

The Amplification Paradox in Recommender Systems

Manoel Horta Ribeiro, Veniamin Veselovsky, Robert West

EPFL

manoel.hortaribeiro@epfl.ch, veniamin.veselovsky@epfl.ch, robert.west@epfl.ch

Abstract

Automated audits of recommender systems found that blindly following recommendations leads users to increasingly partisan, conspiratorial, or false content. At the same time, studies using real user traces suggest that recommender systems are not the primary driver of attention toward extreme content; on the contrary, such content is mostly reached through other means, e.g., other websites. In this paper, we explain the following apparent paradox: *if the recommendation algorithm favors extreme content, why is it not driving its consumption?* With a simple agent-based model where users attribute different utilities to items in the recommender system, we show through simulations that the collaborative-filtering nature of recommender systems and the nicheness of extreme content can resolve the apparent paradox: although blindly following recommendations would indeed lead users to niche content, users rarely consume niche content when given the option because it is of low utility to them, which can lead the recommender system to deamplify such content. Our results call for a nuanced interpretation of “algorithmic amplification” and highlight the importance of modeling the utility of content to users when auditing recommender systems. Code available: https://github.com/epfl-dlab/amplification_paradox.

1 Introduction

On social media platforms, recommender systems bridge the gap between content creators and regular users. On the one hand, they enable users to navigate through vast content catalogs effortlessly. On the other hand, they help content creators find an audience. As recommender systems become pervasive, scholars (Whittaker et al. 2021), the media (Fischer and Taub 2019), and even the general public (Mozilla Foundation 2019) have criticized the misalignment between what recommender systems optimize for and the goals of users and society. For instance, on YouTube, one of the world’s largest social media platforms, the recommender system is perceived to amplify inappropriate or fringe content (e.g., conspiracy theories). Motivated by the concern of algorithmic amplification, recent studies using sock puppets have audited YouTube’s recommender system, showing that watching videos related to misinformation or pseudoscience causes YouTube to recommend more such content (Hussein, Juneja, and Mitra 2020; Papadamou et al. 2022; Haroon et al. 2022; Brown et al. 2022).

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

However, recent work using real navigation logs complicates this narrative, showing that YouTube’s recommender system is not the primary driver of attention toward extreme content (Hosseini et al. 2021; Chen et al. 2022). On the contrary, extreme content is often reached through other websites and is not frequently present in long algorithmically driven watching sessions. These findings are aligned with the “supply-and-demand” hypothesis for the rise of fringe content on platforms like YouTube (Munger and Phillips 2022): “problematic” content thrives because people want to consume it, and social media affordances (e.g., the ease of distributing videos to niche audiences and monetizing it) allow this demand to be met.

Here, we propose an agent-based model that explains the central paradox emerging from the aforementioned literature, which we name the “amplification paradox:” *if the recommendation algorithm favors extreme content, why is it not driving its consumption?* While our model is simpler than the recommender systems in production on platforms like YouTube, it shows how the collaborative-filtering nature of recommender systems and the nicheness of extreme content can, by themselves, explain the contradicting observations in previous work (i.e., *the algorithm favors extreme content vs. the algorithm does not drive the consumption of extreme content*). The reason is that, although blindly following recommendations would indeed lead users to niche content, users rarely consume niche content when given the option because it is of low utility to them, which can lead the recommender system to deamplify such content.

These results have key implications. First, they suggest that algorithmic audits on recommender systems are of limited utility in determining the prevalence of phenomena like radicalization, rabbit holes, and filter bubbles *if they do not model how users interact with algorithms*. To meaningfully represent reality, algorithmic audits ought to model user preferences, as users do not blindly follow recommendations (Lee et al. 2022). Second, they indicate the dynamics of extreme or harmful content (e.g., QAnon conspiracy) within algorithmically driven platforms may be explained, at least in part, by the nicheness of the content, as our model considers nothing but the popularity and co-consumption patterns of different items. Third, they highlight the need for nuance around the notion of “algorithmic amplification,” which we argue should consider the utility of content towards users.

