IS GOOGLE REPSONSIBLE FOR PROVIDING FAIR AND UNBIASED RESULTS?

Dirk Lewandowski, Hamburg University of Applied Sciences, Germany
dirk.lewandowski@haw-hamburg.de

This is a preprint of an article accepted for publication.

Lewandowski, D.: Is Google responsible for providing fair and unbiased results? In: Floridi, L.; Taddeo, M. (eds.): The Responsibilities of Online Service Providers. Berlin Heidelberg: Springer, 2017. S. 61-77. doi: 10.1007/978-3-319-47852-4_4

ABSTRACT

This chapter discusses the responsibilities of Google as the leading search engine provider to provide fair and unbiased results. In its role, Google has a large influence on what is actually searchable on the Web as well as what results users get to see when they search for information. Google serves billions of queries per month, and users only seldom consider alternatives to this search engine. This market dominance further exacerbates the situation. This leads to questions regarding the responsibility of search engines in general, and Google in particular, for providing fair and balanced results. Areas to consider here are (1) the inclusion of documents in the search engine's databases and (2) results ranking and presentation. I find that, while search engines should at least be held responsible for their practices regarding indexing, results ranking, delivering results from collections built by the search engine provider itself and the presentation of search engine results pages; today's dominant player, Google, argues that there actually is no problem with these issues. Its basic argument here is that "competition is one click away", and, therefore, it should be treated like any other smaller search engine company. I approach the topic from two standpoints: from a technical standpoint, I will discuss techniques and algorithms from information retrieval and how decisions made in the design of the algorithms influence what we as users get to see in search engines. From a societal standpoint, I will discuss what biased search engines mean for knowledge acquisition in society and how we can overcome today's unwanted search monopoly.

INTRODUCTION

In this chapter, I discuss Google's role as the dominant search engine on the market and responsibilities that could derive from this market position. There are many responsibilities

that could be discussed in the context of Google (e.g., whether it has responsibilities deriving from it collecting its users' query data), or Online Service Providers in general (see Taddeo & Floridi 2015). I will focus on the results Google provides. I will discuss these results concerning fairness and biases.

First of all, a search engine in the context of this chapter is defined as a computer system that collects content distributed over the Web through crawling, orders the results to a query by machine-determined relevance, and makes these results available to its users through a user interface.

There is a vast body of research on techniques and technologies to improve search engines, on measuring the quality of results of search engines, on the behaviour of search engine users, and on the role search engines play for knowledge acquisition in society. This research is embedded in the wider context of research on the role of algorithms in knowledge acquisition and search engines as socio-technical systems. It is important to understand the decisions made by search engines through their algorithms, as the algorithmic approach to finding information can be seen as "a particular knowledge logic, one built on specific presumptions about what knowledge is and how one should identify its most relevant components. That we are now turning to algorithms to identify what we need to know is as momentous as having relied on credentialed experts, the scientific method, common sense, or the word of God" (Gillespie 2014, p. 168).

The main argument brought forward in this chapter is that every search engine produces *biased results* in some way, resulting from Web crawling, indexing, and results ranking. As there is no perfect or correct results set or ranking, search engine results are always a product of the algorithmic interpretation of the Web's content by the given search engine. Nevertheless, a search engine can still provide *fair results* when there is no preferential treatment of information objects, neither in the process of indexing nor in the process of ranking.

The remainder of this chapter is structured as follows: first, I will elaborate on Google's role as the world's dominant search engine and how Google as a company sees its responsibility for providing fair and unbiased results. Then, I will define the central concepts used in this chapter, namely fair results and unbiased results. I will further discuss the related concepts. I will then focus on the search engines' databases (the indexes) and show how already in building the index, search engines make decisions on which results they will later be able to produce. Then, I will focus on what I call the "algorithmic interpretation of the Web's contents" and how different forms of interpretation shape the results a user gets to see when

using a given search engine. After that, I will discuss responsibility issues related to indexing and ranking (or, more general, producing results). I will conclude the chapter with a summary and some suggestions for further research.

WHY GOOGLE?

First of all, the importance of search engines for finding information could not be overestimated. Not only are fully-automated search engines like Google the dominant means for finding information on the Web and have made all other approaches to finding content on the Web (like Web directories) nearly obsolete, but more and more information is searched on the Web nowadays instead of other sources outside the Web. While other information sources like social networking sites are sometimes seen as competitors to search engines, as users are directed to information objects through messages displayed there, they do not qualify for *ad hoc* searches, i.e. where a user actually queries an information system in order to find information objects related to his or her information need. Furthermore, when looking at the query volume that search engines process (cf. "Stats: comScore" 2015), we find that search engines not only respond to billions of queries per day, but the query volume is nowhere near declining.

Nearly everybody who uses the Internet also uses search engines (Purcell, Brenner, and Raine 2012). Searching for information is one of the most popular activities on the Internet. On average, European users issue 138 queries per month (comScore 2013). Google's market share is 86% in Europe (comScore 2013) – including eastern European countries, where Yandex has a large market share – with many countries reporting Google's market share well over 90%. Users predominantly relying on one search engine leads to certain problems regarding bias and fairness or at least increases problems resulting from biases inherent in every search engine and from search engine provider's decisions on the fairness of the results and their presentation.

