



TECHNIUM
SOCIAL SCIENCES JOURNAL

Vol. 82/2026
A New Decade for Social Changes



PLUS
COMMUNICATION P



International
Communication & PR

Filter Bubbles in the Age of GenAI - A Literature Review

Emanuel Sanda

Bucharest Academy of Economic Studies

sandaemanuel20@stud.ase.ro

Abstract. The term "filter bubble" refers to the possibility that online content customization, resulting from the use of algorithms, could isolate users from wider or even differing perspectives. Recommender systems, which are implemented in all digital platforms and rely on algorithms to anticipate users' preferences and recommend relevant items, are particularly vulnerable to this phenomenon. The rapid development of generative artificial intelligence (GenAI) has brought in focus how recommender systems narrow what users see and strengthen the filter bubble effect. This literature review examines the recent work on the impact of GenAI on search, social media news feeds and customized recommendations. The most direct evidence suggests that GenAI can amplify selective exposure when users make use of large language model systems in ways that tend to confirm their views and expectations, as results and outputs are personalized to existing preferences. At the same time, GenAI can also be used to interrupt the exposure to more of the same information, by directing users to diverging material and alternative viewpoints. Ultimately, GenAI seems to be less of a cause for filter bubble emergence and more of an amplifier. It can intensify narrowing through personalization and distilled content, but it can also be deployed to diversify exposure if platforms choose to optimize for diversity rather than pure user engagement. The review concludes by identifying further research needed to establish causal effects across social media, search, and e-commerce..

Keywords. Artificial Intelligence, filter bubbles, generative AI, large language models, recommender algorithms

1. Introduction – filter bubbles definition, emergence and role

Filter bubbles and echo chambers are a rather common occurrence in the current digital landscape. Social media platforms like Facebook and others have played a significant role in their formation. Coined by internet activist Eli Pariser, the term "filter bubble" refers to the personalized environment that users experience online, where their content consumption is tailored to match their preferences and beliefs, while filtering out dissenting viewpoints and alternative perspectives. (Pariser, 2011)

It is an established fact that individuals have the tendency to favor information that confirms preexisting beliefs and avoid content that produces cognitive discomfort or uncertainty (Hart et. al., 2009). In our modern digital context, this tendency is not only psychological, but rather created and intensified by technologies that rank, sort, and recommend content based on users' prior activity. As such, filter bubbles arise at the intersection of both technological designs, as well as users' own behavior patterns and preferences. (Mobasher et. al., 2000).

A big concern related to filter bubbles is the reinforcement of existing beliefs and the limited exposure to more diverse viewpoints. This can create an echo chamber in which the user sees more of the same information and opinions. In turn, this can reinforce existing perspectives, leading to polarization and limiting open-mindedness on any given topic. (Bakshy et. al., 2015).

A study on how algorithmic filtering affects the information users receive found that individuals with more diverse media consumption are less likely to encounter polarizing content, suggesting that filter bubbles can lead to political polarization by reducing exposure to broader points of view. (Flaxman et. al., 2016)

Despite the risks associated with filter bubbles and echo chambers, there are also benefits of recommender algorithms. Personalized content can boost user satisfaction by displaying only relevant information based on their interests. It can also save time by reducing information overload and prioritizing content that is more likely to be engaging or useful. Additionally, it can create a sense of community among like-minded individuals, providing them with a platform where they share similar interests and can have meaningful discussions.

However, the risks associated with these recommender algorithms cannot be ignored. One of the major concerns is the potential manipulation of individuals' beliefs and behaviors through targeted content. The notorious Cambridge Analytica case revealed how personalized advertising on Facebook was used to target and influence users during political campaigns, raising serious ethical and privacy concerns.

Filter bubbles can also have unintended consequences related to the tone and content of the public conversation in a democracy. The lack of exposure to diverse perspectives can lead to a polarized society in which individuals are not willing to engage in constructive dialogue or accept others' point of view.

The filter bubble concept could perhaps be best observed in the context of algorithmically driven news feeds on various social media platforms. However, the same type of recommender systems, powered by very similar AI driven algorithms, can be found on video and music streaming platforms, online e-commerce and marketplace sites, and app stores curation. As such, filter bubbles have become part of a broader universe aimed at influencing and guiding consumer behavior. These systems not only respond to user preferences but also shape them by creating feedback loops in which their prior activity becomes the basis for their future experience. Filter bubbles are therefore relevant not only in terms of limiting exposure to alternative information and options, but also by solidifying a narrowed down attention, judgment, and preference over time.

