Evaluating Depolarization-Oriented News Ranking Strategies Using LLM-Generated Articles

Uroš Sergaš[1,*], Marko Tkalčič[1] and Bruce Ferwerda[2]

[1]University of Primorska Faculty of Mathematics, Natural Sciences and Information Technologies Glagoljaška 8, SI-6000 Koper, Slovenia

[2]Jönköping University, Gjuterigatan 5, 553 18 Jönköping, Sweden

Abstract

Algorithmic news ranking and large language models (LLMs) increasingly mediate how citizens receive political information. While personalization is often criticized for reinforcing echo chambers, ranking has also been proposed as a lever for depolarization by shaping exposure to cross-cutting viewpoints. We test this claim in an online experiment (N=100, Prolific) using five LLM-generated news articles on gun legislation spanning pro–gun freedom to pro–gun control. All participants read the same articles; the only manipulation was article order, instantiated as a random baseline and six stance-aware depolarization-oriented strategies (counter-narrative sandwich, balanced alternation, and directional gradients). Pre–post questionnaires measured ideological self-placement (feeling thermometer) and affective evaluations of gun-control and gun-freedom advocates. Across six research questions, we find no statistically reliable evidence that any ranking strategy reduces ideological or affective polarization relative to the baseline, nor that stance-conditioned assignment improves outcomes. These results suggest that, under the constraints of a short, single-session, single-topic exposure study with LLM-generated article stimuli, article ordering alone did not produce detectable depolarization effects. The findings motivate future work to design longer-term, multi-topic, and more interactive solutions for mitigating affective polarization.

Keywords

news recommendation, political polarization, affective polarization, ranking strategies, large language models

1. Introduction

Personalized news ranking systems and large language models (LLMs) increasingly shape how people receive political information online. While personalization and ranking are often criticized for reinforcing echo chambers and political polarization, they are also frequently proposed as tools for mitigating polarization by deliberately shaping exposure to diverse or counter-attitudinal perspectives. However, empirical evidence on whether such ranking strategies can reliably reduce political polarization remains limited and mixed.

Prior work suggests several open questions regarding depolarization-oriented news ranking. Ideological polarization refers to differences in policy positions or self-placement, whereas affective polarization captures emotional evaluations of opposing ideological groups. First, it remains unclear whether changing the order in which users encounter political content can meaningfully shift ideological self-placement, particularly in short-term exposure settings [1, 2]. Second, affective polarization may respond differently to interventions than ideological positions [3, 4]. Third, political-psychology research suggests that ideological and affective polarization are partially independent processes, raising the question of whether ranking strategies that influence one dimension also affect the other [5]. Finally, although several recommender-system designs propose tailoring exposure to a user's prior stance to mitigate selective exposure and confirmation bias [6], it remains an open empirical question whether stance-aware ranking strategies outperform non-personalized or random ordering in reducing polarization [7, 8].

Joint Proceedings of the ACM UMAP Workshops 2026, UMAP 2026, June 8–11, 2026, Gothenburg, Sweden

*Corresponding author.

uros.sergas@upr.si (U. Sergaš); marko.tkalcic@famnit.upr.si (M. Tkalčič); bruce.ferwerda@ju.se (B. Ferwerda)

0009-0007-3587-3628 (U. Sergaš); 0000-0002-0831-5512 (M. Tkalčič); 0000-0003-4344-9986 (B. Ferwerda)

Focusing on ranking order isolates one of the most basic mechanisms through which recommender systems shape exposure, as users' attention and interaction are systematically affected by item position in ranked lists [9]. In this study, we specifically ranked the sequence in which viewpoints are encountered. Even when the set of available articles is fixed, ranking can determine whether users first encounter content aligned with their view, content non-aligned with their view, or a gradual progression across perspectives. This makes article order an easily-tractable intervention for evaluation. At the same time, order alone is a deliberately limited mechanism, and its effectiveness for reducing polarization remains uncertain. We therefore treat ranking order not as a complete depolarization solution, but as a minimal test case for examining whether exposure sequencing can produce detectable short-term changes.

