Did You Find It on the Internet? Ethical Complexities of Search Engine Rankings

Cansu Canca

Abstract Search engines play a crucial role in our access to information. Their search ranking can amplify certain information while making others virtually invisible. Ethical issues arise regarding the criteria that the ranking is based on, the structure of the resulting ranking, and its implications. Critics often put forth a collection of commonly held values and principles, arguing that these provide the needed guidance for ethical search engines. However, these values and principles are often in tension with one another and lead us to incompatible criteria and results, as I show in this short chapter. We need a more rigorous public debate that goes beyond principles and engages with necessary value trade-offs.

1 Introduction

In our digitalized life, search engines function as the gatekeepers and the main interface for information. Ethical aspects of search engines are discussed at length in the academic discourse. One such aspect is the *search engine bias*, an issue that encompasses ethical concerns about search engine rankings being value-laden, favoring certain results, and using non-objective criteria (Tavani 2020). The academic ethics debate has not (yet) converged to a widely accepted resolution of this complex issue. Meanwhile, the mainstream public debate mostly ignored the hard ethical trade-offs, opting instead for a collection of commonly held principles and values such as accuracy, equality, efficiency, and democracy. In this short chapter, I explain why this approach leads to unresolvable conflicts and thus ultimately a dead end.

This chapter builds on the *Mapping* workshop (<https://aiethicslab.com/the-mapping/>) I designed and developed together with Laura Haaber Ihle at AI Ethics Lab. I thank Laura also for her insights and valuable feedback on this chapter.

The original version of this chapter was revised with an update in the affiliation of the author Cansu Canca. A correction to this chapter can be found at https://doi.org/10.1007/978-3-030-86144-5_47

C. Canca (✉)
AI Ethics Lab, Cambridge, MA, USA
e-mail: cansu@aiethicslab.com

2 Value of and Value Within Search Engines

Search engines are invaluable. Without them, it is impossible to navigate the massive amounts of information available in the digital world. They are our main mediator of information (CoE 2018). Every day over 7 billion searches are conducted on Google alone,¹ accounting for over 85% of all searches worldwide.² In addition to that, various searches are conducted on specialized search engines such as Amazon and YouTube. Even academic and scientific research relies on Google Scholar, PubMed, JSTOR, and similar specialized search engines—meaning not only we rely on search engines to access information but we also rely on them while creating further knowledge.

We defer to search engines' ranking of information so much that most people do not even check the second page of the search results.³ On Google, 28.5% of users click the first result, while 2.5% click on the tenth result, and fewer than 1% click results on the second page.⁴ This means that the ranking has a great effect on what people see and learn.

Search engine ranking is never value-neutral or free from human interference. It is optimized to reach a goal. Even if we can agree that this goal is *relevance* to the user, defining what is relevant involves guesswork and value judgments. By definition, in most search queries, users do not know what result they are looking for. On top of that, the search engine has to guess from the search keywords what *kind* of information the user is interested in. For example, a user searching for “corona vaccination” could be looking for vaccine options, vaccine safety information, anti-vaccine opinions, vaccination rates, or celebrities who are vaccinated, and they might be looking for these on a global or local scale. More importantly, they might be equally satisfied with well-explained vaccine safety information or anti-vaccine opinions since they might not have prior reasons to differentiate these two opposing results. Here, the value judgment comes into play in designing the system. Should the system first show vaccine safety information to ensure that the user is well-informed or the anti-vaccine opinions since they are often more intriguing and engaging? Should the system make results depending on user profiles (e.g., being scientifically or conspiracy oriented)? Should it sort the results by click rates globally or locally, by personal preferences of the user, by accuracy of the information, by balancing opinions, or by public safety? Deciding which ranking to present embeds a value judgment into the search engine. And this decision cannot fully rely on evaluating user satisfaction about a search query, because the user does not know the full range

¹<https://www.internetlivestats.com/google-search-statistics/>

²<https://www.statista.com/topics/1710/search-engine-usage/-dossierSummary>

³<https://www.forbes.com/sites/forbesagencycouncil/2017/10/30/the-value-of-search-results-rankings/>

⁴<https://www.sistrix.com/blog/why-almost-everything-you-knew-about-google-ctr-is-no-longer-valid/>

of information they could have been shown. Moreover, user satisfaction might still lead to unethical outcomes.