When looking at public statements made by Google, we find that this search engine has a clear view on what its position on the market is and how it should deal with results ranking and transparency related to the rankings. This position can be summarised as follows:

1. *There is competition on the search market, and users can decide to use another search engine without any problem*. In the words of Amit Singhal, Senior Vice President Search at Google: "the competition is only one click away. […] Using Google is a choice—and there are lots of other choices available to you for getting information: other general-

interest search engines, specialized search engines, direct navigation to websites, mobile applications, social networks, and more"[1].

2. *Google generates its results purely through its algorithms, and does not manually interfere with results generated by these*: "No discussion of Google's ranking would be complete without asking the common - but misguided! :) - question: "Does Google manually edit its results?" Let me just answer that with our third philosophy: no manual intervention".[2] And he gives the following reason: "If we messed with results in a way that didn't serve our users' interests, they would and should simply go elsewhere".[3]
3. *Google does not treat its own content preferentially*: "People often ask how we rank our "own" content, like maps, news or images. In the case of images or news, it's not actually Google's content, but rather snippets and links to content offered by publishers. We're merely grouping particular types of content together to make things easier for users. In other cases, we might show you a Google Map when you search for an address. But our users expect that, and we make a point of including competing map services in our search results".[4]
4. *Google is as transparent as possible on how its results are generated*: "Be transparent. We share more information about how our rankings work than any other search engine, through our Webmaster Central site, blog, diagnostic tools, support forum, and YouTube".[5] In another blog post, Singhal says that, "Google's search algorithm is actually one of the world's worst kept secrets".[6] On the other hand, Udi Manber, then Vice President Engineering, Search Quality, said in a blog post: "For something that is used so often by so many people, surprisingly little is known about ranking at Google. This is entirely our fault, and it is by design. We are, to be honest, quite secretive about what we do. There are two reasons for it: competition and abuse".[7]

Much has been written about Google's actual practices, and so I will only summarise some of the findings on Google's practices regarding these statements.

[1] http://googlepublicpolicy.blogspot.de/2011/06/supporting-choice-ensuring-economic.html
[2] https://googleblog.blogspot.de/2008/07/introduction-to-google-ranking.html
[3] http://googlepublicpolicy.blogspot.de/2010/09/competition-in-instant.html
[4] http://googlepublicpolicy.blogspot.de/2010/09/competition-in-instant.html
[5] http://googlepublicpolicy.blogspot.de/2011/06/supporting-choice-ensuring-economic.html
[6] http://googlepublicpolicy.blogspot.de/2010/09/competition-in-instant.html
[7] https://googleblog.blogspot.de/2008/05/introduction-to-google-search-quality.html

Considering the competition on the search engine market (cf. Lewandowski, forthcoming), we can firstly see that Google has been the dominant player on the market for years and that it only faces competition in general-interest search engines from Microsoft's Bing. Either other search engines do not have a large enough index to compete with these two engines or they do not provide their own database at all but instead rely on partnerships with either Google or Bing. Regarding vertical search engines, these are to a large degree accessed through general-interest search engines and thus rely on being ranked high in Google.

Google says that it does not manually interfere with its results. While it is true that there is no simple manipulation in the way that Google would manually adjust the organic results for certain queries (although there has been some doubt in the past; see, Edelman 2010), it does exclude certain results due to law or deliberate choice, or it penalizes information objects for not confirming to its self-defined rules. As these information objects are not considered for ranking, they will not have the chance to be discovered through a Google search. Examples for the exclusion of documents from being found include:

- Exclusion of SPAM documents: Some documents are excluded from Google's index due to being irrelevant and classified as SPAM. While it is necessary for any search engine to take action against spamming, the criteria to qualify a document or a website as SPAM are not transparent.
- Penalties for gaming the system: Google reserves the right to penalise certain documents or websites if it finds that the owners of these were trying to "game the system", i.e., trying to achieve better rankings, e.g., by buying links to their documents. Such penalties have nothing to do with the actual quality of the documents' contents.
- Deliberate choice on how to process "Right to Be Forgotten" (RTBF) queries: There is no clearly defined and transparent process on when documents are not shown due to RTBF requests. One has to admit, however, that the RTBF is relatively new and that it may take some time to establish such a process.
- Exclusion of certain sites from vertical searches: Some vertical search engines are built through focussed crawling, i.e., a process where only content from pre-selected sources is considered for inclusion in the search index. This approach is fundamentally different from Web indexing in general, where a search engine basically crawls all the contents from the Web without humans manually excluding some websites.[8] An example of a vertical search engine that uses focussed crawling is Google News, where humans pre-

[8] It should be noted, however, that there are certain quality thresholds for the inclusion of websites, although whether a website is below such a threshold is determined automatically.

select news sources that are then regularly crawled for new content. This means that if a website is not considered a news source by Google, its documents do not have a chance to be found through a news search (or in the news box in Google Web search).

