich, J., 2017. Measuring fairness in ranked outputs. In: Proc. of the 29th International Conference on Scientific and Statistical Database Management. pp. 1–6.
 - [85] Yu, H.-T., Jatowt, A., Blanco, R., Joho, H., Jose, J. M., Chen, L., Yuan, F., 2018. Revisiting the cluster-based paradigm for implicit search result diversification. *Information Processing & Management* 54 (4), 507–528.
 - [86] Zehlike, M., Bonchi, F., Castillo, C., Hajian, S., Megahed, M., Baeza-Yates, R., 2017. Fa* ir: A fair top-k ranking algorithm. In: Proc. of CIKM 2017. pp. 1569–1578.
 - [87] Zhai, C., 2008. *Statistical Language Models for Information Retrieval. A Critical Review*. Foundations and Trends in Information Retrieval (FnTIR) 2 (3), 137–213.
 - [88] Zhang, X., Ghorbani, A. A., 2020. An overview of online fake news: Characterization, detection, and discussion. *Information Processing & Management* 57 (2), 102025.
 - [89] Zhao, J., Zhou, Y., Li, Z., Wang, W., Chang, K.-W., 2018. Learning gender-neutral word embeddings. In: Proc. of EMNLP 2018. ACL, pp. 4847–4853.
 - [90] Zielinski, K., Nielek, R., Wierzbicki, A., Jatowt, A., 2018. Computing controversy: Formal model and algorithms for detecting controversy on wikipedia and in search queries. *Information Processing & Management* 54 (1), 14 – 36.

Appendix A. Traits and terms for stereotypical associations

agency	communion
aggressive	affectionate
ambitious	compassionate
arrogant	emotional
confident	generous
corageous	honest
critical	nurturing
decisive	outgoing
demanding	patient
hardworking	polite
independent	romantic
possessive	sensitive
proud	unselfish
selfish	
strong	
stubborn	

Table A.4: *agency vs communion*: adjectives associated to each construct [20].

science	arts
astronomy	art
chemistry	dance
Einstein	drama
experiment	literature
NASA	novel
physics	poetry
science	Shakespeare
technology	symphony

Table A.5: *science vs arts*: associated attributes [59].

career	family
business	children
career	cousin
corporation	family
executive	home
management	marriage
office	parents
professional	relatives
salary	wedding

Table A.6: *career vs family*: associated attributes [60].

Predominantly male			Predominantly female		
occupation	%F	%M	occupation	%F	%M
stonemason	0.7	99.3	hygienist	96.0	4.0
roofer	1.9	98.1	secretary	93.2	6.8
electrician	2.2	97.8	hairdresser	92.3	7.7
plumber	2.7	97.3	dietician	92.1	7.9
carpenter	2.8	97.2	paralegal	89.6	10.4
firefighter	3.3	96.7	receptionist	89.3	10.7
millwright	5.0	95.0	phlebotomist	89.3	10.7
welder	5.3	94.7	maid	89.0	11.0
machinist	5.6	94.4	nurse	88.9	11.1
driver	6.7	93.3	typist	86.0	14.0

Table A.7: *jobs_m vs jobs_f*: occupations with highest gender gap in representation [10].

Appendix B. Gendered entities

The following are used in section 4.2.5 to detect mentions of intrinsically gendered entities.

Words associated with male entities:

actor, actors, bachelor, bachelors, bloke, blokes, boy, boys, boyfriend, boyfriends, brother, brothers, brethren, businessman, businessmen, chairman, chairmen, chap, chaps, congressman, congressmen, councilman, councilmen, dad, daddy, dads, dude, dudes, ex-boyfriend, ex-boyfriends, exboyfriend, exboyfriends, father, fathers, fella, fellas, gentleman, gentlemen, godfather, godfathers, grandfather, grandfathers, grandpa, grandson, grandsons, guy, guys, handyman, handymen, he, him, himself, his, husband, husbands, king, kings, lad, lads, male, males, man, men, monk, monks, mr, nephew, nephews, pa, prince, princes, salesman, salesmen, schoolboy, schoolboys, son, sons, spokesman, spokesmen, statesman, statesmen, stepfather, stepfathers, stepson, stepsons, uncle, uncles, waiter, waiters.

Words associated with female entities:

actress, actresses, aunt, aunts, ballerina, ballerinas, bride, brides, businesswoman, businesswomen, chairwoman, chairwomen, congresswoman, congresswomen, councilwoman, councilwomen, daughter, daughters, exgirlfriend, ex-girlfriends, ex-girlfriend, ex-girlfriends, female, females, gal, gals, girl, girls, girlfriend, girlfriends, godmother, godmothers, granddaughter, granddaughters, grandma, grandmas, grandmother, grandmothers, her, hers, herself, hostess, hostesses, housewife, housewives, lady, ladies, ma, maid, maiden, maids, mama, mom, mommy, moms, mother, mothers, ms, mrs, niece, nieces, nun, nuns, princess, princesses, queen, queens, schoolgirl, schoolgirls, she, sister, sisters, spokeswoman, spokeswomen, stepdaughter, stepmother, waitress, waitresses, wife, wives, woman, women.