2 Agent-Based Model

Our model captures three key ingredients present in online platforms like YouTube: 1) the recommender systems suggest items that similar users have consumed; 2) different topics appeal to different audiences; and 3) users consume content according to their internal preferences.

User preferences. We consider the scenario commonly used in the literature [e.g., Haroon et al. (2022), Hosseinmardi et al. (2021)], with five topics: *Far Left (FL)*, *Left (L)*, *Center (C)*, *Right (R)*, *Far Right (FR)* that appeal differently to individuals across the political spectrum. We illustrate our *desiderata* in Fig. 1 (left). Considering users that range from the most left-leaning (user 1 in the figure) to very right-leaning (user 100), items from a topic are high-utility to users whose political views are well-aligned, e.g., items from the *Far Left* topic are high-utility to low-index users and low-utility to high-index users. Further, the more extreme a topic, the more its utility distribution is concentrated among a few users. Last, items belonging to the same topic are indistinguishable, i.e., each item of a given topic (e.g., *Left*) has the same utility for a given user.

We operationalize this scenario by constructing a matrix M of dimensions $|U| \times |C|$ capturing user preferences, i.e., each element m_{ij} captures the utility of item j to user i . To flexibly model m_{ij} , we use the (scaled) probability mass function of the beta-binomial distribution:

$$m_{ij} = \gamma_j \left(\frac{|U|}{i} \right) \frac{\text{Beta-binomial PMF}}{\text{B}(\alpha_j, \beta_j)}, \quad (1)$$

where B is the beta function, α_j and β_j are *concentration parameters* that control the shape of the curve, and γ_j is a *scale parameter* that determines the area under the curve (when γ_j equals 1, so does the area under the curve). For each topic, we consider items that all share the same parameters, e.g., the topic *Left* has η_L items each associated with parameters α_L , β_L , and γ_L . We illustrate an M matrix constructed as described above in Fig. 1 (right), where each dot represents the utility m_{ij} associated with a user–item pair.

Recommender system. Let a collaborative-filtering recommender system (CFRS) serve items from an item catalog C to users U . The CFRS shows items to users and records which items users consume. We represent the input of the CFRS as a $|U| \times |C|$ matrix S , where each element s_{ij} equals 1 if user i has consumed item j in the past, and 0 otherwise. Let $N_U^w(i)$ be the set of the w users most similar to user i in matrix S according to the cosine similarity (\cos). We estimate a score \hat{s}_{ij} for the user–item pair $\langle i, j \rangle$ as

$$\hat{s}_{ij} = \frac{\sum_{k \in N_U^w(i)} \cos(s_{i*}, s_{k*}) s_{kj}}{\sum_{k \in N_U^w(i)} \cos(s_{i*}, s_{k*})}, \quad (2)$$

where s_{i*} is the i -th row of S , containing user i 's records.

Interaction. Let $z_i \sim \text{Poisson}(\lambda)$ be the number of interaction rounds of user i with the recommender system. At each round, the recommender system provides the user with a set of v items that the user has not yet consumed. We consider two ways in which users select an item among recommendations given to them. Either they choose items uniformly at




Figure 1: On the left, we depict the scenario considered: five topics, each corresponding to a political position that appeals differently to users. On the right, we depict the user preference matrix M , where the size of the dots represents the utility of a user–item pair. Users are ordered from the highest index (top; most right-leaning) to the lowest (bottom; most left-leaning), and items are ordered according to their topic.

random, i.e., disregarding the utility of the items (“random selection”), or they select an item j at random with a probability proportional to the item’s utility m_{ij} (“utility-informed selection”).

Relative utility. Let the relative utility r_{iq} be the percentage of content belonging to topic q that user i would consume if they choose items from the whole catalog at random with probability proportional to each item’s utility, i.e.,

$$r_{iq} = \frac{\sum_{j \in q} m_{ij}}{\sum_{j \in C} m_{ij}}. \quad (3)$$

This is similar to Chang and Ugander’s 2022 “organic model,” a counterfactual that simulates consumption without a recommender system. A topic is “amplified” [“deamplified”] by the recommender system for a user if user consumption of the topic is above [below] its relative utility; e.g., if the relative utility of the topic *Left* for user i equals 25%, but 50% of the items i consumed were from the topic, the topic is said to be amplified by the recommender system.