A further example where Google does not exactly interfere with the results themselves, but with the process that leads a user to the results, is its interference with the autocomplete function. While Google claims that query suggestions are solely based on other users' past queries and determined automatically[9], there are examples where one can easily see that for certain queries, humans at Google have decided that no suggestions should be made or that suggestions should be filtered (Diakopoulos 2013).

These examples show that Google does not function solely on algorithms and that there are human decisions, not only in the design of the algorithms but also in maintaining the search engine. It is a myth that Google does not manually interfere with the results. However, it is in the interest of Google to prolong this myth because, otherwise, information providers whose content is ranked low within Google's results could argue for them to be ranked higher based on an *assumed* better quality. Google will try not to raise any discussions on the actual quality of its results (apart from it being produced by an algorithm that treats each document the same).

Regarding the question of whether Google gives its own content preferential treatment on its *Web search* results pages (and therefore, using its dominant position in Web search to promote its own content or the results from its vertical search engines, respectively), we can see that giving these results a different (and more attractive) layout than results from competitors alone constitutes preferential treatment. Users are not only attracted by the position of a result but also by its graphical design. For instance, if a news box with a result including an image is presented above the fold on a search engine results page, users will be attracted to it to a large degree (see Lewandowski and Sünkler 2013). Therefore, the if-question may be the wrong one. Instead, one should ask whether search engines have a moral responsibility when it comes to their own content. I will discuss this further below.

This brief comparison of Google's statements with its actual practices shows that Google operates on statements that are at least in part contrary to their actual practices. One could simply qualify these statements as public relations, but the point is that in many cases, even scholars use these arguments when discussing search engines and the role of Google. Moving

[9] https://support.google.com/websearch/answer/106230?hl=en

away from current practices, in the next sections, I will define what fair and unbiased results are and whether search engines are able to provide such results.

WHAT ARE FAIR AND UNBIASED RESULTS?

The Oxford English Dictionary gives several definitions for the term "fair", depending on the context:

> "Of conduct, actions, methods, arguments, etc.: free from bias, fraud, or injustice; equitable; legitimate, valid, sound."
>
> "Of conditions, circumstances, etc.: providing an equal chance of success to all; not unduly favourable or adverse to anyone."
>
> "Of remuneration, reward, or recompense: that adequately reflects the work done, service rendered, or injury received."[10]

For our purposes, we can define fair search results as results that are produced in a way where every document on the Web is treated in the same way by the search engine and, therefore, has the same chance of being found and ranked by that search engine and that there are no human interferences with algorithmic decisions on crawling, indexing and ranking.

Bias, then, is,

> "An inclination, leaning, tendency, bent; a preponderating disposition or propensity; predisposition *towards*; predilection; prejudice."[11]

Search engine bias is the tendency of a search engine to prefer certain results through the assumptions inherent in its algorithms. This means that every search engine is biased, as it is impossible to design algorithms without human assumptions. Therefore, search engine bias does not mean that search results are deliberately manipulated by the search engine vendor but simply that results are ordered in a certain way that is determined by assumptions of what constitutes a good or relevant result in response to queries. It is even at the core of every idea of ranking, based on certain technically mediated assumptions, that certain items are preferred over others.

Yet it should be mentioned that there are other definitions of search engine bias that do not define search engines as biased *per se* but focus on the deliberate preference for certain results (as in the case of Google discussed above, when it favours its own content over content from

[10] http://www.oed.com/view/Entry/67704?rskey=FgCXaJ&result=2#eid
[11] http://www.oed.com/view/Entry/18564?rskey=dJWCZ3&result=1#eid

its competitors. Tavani (2012) summarises the three concerns underlying the definitions of search engine bias:

> "(1) search-engine technology is not neutral, but instead has embedded features in its design that favor some values over others; (2) major search engines systematically favor some sites (and some kind of sites) over others in the lists of results they return in response to user search queries; and (3) search algorithms do not use objective criteria in generating their lists of results for search queries." (Tavani 2012)

So, is the question on whether Google is responsible for fair and biased results put the wrong way? If there is no such thing as an unbiased search engine, Google cannot be made responsible for being biased. In my opinion, the bias inherent in search algorithms in fact leads to the severe need for more search engines rather than the demand for Google to reduce or even erase the bias in its search results. I argued for a public search infrastructure (as opposed to alternative search engines) elsewhere (Lewandowski 2014a) and see this as the only solution for dealing with the problem of every search engine being biased by design.

Even if we consider that every search engine is biased by design, Google could still be made responsible for providing fair results. Fair in this sense would be that every information object on the Web has the same chance of being included in Google's database (index) and that every information object in the index has the same chance of being found, solely on the basis of algorithms that treats every information object in the database in the same way.

However, there is an important restriction to this: As search engines not only provide textual documents, but also images, videos etc., they need to treat these different kinds of information objects differently, if only for their different properties. For instance, it is not possible to treat images to be found in the same way by the same algorithm as textual documents. Therefore, it is misleading to speak of *the* index of a search engine, as search engines have multiple indexes for each type of content.