3 Ethical Importance of Search Engine Rankings

Imagine basing your decision whether to get vaccinated on false information because that is what came up on your search (Ghezzi et al. 2020; Johnson et al. 2020). Imagine deciding whom to vote for based on conspiracy theories (Epstein and Robertson 2015). Imagine having your perception of other races and genders pushed to negative extremes because that is the stereotype you are presented with online (Noble 2018). Imagine searching for a job but not seeing any high pay and high power positions because of the search engine's knowledge or estimate of your race, gender, or socio-economic background.⁵ Imagine having your CV never appear to the employers for those jobs that you easily qualify for because of search engine profiling (Deshpande et al. 2020). These are all ethical issues. They stem from value judgments embedded in search engine processing and, as a result, impacting individual autonomy, well-being, and social justice.

We base our decisions on what we know. By selecting which information to present in what order, search engines affect and shape our perception, knowledge, decision, and behavior both in the digital and physical sphere. As a result, they can manipulate or nudge individuals' decisions, interfere with their autonomy, and affect their well-being.

By sorting and ranking what is out there in the digital world, search engines also impact how benefits and burdens are distributed within the society. In a world where we search online for jobs, employees, bank credits, schools, houses, plane tickets, and products, search engines play a significant role in the distribution of opportunities and resources. When certain goods or opportunities are systematically concealed from some of us, there is an obvious problem of equality, equal opportunity and resources, and, more generally, fairness. More to the point, once certain information is not accessible to some, they often do not even know that it exists. If we cannot realize the injustice, we also cannot demand a fair treatment.

4 Do You See Female Professors?

A running example for search engine bias has been the image search results for the term "professor."⁶ When searching for "professor," search engines present a set of overwhelmingly white male images. In the USA, for instance, women make up 42%

⁵<https://www.nytimes.com/2018/09/18/business/economy/facebook-job-ads.html>

⁶<https://therepresentationproject.org/search-engine-bias/>

of all tenured or tenure-track faculty.⁷ In search engine results, the ratio of female images has been about 10–15% and only recently went up to 20–22%.⁸ When searching specifically for “female professors,” the image results are accompanied by unflattering suggestions: Google’s first suggested term is “clipart” (Fig. 1), whereas Bing’s top four suggestions include “crazy,” “cartoon,” and “clipart” female professors (Fig. 2).⁹

Why is this an ethical problem? Studies show that girls are more likely to choose a field of study if they have female role models (Lockwood 2006; Porter and Serra 2020). Studies also show that gender stereotypes have a negative effect on hiring women for positions that are traditionally held by men (Isaac et al. 2009; Rice and Barth 2016). By amplifying existing stereotypes, search engine results contribute to the existing gender imbalance in high-powered positions. It is reasonable to think that this gender imbalance in real life has its roots in unjustified discrimination against women in the workplace as well as discriminatory social structures, both of which do not allow female talent to climb the career ladder. The search engine bias can contribute to the perpetuation of this gender imbalance.

In fact, this is not special to image search results for “professor.” Women are underrepresented across job images and especially in high-powered positions (Kay et al. 2015; Lam et al. 2018). Take one step further in this problem, and we end up with issues such as LinkedIn—a platform for professional networking—autocorrecting female names to male ones in its search function¹⁰ and Google showing much fewer prestigious job ads to women than to men.¹¹

Most mainstream reactions criticize search engine rankings from commonly held values and principles: search engines should reflect the truth and be fair; they should promote equality and be useful; they should allow users to exercise agency and prevent harm; and so on. On close inspection, however, these commonly held values and principles fail to provide guidance and may even conflict with one another. In the next paragraphs, I briefly go over three values—accuracy, equality, and agency—to show how such simple guidance is inadequate for responding to this complex problem.

Accuracy One could argue that search engines, being a platform for information, should accurately reflect the world as it is. This would imply that the image results

⁷<https://www.aaup.org/news/data-snapshot-full-time-women-faculty-and-faculty-color>.

⁸Calculated from search engine results of Google, Bing, and DuckDuckGo in May 2021. Note that these numbers fluctuate depending on various factors about the user such as their region and prior search results.

⁹In comparison, Bing’s suggestions when searching for “male professor” remain within the professional realm: “black professor,” “old male teacher,” “black female professor,” “black woman professor,” “professor student,” and such. However, Google fails equally badly in its suggestions for “male professor.” Its first two suggestions are “stock photo” and “model.”

¹⁰<https://phys.org/news/2016-09-linkedin-gender-bias.html>

¹¹<https://www.washingtonpost.com/news/the-intersect/wp/2015/07/06/googles-algorithm-shows-prestigious-job-ads-to-men-but-not-to-women-heres-why-that-should-worry-you/>

Fig. 1 Female professors and Google

Fig. 2 Female professors and Bing

should be revised and continuously updated to reflect the real-life gender ratio of professors. In contrast, the current search results portray the social perception about gender roles and prestigious jobs. Note that implementing accuracy in search results would require determining the scope of information: Should the results accurately reflect local or global ratio? And what other variables—such as race, ability, age, and socio-economic background—and their ratio should the results reflect accurately?