Simulation procedure. First, in the *burn-in phase*, we populate the matrix S . Until all users have carried out all interactions with the recommender, we 1) sample a random user i who has interacted with the recommender system fewer than z_i times, and 2) let i interact with the recommender system with utility-informed selection. Second, in the *measurement phase*, we quantify how new users would receive recommendations given an already-populated matrix S . We 1) sample a random user i and temporarily erase their corresponding row in the CFRS matrix, i.e., we set $s_{i*} = \mathbf{0}$, 2) add one item to the user vector, creating a starting condition that varies depending on the simulation, 3) let i interact with the recommender system (the selection procedure is either random or utility-informed, depending on the simulation).

Parameter summary. Altogether, our model has the following parameters: the number of users $|U|$ and items $|C|$, the number of recommendations v given and nearest neighbors w used by the recommender system, the parameter λ governing the number of times each user interacts with the recommender system, and the number of topics $|T|$ and, for each topic q , parameters α_q , β_q , γ_q , and η_q .

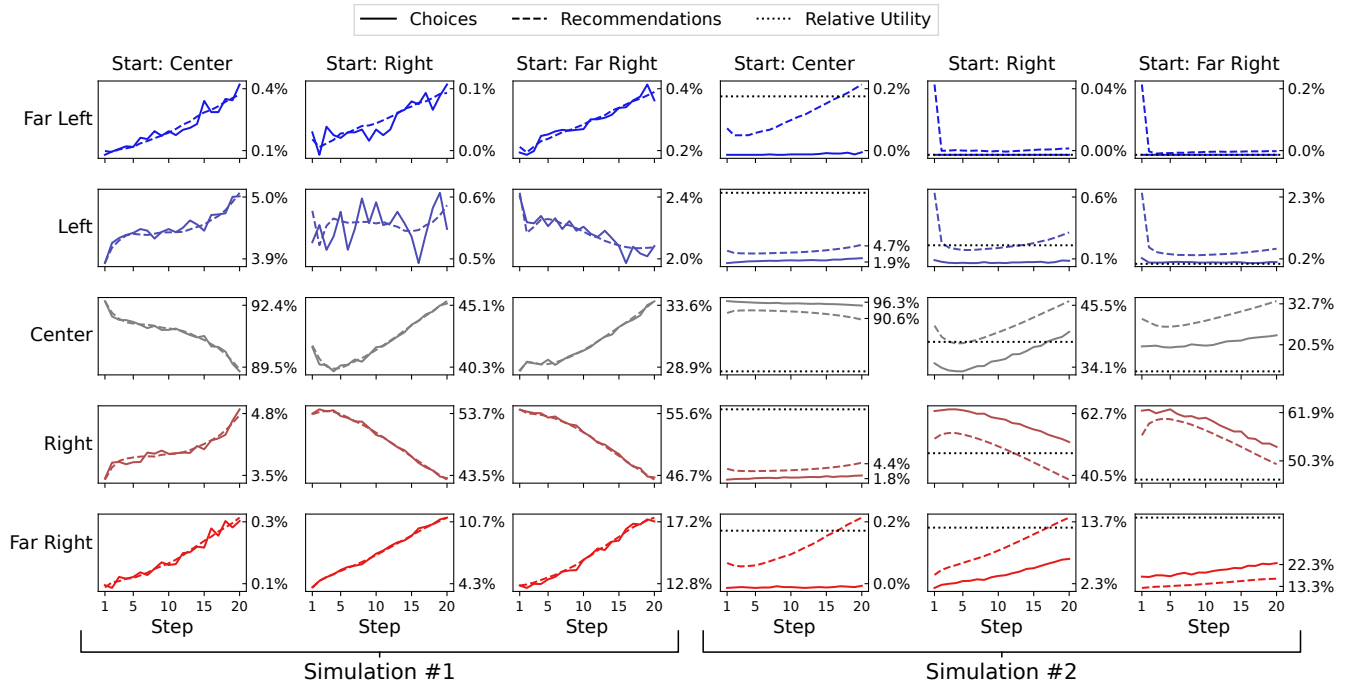


Figure 2: Simulation results. The y-axes in the plots show the percentage of times users chose (solid line) or were recommended (dashed line) an item of a specific topic (one per row) in different starting conditions (i.e., what video we initialize the user history with; one per column). The x-axis depicts the number of steps in the simulation. For the second simulation, we also show each topic’s relative utility (dotted line), a counterfactual estimate simulating consumption without a recommended system, i.e., if users choose from the whole catalog of items with probability proportional to each item’s utility to them. We omit starting conditions *Left/Far Left* as they are symmetrical to *Right/Far Right*, y scales differ per subplot.

$ T $:	5	λ :	60	α :	1, 1, 1.3, 5, 16
$ U $:	600	v :	20	γ :	1, 1.2, 1.5, 1.2, 1
$ C $:	600	w :	10	η :	75, 125, 200, 125, 75

Table 1: Parameters used in our simulation, note that for α , γ , and η , we list the parameter associated with each topic (L, CL, C, CR, R). We omit β as the parameters used are symmetrical, e.g., $\alpha_L = \beta_R$, $\alpha_{CL} = \beta_{CR}$.

3 Simulations

We conduct two simulations that attempt to explain the amplification paradox. Both share the same parameters (see Tab. 1; discussed in Sec. 4) and the same burn-in phase.

Simulation #1 examines what is recommended after users consume items from a topic and then blindly follow recommendations, similar to how recent studies audit recommender systems [e.g., Brown et al. (2022)]. If we observe that the algorithm favors niche content, we will have reached similar conclusions to previous work [assuming “extreme” content is niche, which previous work supports, see Horta Ribeiro et al. (2020)]. In our agent-based model, we operationalize this simulation by, in the measurement phase, adding a random item to the user vector in step #2 and interacting with the recommended through random se-