That being said, we must distinguish between fairness in including documents in the indexes and fairness in ranking the results in the indexes in response to a query.

PROVIDING A COMPREHENSIVE AND FRESH INDEX

Search engines make decisions about what to include in their indexes and how often to refresh already known documents. As there is an overwhelming amount of mostly automatically generated content that can be considered SPAM, search engines require technical filters that

allow them to not even consider such documents in the crawling/indexing process, which can lead to unwanted exclusion of documents from the search engine. In addition to filtering SPAM content, search engines also apply filters based on – often country-specific – laws (e.g., the "right to be forgotten" in the European Union) and based on the deliberate choices of the search engine provider (mainly in the context of self-defined rules for the protection of children and young persons and takedown notices from copyright holders).

As the Web is of an immense size and continuously changes (Ntoulas, Cho & Olston, 2004), building and maintaining a comprehensive index is a huge challenge (Patterson 2004). Related to these challenges are issues with comprehensiveness, freshness and the deliberate choice of search engines to exclude certain documents from their indexes. The latter can either be deciding not to index certain documents at all or excluding documents after indexing, i.e., not making them available to users in certain countries or regions.

Issues with comprehensiveness

The first and arguably the most important issue with comprehensive search engine indexes is the size of the Web. While we know that the Web consists of many billions of documents (some years ago, Google even claimed that it knew more than one trillion different URLs[12]), we do not know the exact number, as there is no central registration for URLs on the Web. The best estimates we have are actually from numbers derived through Web crawling (i.e., finding content on the Web through following links), which on a large scale is mainly done by search engines.

Some years ago, search engine companies stopped reporting index sizes. This can be seen as a consequent move, as index sizes do not say much about the quality of a search engine. As there are vast amounts of documents on the Web that a search engine will not want to index (such as copies of the ever-same content and SPAM pages), a search engine having a lot of these documents in its index would surely increase its index size but for no one's benefit.

Then, there is the problem with defining what a document on the Web actually is. One could argue that everything that has an URL should be considered to be a document. However, as such documents can be easily created automatically, and can be built by combining elements from other documents, a lot of documents without any benefit could be (and are) built. This does not have to do with spamming search engines. Consider, for instance, blogs where different kinds of overview pages (such as teasers for all articles from a certain month, teasers

[12] https://googleblog.blogspot.de/2008/07/we-knew-web-was-big.html

for articles tagged with a certain keyword etc.) are generated. We can question whether a search engine should index all these "documents".

However, while this seems to be a purely technical problem, it still comprises decisions about what is worth indexing, and there is no guarantee that no potentially relevant document will slip through. Presumably due to the structure of the Web, certain content, e.g., from certain countries, is not as well represented as content from other countries (cf. Vaughan and Zhang 2007).

Furthermore, search engines have technical and financial limitations regarding index sizes: Even if a search engine wanted to build a complete index of the Web, it would still face limitations due to its technical possibilities and financial resources. No search engine is able to index the Web in its entirety. The problem arising from that, however, may not be the lack of completeness but the lack of transparency regarding the criteria that lead a search engine to index certain documents and exclude others.

Issues with freshness

Apart from the issues of building a comprehensive index of the Web, search engines must also keep up with the ever-changing Web. New documents are created, existing documents are changed and documents are deleted. The issue related to these Web dynamics is twofold: Firstly, search engines have to make sure to index documents afresh and, secondly, provide fresh results through ranking. It would be a bad idea for a search engine to present a user with a result description that points him or her to a page that no longer exists.

The issue with freshness is that no search engine can keep all the documents in its index current. On the one hand, no search engine could afford to crawl every document every second. On the other hand, even if a search engine would be able to do so, this would account for too much bandwidth and would send too many requests to Web servers. The approach that search engines take is to decide which documents to revisit when based on popularity and on the refreshment rate of each document in the past (Risvik and Michelsen 2002). On the one hand, this leads to a technically feasible solution. On the other hand, decisions about freshness (i.e., which documents to index more frequently) may lead to fairness issues. Preferring popular and/or often refreshed content is a decision that could lead to the oppression of other documents in the results sets.

Issues with deliberate choices made by the search engines

Search engine providers also make deliberate choices about documents to exclude from their indexes, sometimes not actually excluding them, but removing them from the results in certain regions or countries. The prime example for this is the "Right to Be Forgotten" (RTBF) in European legislation. Persons can request for certain results to be removed from search results if these results refer to the person's past that is no longer occurring. A problem with the RTBF may lie in it not being precise as to when such data should be removed. However, the issue with search engines and the RTBF lies in that the procedure for having content removed is not transparent. While Google provides information about how many requests it received and how it decided as well as established an advisory council on the topic, there are still no transparent rules on how these requests are treated. So, some documents may have been removed even though they do not fall under the RTBF, while others may have been removed without actually falling under the RTBF.

Very similar is the case of takedown notices by copyright holders. Here also, Google must process a large number of requests, but this time, it mainly processes them automatically. This could lead to documents taken down erroneously, simply due to the sheer volume of these requests and standard procedures to treat them (Karaganis and Urban 2015).

