AH ANALECTA HERMENEUTICA
INTERNATIONAL INSTITUTE FOR HERMENEUTICS / INSTITUT INTERNATIONAL D'HERMÉNEUTIQUE

ISSN 1918-7351
Volume 15.1 (2023)

Everything but the Truth: On the Relevance of Algorithms

Sebastian Rosengrün
CODE University of Applied Sciences, Berlin
ORCID: 0000-0002-0747-8424

Abstract
This paper defends two theses: (1) Humans aim for relevance, not for truth. (2) Algorithms provide relevance, not truth. While the first thesis builds upon research in cognitive linguistics and neuroscience, the second thesis will be based upon an analysis of algorithm functionality in digital platforms like Google, Amazon, and Social Media, but also in most recent developments in the field of Generative AI. Based on these two theses, this paper outlines asymmetric power structures in digital capitalism and how commercial interests undermine the democratic discourse by spreading fake news and conspiracy theories.

Keywords: artificial intelligence, relevance, algorithms, digital capitalism, fake news

Introduction

Algorithm-based recommendation systems have gained increasing influence over how humans perceive and interact with the world. These systems play a pivotal role in shaping various aspects of our lives, both online and offline. For instance, Google's search algorithm determines the information users find and the sources they consider during their research endeavors. Social media algorithms curate the news users consume, dictate the topics they encounter, and influence how these topics are presented. Similarly, Amazon's algorithm influences the products buyers discover and the prices they pay, thereby impacting market dynamics and the chances of success for different companies. Importantly, the reach of algorithm-based recommendation systems extends beyond the confines of the online world, permeating our daily lives in ways that often go unnoticed. From the restaurants we choose to dine in and the vacation destinations we select to the potential sex partners we encounter through dating apps, algorithms subtly shape our normal lives. In essence, algorithm-based recommendation systems can be seen as a form of regulation on human behavior, akin to what legal scholar Lawrence Lessig termed "Code is Law."[1]

The objective of this paper is twofold: Firstly, it aims to provide an answer to the question of why algorithm-based recommendation systems have achieved such remarkable success. Secondly, it seeks to contextualize some of the most problematic effects arising from these systems. Of particular importance is the rising concern of fake news and conspiracy theories, where the asymmetric power structures inherent in algorithm-based recommendation systems play a crucial role. To comprehend this phenomenon, this paper will explore the concept of relevance—a fundamental idea within cognitive linguistics and philosophy of language that must be distinguished from the concept of truth. In this context, "relevance" refers to the usefulness of an assertion (a recommendation, or an explanation) to a speaker, which may not always align with factual accuracy.

In the first section, this paper will delve into this concept of relevance and its significance in human cognition, highlighting that humans prioritize relevance over truth. This understanding will lay the groundwork for the next section, where it will be demonstrated how algorithm-based recommendation systems excel at providing relevance to users compared to human-based recommendations. Leveraging their algorithms, these systems offer a significant advantage in satisfying users' cognitive needs efficiently. Thereafter, this paper will shed light on some of the problematic aspects of algorithm-based relevance in various facets of our social life, particularly considering the asymmetric power structures prevalent in digital capitalism and how algorithm-based relevance contributes to the rise of fake news and conspiracy theories.

---

[1] Lawrence Lessig, *Code: Version 2.0* (New York: Basic Books, 2006). For a detailed analysis of this broader picture of 'regulation by AI', cf. Sebastian Rosengrün, "Why AI Is a Threat to the Rule of Law," *Digital Society* 1, no. 2 (2022): 10, https://doi.org/10.1007/s44206-022-00011-5.

Finally, this paper will briefly talk about recent developments in generative AI and about their potential implications for algorithm-based recommendation systems.

The *Relevance* of Relevance

Relevance theory is a major theme in cognitive linguistics. Research in this field indicates "the search for relevance is a basic feature of human cognition" and:

> [T]he spontaneous working of our perceptual mechanisms tends to pick out the most relevant potential inputs, the spontaneous working of our memory retrieval mechanisms tends to activate the most relevant potential contextual assumptions, and the spontaneous working of our inferential mechanisms tends to yield the most relevant conclusions.[2]

Similarly, studies from neuroscience further suggest innate processes of selective attention "that allow an individual to select and focus on particular input for further processing while simultaneously suppressing irrelevant or distracting information."[3] There is also a strong tradition of philosophical pragmatism building on the assumption that, at least in everyday situations, the factuality ('truth') of an assertion is significantly less important than whether an assertion is useful to fulfill a particular goal within a specific context, hinting already at a minimal definition of the term 'relevance': *Def.* An assertion is relevant if and only if it is useful in a concrete situation.[4] Furniture assembly instructions (which are assertions, too), for example, are relevant for me if and only if they lead to a successful practice, i.e., they help me connect the various parts so that the bookcase will hold together. One crucial feature of relevance, therefore, is that it often only manifests itself in retrospect. Whether instructions are relevant, I will most likely only find out after having followed them through without the bookcase falling apart.