These open questions motivate the following research questions:

- **RQ1:** Do different news ranking strategies produce different levels of ideological depolarization, as measured by changes in a feeling thermometer on gun legislation?
- **RQ2:** Do ranking strategies influence affective polarization toward gun-control and gun-freedom advocates?
- **RQ3:** Do changes in affect toward one side co-occur with changes toward the opposing side, indicating cross-side affective convergence?
- **RQ4:** Are ideological shifts associated with affective depolarization, or do these processes unfold independently?
- **RQ5:** Does the effectiveness of ranking strategies depend on a participant's initial stance on the issue?
- **RQ6:** Does any depolarization-oriented ranking strategy outperform a random-order baseline?

To answer these research questions, we conducted a controlled online experiment in which we manipulated the order of news articles (ranking strategy) while holding content constant, and measured pre-post changes in ideological and affective polarization. All participants read the same five LLM-generated news articles on gun legislation, spanning perspectives from strict gun freedom to strict gun control. The only experimental manipulation was the order in which these articles were presented. This design does not evaluate a full news personalization system. Instead, it utilizes exposure order as a single design variable and tests whether order alone can produce detectable short-term depolarization effects, when the order is applied to a fixed set of ideologically varied articles and conditioned on prior stance.

Participants completed ideological self-placement and affective evaluations toward both sides of the issue before and after exposure. Based on their initial stance, participants were assigned either to a random-order baseline or to one of six depolarization-oriented ranking strategies. We then compared pre-post changes in ideological stance, affective attitudes toward each side, and the affective gap between them across conditions.

The contributions of this work are threefold. First, we introduce an experimental protocol and analysis procedure for evaluating depolarization-oriented news ranking strategies over LLM-generated content, jointly measuring ideological and affective outcomes. Second, we provide an initial empirical test of ranking-only depolarization interventions and show that, in this short, single-topic exposure setting, we did not observe measurable reductions in political polarization. Third, we derive design implications for human-centered recommender systems that aim to mitigate echo chambers and manipulation by explicitly targeting depolarization rather than engagement alone.

2. Related Work

News recommendation, diversity, and polarization: A central concern in personalized news recommendation is that systems optimized for user engagement tend to repeatedly surface content that aligns with users' existing preferences. Over time, this can limit exposure to alternative viewpoints, reduce viewpoint diversity, and potentially reinforce political polarization. In response, recent work has argued for normative objectives in recommender systems that explicitly promote exposure to

multiple perspectives. Heitz et al. [8] introduced deliberative diversity as a design goal and empirically examined its trade-offs with user satisfaction and perceived quality. Vrijenhoek et al. [7] proposed RADio, a rank-aware divergence framework that accounts for position bias when evaluating diversity in ranked news lists. Broader reviews emphasize that existing diversity mechanisms often remain weakly connected to societal outcomes such as polarization reduction, calling for clearer value models and outcome-driven evaluation [10]. These works provide methods for increasing diversity in ranked lists, but less direct evidence on whether such ranking interventions reduce polarization. Motivated by this gap, we evaluate whether concrete stance-aware ranking strategies contribute to measurable depolarization outcomes in a controlled setting.

Human-centered interventions and perspective exposure: Prior HCI research shows that interface design can shape how users interpret political content, for example by supporting comparison, contextualization, and reflective reading across perspectives [11]. Our study complements this work by holding the interface constant and isolating ranking order as the sole design manipulation, rather than testing broader interface or interaction effects.

Cross-cutting exposure and depolarization effects: Empirical evidence on whether exposure to counter-attitudinal content reduces polarization remains mixed. Field and experimental studies show that such exposure can sometimes backfire, increasing polarization or reinforcing motivated reasoning. Bail et al. [12] demonstrated that exposure to opposing political views on social media increased polarization among some users, while Garrett et al. [4] highlighted the complex role of selective exposure in shaping affective responses. Large-scale platform studies suggest that short-term changes in feed ranking have limited effects on political attitudes, even when exposure patterns shift [2]. Experimental work on counter-attitudinal media exposure similarly reports heterogeneous and context-dependent effects [13]. These findings provide the theoretical basis for treating ordered exposure as a plausible but uncertain intervention. Given these mixed findings, we test whether rank-ordering alone produces detectable effects when content is held fixed.