A third area of concern is the protection of children and young persons. At the request of authorities (e.g., the *Bundesprüfstelle für jugendgefährdende Medien* in Germany), Google removes websites from its search results no matter if a child or adult is searching for that content. Furthermore, adjusting rankings in a way that prefers non-offensive content is also a decision on what constitutes a document potentially relevant in response to a query (see the section in rankings below).

From this brief discussion of the RTBF, takedown notices and the protection of minors, we can see how Google and other search engines make decisions about which documents to include in their results sets that are opaque to their users. In contrast to the technical and financial issues related to index comprehensiveness and index freshness, these decisions are, even though they are founded on law, deliberate decisions made by the search engines and, therefore, moral decisions. Mainly due to technical reasons, we cannot expect a search engine to provide a complete and fresh copy of the Web. However, what we could expect from a search engine is to make transparent how its index is built, what is left out and for what reasons.

SEARCH ENGINES' ALGORITHMIC INTERPRETATION OF THE WEB'S CONTENT

For every query, a search engine – through its algorithms – presents certain results in a certain way. We can call this an "algorithmic interpretation" of the Web's content (Lewandowski 2015a), and users tend to follow this interpretation uncritically (e.g., Pan et al. 2007; Purcell, Brenner, and Raine 2012). Algorithmic interpretation does not only consider the ranking of the results lists (Pan et al. 2007) but also the positioning of elements on the search engine results pages (Lewandowski and Sünkler 2013; Liu et al. 2015), the correctness of results (e.g., (White and Horvitz 2009), the labelling of advertisements and the diversity within the top results (Denecke 2012).

Due to search engines presenting different kinds of results from different collections within search engine results pages (so-called "Universal Search"), the ranking of these results is not only list-based anymore but consists of at least three different ranking functions: (1) Ranking of the results lists from the Web index, (2) ranking of vertical results within collections such as news, images, etc. and (3) ranking of Universal Search containers (i.e., boxes presenting top results from vertical search engines) within the general SERP or the list of ranked Web search results, respectively.

Ranking within collections (whether considering the Web collection or vertical collections) is based on groups of ranking factors as follows (cf. Lewandowski 2015b, p. 91-92):

(1) Text statistics: This is where the query is matched with the representation of the information objects and statistics are applied to rank documents according to their fit with the query. As queries on the Web are usually very short and there is no standard quality for documents found on the Web (opposed to documents in a curated database, like the electronic archive of a newspaper), text statistics alone are not sufficient for ranking Web documents. They must be accompanied by so-called quality factors.

(2) Popularity: Link-based ranking algorithms (e.g., PageRank) as well as click-based algorithms assign popularity scores to documents and use these scores for a ranking based on the assumption that what has been useful to others will also be useful to a given user. Popularity-based algorithms can either take all users into account or only a certain group of users.

(3) Freshness: As new content is produced in large amounts on the Web, the freshness of information objects is also considered in ranking algorithms.

(4) Locality: Information objects are matched to the geographical location of the user.

(5) Personalisation: Information objects are matched to the interests of an individual user, mainly based on his past behaviour (e.g., queries entered, results viewed).

(6) Technical ranking factors: These factors are mainly used to determine how reliable Web servers are in providing results to the user. As a search engine in most cases only links to information objects from external sources, a user clicking on a result on the search engine results page (SERP) will have to wait for the information object to be produced by the server. Search engines take into account how fast a server is able to process requests and how reliable it is (concerning downtimes). Technical ranking factors are an interesting case, as they judge where an information object should be displayed in a results set, not based on the assumed quality of the content of the information object but rather on the convenience for the user to get to the information object.

To understand what types of results search engines prefer, it is important to consider that the popularity group is considered one of the most important to determine the quality of individual documents. This means that while search engines try to measure such things as credibility, this can only be simulated through measuring popularity (cf. Lewandowski 2012).

The ranking of search results is often misinterpreted as either correct or wrong. I argue that this is mainly due to users having success with navigational queries (where they search for a certain website and the aim is to navigate to that website), with transactional queries where they already have a certain website in mind and with informational queries where they either only search for trivia that can be found on a multitude of websites or they already have a website providing that information in mind (such as Wikipedia). Based on their experience of success for these queries, they assume that their favourite search engine will also produce "correct" results for other types of – mainly informational – queries where there may be many relevant results, and, therefore, no single correct ordering of these results (Lewandowski 2014b).

The basis for algorithmic interpretation is the assumptions that search engine engineers put into these algorithms. Little research has been conducted on the motivations, beliefs and assumptions of this group of people. However, the research already done shows that engineers (and other search engine employees) see search engines as rather purely technical systems, and they conform to a capitalistic way of looking at them (van Couvering 2007; Mager 2012).