2. The Rise of LLMs and Generative AI (GenAI)

The emergence and quick adoption of large language models (LLMs) and other forms of generative artificial intelligence (GenAI) in recent years has completely changed the information landscape users are exposed to. These changes may have a direct impact on the already formed filter bubble reality of our digital lives.

While previous technologies worked by sorting, ranking and prioritizing information which was deemed relevant to the user based on previous behavior, GenAI systems have the added ability to summarize and generate content as a response to user prompts and actions. This is a major development as it completely transforms how users interact with information. Rather than accessing multiple sources, users can now tap a single interface that responds quickly, fluently and bears the appearance of authority and knowledgeable.. in A look at ChatGPT from the perspectives of human-computer interaction and psychology (Liu, 2024), shows that

conversational AI changes what users expect in terms of convenience and personalization. This new technology may further reduce the burden associated with users' search for information, but it may also further narrow their exposure to a more diverse landscape that allows them to compare options. When users receive one curated answer rather than multiple possibilities and sources to compare, they are much more prone to a narrowing of the information at their disposal.

GenAI systems also have the capacity for highly relevant personalization. While earlier recommender systems defined relevance based on users' prior clicks, views, and purchases, GenAI systems can be more adaptable, as they take into account the user context, their input / prompt, as well as previous preferences. These systems may therefore become increasingly responsive to what users are likely to find agreeable, useful, and fit to their emotional state at any given time. While this can lead to higher user satisfaction, it also raises a serious concern: the more accurately a system can match user expectations, the more likely it is to reinforce existing biases rather than challenge them.

Perhaps even more relevant for our context is the fact that GenAI can produce personalized content extremely rapidly and at scale. As a result, social media platforms have become even more complex for the average user as they can now produce realistic, customized, and highly engaging content specifically tailored to the respective user. (Pan et al., 2024). This new reality adds to the problem selective filtering that of selective content generation. While previous systems determined what content to present to users, GenAI can determine what content will be created for that particular user in the first place.

3. GenAI impact on Social Media News Feeds

One of the most important areas in which GenAI may impact the filter bubble effect is the social media news feed. These platforms have been known for quite some time for employing algorithms that favor user engagement rather than exposure to diverse content and information. The use of GenAI in these environments may amplify this tendency in different ways. It may increase the accuracy with which platforms profile their users and their preferences, while also creating and expanding the content available to users.

GenAI enables real time generation of content that can be displayed together with human created material, thus further increasing the complexity of the platform for users (Pan et al., 2024). As such, the bubble effect no longer depends on curating existing content but can be further amplified by the creation of new content specifically tailored to a particular user, community, audience or group.

This is an important development because previously the filter bubble concerns were mostly related to excluding more diverse content and information, while GenAI brings another layer, that of saturating users with one-sided content. Users may now encounter repeated narratives exposure to same or very similar content adapted to capture and maintain their attention and reinforce existing habits. This new environment not only makes it easier to 'find' consistently similar content, but this is now served on a continuous and in a pretty convincing manner.

GenAI can significantly increase the spread of misleading or false content by increasing the volume and the personalization of content that is presented through already narrow information channels (Jain et. al., 2025). Although not directly connected to the filter bubble effect, it can be inferred that when users are present in an environment where exposure is already

filtered to their preferences, the additional generation of tailored content may only increase the persuasiveness of whatever content is present in those environments.

The impact of generating personalized content is even bigger when put in the context of exhibiting trust and persuasion in relation to the user. When LLM models such as ChatGPT are given access to a user's personal profile information, the dialogue they can engage in can be as persuasive as that with a human counterpart (Salvi et al., 2025). This is of particular interest for the study of filter bubbles because it suggests that the issue is not only what users see, but of how effectively what they see can influence their beliefs and attitudes. A bubble populated by highly customized and persuasive content may be more difficult to escape from than one simply generated by an algorithm.

Extrapolating this to e-commerce, tailored GenAI content may narrow product discovery by keeping users trapped in a universe of familiar brands, styles, or price categories. Consequently, the filter bubble problem in the age of GenAI may be broader than simply information and idea debates, extending to narrowing of user horizons across many facets of our online lives.

Still, the currently available literature on this subject is quite limited. While the workings of generating customized content is quite clear, we are yet to perform the kind of long-term studies demonstrating how such systems affect users' beliefs, preferences, and exposure over extended periods. Much of the existing work is experimental, simulated, or inferential. The evidence is significant, but it remains incomplete.