Equality Contesting prioritization of accuracy, one could argue that search engines should promote equality because they shape perception while presenting information, and simply showing the status quo despite the embedded inequalities would be unfair. If we interpret equality as equal representation, this would imply showing equal number of male and female professors. Implementing equal representation in search results would also require taking into account other abovementioned variables—such as race, ability, age, and socio-economic background—and equally representing all their combinations. A crucial question would then be, what would the search results look like if all possible identities are represented and would these results still be relevant, useful, or informative for the user?

Agency Contesting both accuracy and equality, one could argue that the system should prioritize user agency and choice by ranking the most clicked results at the top. This is not a strong argument. When conducting a search, users do not convey an informed and intentional choice through their various clicks. One could, however, incorporate user settings to the search engine interface to encourage user agency and provide them with a catalogue of setting options for ranking. Yet, since most people have psychological inertia and status quo bias, most users would still end up using the default (or easiest) setting—which brings us back to the initial question: What should the default ranking be?

An additional consideration must be the content of web pages that these images come from. It is not sufficient to have an ethically justifiable combination of images in the search results. It is also important that the web pages that link to these images follow a similar ethical framework. If, for example, search results show equal number of female and male images but the pages with female images dispute gender equality, this would not be a triumph of the principle of equality.

We could continue with other values such as well-being, efficiency, and democracy. They would yield even more different and conflicting ranking outcomes. Therein lies the problem. These important and commonly held values do not provide a straightforward answer. They are often in tension with each other. We all want to promote our commonly cherished and widely agreed-upon values. This is apparent from the Universal Declaration of Human Rights, from the principlism framework, and from the common threads within various sets of AI principles published around the world.¹² But simply putting forth some or all of these values and principles and

¹²<https://aiethicslab.com/big-picture/>

demanding them to be fulfilled is an unreasonable and impossible request, which do not get us very far in most cases (Canca 2019, 2020).

5 The Process and the End Product

Thus far I focused on the end product. What is the ethically justified composition for search engine results? But we also need to focus on the process: how did we end up with the current results, and what changes can or should be implemented?

Search engines use a combination of proxies for “relevance.” In addition to the keyword match, this might include, for example, click rates, tags and indexing, page date, placing of the term on the website, page links, expertise of the source, and user location and settings.¹³ Search is a complex task, and the way that these proxies fit together changes all the time. Google reports that in 2020 they conducted over 600,000 experiments to improve their search engine and made more than 4500 changes.¹⁴ Since search engine companies compete in providing the most satisfactory user experience, they do not disclose their algorithms.¹⁵ Returning to our example, when we compare image search results for “professor” in Google, Bing, and DuckDuckGo, we see that while Google prioritizes Wikipedia image as the top result, Bing and DuckDuckGo refer to news outlet and blog images for their first ten results, excluding the image from the Wikipedia page.¹⁶

Value judgments occur in deciding which proxies to use and how to weigh them. Ensuring that the algorithm does not fall into ethical pitfalls requires navigating existing and expected ethical issues. Going back to our example, content uploaded and tagged by users or developers is likely to carry their implicit gender biases. Therefore, it is reasonable to assume that the pool of images tagged as “professor” would be highly white-male oriented to start with. Without any intervention, this imbalance is likely to get worse when users click on the results that match their implicit bias and/or when an algorithm tries predicting user choice and, thereby, user bias.

¹³<https://www.google.com/search/howsearchworks/algorithms/>; <https://www.bing.com/webmasters/help/webmasters-guidelines-30fba23a>; <https://dl.acm.org/doi/abs/10.1145/3361570.3361578>

¹⁴<https://www.google.com/search/howsearchworks/mission/users/>

¹⁵Even if they did disclose their algorithms, this would likely be both extremely inefficient and ethically problematic. See Grimmelmann (2010) for a more detailed discussion of transparency in search engines.

¹⁶At the height of these discussions in 2019, the image in the Wikipedia entry for “professor” was switched to Toni Morrison, a black female professor, thereby making this the first image to appear on Google image search results. Currently (May 2021), the Wikipedia image has been changed to a picture of Einstein.