The effects of algorithmic interpretation can be seen on different levels. Most obviously, algorithmic interpretation affects the ranking of the organic results, i.e., the results produced

from the basic Web index. Here, every information object included in the index is treated in the same way by the same ranking algorithms, i.e., the results are ranked in a fair way, although not without bias towards the information objects that fulfil assumptions inherent in the ranking algorithms. The effect of algorithmic interpretation for organic results can best be seen when considering drastic examples like the infamous martinlutherking.org website (Piper 2000) that, even in 2015, still ranks very well in Google, and the results produced by Google in response to queries related to race and gender (Noble 2013; Noble 2012). These examples also show that a good ranking position in Google does not necessarily conform with a result being credible or trustworthy (Lewandowski 2012).

Applying certain algorithms can also lead to search engines presenting one side of an argument or only the results of a certain type or tendency. Some algorithms not only try to rank results according to relevance but also mix different result types within the top results to achieve diversity (Giunchiglia et al. 2009).

Algorithmic interpretation also affects the composition of results pages from different indexes ("Universal Search"). For instance, in addition to results from the Web index, results from vertical indexes like news, images and videos can be included in the results pages. This leads to a manifold ranking: Firstly, the results within each vertical index and the main Web index must be ranked. Then, the top results from these indexes must be incorporated into one search engine results page, where results from the vertical searches are to be positioned, and which is another form of ranking. As the presentation of results on the SERP (and even more importantly, in the area "above the fold" of the SERP) heavily influences users' decision on what results to select, search engines are able to lead users to certain types of results merely through results presentation. An important example for this is Google presenting results from its own vertical search engines (such as Google News, Google Scholar and Google Maps) as attractive boxes within the SERP, which then preferably leads to users clicking on them (Lewandowski and Sünkler 2013).

The personalisation of search results is another form of algorithmic interpretation, this time also related to a user's preferences and interests. Results are then produced according to these assumed preferences, mostly without the user knowing what data about him or her is actually collected and how the use of this data affects his or her results. In extreme cases, personalisation can lead to what Eli Pariser termed the "filter bubble" (Pariser 2011), where information objects presenting contradicting views and beliefs from the users are oppressed, and the users only receive results confirming their already established opinions.

Last but not least, search engines present text-based, contextual advertisements on the SERPs. These can be seen as a distinct type of result, and the often-used term "sponsored link" may describe them best: They are a type of result but different from organic results in that they are paid for. Studies lead to the conclusion that users are not able to properly distinguish between organic results and advertisements (Bundesverband Digitale Wirtschaft 2009; Filistrucchi et al. 2012), and, in the case of Google, ad labelling has not become clearer in recent years (Edelman 2014).

It should also be mentioned that apart from the assumptions underlying search engines' algorithms, there is also external influence on the results, namely in the form of search engine optimisation (SEO). The aim of SEO is to optimise information objects in a way that leads to optimal findability through search engines, mainly Google. Search engine optimisation has grown to be a billion-dollar industry, and, at least for queries assumed to have a commercial intent, it will be difficult to find top results in Google that have not been optimised.

While it is common knowledge in the industry and academia that search results are heavily influenced by search engine optimisers, users' knowledge about these practices seems to be low. Furthermore, we see that users generally know little about search engines' workings in general (see, e.g., Purcell, Brenner, and Raine 2012). They often have misconceptions about how a search engine actually works (e.g., Hendry and Efthimiadis 2008), they are not good at searching and they lack knowledge about search engines' ways of making money. On the other hand, they trust in Google's rankings when it comes to results quality (Keane, O'Brien, and Smyth 2008; Bar-Ilan et al. 2009), sometimes even more than their own judgments (Pan et al. 2007).

RESPONSIBILITIES

As we can see from the discussion above, there are multiple areas where we can ask for the responsibility of search engines, especially Google as the dominant player on the market. While we cannot expect Google to provide unbiased results, since we can see that search engine rankings are biased *per se*, we can expect Google to give every information object in its index a fair chance of being ranked in response to a query. "Fair results" here would mean that every information object is treated in the same way. This leads to the conclusion that Universal Search is an unfair treatment to certain results as soon as Google presents results from its own offerings preferentially.

We can also demand for Google to be transparent about its practices, be it the sources its vertical results come from and why they are given preferential treatment or the labelling of its

advertisements. While information on both can be found on Google's help pages, we can see in practice that users do not understand – or are not interested in – the workings behind the composition of search engine results pages. This may be seen as the users' own fault, but in its current practices, Google at least accepts that users are deceived about the true reasons for the display of certain results (types).

CONCLUSION

We cannot expect a search engine to provide fair and unbiased results. Every search engine is per definition biased in that it is not able to provide sets of correct results, as separated from "incorrect" results. Correct results can only be provided for a subset of queries, mostly navigational queries (Broder 2002; Lewandowski 2014b). With informational queries, search engines can provide *relevant* results. However, as relevance always refers to a given user in a given context (Saracevic 2015), a search engine can only make more or less good guesses to what a user may find relevant in his or her current context.

Before producing a ranked results set in response to a query, a search engine must build an index from content found on the Web. A problem here is the size of the Web and its dynamics. Due to the vast amount of information objects on the Web, search engines produce more results than a user is able to consider for most queries. This means that users must trust the ranking provided by the search engine. However, this does not mean that there may be no additional relevant results (or even results being more relevant to a given user in a given context) on lower positions of the results lists.