lections in step #3. We then analyze the percentage of times items from each topic are recommended/chosen depending on the topic of the item topic added to the user history.

Simulation #2 examines how topics are recommended and consumed when users follow their preferences. If the algorithm does not drive the consumption of extreme topics, as indicated by studies analyzing real user traces (Hosseiniardi et al. 2021; Chen et al. 2022), we would expect that these topics are not systematically amplified. In our agent-based model, we operationalize this simulation by, in the measurement step, adding an item of the topic of the highest utility to the randomly selected user in step #2 and interacting with the recommended through utility-informed selections in step #3. Again, we estimate the percentage of times each topic is recommended/chosen.

3.1 Results

We present the results of simulations #1 and #2 in Fig. 2 (left and right, respectively). The figure reads like a table. Each row shows the percentage of times a topic was recommended and chosen by users that were initialized with items from different topics, each in a column. We show only three initial conditions (*Center*, *Right*, and *Far Right*) as topics are symmetrical: e.g., the occurrence of *Right* items under the starting condition *Far Left* equals the occurrence of *Left* items under the starting condition *Far Right*, etc.

In simulation #1, we find that no matter where users start, they become increasingly exposed to content in the *Far Right* and the *Far Left*, the most niche and “extreme” of topics, e.g., users that start with one *Far Right* video in their history (third column) go from having roughly 13% of recommended videos belonging to the *Far Right* topic when interacting with the recommender system for the first time in step 1 to having around 17% in step 20. This is similar to what recent studies found when auditing the YouTube recommender system (Haroon et al. 2022; Brown et al. 2022) with bots. Note that as the selection here is random, the fraction of topics recommended and chosen are very similar.

In simulation #2, we also depict the *relative utility* of each topic to users in each initial condition as a horizontal dotted line in each plot. We find that the *Far Left* and *Far Right* topics (in the first and the fifth row, respectively) are rarely recommended to, and chosen by, users who start in the *Center* initial condition (fourth column). Considering users that start on the *Far Right* initial condition (sixth column), we see that *Far Right* content is not recommended or chosen substantially more than in Simulation #1, and *Left* items are seldom recommended and never chosen. Most important, across all starting conditions, extreme content is never chosen above the relative utility of the items to users in the starting condition. As the users are randomly sampled in the experiment, we more generally state that, on average, *Far Right* and *Far Left* items are deamplified by the recommender system. This is in accordance with the analyses of real user traces from previous work, e.g., Hosseinmardi et al. (2021) have found that consumers of extreme content do not consume more extreme content deep into long algorithmically driven viewing sessions.

