

This item is the archived peer-reviewed author-version of:

What are filter bubbles really? A review of the conceptual and empirical work

Reference:

Michiels Lien, Leysen Jens, Smets Annelien, Goethals Bart.- What are filter bubbles really? A review of the conceptual and empirical work
Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization - ISBN 978-1-4503-9232-7 - New York, N.Y., Association
for Computing Machinery, 2022, p. 274-279
Full text (Publisher's DOI): <https://doi.org/10.1145/3511047.3538028>
To cite this reference: <https://hdl.handle.net/10067/1904900151162165141>

What Are Filter Bubbles Really? A Review of the Conceptual and Empirical Work

Jens Leysen*

Lien Michiels*

jens.leysen@student.uantwerpen.be

lien@froomle.com

Froomle

Belgium

University of Antwerp

Belgium

Bart Goethals

bart.goethals@uantwerpen.be

Froomle

Belgium

University of Antwerp

Belgium

Monash University

Australia

Annelien Smets

annelien.smets@vub.be

imec-SMIT, Vrije Universiteit Brussel

Belgium

ABSTRACT

The original filter bubble thesis states that the use of personalization algorithms results in a unique universe of information for each of us, with far-reaching individual and societal consequences. The ambiguity of the original thesis has prompted both a conceptual debate regarding its definition and has forced empirical researchers to consider their own interpretations. This has led to contrasting empirical results and minimal generalizability across studies. To reliably answer the question of whether filter bubbles exists, on what platforms, and what caused them, we need a systematically and empirically verifiable definition of the filter bubble that can be used to develop rigorous tests for the existence and strength of a filter bubble. In this paper, we propose an operationalized definition of the (technological) filter bubble and interpret previous empirical work in light of this new definition.

CCS CONCEPTS

• Human-centered computing → HCI theory, concepts and models.

KEYWORDS

filter bubble, diversity, recommendation, personalization, recommender system

ACM Reference Format:

Jens Leysen, Lien Michiels, Bart Goethals, and Annelien Smets. 2025. What Are Filter Bubbles Really? A Review of the Conceptual and Empirical Work. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

The idea of ‘filter bubbles’ originated from internet activist Eli Pariser [28]. According to Pariser’s original thesis, a filter bubble is an environment, *created by a personalization algorithm*, in which a person only encounters familiar information or opinions. Additionally, he argues that these filter bubbles “close us off to new ideas, subjects, and important information” which harms the democratic process and leads to increased polarization, among other things [28]. The term quickly found its way into the common discourse with

‘endorsements’ by prominent figures such as former U.S. president Barack Obama [10] and former Microsoft CEO Bill Gates [13]. All the while, the filter bubble has been under academic scrutiny both from a theoretical and empirical standpoint. Legal and media scholars have criticized the concept for its lack of a clear definition, its unsubstantiated leap from individual observations to societal effects, its changeability, and the absence of irrefutable empirical evidence for its existence [6–8, 12, 33]. Despite their criticism of Pariser’s original thesis, these same scholars (and others) have emphasized that the study of the effect of personalization technology on diversity is highly relevant as more and more content is being filtered by algorithms [12, 14, 24, 33]. This lack of a clear definition has forced empirical researchers to adopt their own interpretations of the filter bubble, which vary wildly. Filter bubbles have, as such, become a moving target: Because the theory is constantly adapted, it cannot be criticized or falsified [12]. This concern of falsifiability was recently addressed by Dahlgren, who proposed an improved conceptual definition of the filter bubble [12]. However, as we argue in this paper, this conceptual definition of a technological filter bubble could still benefit from further operationalization. In this work we build on Dahlgren’s definition of the filter bubble and propose an operationalized definition of the technological filter bubble. To the best of our knowledge, this work is the first to provide such an operationalized, systematically and empirically verifiable definition of the technological filter bubble. It is our hope that this definition can henceforth be used to develop “methodologically sound and empirically rigorous tests for the existence and strength of filter bubbles” [7].