4. GenAI impact on search

An area of direct impact of GenAI is the recommender systems embedded in search. These are essential to our digital lives as they span domains such as entertainment, shopping, news, and general information search. The application of GenAI to these systems has changed not only their capabilities but also the way in which users interact with them. Recommendations are increasingly becoming part of conversations between users and chatbots, as search turns into a dialogue rather than a listing of multiple links.

In a 2024 study, Sharma et al. provide direct evidence that this shift in how search is conducted may amplify the filter bubble effect. Comparing 'classical' search with LLM-powered conversational search, the study found that users' behavior is more biased when interacting with conversational systems, especially when those systems produce responses that align with users' existing preferences. This is particularly relevant because it identifies the new conversational style search as a possible cause of creating narrower perspectives for users. A single fluent answer may discourage comparison and validate the existing bias of the users' approach to the initial search.

The implications of these findings can be quite big, especially since search engines performed the role of exposing users to multiple sources and viewpoints through ranked results. Conversational AI systems distill information into a single structured response which can be accurate and balanced but still reduce a user's exposure to multiple sources, allowing them to exercise critical thinking relative to these. Moreover, such responses are usually highly personalized, running the risk of only confirming users' preexisting biases.

Recommender systems function in a similar manner. Zhang et al. (2024), in their work on generative agents in recommendation, developed a framework for simulating how repeated interactions with recommendation systems can lead to a pattern that exposes the user to the same kind of choices repeatedly.

There is also research that focuses on the positive effects as well. Gao et al. (2023) show that recommender systems can be designed to prevent the creation of filter bubbles by incorporating mechanisms that can lead to diversity, including content diverging from what might be users' expectations and biases. Zhang et al. (2024), show that LLM-based systems can be designed to help users engage with opposing viewpoints in reflective ways. These studies prove that GenAI is not only reinforcing existing bubbles but the effect depends on how it is integrated into recommender systems and what goals these are designed to accomplish.

5. Reducing the Effects of GenAI on Filter Bubbles

If GenAI has the capacity to impact the filter bubble phenomenon, the question is what can be done to prevent further worsening of this effect. One of the most discussed solution is to include specifications that can ensure diversity even from the design phase. Platforms should not code only for engagement, but also for offering users a balanced exposure to multiple sources and views in how they filter and rank the news feed. The goal would not be to completely eliminate customization, but prevent the risk of creating an echo chamber around users.

Another possibility could be offering more transparency to users. Instead of well polished answers usually provided by GenAI chatbots, these could potentially offer multiple views as well as source, leaving room for the fact that not everything is clear cut and ambiguities are possible. Such provisions may help users retain the sense that the information provided is open to examination, interpretation before drawing their own conclusions.

At the same time, GenAI itself may be involved as part of the solution. LLM systems can be designed to encourage the comparison of alternative viewpoints. (Zhang et al., 2024) Rather than simply reinforcing what users already believe, AI systems could raise questions and counterarguments to challenge users and provoke them to deeper search and exercising critical thinking. In the same manner, recommender systems could integrate mechanisms that can balance users satisfaction with diversity of options presented. (Gao et al., 2023)

Some forms of auditing the development and implementation of such systems should also be part of the equation. Profit seeking corporations will always seek commercial goals which are easily achieved through customization and engagement, while diversity is never a top priority. As such, access for independent researchers, clearer standards for algorithm design and governmental regulation for these AI powered systems would only contribute to a more transparent digital environment.

6. Directions for Further Research

The recent work on the GenAI and filter bubbles has aimed to broaden the approach of the topic from a multidimensional point of view. Rather than treating filter bubbles as primarily related to political news, the research has looked also at how conversational interfaces (chatbots), recommender systems, AI generated media, and customization interact with each other. This is essential in understanding the role of GenAI across the multitude of digital platforms on which users are present at any given time. Still, the research is still quite concentrated on the realm of political information, misinformation, and social media, while less attention has been paid to e-commerce, entertainment, education, and health, even though these are the domains in which personalized GenAI systems may become deeply integrated into routine decision-making.

An aspect which is commendable is the fact that the recent research does not simply assume that GenAI will only exacerbate the filter bubble effect. Some studies show that the same systems can be used towards promoting diversity and varying viewpoints. This approach can encourage constructive discussions when it comes design choices, as well as private and public incentives which are at play when such choices need to be made.

Still, important weaknesses remain. Most research is still indirect or conducted under controlled conditions that may not capture actual user behavior. Lab studies, simulations, and prototype evaluations are important tools, but they do not fully substitute for natural longitudinal studies. We still need empirical studies that can show how repeated use of GenAI systems impact diversity, trust, beliefs and behaviors over time.