So, even if Google treated all information objects in a fair manner, users would still see only a fraction of the relevant results available. And as all ranking algorithms organise results in a certain order based on assumptions about what is relevant to users, results from different search engines could differ considerably (or may not even overlap at all) without the results from one search engine being less relevant than the others.

This leads to the conclusion that to release us from only one (or considering the current competition, from a few) of many possible algorithmic interpretations of the Web's content, we need more search engines. With "more", I do not mean just one or two more search engines but a considerable amount of them. One way to achieve this is to view Web indexing as a public service to be provided for the good of all and then have services built upon that infrastructure (Lewandowski 2014a).

Further research is needed on the actual differences of the algorithmic interpretation by different search engines. While some empirical studies already determined overlaps between

search engine results (e.g., Spink, Jansen, Blakely and Koshman, 2006), they do not deal with the actual content found, but from a technical viewpoint with URL overlap only. Furthermore, more research is needed on the types of results and the beliefs reproduced through search engine algorithms.

REFERENCES

Bar-Ilan, Judit, Kevin Keenoy, Mark Levene, and Eti Yaari. 2009. Presentation Bias Is Significant in Determining User Preference for Search results—A User Study. *Journal of the American Society for Information Science and Technology* 60 (1): 135–149.

Broder, Andrei. 2002. A Taxonomy of Web Search. *ACM Sigir Forum* 36 (2): 3–10.

Bundesverband Digitale Wirtschaft. 2009. *Nutzerverhalten Auf Google-Suchergebnisseiten: Eine Eyetracking-Studie Im Auftrag Des Arbeitskreises Suchmaschinen-Marketing Des Bundesverbandes Digitale Wirtschaft (BVDW) e.V*.

comScore. 2013. *Europe Digital Future in Focus: Key Insights from 2012 and What They Mean for the Coming Year*.

Denecke, Kerstin. 2012. Diversity-Aware Search : New Possibilities and Challenges for Web Search. In *Web Search Engine Research*, ed. Dirk Lewandowski, 139–162. Bingley: Emerald Group Publishing Ltd. doi:10.1108/S1876-0562(2012)002012a008.

Diakopoulos, Nicholas. (2013). Sex, Violence, and Autocomplete Algorithms: What words do Bing and Google censor from their suggestions? Slate. http://www.slate.com/articles/technology/future_tense/2013/08/words_banned_from_bing_and_google_s_autocomplete_algorithms.single.html [Accessed 12 May 2016.]

Edelman, Benjamin. 2010. Hard-Coding Bias in Google 'Algorithmic' Search Results. Benedelman.org. http://www.benedelman.org/hardcoding/. [Accessed 12 May 2016.]

Edelman, Benjamin . 2014. Google's Advertisement Labeling in 2014. Benedelman.org. http://www.benedelman.org/adlabeling/google-colors-oct2014.html. [Accessed 12 May 2016.]

Filistrucchi, Lapo, Catherine Tucker, Benjamin Edelman, and Duncan S. Gilchrist. 2012. Advertising Disclosures: Measuring Labeling Alternatives in Internet Search Engines. *Information Economics and Policy* 24 (1): 75–89.

Gillespie, Tarleton. 2014. The Relevance of Algorithms. In *Media Technologies*, ed. Tarleton Gillespie, Pablo Boczkowski, and Kirsten Foot, 167–193. Cambridge, MA: MIT Press.

Giunchiglia, Fausto, Vincenzo Maltese, Devika Madalli, Anthony Baldry, Cornelia Wallner, Paul Lewis, Kerstin Denecke, Dimitris Skoutas, and Ivana Marenzi. 2009. Foundations for the Representation of Diversity, Evolution, Opinion and Bias.

Hendry, D.G., and E.N. Efthimiadis. 2008. Conceptual Models for Search Engines. In *Web Searching : Multidisciplinary Perspectives*, ed. Amanda Spink and Michael Zimmer, 277–308. Berlin: Springer.

Karaganis, Joe, and Jennifer Urban. 2015. The Rise of the Robo Notice. *Communications of the ACM* 58 (9): 28–30. doi:10.1145/2804244.

Keane, Mark T., Maeve O'Brien, and Barry Smyth. 2008. Are People Biased in Their Use of Search Engines? *Communications of the ACM* 51 (2): 49–52.

Lewandowski, Dirk. 2012. Credibility in Web Search Engines. In *Online Credibility and Digital Ethos: Evaluating Computer-Mediated Communication*, ed. Moe Folk and Shawn Apostel, 131–146. Hershey, PA: IGI Global.

Lewandowski, Dirk. 2014a. Why We Need an Independent Index of the Web. Information Retrieval; Digital Libraries. In *Society of the Query Reader: Reflections on Web Search*, ed. René König and Miriam Rasch, 49–58. Amsterdam: Institute of Network Culture.