This paper proceeds as follows: Sections 2 and 3 discuss Pariser’s thesis and Dahlgren’s definition respectively. In Section 4, we build upon the definition of Dahlgren and introduce our operationalized definition of the technological filter bubble. We then reinterpret prior empirical work in light of this definition in Section 5. In Section 6, we briefly discuss the main issues that research on filter bubbles has faced. We conclude the paper in Section 7.

2 THE ORIGINAL FILTER BUBBLE

Eli Pariser introduced the term and described it as follows:

“The new generation of Internet filters looks at the things you seem to like—the actual things you’ve done, or the things people like you like—and tries to extrapolate. They are prediction engines, constantly creating

*Both authors contributed equally to this research.

and refining a theory of who you are and what you'll do and want next. Together, these engines create a unique universe of information for each of us—what I've come to call a filter bubble—which fundamentally alters the way we encounter ideas and information.” [28, p. 10]

Pariser's description of a “unique universe of information for each of us” lays the blame for filter bubbles squarely with personalization algorithms, which he terms ‘Internet filters’. According to Pariser, this filter bubble has far-reaching personal and societal costs. Two of the most far-reaching consequences he mentions are “auto propaganda that indoctrinates us with our own ideas” leading to increased polarization and “obstruction of information flows” which harms the democratic process [28].

3 THE TECHNOLOGICAL FILTER BUBBLE

Following the widespread criticism of Pariser's thesis [6–8, 33], Dahlgren explicitly sets out to interpret Pariser's original thesis and clear up any ambiguities [12]. Importantly, Dahlgren argues that Pariser's filter bubble thesis is an argument consisting of two parts: a technological and societal filter bubble. At the technological level, we have that “any single choice affects the content recommended by personalization algorithms, thereby narrowing the type of content available over time” [12]. The societal level, on the other hand, is concerned with “the causes and consequences of these choices and technologies for humans and society, and, more importantly, for the political process and democracy over time” [12]. Pariser's thesis implicitly assumes a technological filter bubble will inevitably lead to societal issues, but this leap is as yet unsupported by empirical evidence [12]. In the remainder of this work, we limit our discussion to the technological filter bubble, which is concerned with personalization technology. For further discussion on the potential societal consequences of filter bubbles, we refer the interested reader to the following related work [6–8, 12, 32, 33].

Dahlgren's conceptualization of the technological filter bubble addresses two sources of ambiguity found in the original thesis. First, he makes explicit that a technological filter bubble is not a state in time but instead is created ‘over time’. He quotes Pariser saying: “personalization algorithms can cause identity loops, in which what the code knows about you constructs your media environment, and your media environment helps to shape your future preferences” [12, 28]. Second, he specifies that a technological filter bubble requires ‘narrowing the type of content available’. We can draw this conclusion from another direct quotation of Pariser's book: “You click on a link, which signals an interest in something, which means you're more likely to see articles about that topic in the future, which in turn prime the topic for you.” [12, 28].

4 OPERATIONALIZING THE TECHNOLOGICAL FILTER BUBBLE

Although Dahlgren's effort to bring some clarity is laudable, looking at it through the lens of empirical research, his definition still leaves important properties of a technological filter bubble open to interpretation. It is unclear what is meant by “narrowing the type of content available” and exactly whose choices are implied with his use of “any single choice”. We believe a more systematically and

empirically verifiable definition is required to move towards comparable and generalizable empirical works. For that reason, we now propose an operationalized version of Dahlgren's conceptual definition. With operationalization, we refer to the translation of abstract concepts to measurable observations. Our operationalized definition uses the terminology commonly found in the recommender systems literature. We posit that:

A technological filter bubble is a decrease in the diversity of a user's recommendations over time, in any dimension of diversity, resulting from the choices made by different recommendation stakeholders.

Our definition is built around four foundational elements: diversity, recommendations, time, and recommendation stakeholders.