4 Discussion

Our first simulation shows that the most extreme topics (*Far Left* and *Far Right*) are increasingly recommended when blindly following recommendations, similar to what Haroon et al. (2022) and other recent audits observe. However, when users choose items based on their preferences, as in our second simulation, we find that extreme topics are *deamplified* by the recommender system, i.e., users consume these topics less than they would have in the absence of a recommender system. This is aligned with empirical studies with real navigation logs (Hosseinmardi et al. 2021; Chen et al. 2022) that have found that the recommender system is not a key driver of extreme content. Since users do not meet their demand for this content through recommendations, it is only natural that they resort to subscriptions or other websites to find it. Thus, we provide a simple potential explanation for the amplification paradox: although blindly following recommendations would indeed lead users to niche content, users rarely consume niche content when given the option because it is of low utility to them, which can lead the recommender system to deamplify such content. Importantly, our findings have nothing to do with how “ideologically extreme” a topic is *but how niche it is*. Thus, we might observe this same behavior with harmless niche content (e.g., Japanese carpentry), which may appeal to a specific group.

Metaxa et al. (2021) define an algorithm audit as “a method of repeatedly and systematically querying an algorithm with inputs and observing the corresponding outputs in order to draw inferences about its opaque inner workings.” This methodology is appropriate to audit “single-round” interactions between humans and algorithms, e.g., when Buolamwini and Gebu (2018) show how commercial gender classification algorithms systematically misclassify darker-skinned women. However, recent audits of the YouTube recommender system try to uncover phenomena that, like “echo chambers,” arise from multiple interactions between humans and algorithms *without realistically modeling the human side of the interaction* (Haroon et al. 2022; Brown et al. 2022). Our agent-based model illustrates how factoring in user preferences can yield substantially different results, and, therefore, it follows that audits on YouTube are of limited utility in determining the prevalence of phenomena like radicalization, echo chambers, etc., insofar as they do not realistically model how users interact with recommender systems [see Lee et al. (2022) and Shin (2020)].

The limitations of algorithmic audits on YouTube reflect a broader issue with the notion of “algorithmic amplification.” While the term is increasingly present in the regulatory debate [see Whittaker et al. (2021)], experts have pointed out that it is ambiguous (Thorburn, Stray, and Bengani 2022) and that enforcing laws around it is challenging (Keller 2021). In our model, we adopt a “utility-based” notion of algorithmic amplification; we consider that a topic is amplified if it is systematically consumed by users attributing low utility to it. This perspective, currently not present in the regulatory debate (Keller 2021), can help stakeholders more clearly understand recommender systems.

A possible criticism of the work at hand is that this is an exceedingly simple model and that evaluations were not thorough (e.g., we did not examine the model with various parameters). We argue that these flaws do not undermine our results, as the purpose of the model is to provide a possible explanation for seemingly contradictory results in the existing literature and not to create a realistic model of how users interact with the YouTube recommender system. In a sense, this paper is analogous to an “existence proof,” showing that *there exists* a simple model that, parametrized a certain way (which we argue is reasonable), can explain the results in the literature. Nonetheless, extending the present model to be more realistic may be a worthy pursuit. Similar to how we can reason about possible answers to the “amplification paradox” given our simple model, other models that take into account how user preferences are shaped by the recommender system (Ben-Porat and Tennenholtz 2018; Cotter 2019) or how the recommender system creates incentives to produce specific kinds of content (Kalimeris et al. 2021) may help guide empirical work trying to understand the impact of recommender systems on society.

Ethical Considerations

We do not foresee a negative societal impact coming from this research, which, on the contrary, may help improve algorithmic audits of recommender systems like YouTube.