While this paper does not presuppose a specific definition of truth, there is one significant difference between the concept of truth (both according to correspondence and coherence theories) and the concept of relevance: Truth is

---

[2] Deirdre Wilson, "Relevance Theory," in *The Pragmatics Encyclopedia*, ed. Louise Cummings (New York: Routledge, 2010), 395; Dan Sperber and Deirdre Wilson, "Relevance Theory," in *The Handbook of Pragmatics*, ed. Laurence Horn and Gregory Ward (Oxford: Blackwell, 2006), 608.

[3] Courtney Stevens and Daphne Bavelier, "The Role of Selective Attention on Academic Foundations: A Cognitive Neuroscience Perspective," *Developmental Cognitive Neuroscience* 2 (2012): 30, https://doi.org/10.1016/j.dcn.2011.11.001; cf. Wolf Singer, *Ein neues Menschenbild? Gespräche über Hirnforschung*, Suhrkamp-Taschenbuch Wissenschaft 1596 (Frankfurt am Main: Suhrkamp, 2003).

[4] For this definition and the following explanations, cf. Thomas Becker, "Is Truth Relevant? On the Relevance of Relevance," *Etica & Politica / Ethics & Politics* XVI, no. 2 (2014): 595–618; Sebastian Krebs, "Does Truth Really Matter? On the Irrelevance of Truth," in *Practical Rationality in Political Contexts. Facing Diversity in Contemporary Multicultural Europe*, ed. Gabriele De Anna and Riccardo Martinelli (Trieste: EUT Edizioni Università di Trieste, 2016), 31–58.

objective, relevance is not. If the assertions "The 12th decimal place of π is 9" and "The last word of this paper is 'relevant'" are true, then they are true for everyone at all times. They would even be true if no one had ever calculated π or would read this paper to the last paragraph. However, the 12th decimal place of π is relevant only for a small group of people, and whatever the last word of this paper shall be, is, at this point, still irrelevant. Those examples also indicate that not all true assertions are relevant. With Bernard Bolzano, drawing from his *Wissenschaftslehre* from 1837, it can be stated that science is not (only) about finding truth but (also) about selecting from an infinite number of truths those that have practical use.[5] This idea is exemplified by Thomas Becker's thought experiment of the Library of Baghdad, an adaptation from Argentinian writer Jorge Luis Borges' Library of Babel:

> [In this fictive library,]: books contain only true sentences (not a single false one) in impeccable English, without a single misprint. It contains, just like the Library of Babel, an infinite set of true sentences derived logically or by other recursive definitions from a basis of true and known sentences compiled by a large committee of scholars. All the sentences differ from each other, not a single sentence is recorded twice, and all sentences are of finite length. Nevertheless, it is as useless as the Library of Babel, because you have virtually no chance to find a single interesting sentence among the infinite number of true and irrelevant ones.[6]

Becker claims that this library is as useless as Borges' complete library, as this library is infinitely extensive due to multiple reasons like the recursiveness of natural languages and the formal logic of adding a true disjunct to an otherwise false sentence yielding a true sentence.[7] He concludes that "the point of assertion is to pick out the most relevant proposition of an infinite number of true, known and justifiable ones."[8]

The most influential definition of relevance can be attributed to Dan Sperber and Deirdre Wilson, who established 'relevance theory' as an independent field of research in cognitive linguistics. According to them, relevance is a function determined by two factors, (a) cognitive effects and (b) processing effort. For individuals, what is considered relevant is typically the assertion that yields the highest cognitive effects (cognitive reward) with the least processing effort (cognitive costs).[9] This function is

---

[5] Bernard Bolzano, *Grundlegung Der Logik. Ausgewählte Paragraphen Aus Der Wissenschaftslehre*, ed. Friedrich Kambartel, 2., durchges. Aufl, vol. 259, Philosophische Bibliothek (Hamburg: Meiner, 1978), 3.