4.1 Diversity

First of all, we define a filter bubble as a decrease in *any dimension of diversity*, whereas Dahlgren refers to a *narrowing of the type of content available*. Empirical research on filter bubbles is mainly concerned with the three core dimensions of media diversity, which find their origin in the early works of Napoli [25], Voakes et al. [31] and McQuail [23]. However, many more dimensions of diversity can be identified [19, 21]. Various terms have been used to describe these three core dimensions of diversity over the years, adding to the ambiguity in the field. We follow the naming convention of Loecherbach et al. and refer to these dimensions as structural, topic, and viewpoint diversity [21]. Each of these dimensions is considered essential to a user's media diet in their own right [15, 33]. To avoid any confusion that could arise due to conflicting definitions, we give explicit definitions for each of the dimensions.

Structural Diversity. Napoli first defined structural diversity as diversity in news outlets but named it source diversity [25]. As the media landscape has evolved beyond news outlets alone, Loecherbach et al. expanded this definition to “diversity in who supplies the information,” which we also adopt here [21]. Examples of these information suppliers are ‘other users’ on social media platforms, or news brands on social media and news aggregators.

Topic Diversity. Topic diversity relates to any diversity in topics or subjects. What constitutes a topic is open for debate. At the highest level, one can consider the sections of a newspaper to be topics: politics, sports, health, business, science. However, it is equally valid to argue that a news article on volleyball does not have the same subject as an article on basketball or that they share only some limited similarity, since both are ball games.

Viewpoint Diversity. Viewpoint diversity is the dimension that pertains to the variety of stances that can be taken on a topic. One could, for example, think of proponents of Brexit and those against it. Some empirical works on filter bubbles study misinformation, which we consider a special case of viewpoint diversity in which the viewpoints represented are provably false.

4.2 Recommendations

Secondly, we make explicit that a technological filter bubble is a decrease in the diversity of *a user's recommendations*. Pariser's thesis is clear that personalization algorithms, or ‘internet filters’,

are to blame for filter bubbles. Additionally, he states that they “create a unique universe of information for each of us”, which indicates that we are all alone in our respective filter bubbles [28]. Therefore a filter bubble should be measured as an effect in every individual user’s recommendations, i.e. what a user is exposed to as a result of ‘internet filters’.

4.3 Time

Thirdly, both in Pariser’s original thesis and Dahlgren’s later definition, the filter bubble is positioned as a longitudinal effect of a recommender system that arises while it “refines its theory of who you are and what you’ll do and want next” [28]. This implies that ideally, one starts measuring before the recommender system has had the chance to start ‘refining its theory of you’.

4.4 Recommendation Stakeholders

Finally, we emphasize that there are multiple *recommendation stakeholders* each of which can affect the diversity of content recommended [1]. A recommendation stakeholder is defined as “any group or individual that can affect, or is affected by, the delivery of recommendations to users” and three key groups of stakeholders have been identified: users, providers, and the system [1]. This clarification is important, as today’s recommender systems are not the prediction engines - that focus only on refining their model of your preferences - Pariser envisioned [28]. Users are arguably the most important stakeholder: They are affected by and affect the diversity of recommendations through their consumption behavior. However, other recommendation stakeholders can also play an important role. Recommender systems can be (and are) deployed with multiple system or business objectives in mind [1, 18, 29]. For example, product owners can decide recommendations should meet a minimum threshold of diversity and post-process initial recommendations accordingly [9]. The diversity of content that is written, promoted, relevant and available at any moment in time is also an important factor. If this supplied diversity is temporarily lowered, e.g. because most articles or posts are on the topic of presidential elections, recommendation diversity must inevitably follow. We believe this multi-stakeholder perspective is essential, as a fundamental understanding of filter bubbles will not come from observing their existence but from explaining why these filter bubbles came to be in the first place.

5 REINTERPRETING PRIOR EMPIRICAL STUDIES

To move toward greater comparability and generalizability between empirical studies the field would benefit from a clear distinction between studies that do and do not research a technological filter bubble [6]. In this section, we reinterpret prior empirical work in light of our operationalized definition of the technological filter bubble. We relied on three search strategies to identify these works. First, we conducted a Google Scholar search with the terms ‘filter bubble(s)’, sorted by relevance, and retained the first 150 results for each query. Second, we reviewed the proceedings of the FAccT, WWW, RecSys, WSDM and UMAP conferences in the last three years. Finally, we mined the citations of the conceptual works we discussed earlier in this work. We then obtained our final selection

by retaining only peer-reviewed empirical studies related to filter bubbles. A reference check showed high saturation [16], although we do not claim to be exhaustive.