References

- Ben-Porat, O.; and Tennenholtz, M. 2018. A game-theoretic approach to recommendation systems with strategic content providers. *Advances in Neural Information Processing Systems*, 31.
- Brown, M. A.; Bisbee, J.; Lai, A.; Bonneau, R.; Nagler, J.; and Tucker, J. A. 2022. Echo Chambers, Rabbit Holes, and Algorithmic Bias: How YouTube Recommends Content to Real Users. Available at SSRN 4114905.
- Buolamwini, J.; and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, 77–91. PMLR.
- Chang, S.; and Ugander, J. 2022. To Recommend or Not? A Model-Based Comparison of Item-Matching Processes. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, 55–66.
- Chen, A. Y.; Nyhan, B.; Reifler, J.; Robertson, R. E.; and Wilson, C. 2022. Subscriptions and external links help drive resentful users to alternative and extremist YouTube videos. *arXiv preprint arXiv:2204.10921*.
- Cotter, K. 2019. Playing the visibility game: How digital influencers and algorithms negotiate influence on Instagram. *New Media & Society*, 21(4): 895–913.
- Fischer, M.; and Taub, A. 2019. How YouTube Radicalized Brazil. <https://www.nytimes.com/2019/08/11/world/americas/youtube-brazil.html>. [Online; accessed 20-Aug-2022].
- Haroon, M.; Chhabra, A.; Liu, X.; Mohapatra, P.; Shafiq, Z.; and Wojcieszak, M. 2022. YouTube, The Great Radicalizer? Auditing and Mitigating Ideological Biases in YouTube Recommendations. *arXiv preprint arXiv:2203.10666*.
- Horta Ribeiro, M.; Ottoni, R.; West, R.; Almeida, V. A.; and Meira Jr, W. 2020. Auditing radicalization pathways on YouTube. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 131–141.
- Hosseinmardi, H.; Ghasemian, A.; Clauset, A.; Mobius, M.; Rothschild, D. M.; and Watts, D. J. 2021. Examining the consumption of radical content on YouTube. *Proceedings of the National Academy of Sciences*, 118(32): e2101967118.
- Hussein, E.; Juneja, P.; and Mitra, T. 2020. Measuring misinformation in video search platforms: An audit study on YouTube. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW1): 1–27.
- Kalimeris, D.; Bhagat, S.; Kalyanaraman, S.; and Weinsberg, U. 2021. Preference amplification in recommender systems. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 805–815.
- Keller, D. 2021. Amplification and its discontents: Why regulating the reach of online content is hard. *J. FREE SPEECH L.*, 1: 227–268.
- Lee, A. Y.; Mieczkowski, H.; Ellison, N. B.; and Hancock, J. T. 2022. The Algorithmic Crystal: Conceptualizing the Self through Algorithmic Personalization on TikTok. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2): 1–22.
- Metaxa, D.; Park, J. S.; Robertson, R. E.; Karahalios, K.; Wilson, C.; Hancock, J.; Sandvig, C.; et al. 2021. Auditing algorithms: Understanding algorithmic systems from the outside in. *Foundations and Trends® in Human-Computer Interaction*, 14(4): 272–344.
- Mozilla Foundation. 2019. YouTube Regrets. <https://foundation.mozilla.org/en/youtube/findings/>. [Online; accessed 20-Aug-2022].
- Munger, K.; and Phillips, J. 2022. Right-wing YouTube: A supply and demand perspective. *The International Journal of Press/Politics*, 27(1): 186–219.
- Papadamou, K.; Zannettou, S.; Blackburn, J.; De Cristofaro, E.; Stringhini, G.; and Sirivianos, M. 2022. “It is just a flu”: Assessing the Effect of Watch History on YouTube’s Pseudoscientific Video Recommendations. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, 723–734.
- Shin, D. 2020. How do users interact with algorithm recommender systems? The interaction of users, algorithms, and performance. *Computers in Human Behavior*, 109: 106344.
- Thorburn, L.; Stray, J.; and Bengani, P. 2022. What Will “Amplification” Mean in Court? <https://techpolicy.press/what-will-amplification-mean-in-court/>. [Online; accessed 21-Nov-2022].
- Whittaker, J.; Looney, S.; Reed, A.; and Votta, F. 2021. Recommender systems and the amplification of extremist content. *Internet Policy Review*, 10(2): 1–29.