6 Conclusions

A comprehensive ethical approach to search engine rankings must take into account search engines' impact on individuals and society. The question is how to mitigate existing ethical issues in search engine processing and prevent amplifying them or creating new ones through user behavior and/or system structure. In doing so, values and principles can help us formulate and clarify the ethical problems at hand, but they cannot guide us to a solution. For that, we need to engage in deeper ethics analyses which provide insights to the value trade-offs and competing demands that we must navigate to implement any solutions to these problems. These ethics analyses should then feed into the public debate so that the discussion can go beyond initial arguments, reveal society's value trade-offs, and be informative for decision-makers.

References

- Canca, C. (2019) 'Human Rights and AI Ethics – Why Ethics Cannot be Replaced by the UDHR', in *AI & Global Governance*, United Nations University Centre for Policy Research [online]. Available at: <https://cpr.unu.edu/publications/articles/ai-global-governance-human-rights-and-ai-ethics-why-ethics-cannot-be-replaced-by-the-udhr.html> (Accessed: 1 May 2021).
- Canca, C. (2020) 'Operationalizing AI Ethics Principles', *Communications of the ACM*, 63(12), pp.18–21. <https://doi.org/10.1145/3430368>
- Council of Europe (CoE) (2018) *Algorithms and Human Rights* [online]. Available at: <https://rm.coe.int/algorithms-and-human-rights-en-rev/16807956b5> (Accessed: 1 May 2021)
- Deshpande, K.V., Shimei, P. and Foulds, J.R. (2020) 'Mitigating Demographic Bias in AI-based Resume Filtering', in *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization*. <https://doi.org/10.1145/3386392.3399569>
- Epstein, R. and Robertson R.E. (2015) 'The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections', *Proceedings of the National Academy of Sciences*, 112(33). <https://doi.org/10.1073/pnas.1419828112>
- Grimmelmann, J. (2010) 'Some Skepticism About Search Neutrality' [online]. Available at: <http://james.grimmelmann.net/essays/SearchNeutrality> (Accessed: 1 May 2021)
- Ghezzi, P., Bannister, P.G., Casino, G., Catalani, A., Goldman, M., Morley, J., Neunez, M., Prados-Bo, A., Smeesters, P.R., Taddeo, M., Vanzolini, T. and Floridi, L. (2020) 'Online Information of Vaccines: Information Quality, Not Only Privacy, Is an Ethical Responsibility of Search Engines', *Frontiers in Medicine*, 7(400). <https://doi.org/10.3389/fmed.2020.00400>
- Isaac, C., Lee, B. and Carnes, M. (2009) 'Interventions that affect gender bias in hiring: a systematic review', *Academic medicine: Journal of the Association of American Medical Colleges*, 84(10), pp.1440–1446. <https://doi.org/10.1097/ACM.0b013e3181b6ba00>
- Johnson, N.F., Velásquez, N., Restrepo, N.J., Leahy, R., Gabriel, N., El Oud, S., Zheng, M., Manrique, P., Wuchty, S. and Lupu, Y. (2020) 'The online competition between pro- and anti-vaccination views', *Nature*, 582, pp.230–233. <https://doi.org/10.1038/s41586-020-2281-1>
- Kay, M., Matuszek, C. and 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*. <https://doi.org/10.1145/2702123.2702520>

- Lam, O., Wojcik, S., Broderick, B. and Hughes, A. (2018) 'Gender and Jobs in Online Image Searches', *Pew Research Center* [online]. Available at: <https://www.pewresearch.org/social-trends/2018/12/17/gender-and-jobs-in-online-image-searches/> (Accessed: 1 May 2021)
- Lockwood, P. (2006) 'Someone Like Me can be Successful: Do College Students Need Same-Gender Role Models?', *Psychology of Women Quarterly*, 30(1), pp.36–46. <https://doi.org/10.1111/j.1471-6402.2006.00260.x>
- Noble, S.U. (2018) *Algorithms of Oppression: How Search Engines Reinforce Racism*, New York: New York University Press.
- Porter, C. and Serra, D. (2020) 'Gender Differences in the Choice of Major: The Importance of Female Role Models', *American Economic Journal: Applied Economics*, 12(3), pp.226–254. <https://doi.org/10.1257/app.20180426>
- Rice, L. and Barth, J.M. (2016) 'Hiring Decisions: The Effect of Evaluator Gender and Gender Stereotype Characteristics on the Evaluation of Job Applicants', *Gender Issues* 33, pp.1–21. <https://doi.org/10.1007/s12147-015-9143-4>
- Tavani, H. (2020) 'Search Engines and Ethics', *The Stanford Encyclopedia of Philosophy*, Fall 2020 Edition in Zalta, E.N. (ed.) [online]. Available at: <https://plato.stanford.edu/archives/fall2020/entries/ethics-search/> (Accessed: 1 May 2021)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.