Based on our definition, a study on the technological filter bubble must consider the *diversity of recommendations* and measure a decrease in diversity over *time*.

We find that each of the studies reviewed considers *diversity*, albeit in different dimensions. Viewpoint diversity has been studied in many different contexts, ranging from video platforms (YouTube) [17, 30], social media platforms [4, 11] to news [20]. Adequate topic diversity is also a concern in many different contexts, though most often studied in the context of online news [5, 15, 22, 24, 27]. Structural diversity is a common concern in the context of news aggregators, such as Google News [15, 26].

We now turn to the second criterion of our definition: *time*. A large share of studies does not consider a decrease over time [2, 4, 11, 15, 17, 20, 24, 26]. Broadly speaking, prior empirical work that does not consider a change in diversity over time can be classified into studies that compare diversity between (groups of) users and studies that compare diversity between algorithms. Both are a form of *recommendation stakeholder* analysis. An example of the former is the work by Haim et al. who conducted two audit studies of Google News [15]. In the first study, they measured the effect of expressing a preference for a certain topic on topic and structural diversity. They found that expressing a preference indeed makes this topic appear more often on your Google News homepage, and consequently, they observed differences in structural diversity as well. In their second study, they created four agents modeled after a media archetype and one control agent. Over the course of a week, termed the training phase, research assistants impersonated the four archetypical users on Facebook, Google+, Google Search, Amazon, and online news websites. At the end of the training phase, all five accounts queried Google News using three search terms. They found only minor differences in recommended content across the four accounts, indicating limited effects of the training on search results. It bears mentioning that the study makes no mention of interactions with Google News during the training phase. This could very well explain the absence of any differences between users, as neither of the platforms used during training is guaranteed to share user data with the purpose of personalizing Google News. In other words, it cannot be determined whether the training phase was sufficient to obtain an adequate degree of personalization, which the authors point out, albeit for different reasons. Similarly, two other highly cited research articles also researched differences between (groups of) users. Nechustai et al. also studied personalization on Google News and found virtually no differences in content recommended and, therefore, no difference in topic diversity and structural diversity [26]. Bakshy et al. on the other hand, researched to what extent ideologically different groups of users are exposed to ideologically diverse content on Facebook [4]. The latter found considerable polarization among hard news shared by users, but also that most users are exposed to ideologies different from their own. Hussein et al. investigated whether specific attributes of users (age, gender, geolocation, and watch history) lead to more (or less) misinformation in recommendations and search results on YouTube [17]. They found that for specific combinations of attributes, some misinformation topics are shown to users who have not previously

shown interest in videos promoting conspiracies. Möller et al. [24] and Liu et al. [20] presented differences in the diversity of recommendations between algorithms. The former found that there are considerable differences between algorithms in the diversity of content recommended (in all of five dimensions) and that algorithms recommend more diversely than editors would have for a given ‘source article’ [24]. The latter found markedly different political diversity between recommendations made by collaborative filtering and content-based algorithms [20]. In all of the above, the authors compared recommendation diversity between users or algorithms for a window in time, but not how that diversity changed over time. Both Hussein et al. and Möller et al. did acknowledge the need to study the longitudinal effects of personalization as part of their future work [17, 24].

Yet another study by Chitra et al. simulated a social network and studied the network’s polarization before and after introducing a network administrator who filters content based on its agreement with the user’s prior beliefs [11]. This study does not meet our definition as they do not consider the *diversity of recommendations* nor a decrease in diversity over time. Anderson et al. studied the diversity of music consumption on Spotify [2]. They found that organic consumption is more diverse than algorithmically-driven consumption and that users who increase their diversity do so by increasing their organic consumption. While they do consider an aspect of time, they consider the diversity of consumption, not recommendations.

