

bias goggles: Exploring the Bias of Web Domains through the Eyes of Users

Giannis Konstantakis¹, Gianins Promponas¹, Manthos Dretakis¹ and Panagiotis Papadakos^{1,2}[0000-0001-8926-4229]

¹ Computer Science Department, University of Crete, Greece

² Institute of Computer Science, FORTH-ICS, Greece

{jkonstan,csd3522,csd3357}@csd.uoc.gr, papadako@ics.forth.gr

Abstract. Ethical issues, along with transparency, disinformation and bias are in the focus of our information society. In this demo, we will present the *bias goggles* system, that based on the web graph computes the bias characteristics of web domains to user-defined concepts. Our approach uses adaptations of propagation models and a variation of the pagerank algorithm named **Biased-PR**, that models various behaviours of biased surfers. Currently, the system runs over a subset of the greek web graph. We have developed cross-browser plugins that let users explore the bias characteristics of domains, define their own biased concepts and monitor the consumption of biased information through some analytics.

Keywords: bias · web graph · propagations models · biased pagerank · browser plugins

1 Introduction

There is an increasing concern about the potential risks in the consumption of the abundant biased information in online platforms like Web Search Engines (WSEs) and social networks. Terms like echo chambers and filter-bubbles [2] depict the isolation of groups of people and its aftereffects that result from the selective and restricted exposure to information. This restriction can be the result of helpful personalized algorithms that suggest user connections or rank high information relevant to the users' profile. Yet, this isolation might inhibit the growth of informed and responsible humans / citizens / consumers, and can also be the result of malicious algorithms that promote and resurrect social, religious, ethnic, and other kinds of discriminations and stereotypes.

Currently, the community focus is towards the transparency, fairness and accountability of mostly machine learning algorithms for decision making, classification, and recommendation in social platforms like twitter. However, social platforms and WSEs mainly act as gateways to information published on the web as common web pages (e.g., blogs and news). But, users are unaware of the bias characteristics of these pages, except from obvious facts like that a page in a political party's web site will be biased towards this party.

bias goggles is a system that allows users to explore the bias characteristics of web domains for a specific concept (i.e., a bias goggle). Since there is no objective definition about what bias and biased concepts are [3], we let users define them. For these concepts, the system computes the *support* and the *bias score* of the domains, using the *support* of this domain for each aspect of the biased concept. These *support* scores are calculated by graph-based algorithms that exploit the structure of the web graph and a set of user-defined seeds for each aspect of bias. As an example consider the bias goggle of "politics in Greece", which is defined using as aspects the 9 most popular greek political parties, and their homepages as seeds. More information about the model, its assumptions, the computation of *support* and *bias* scores, the propagation and **Biased-PR** algorithms, and an experimental evaluation, is provided in [1].

In this demo we showcase two front-ends that allow users to search the bias characteristics of web domains. The first one is a simple web-page that can be used for searching web domains over predefined biased concepts. The second one is a cross-browser plugin that allows users to define or search the available biased concepts and aspects of bias, explore and compare the bias characteristics of the web domains as they visit or show interest for specific web pages, and finally to monitor their behavior regarding the bias characteristics of the sites and the information they consume.

2 Implementation Discussion

2.1 Back-end

The developed prototype of the *bias goggles* system allows the exploration of biased concepts over a set of mainly greek domains. For gathering pages from the web, we have developed a crawler based on the open-source Java crawler4j³ project, which is a multi-threaded but unfortunately non-distributed crawler. The rest of the system is also implemented in Java. Specifically, we have developed two propagation models and the **Biased-PR** algorithm, which is a variation of the PageRank algorithm that models a biased surfer. These algorithms exploit the structure of the web graph to compute the bias scores of the domains and can be used on demand. Further, we have developed highly efficient inverted-file indexes, that provide fast access to the bias scores of the crawled domains, and to any information related to the available biased concepts and aspects of bias. Such functionality includes searching specific biased concepts or aspects of bias, either over their descriptions or the names of the domains over which they are defined. On top of all this functionality we have implemented a REST-API based on the spark-java micro-services library⁴, which is available at <http://pangaia.ics.forth.gr:4567/bias-goggles-api/>, that allows users to query and insert relevant information.

³ <https://github.com/yasserg/crawler4j>

⁴ <http://sparkjava.com/>

2.2 Front-ends

Currently, we offer two different front-ends that provide access to the underlying information. The first one is a simple web-page that allows users to search the bias characteristics of web domains for two biased concepts: the “greek political parties” and the “greek football teams”. Also, this page helps users to explore the currently crawled domains in descending order according to their bias scores for the two previously mentioned biased concepts. A screenshot of our web front-end that depicts the top-k biased results for the “greek political parties” goggle is given in Fig. 1. The web site is publicly accesible at <http://pangaia.ics.forth.gr/bias-goggles>.