[6] Becker, "Is Truth Relevant?," 602.

[7] Cf. Jorge Luis Borges, *Collected Fictions.*, trans. Andrew Hurley (London: Penguin, 1999); Jorge Luis Borges, *The Total Library: Non-Fiction 1922-1986*, trans. Esther Allen, Suzanne Jill Levine, and Eliot Weinberger (London: Penguin, 2001).

[8] Becker, 603.

[9] Cf. Dan Sperber and Deirdre Wilson, *Relevance: Communication and Cognition*, 2nd ed. (New York: Blackwell, 1995); Sperber and Wilson, "Relevance Theory"; Yan Huang, *Pragmatics* (Oxford: Oxford University Press, 2007).

crucial for the following analysis of relevance in algorithms, as it adds another essential aspect to the understanding of usefulness from above: An assertion is more useful to a hearer when the conveyed information requires less effort to process. For example, while a 200-page documentation on using various screws and tools may effectively prepare me for assembling a bookcase, too, a concise and visually appealing four-page comic-like guide offers me an equally effective cognitive while incurring significantly fewer cognitive costs.

Relevance in Algorithms

Tech companies have embodied this cognitive principle of relevance, which is, as this paper claims, a crucial aspect of the economic success of algorithm-driven business models. In 1995, Larry Page and Sergey Brin developed the Google precursor Backrub. It was based on the idea that the value of a website should be determined by the number of backlinks, i.e., the more often a website was linked by other websites, the higher it was ranked at Google. In other words: They understood that the quality of content is less important for the user than how popular it is among other people. This approach was revolutionary because it prioritized the popularity of a website over its content quality, which was a departure from previous search engines. By relying on backlinks to rank websites, Google could provide more relevant results to users, as websites frequently linked to by other reputable sources were deemed more valuable. This approach has since been widely adopted by other tech companies, who use similar algorithms to provide personalized recommendations to users based on their past behaviors and preferences.

Interestingly, Page and Brin took their idea from academia, where a researcher's reputation is mainly determined not by the quality of their research but rather by how often their papers are quoted by others (or rather, that the number of citations is the most important criteria for the quality of research). This paper, for example, becomes a relevant contribution to the philosophy of digital technologies if and only if it is quoted in many academic publications.[10] The quality of its arguments matters only insofar as there is a common agreement in academia that one ought only quote papers with argumentative quality (and reviewers typically ensure a certain standard within the publication process).

While there are neither peer-reviewers nor silent agreements in website publishing, the idea of Backrub was still a success—compared to earlier search engines that solely analyzed a website's content. To have a website placed among the top search results for, e.g., "how to assemble a bookshelf," a publisher only had to ensure to use words like "bookshelf" and "assembly" more often than their competitors within their

[10] For a performative criticism of this common academic practice, see, Sebastian Rosengrün, "Everything but the Truth: On the Relevance of Algorithms," *Analecta Hermeneutica* 15 (2023).

HTML documents. The idea of Backrub required a website to be linked back by others to achieve a high PageRank—a concept introduced by Page and Brin. From a user perspective, search results suddenly became far more relevant, and Google (as Backrub was re-named before becoming successful) soon became the gold standard of internet search engines (which it still is, even though Generative AI might pose a significant threat to its market position, see section 5).

While relevance should be defined as leading to successful practice (see the previous of this paper), the Google algorithm defines what leads to a successful practice by what other people think is leading to successful practice. Algorithms have been optimized based on such an understanding of relevance, and search engines are just one of many similar examples. Spotify, Netflix, and YouTube make algorithm-based recommendations of what to watch/listen next dependent on what other people with a similar background (age, gender, hobbies, favorite band, etc.) like to watch/listen. Amazon optimizes their product placement algorithms based on what other similar shoppers tend to buy, and the whole business model of social media like Facebook, Instagram and TikTok relies on the quality of their algorithms to select relevant content for their users. The same principle applies, more and more, to news websites, political campaigning and commercial advertisement in general, but also to dating apps, restaurant recommendations and travel planning. Algorithms have been optimized to display what 'similar' users seem to enjoy, as this seems to be what brings any individual the highest cognitive reward for as little processing costs as possible.