Next, we discuss four empirical studies that do meet all of our criteria to be considered a study of the technological filter bubble [5, 22, 27, 30]. Nguyen et al. studied the topic diversity of recommendations on a movie recommendation platform (MovieLens) [27]. They found a decrease in topic diversity of recommendations for all users, indicating that a topical technological filter bubble existed on the platform. Additionally, they found that users who use recommendations rate movies more positively. Aridor et al. later described a user model that may explain the results found by Nguyen et al. [3]. Bountouridis et al. presented a simulation framework for news recommendations (SIREN) [5]. The framework simulates an online news website and its users for a fixed number of news cycles and then visualizes the evolution of diversity over time. They conducted a case study using a dataset of BBC news articles in which they simulated recommendations using five different recommendation algorithms: ItemKNN, MostPopular, Random, UserKNN, and WeightedBPRMF. A decrease in *unexpectedness* was observed for all algorithms except MostPopular and UserKNN, the latter of which gave rise to an apparent increase in unexpectedness. Tomlein et al. studied the effects of watching promoting and debunking content on misinformation in the search results and recommendations on YouTube [30]. In essence, their study had both an observational and interventional design. They concluded that even users that have previously watched videos promoting misinformation were still exposed to alternative (correct) viewpoints in search results. However, an increase in misinformation was observed in ‘What to watch next’ recommendations, showing evidence of a technological filter bubble in this location. In both places, watching debunking content ameliorated results. Both Tomlein et al. and Bountouridis et al. perform an analysis of recommendation stakeholders. The former studies the impact of algorithms (the system), whereas the latter studies the

impact of user behavior (the consumer). Lastly, Lu et al. conducted an interventional study on a news website [22]. Participants first received recommendations from the pre-existing production recommender system and afterwards, from an improved, more diverse recommender system. They found that this intervention resulted in significant increases in topic diversity of recommendations. Interestingly, none of the studies investigate the impact of the providers as a recommendation stakeholder.

6 DISCUSSION

Considering the prior conceptual and empirical work, we argue that filter bubble research has suffered from three main problems that can be (partially) resolved by our proposed operationalized definition of the technological filter bubble. We now discuss each of these issues in greater detail.

Ambiguity. The vagueness and changeability of the filter bubble concept have led to many different interpretations in the empirical works. To further add to this ambiguity, most empirical works do not give a definition or even explicitly state their interpretation of the filter bubble. The only constant we are able to identify across all empirical studies discussed in Section 5 is the notion of diversity. Many empirical studies disregard the longitudinal nature of the filter bubble. Instead, they often ask whether some (groups of) users receive more diverse recommendations than others, whereas Pariser’s original thesis assumes filter bubbles affect all of us. Other works make comparisons between different algorithms, whereas Pariser was unconcerned with specific algorithms and broadly termed all of them ‘prediction engines’. Although the study of differences in recommendation diversity between users or algorithms is interesting in its own right, these studies should not be used as evidence of the existence of a (technological) filter bubble or lack thereof. A clear, systematically, and empirically verifiable definition of the technological filter bubble is an essential tool in clearing up this ambiguity and allowing rigorous testing of the filter bubble thesis.

Negativity. Pariser’s original thesis linked the existence of a technological filter bubble, i.e. a decrease in diversity over time, to dire consequences for society [28]. As a result, the term filter bubble has a negative connotation. However, to the best of our knowledge, this leap from the technological to the societal filter bubble has not been substantiated, except in simulations where extreme content filtering is applied [11]. It is unlikely that a small decrease in recommendation diversity will lead to undesirable outcomes for society. Instead, some amount of the recommender system ‘refining its theory of you’ may be acceptable and even desired in some contexts. Dahlgren’s conceptual, and consequently our operationalized, definition allows us to reframe this discussion. The technological filter bubble is intended to be an inherently neutral concept. The identification of technological filter bubbles can then lead to future research on when (or if) they indeed lead to societal filter bubbles.

Filter Bubble Fatigue. Recently, some scholars have called for abandoning the term ‘filter bubble’ altogether. They advise to study the effects of personalization algorithms without ‘invoking the theoretical baggage of filter bubbles’ [12]. Others have started to refer to ‘filter bubbles’ and ‘echo chambers’ as ‘metaphors that have failed us’ [7]. While the term ‘filter bubble’ has been used without

discretion in the past, abandoning it altogether creates a disconnect between the common discourse and academia. Without this shared vocabulary, it is also much more difficult for researchers from different fields to find what their peers have written on the subject, hindering interdisciplinary collaboration. Instead, we argue that the concept should be clearly defined so that it can be repeatedly and rigorously tested and consequently refuted or confirmed.