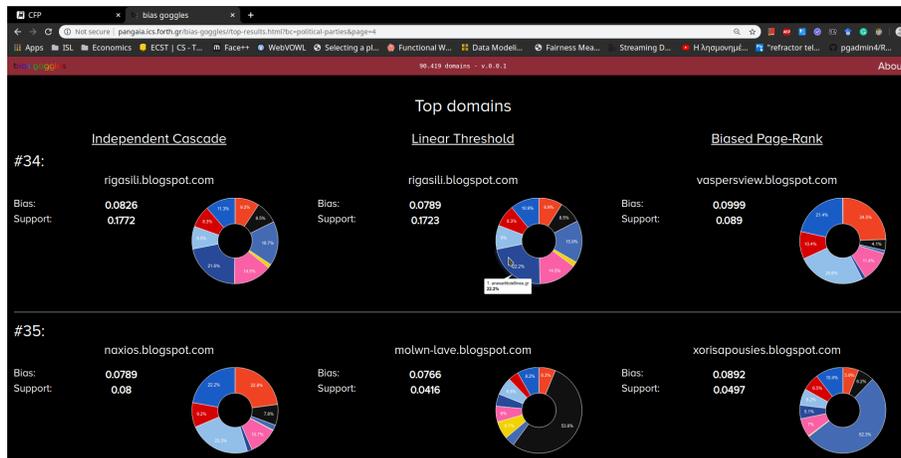


Fig. 1: A screenshot of the web interface, listing the top-biased domains for the Independent Cascade and Linear Threshold propagation models along with the Biased-PR algorithm.

We have also implemented cross-browser plugins⁵ based on the WebExtensions API⁶, along with a user friendlier sidebar extension⁷ for Firefox. These plugins allow users to define their own bias goggles, retrieve the bias characteristics of the currently visited page and its linked pages, and compare the bias scores of domains. Finally, they offer various analytics and graphs about the bias scores of the sites that a user has visited over a period of time for specific bias-goggles (e.g., average bias score, average support score, top-biased domains, etc.). Figure 2 depicts how the bias scores are shown for a specific web site, along with a comparison of the scores of 3 different domains in the Firefox sidebar plugin. Finally, Fig. 3 shows the UIs for searching and creating bias goggles, along

⁵ The plugins have been tested only on the Firefox and Chrome/Chromium browsers

⁶ <https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions>

⁷ https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions/user_interface/Sidebars

with the analytics information in the Chrome plugin. The plugins will be made available in the prototype’s web site in the near future.

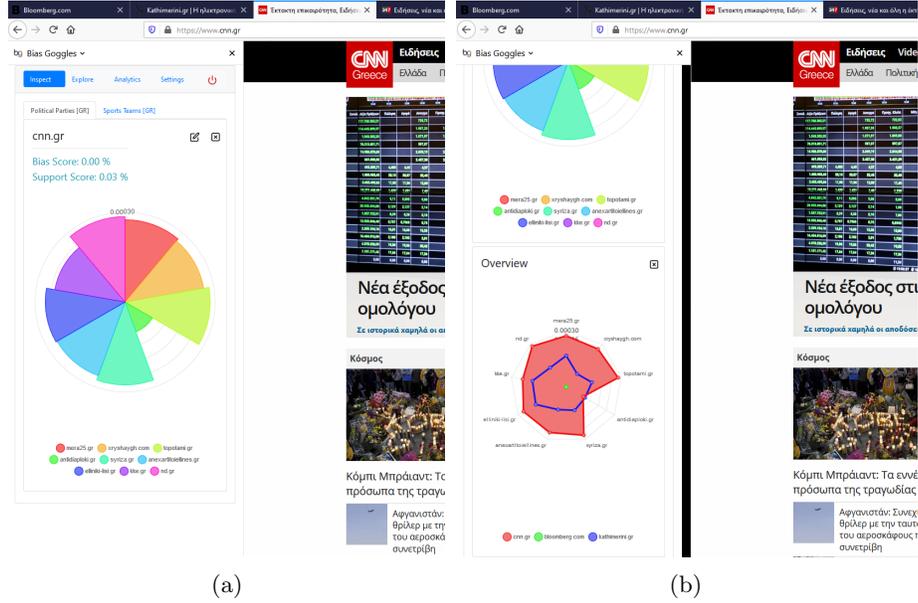


Fig. 2: The Firefox browser plugin. Figure (a) depicts the CNN.gr bias characteristics for the greek politics bias goggle, while Fig. (b) compares CNN.gr to bloomberg.com and the greek news paper kathimerini.gr.

3 Crawled Data

We have crawled a subset of the greek web by running four instances of the crawler: one with 383 sites related to the greek political life, one with 89 sport related greek sites, one with the top-300 popular greek sites according to Alexa, and a final one containing 127 seeds related to big greek industries. We black-listed popular sites like facebook and twitter to control the size of our data and avoid crawling non-greek domains. The crawlers were restricted to depth seven for each domain, and free to follow any link to external domains. In total we downloaded 893,095 pages including 531,296,739 links, which lead to the non-connected **Support Flow Graph (SFG)**⁸ of 90,419 domains, 288,740 links (on average 3.1 links per domain) and with a diameter $k = 7,944$. More information about the crawled pages, the hand-crafted gold standard that we used in

⁸ The SFG graph is the crawled graph, where the nodes are the crawled domains and the links are the corresponding inverse links.