Lewandowski, Dirk. 2014b. Wie Lässt Sich Die Zufriedenheit Der Suchmaschinennutzer Mit Ihren Suchergebnissen Erklären? In *Suchmaschinen (Passauer Schriften Zur Interdisziplinären Medienforschung, Band 4)*, ed. H. Krah and R. Müller-Terpitz, 35–52. Münster: LIT.

Lewandowski, Dirk . 2015a. Living in a World of Biased Search Engines. *Online Information Review* 39 (3): 278–280. doi:10.1108/OIR-03-2015-0089.

Lewandowski, Dirk. 2015b. *Suchmaschinen Verstehen*. Berlin Heidelberg: Springer Vieweg.

Lewandowski, Dirk. Forthcoming. Status Quo und Entwicklungsperspektiven des Suchmaschinenmarkts. In *Hanbuch Medienökonomie*, ed. Tassilo Pellegrini and Jan Krone. Berlin Heidelberg: Springer.

Lewandowski, Dirk, and Sebastian Sünkler. 2013. *Representative Online Study to Evaluate the Revised Commitments Proposed by Google on 21 October 2013 as Part of EU Competition Investigation AT.39740-Google Report for Germany*. Hamburg.

Liu, Zeyang, Yiqun Liu, Ke Zhou, Min Zhang, and Shaoping Ma. 2015. Influence of Vertical Result in Web Search Examination. In *Proceedings of SIGIR'15, August 09 - 13, 2015, Santiago, Chile*. New York: ACM.

Mager, Astrid. 2012. Algorithmic Ideology: How Capitalist Society Shapes Search Engines. *Information, Communication & Society* 15 (5): 769–787. doi: 10.1080/1369118X.2012.676056.

Noble, Safiya Umoja. 2012. Missed Connections: What Search Engines Say about Women. *Bitch Magazine*. (54): 36-41.

Noble, Safiya Umoja. 2013. Google Search: Hyper-Visibility as a Means of Rendering Black Women and Girls Invisible. InVisible Culture: An Electronic Journal for Visual Culture. http://ivc.lib.rochester.edu/google-search-hyper-visibility-as-a-means-of-rendering-black-women-and-girls-invisible/. [Accessed 12 May 2016.]

Ntoulas, A., Cho, J., & Olston, C. (2004). What's new on the web?: the evolution of the web from a search engine perspective. In *Proceedings of the 13th international conference on World Wide Web*, 1–12. New York: ACM.

Pan, B., H. Hembrooke, T. Joachims, L. Lorigo, G. Gay, and L. Granka. 2007. In Google We Trust: Users' Decisions on Rank, Position, and Relevance. *Journal of Computer-Mediated Communication* 12 (3): 801–823.

Pariser, Eli. 2011. *The Filter Bubble: What The Internet Is Hiding From You*. London: Viking.

Patterson, Anna. 2004. Why Writing Your Own Search Engine Is Hard. *Queue* 2 (2): 49–53.

Piper, Paul S. 2000. Better Read That Again: Web Hoaxes and Misinformation. Searcher. *Searcher* 8 (8): 40.

Purcell, Kristen, Joanna Brenner, and Lee Raine. 2012. Search Engine Use 2012. *Pew Internet & American Life Project*. Washington, DC.

Risvik, Knut Magne, and Rolf Michelsen. 2002. Search Engines and Web Dynamics. *Computer Networks* 39 (3): 289–302.

Saracevic, Tefko. 2015. Why Is Relevance Still the Basic Notion in Information Science ? In *Re:inventing Information Science in the Networked Society. Proceedings of the 14th International Symposium on Information Science (ISI 2015), Zadar, Croatia, 19th—21st May 2015*, ed. Franjo Pehar, Christian Schlögl, and Christian Wolff, 26–35. Glückstadt: Verlag Werner Hülsbusch.

Spink, A., Jansen, B. J., Blakely, C., & Koshman, S. (2006). A study of results overlap and uniqueness among major web search engines. Information Processing & Management, 42(5), 1379–1391.

Stats: comScore. 2015. Search Engine Land. http://searchengineland.com/library/stats/stats-comscore. [Accessed 12 May 2016.]

Taddeo, M., and L. Floridi (2015). The Debate on the Moral Responsibilities of Online Service Providers. *Science and Engineering Ethics.* doi:10.1007/s11948-015-9734-1.

Tavani, Herman. 2012. Search Engines and Ethics.. [Plattformname.]http://plato.stanford.edu/entries/ethics-search/. [Accessed 12 May 2016.]

van Couvering, Elizabeth. 2007. Is Relevance Relevant? Market, Science, and War: Discourses of Search Engine Quality. *Journal of Computer-Mediated Communication* 12 (3): 866–887.

Vaughan, Liwen, and Yanjun Zhang. 2007. Equal Representation by Search Engines? A Comparison of Websites across Countries and Domains. *Journal of Computer-Mediated Communication* 12 (3): 888–909.

White, Ryen W., and Eric Horvitz. 2009. Cyberchondria. *ACM Transactions on Information Systems* 27 (4): Article No. 23. doi:10.1145/1629096.1629101.