7 CONCLUSION

The study of filter bubbles is plagued by ambiguity, negativity, and as a result, filter bubble fatigue. This has contributed to contrasting empirical results and minimal generalizability across empirical studies. As a result, the answers to whether or not technological or societal filter bubbles exist, on which platforms they do, and what causes them are as yet unknown. This work builds upon Dahlgren's conceptual definition of Pariser's original thesis and proposes an operationalized, systematically and empirically verifiable definition of the technological filter bubble. We state that a technological filter bubble is a 'decrease in the diversity of a user's recommendations over time, in any dimension of diversity, resulting from the choices made by different recommendation stakeholders'. We revisit prior empirical works on the filter bubble and find that many different (implicit) interpretations of a filter bubble are used. The studies that meet the technological filter bubble criteria give reason for both concern and optimism. Some observational and simulation studies have found that a decrease in the diversity of recommendations over time can indeed occur. Interventional research designs have found that the diversity of recommendations can be increased through intervention of the system stakeholder. Finally, we list what we believe are the main issues filter bubble research has faced. We believe this operationalized definition and reinterpretation of prior empirical works can bring some much needed clarity to the field. It can be used as a stepping stone towards a universally agreed-upon definition of the filter bubble and standardized methodologies for measuring them. Eventually, repeated rigorous tests using this definition and standardized methodologies should allow us to answer whether or not technological filter bubbles exist, if they lead to societal filter bubbles, and if they do, how we can combat them.