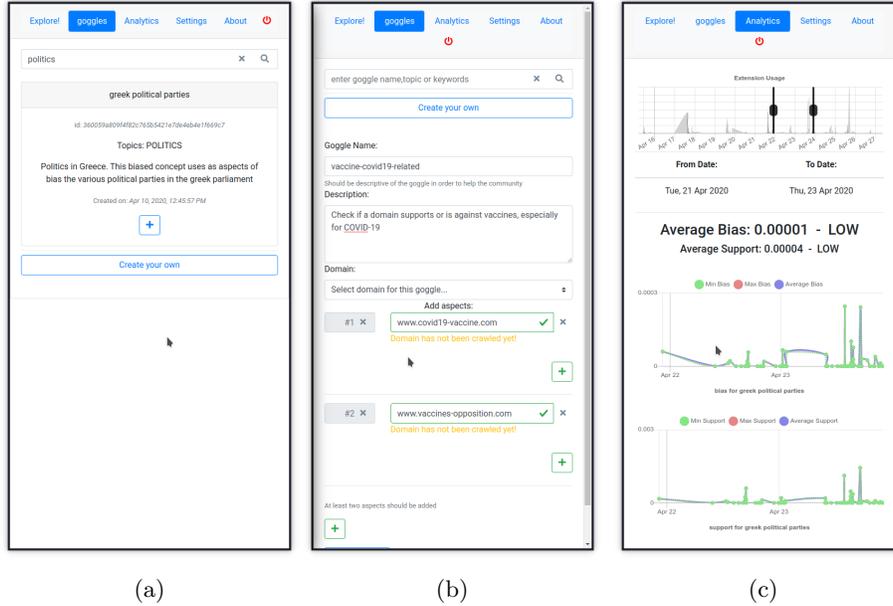


Fig. 3: The chrome browser plugin. Figure (a) depicts the results of searching available bias goggles for the query ‘politics’, Fig. (b) the UI for creating a new bias goggle related to COVID-19 and vaccines, and finally Fig. (c) depicts some analytics of the average bias and support scores over the consumed information for a specific period of time.

our evaluation in [1], the SFG graph and the graph itself are available in the prototype’s web site.

4 Performance Discussion

The most expensive part of our approach is the computation of the *support* scores of the domains for a specific bias goggle, due to the huge scale of the web graph. What is encouraging though is that the algorithms are applied to the much more compact SFG graph, that contains only the Second and Third Level Domains (SLDs & TLDs)⁹ of the urls of the pages and their corresponding links, instead of the original web graph. In our experiments (reported in [1]), we have shown that one propagation model can be used in real-time to compute the *support* scores. On the other hand the whole process of computing the *support* scores of the domains for a specific bias goggle, can be considered an offline process. The user can submit his/her own bias concepts to the system and then query the results after being notified that the system has computed the required *support*

⁹ We follow the standard URL normalization method (see https://en.wikipedia.org/wiki/URI_normalization) and get the SLD/TLD of an url.

scores. However, what is crucial, is the ability to let users explore in real-time the **bias** scores of the domains space for any precomputed and commonly used bias goggle. This can be easily supported by providing efficient ways to store and retrieve the computed *support* scores of domains, which we have implemented through appropriate indexes that allow the fast retrieval of offsets in files, where the *support* scores and the related metadata are stored. Given the above, the computation of the *bias score* of a domain for a bias goggle is rather fast.

5 Future Work

In the future, we plan to explore variations of the proposed approach where our assumptions do not hold. For example, currently we are trying to exploit the supportive, neutral or opposite nature of the available links, as identified by sentiment analysis methods, along with the importance of the web pages they appear in. Content-based and hybrid approaches for computing the *support* scores of domains are also in our focus, as well as the exploitation of other available graphs, like the graph of friends, retweets, etc. In addition interesting aspects include how the *support* and *bias scores* of multiple biased concepts can be composed, providing interesting insights about possible correlations, as well as how the *bias* scores of domains change over time. Finally, our vision is to integrate the approach in a large scale WSE/social platform, in order to study how users define bias, create a globally accepted gold standard of biased concepts, and explore how such tools can affect the consumption of biased information. In this way, we will be able to evaluate and tune our approach to real-life scenarios, and mitigate any performance issues. Also, by conducting a large scale user study and releasing the plugins to the public, we will gather useful insights about how various users define bias and biased concepts. At the same time, based on the gathered user analytics we will be able to study how such tools affect online user behavior and the consumption of biased information.

References

1. P. Papadakos and G. Konstantakis. bias goggles - Graph-based Computation of the Bias of Web Domains through the Eyes of Users. *ECIR: 42th European Conference on Information Retrieval*, 2020.
2. E. Pariser. *The filter bubble: What the Internet is hiding from you*. Penguin UK, 2011.
3. E. Pitoura, P. Tsaparas, G. Flouris, I. Fundulaki, P. Papadakos, S. Abiteboul, and G. Weikum. On measuring bias in online information. *ACM SIGMOD Record*, 46(4):16–21, 2018.