In order to train those algorithms, tech companies like Google heavily rely on the data they collect from their users, leading to a phenomenon that Shoshana Zuboff, in her well-received study, describes as "surveillance capitalism," a "new instrumentarian power that asserts dominance over society and presents startling challenges to market democracy."[11] According to Zuboff, companies increasingly use algorithms to control human behavior by predicting precisely how humans behave and how manipulating the input variables for human behavior will affect behavioral outcomes. Zuboff's analysis offers profound insights into the historical development and business models of big tech corporations, but also a blunt criticism of surveillance capitalism which she describes as "parasitic and self-referential" economic order that "feeds on every aspect of every human's experience."[12] With algorithm-based behavior prediction and control, tech corporations endanger every human's right to make their own life choices—which is a core value of any democratic society.[13]

However, it is highly questionable whether algorithms trained with machine learning can be accurate enough to achieve such an ultimate behavior prediction. The effort (the cognitive costs) to perform a Google search (or following Amazon's or

---

[11] See, Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, First edition (New York: PublicAffairs, 2019).
[12] Zuboff, *The Age of Surveillance Capitalism*, 9.
[13] For a more detailed analysis, See, Rosengrün, Sebastian, "Why AI is a threat to the rule of law," *Digital Society* 1, no. 2 (2022).

Netflix's recommendations), however, is very little compared to asking friends, colleagues, and experts for help and recommendations. Also, their output is optimized in a way that makes it very easy to process: a brief list of 10 suggestions of what product to buy/read/watch/invest in is far easier to deal with than extracting valuable information out of the monumental speeches by our enthusiastic librarians, geeky little nieces, or dodgy insurance brokers. In fact, all tech companies spend tremendous research efforts in order to 'optimize' their user experience by introducing new design patterns, color schemes, or features (like Amazon's infamous One-Click-Buying, which they even patented, and which reduces the cognitive costs of online shopping to the bare minimum).[14]

Therefore, their suggestions automatically have a tremendous advantage in terms of relevance over 'traditional' search methods in terms of how relevant they are for an individual. When one assumes that whatever an algorithm suggests is accurate (and people do not question those suggestions because of the relatively little processing effort), it is easy to conclude that algorithmic suggestions are relevant. A user being amazed by algorithmic predictions (like Google's search results or Netflix's movie predictions) is comparable to a tiny baby being fascinated by observing her reflection in a mirror without realizing it is herself she is watching. This observation is also (and even more) accurate for generative AI (like ChatGPT, Google Bard, Stable Diffusion, DALL-E etc.) that often surprises their users with highly relevant outputs to their initial prompts. However, the reason for the relevance of their outputs is, at least partially, to be explained by the little processing costs they mean for their users, while the cognitive effects are simply a result of a rearrangement of an extremely large set of language tokens (i.e., language used by other human beings) with the help of stochastic means.

Especially in cases in which factuality and objectivity do not (seem to) exist, it seems that the relevance of algorithmic predictions must be explained mainly by their low cognitive costs, not by their invaluable cognitive effects: What book to buy next, what movie to watch next, and—what party/candidate to vote next for in the upcoming elections. Those are rather questions of relevance rather than questions of truth. Those question even presuppose that buying yet another book, watching yet another movie and participating in an election are the only feasible options (see section 4). They have a tremendous social impact, however, when their answers are calculated by algorithms within the asymmetric power structures of digital capitalism.

[14] Within academia, judging a researcher's work solely based on their h-index also requires much less processing effort than reading through their publications.

Relevance in Asymmetric Power Structures

What has been said in the two sections above can be summed up as follows: Algorithms control what appears to be most relevant for individual users. Keeping in mind that relevance is determined by the highest cognitive effects for the lowest cognitive costs, however, the relevance of algorithmic suggestions is not necessarily to be explained by high cognitive effects for the user but rather because of the small cognitive costs. Given that those algorithms are controlled by a handful of tech corporations, it is crucial to shed light on the concept of relevance within those asymmetric power structures created by what scholars call 'digital capitalism.'[15]

Given the monopolistic tendencies within digital capitalism, tech corporations control what is relevant. In the search engine market, Google has a global market share of 93.11 percent (May, 2023), with Microsoft's Bing being the only noteworthy alternative (at least within the so-called Western context, leaving Russian and Chinese search engines aside).[16] Similarly, companies and services like Amazon, Netflix, Twitter, Instagram, and YouTube have become dominant players with their algorithm-based recommendations, influencing, for example, what people buy, read, watch, or listen to next. It is important to note that those companies hold this power—whether they want it to or not. Not making an active decision on what to recommend someone, is also a decision—especially if the person looking for a recommendation looks at you as the sole source of truth.[17] While most tech companies actively shy away from the responsibility that comes along with this power, they actively make use of this power: By controlling those algorithms, they influence the way people perceive and interact with the world.[18] It is important to note here that recommendation algorithms do not only offer relevant answers, but the way those algorithms have been designed, also presupposes that what to buy, read, watch, or listen to next is even a relevant question for their users. *That* everyone wants to buy, read, watch or listen to something else, is a decision already made for the users, and it seems that, to many people, this is a more relevant option (i.e., a question that takes less processing costs) than asking the (admittedly, more complex) question of what else one could do with their time. While