REFERENCES

- [1] Himan Abdollahpour, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Augusto Pizzato. 2020. Multistakeholder recommendation: Survey and research directions. *User Model. User Adapt. Interact.* 30, 1 (2020), 127–158. <https://doi.org/10.1007/s11257-019-09256-1>
- [2] Ashton Anderson, Lucas Maystre, Ian Anderson, Rishabh Mehrotra, and Mounia Lalmas. 2020. Algorithmic Effects on the Diversity of Consumption on Spotify. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) (WWW '20). Association for Computing Machinery, New York, NY, USA, 2155–2165. <https://doi.org/10.1145/3366423.3380281>
- [3] Guy Aridor, Duarte Goncalves, and Shan Sikdar. 2020. Deconstructing the Filter Bubble: User Decision-Making And Recommender Systems. In *Proceedings of the Fourteenth ACM Conference on Recommender Systems*. Association for Computing Machinery, New York, NY, USA, 82–91. <https://doi.org/10.1145/3383313.3412246>
- [4] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132. <https://doi.org/10.1126/science.aaa1160> arXiv:<https://www.science.org/doi/pdf/10.1126/science.aaa1160>
- [5] Dimitrios Bountouridis, Jaron Harambam, Mykola Makhortykh, Mónica Marrero, Nava Tintarev, and Claudia Hauff. 2019. SIREN: A Simulation Framework for Understanding the Effects of Recommender Systems in Online News Environments. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA) (FAT* '19). Association for Computing Machinery, New York, NY, USA, 150–159. <https://doi.org/10.1145/3287560.3287583>
- [6] Axel Bruns. 2019. Filter Bubble. *Internet Policy Review* 8, 4 (Nov. 2019). <https://doi.org/10.14763/2019.4.1426>
- [7] Axel Bruns. 2019. It's not the technology, stupid: How the 'Echo Chamber' and 'Filter Bubble' metaphors have failed us. (2019). <https://eprints.qut.edu.au/131675/>
- [8] Axel Bruns. 2022. Echo Chambers? Filter Bubbles? The Misleading Metaphors That Obscure the Real Problem. In *Hate Speech and Polarization in Participatory Society*, Marta Pérez-Escobar and José Manuel Noguera-Vivo (Eds.). Routledge, Abingdon, Oxon, 33–48. <https://doi.org/10.4324/9781003109891-4>
- [9] Jaime Carbonell and Jade Goldstein. 1998. The Use of MMR, Diversity-Based Reranking for Reordering Documents and Producing Summaries. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Melbourne, Australia) (SIGIR '98). Association for Computing Machinery, New York, NY, USA, 335–336. <https://doi.org/10.1145/290941.291025>
- [10] Nyshka Chandran. 2018. Obama to David Letterman: Media is dividing Americans. <https://www.cnn.com/2018/01/12/former-president-barack-obama-warns-on-polarizing-media-us-electoral-system.html> Accessed: 2022-05-16.
- [11] Uthsav Chitra and Christopher Musco. 2020. Analyzing the Impact of Filter Bubbles on Social Network Polarization. In *Proceedings of the 13th International Conference on Web Search and Data Mining* (Houston, TX, USA) (WSDM '20). Association for Computing Machinery, New York, NY, USA, 115–123. <https://doi.org/10.1145/3336191.3371825>
- [12] Peter Dahlgren. 2021. A critical review of filter bubbles and a comparison with selective exposure. *Nordicom Review* 42 (01 2021), 15–33. <https://doi.org/10.2478/nor-2021-0002>
- [13] Kevin J. Delaney. 2017. Filter bubbles are a serious problem with news, says Bill Gates. <https://qz.com/913114/bill-gates-says-filter-bubbles-are-a-serious-problem-with-news/> Accessed: 2022-05-16.
- [14] Mehdi Elahi, Dietmar Jannach, Lars Skjærven, Erik Knudsen, Helle Sjøvaag, Kristian Tolonen, Øyvind Holmstad, Igor Pipkin, Eivind Thordsen, Agnes Stenbom, Eivind Fiskerud, Adrian Oesch, Loek Vredenberg, and Christoph Trattner. 2021. Towards responsible media recommendation. <https://doi.org/10.1007/s43681-021-00107-7>
- [15] Mario Haim, Andreas Graefe, and Hans-Bernd Brosius. 2018. Burst of the Filter Bubble?: Effects of personalization on the diversity of Google News. *Digital Journalism* 6 (03 2018), 330–343. <https://doi.org/10.1080/21670811.2017.1338145>
- [16] Julian PT Higgins, James Thomas, Jacqueline Chandler, Miranda Cumpston, Tianjing Li, Matthew J Page, and Vivian A Welch. 2019. *Cochrane Handbook for Systematic Reviews of Interventions*. John Wiley & Sons.
- [17] Eslam Hussein, Prerna Juneja, and Tanushree Mitra. 2020. Measuring Misinformation in Video Search Platforms: An Audit Study on YouTube. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 048 (may 2020), 27 pages. <https://doi.org/10.1145/3392854>
- [18] Dietmar Jannach and Gediminas Adomavicius. 2016. Recommendations with a Purpose. In *Proceedings of the 10th ACM Conference on Recommender Systems* (Boston, Massachusetts, USA) (RecSys '16). Association for Computing Machinery, New York, NY, USA, 7–10. <https://doi.org/10.1145/2959100.2959186>
- [19] Glen Joris, Frederik De Grove, Kristin Van Damme, and Lieven De Marez. 2020. News Diversity Reconsidered: A Systematic Literature Review Unraveling the Diversity in Conceptualizations. *Journalism Studies* 21, 13 (2020), 1893–1912. <https://doi.org/10.1080/1461670X.2020.1797527> arXiv:<https://doi.org/10.1080/1461670X.2020.1797527>
- [20] Ping Liu, Karthik Shivaram, Aron Culotta, Matthew A. Shapiro, and Mustafa Bilgic. 2021. The Interaction between Political Typology and Filter Bubbles in News Recommendation Algorithms. In *Proceedings of the Web Conference 2021*. Association for Computing Machinery, New York, NY, USA, 3791–3801. <https://doi.org/10.1145/3442381.3450113>
- [21] Felicia Loecherbach, Judith Moeller, Damian Trilling, and Wouter van Atteveldt. 2020. The Unified Framework of Media Diversity: A Systematic Literature Review. *Digital Journalism* 8, 5 (2020), 605–642. <https://doi.org/10.1080/21670811.2020.1764374> arXiv:<https://doi.org/10.1080/21670811.2020.1764374>
- [22] Feng Lu, Anca Dumitrache, and David Graus. 2020. Beyond Optimizing for Clicks: Incorporating Editorial Values in News Recommendation. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*. Association for Computing Machinery, New York, NY, USA, 145–153. <https://doi.org/10.1145/3340631.3394864>
- [23] Denis McQuail. 1992. *Media Performance: Mass Communication and the Public Interest*. SAGE Publications Ltd.
- [24] Judith Möller, Damian Trilling, Natali Helberger, and Bram van Es. 2018. Do not blame it on the algorithm: an empirical assessment of multiple recommender systems and their impact on content diversity. *Information, Communication & Society* 21, 7 (2018), 959–977. <https://doi.org/10.1080/1369118X.2018.1444076> arXiv:<https://doi.org/10.1080/1369118X.2018.1444076>
- [25] PM Napoli. 1999. Deconstructing the diversity principle. *Journal of Communication* 49, 4 (1999), 7–34. <https://doi.org/10.1111/j.1460->