To establish more decisively how GenAI affects users in relation to filter bubbles, several hypotheses could be prioritized. First, longitudinal studies are urgently needed. Researchers must examine not only short term reactions to isolated GenAI interactions, but also how repeated use over weeks or months shapes informational exposure, habits, and susceptibility to reinforcement.

Future research should also expand the scope to a cross platform perspective. Users move easily across social media, search engines, online shopping, messaging applications, and AI assistants. The cumulative effect of these environments may be significantly greater than the effects of any one platform taken in isolation. A user who receives similar recommendations from multiple platforms may experience a stronger bubble effect that current single platform studies do not reflect.

Future research should also focus more explicitly on outcomes. The question is not simply whether users see narrower content, but whether such narrowing affects confidence, trust in expertise, critical reasoning, emotional wellbeing, buying behavior, or political choices. The social consequences of filter bubbles cannot be fully assessed without attention to these combined effects.

7. Conclusion

In the short time span since the emergence of LLMs and GenAI, the available research seems to suggest that GenAI acts mostly as an amplifier of existing filter bubbles. Older research, predating the rise of GenAI, had already demonstrated that digital environments can narrow information exposure for users through algorithmic curation. GenAI has impacted this environment by making the presentation of information more conversational, more personalized, and more capable of producing distilled, well organized content tailored to users' preferences.

The existing literature also points out that such outcomes are not technologically predetermined. GenAI can also be used towards creating diversity and encouraging critical thinking, thus leading to a disruption of self reinforcing recommendation loops.

The risk that GenAI can only deepen filter bubbles when used in conjunction with systems already optimized for user engagement is quite real. But GenAI can also weaken the bubble effect when integrated into design choices aimed to promote diversity of exposure, transparency, and user agency. The significance of GenAI for our online lives will depend not only on technical aspects, but also on the institutional standards and policies that govern how those capabilities are used.

References

- [1] Pariser, E., 2011. *The Filter Bubble: What the Internet is Hiding from You*. Penguin, London.
- [2] Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling Validated Versus Being Correct: A Meta-Analysis of Selective Exposure to Information. *Psychological Bulletin* <https://doi.org/10.1037/a0015701>
- [3] Mobasher B, Cooley R, Srivastava J. 2000. Automatic personalization based on web usage mining. *Communications of the ACM* 43. <https://doi.org/10.1145/345124.345169>
- [4] Bakshy E, Messing S, Adamic L. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* <https://doi.org/10.1126/science.aaa1160>
- [5] Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly* <https://doi.org/10.1093/poq/nfw006>
- [6] Liu, J. L. (2024). ChatGPT: Perspectives from human-computer interaction and psychology. *Frontiers in Artificial Intelligence* <https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2024.1418869/full>
- [7] Pan, L., Wang, C.-Y., Zhou, F., & Lu, L. (2024). Complexity of social media in the era of generative AI. *National Science Review* <https://doi.org/10.1093/nsr/nwae323>
- [8] Jain, K., & Achuthan, K. (2025). Modeling the dynamics of misinformation spread: A multi-scenario analysis incorporating user awareness and generative AI impact. *Frontiers in Computer Science* <https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2025.1570085/full>
- [9] Salvi, F., Horta Ribeiro, M., Gallotti, R., & West, R. (2025). On the conversational persuasiveness of GPT-4. *Nature Human Behaviour* <https://doi.org/10.1038/s41562-025-02194-6>
- [10] Sharma, N., Liao, Q. V., & Xiao, Z. (2024). Generative echo chamber? Effects of LLM-powered search systems on diverse information seeking. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* <https://doi.org/10.1145/3613904.3642459>
- [11] Zhang, A., Chen, Y., Sheng, L., Wang, X., & Chua, T.-S. (2024). On generative agents in recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* <https://doi.org/10.1145/3626772.3657844>
- [12] Gao, C., Wang, S., Li, S., Chen, J., He, X., Lei, W., Li, B., Zhang, Y., & Jiang, P. (2023). CIRS: Bursting filter bubbles by counterfactual interactive recommender system. *ACM Transactions on Information Systems* <https://doi.org/10.1145/3594871>
- [13] Zhang, Y., Sun, J., Feng, L., Yao, C., Fan, M., Zhang, L., Wang, Q., Geng, X., & Rui, Y. (2024). See widely, think wisely: Toward designing a generative multi-agent system to burst filter bubbles. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* <https://doi.org/10.1145/3613904.3642545>