---

[15] Cf. Michael Betancourt, *The Critique of Digital Capitalism: An Analysis of the Political Economy of Digital Culture and Technology* (Brooklyn, NY: punctum books, 2015); Dan Schiller, *Digital Capitalism: Networking the Global Market System* (Cambridge, Mass: MIT Press, 1999); Dan Schiller, *Digital Depression: Information Technology and Economic Crisis*, The Geopolitics of Information (Urbana, Chicago: University of Illinois Press, 2014); Philipp Staab, *Digitaler Kapitalismus. Markt Und Herrschaft in Der Ökonomie Der Unknappheit* (Suhrkamp, 2019); Amy Webb, *The Big Nine: How the Tech Titans and Their Thinking Machines Could Warp Humanity* (New York: PublicAffairs, 2020); Rosengrün, "Why AI Is a Threat to the Rule of Law."

[16] StatCounter Global Stats, accessed June 21, 2023, https://gs.statcounter.com/search-engine-market-share. Privacy-aware people often prefer meta search engines such as DuckDuckGo and Ecosia, which, however, rely in their search results mainly on the algorithms of Google or Bing.

[17] Cf. Sebastian Rosengrün, *Künstliche Intelligenz zur Einführung*, Zur Einführung (Hamburg: Junius, 2021).

[18] Cf. Adrian Daub, *What Tech Calls Thinking* (New York: Farrar, Straus and Giroux, 2020).

the relevance of algorithmic recommendations in our daily life causes enormous social problem, this paper does not suggest that the companies behind that follow a broader political agenda. There mere goal is capitalist profit,[19] and that's exactly why it is problematic that their power over what is presented as relevant for people takes place beyond democratic discourse.

This can be, for example, illustrated by focusing on the issue of so-called filter bubbles and how they lead to the spread of fake news and conspiracy theories.[20] As this paper suggests, those result from the monopolistic control of social media companies over what is relevant for their users. By filter bubbles, this paper refers to the phenomenon where individuals are exposed only to information that confirms their existing beliefs, resulting in a narrowing of perspectives and an echo chamber effect. Tech corporations play a central role in creating these filter bubbles through their algorithms, which are designed to show users content that is most likely to keep them engaged on their platforms. Algorithmic recommendations limit exposure to diverse perspectives by prioritizing content that confirms users' existing beliefs and biases. In addition, the spread of fake news is facilitated by the ability of tech corporations to amplify and distribute information at unprecedented speeds. Without robust fact-checking mechanisms, misinformation can quickly spread through social media networks, further contributing to the creation of filter bubbles and the erosion of trust in traditional sources of information.

In this landscape, traditional media organizations have been confronted with the following dilemma of staying relevant: Either they adjust themselves by lowering the cognitive costs for their audience or they try to focus on the quality of their information (and increase the cognitive rewards), for which, however it is increasingly difficult to find a viable business model given how easy it is to find equally relevant information that might not have the same cognitive reward but comes with much lower cognitive costs. Quality media do not only have higher cognitive costs for their users, but also significantly higher convenience and financial costs: While social media (or the internet, in general) is full of fake news and conspiracy theories that are freely accessible, quality media are 'hiding' their content more and more behind paywalls, opaque subscription models and premium accesses.[21] While this approach seems to be economically necessary to generate revenue and pay for journalistic endeavors, it also restricts access to valuable information and excludes those who cannot afford or are unwilling to pay for digital content (or those who are unwilling to give away their personal data and manage numerous online accounts). Especially people who are

---

[19] Rosengrün, "Why AI Is a Threat to the Rule of Law," 10.

[20] See, Seth Flaxman, Sharad Goel, and Justin M. Rao, "Filter Bubbles, Echo Chambers, and Online News Consumption," *Public Opinion Quarterly* 80, no. S1 (2016): 298–320, https://doi.org/10.1093/poq/nfw006; Eli Pariser, *The Filter Bubble: What the Internet Is Hiding from You* (London: Penguin Books, 2012); Michael Butter, *The Nature of Conspiracy Theories* (Cambridge: Polity Press, 2020).