- 2466.1999.tb02815.x arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1460-2466.1999.tb02815.x>
- [26] Efrat Nechushtai and Seth C. Lewis. 2019. What kind of news gatekeepers do we want machines to be? Filter bubbles, fragmentation, and the normative dimensions of algorithmic recommendations. *Computers in Human Behavior* 90 (2019), 298–307. <https://doi.org/10.1016/j.chb.2018.07.043>
 - [27] Tien T. Nguyen, Pik-Mai Hui, F. Maxwell Harper, Loren Terveen, and Joseph A. Konstan. 2014. Exploring the Filter Bubble: The Effect of Using Recommender Systems on Content Diversity. In *Proceedings of the 23rd International Conference on World Wide Web (Seoul, Korea) (WWW '14)*. Association for Computing Machinery, New York, NY, USA, 677–686. <https://doi.org/10.1145/2566486.2568012>
 - [28] E. Pariser. 2011. *The Filter Bubble: What The Internet Is Hiding From You*. Penguin Books Limited. <https://books.google.be/books?id=-FWO0puw3nYC>
 - [29] Annelien Smets, Jonathan Hendrickx, and Pieter Ballon. 2022. We're in This Together: A Multi-Stakeholder Approach for News Recommenders. *Digital Journalism* (2022), 1–19. <https://doi.org/10.1080/21670811.2021.2024079>
 - [30] Matus Tomlein, Branislav Pecher, Jakub Simko, Ivan Srba, Robert Moro, Elena Stefanova, Michal Kompan, Andrea Hrcakova, Juraj Podrouzek, and Maria Bielikova. 2021. An Audit of Misinformation Filter Bubbles on YouTube: Bubble Bursting and Recent Behavior Changes. In *Fifteenth ACM Conference on Recommender Systems*. 1–11.
 - [31] Paul S. Voakes, Jack Kapfer, David Kurpius, and David Shano-Yeon Chern. 1996. Diversity in the News: A Conceptual and Methodological Framework. *Journalism & Mass Communication Quarterly* 73, 3 (1996), 582–593. <https://doi.org/10.1177/107769909607300306> arXiv:<https://doi.org/10.1177/107769909607300306>
 - [32] Sanne Vrijenhoek, Mesut Kaya, Nadia Metoui, Judith Möller, Daan Odijk, and Natali Helberger. 2021. Recommenders with a Mission: Assessing Diversity in News Recommendations. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval (Canberra ACT, Australia) (CHIIR '21)*. Association for Computing Machinery, New York, NY, USA, 173–183. <https://doi.org/10.1145/3406522.3446019>
 - [33] Frederik J. Zuiderveen Borgesius, Damian Trilling, Judith Möller, Balázs Bodó, Claes H. de Vreese, and Natali Helberger. 2016. Should we worry about filter bubbles? *Internet Policy Review* 5, 1 (March 2016). <https://doi.org/10.14763/2016.1.401>