[21] Unfortunately, the same must be said for many academic publications.

already affected by their individual filter bubbles and echo chambers will be excluded even more from reliable information, given that what seems relevant to them is already provided by social media who mostly care about low processing costs for their users.

Abuse of asymmetric power structures within digital capitalism goes beyond filter bubbles and echo chambers. Users have been abused by tech companies which confront them with seemingly relevant algorithmic recommendations for many years. An extreme example is the infamous 'Emotional Contagion Experiment,' conducted by Facebook in 2014.[22] In this study, Facebook manipulated the News Feed algorithm for a subset of users by selectively filtering out positive or negative posts for a given period. This alteration in the content was done to observe if it would lead to changes in users' own emotional expressions in their subsequent posts. While this study has been controversially discussed (not only for the obvious ethical, but also for methodological reasons), its results suggest that emotional contagion could occur through social media platforms, as user emotions seemed to be influenced by the emotional content they were exposed to. Of course, the main intention behind such commercial experiments is to turn their results into a business case, knowing that a user's emotional status not only heavily affects how they interact with other users on a platform, but also their media consumption and (online) shopping behavior. Another extreme case highlighting the implications of algorithm-based recommendation systems is the Cambridge Analytica scandal.[23] This notorious incident revealed the potential misuse of personal data by a political consulting firm, which utilized algorithms to target and manipulate users with tailored political content to the effect of influencing the outcome of democratic elections.

Summary and Outlook

Algorithm-based recommendation systems have become integral to our modern lives, shaping our behavior and influencing the information we consume. This paper has explored the concept of relevance from both a philosophical and linguistic perspective. Relevance, defined as practical usefulness and determined by cognitive costs and rewards, was presented as a key concept to explain the success of those algorithmic systems. By prioritizing relevance over truth, algorithm-based recommendations cater to our cognitive needs effectively. Given that many users already assume that algorithmic recommendations are more relevant than traditional systems, the

---

[22] Cf. Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock, "Experimental Evidence of Massive-Scale Emotional Contagion through Social Networks," *Proceedings of the National Academy of Sciences* 111, no. 24 (2014): 8788–90, https://doi.org/10.1073/pnas.1320040111; Jukka Jouhki et al., "Facebook's Emotional Contagion Experiment as a Challenge to Research Ethics," *Media and Communication* 4, no. 4 (2016): 75–85, https://doi.org/10.17645/mac.v4i4.579.

[23] Cf. Christopher Wylie, *Mindf*ck: Cambridge Analytica and the Plot to Break America* (New York: Random House, 2019).

companies behind those algorithms gain power over shaping the way people perceive and interact with the world. This comes with inherent risks and challenges. The rise of fake news and conspiracy theories as well as other threats to a free and open society underscores the need to critically examine the asymmetric power structures embedded in these systems.

A gamechanger with regards to algorithm-based recommendations systems is the rather recent development in generative AI (including Chat-GPT), which has revolutionized the capabilities of these systems in understanding and responding to user preferences and needs. Generative AI models, such as Chat-GPT, possess the ability to generate human-like text, engage in dynamic conversations, and adapt their recommendations based on user interactions. While business cases building on generative AI are still in the process of being fully explored, the impact of this technology on search engines is already poised to be profound. A notable example of this is Microsoft's recent collaboration with OpenAI, leading to the integration of Chat-GPT into the Bing search engine. This partnership demonstrates the potential of generative AI in transforming the way search engines function and the experiences they offer to users. While most generative AI tools are currently designed in a way to protect their users from conspiracy theories and even very hesitant to give recommendations (e.g., on whom to vote for in the upcoming election), there is, without doubt, both commercial interest in the companies providing those tools and organizations willing to pay for influencing algorithmic recommendations towards their own interest (see the Cambridge Analytica example outlined in the previous section). Considering the extensive media coverage and enthusiastic celebration surrounding the technology underpinning Chat-GPT, the answers and recommendations generated by generative AI systems are likely to be perceived as highly relevant by their users. Nevertheless, it is crucial to acknowledge that the decision-making process governing these algorithms rests not in the realm of democratic discourse but rather within the realm of commercial interests controlled by a small number of major tech corporations. For truth to prevail over falsehoods, including the spread of fake news and conspiracy theories, within a democratic discourse, it is imperative for any free and open society to guarantee that truth remains relevant.