Affective polarization as a distinct outcome: Affective polarization (emotional evaluations of ideological groups) is analytically distinct from ideological distance, and shifts in policy preferences do not necessarily reduce intergroup hostility. Prior work on cross-partisan dialogue and perspective-getting suggests that affective polarization may be more responsive to sustained, interactive interventions than to passive information exposure alone [14, 15]. These insights motivate evaluating ideological and affective outcomes jointly when assessing depolarization-oriented systems. Accordingly, we measure both ideological shifts and directionally consistent affective attitudes toward each side of a single issue. We do however, acknowledge that polarization is not only an individual response to information, but also a long-term social process shaped by identity, group norms, repeated interaction, and community structure. Ranking interventions may change immediate exposure patterns, but they do not directly alter the social contexts in which political attitudes are formed and reinforced.

3. Methodology

We conducted a controlled online experiment in which we manipulated the order of news articles (ranking strategy) while holding article content constant. Participants were recruited through the Prolific online platform, where no demographic or geographic recruitment filters were applied beyond Prolific platform eligibility and English-language proficiency. Moreover, their demographic information was not stored. The participants first completed a pre-exposure questionnaire, then read five ideologically distinct news articles on gun legislation, and finally completed a post-exposure questionnaire; after each article, they answered a brief comprehension question. Anyone who answered fewer than four of the five comprehension questions correctly (that is, failed to meet an 80% comprehension threshold) was excluded from analysis, resulting in the exclusion of ten participants. The primary outcomes were pre-post changes in ideological self-placement, affective attitudes toward each side, and the affective gap between them. Participants provided informed consent, were fully debriefed after the study, and the protocol was approved by the institution's review board.

Article stimuli: The five news articles used as experimental stimuli were generated by LLM and spanned perspectives from strongly pro-gun freedom to strongly pro-gun control, with a neutral article in between. They were generated using OpenAI's GPT-4 large language model in a single-shot prompting setup. The prompt used was:

> Generate five news articles on the topic of gun control, each up to 500 words in length. The articles should differ in their degree of ideological alignment, ranging from strong support for gun ownership and gun freedom in Article 1 to strong support for gun control in Article 5. The intermediate articles should show a gradual progression between these two positions, with Article 3 being neutral and not expressing clear support for either side. Each article should cover a different incident or event and be structurally similar.

The resulting articles formed a monotonic ideological progression from pro-gun freedom to pro-gun control and were manually reviewed by the authors to verify consistency with the intended stance ordering before inclusion in the study. This verification served as an internal design check rather than as a separate participant-based validation study. The full set of article stimuli is available online [1].

Ranking strategies: Ranking strategies were selected based on prior work arguing that (i) exposure diversity can be shaped through deliberately structured sequences and (ii) rank position (i.e. ordering) can meaningfully affect which viewpoints users attend to [16, 17, 7]. Participants were assigned to one of seven ranking-strategy groups ($G0$-$G6$). The *random ordering* condition ($G0$) served as a non-targeted baseline.

The remaining strategies instantiated three families of structured exposure, each implemented in two mirrored variants based on the stance of the participant. *Counter-narrative sandwich* strategies ($G1, G5$) placed counter-attitudinal content between value-aligned articles. The two variants differed in whether the surrounding aligned content reflected an initially pro-gun freedom or pro-gun control stance. *Balanced alternation* strategies ($G2, G6$) switched between aligned and non-aligned perspectives to avoid contiguous blocks of similar content, again mirrored for pro-gun freedom versus pro-gun control participants. *Directional gradient* strategies ($G3, G4$) presented a monotonic ideological sequence, either from gun freedom to gun control or in the reverse direction, depending on the participant's initial stance. Table 1 summarizes the ranking strategies, article sequences (A1-A5), and ideological stance segments used for group assignment.

Measures: Ideological polarization was measured using a single-item feeling thermometer capturing participants' self-placement on gun legislation (0 = total gun freedom, 100 = strict gun control). Ideological change was operationalized as the pre-post difference in thermometer scores (ΔQ). Affective polarization was measured toward two target groups, advocates of gun freedom and advocates of strict gun control, using directionally consistent composite affective evaluation indices. For each group, participants rated perceived traits (open-mindedness, moderation, morality) and social distance (comfort with a family member, close friend, or coworker from the group) on 10-point Likert scales. These items were combined into A_{GF} and A_{GC} indices, with affective change computed as ΔA_{GF} and ΔA_{GC}. Affective depolarization was further captured by the pre-post change in the affective gap between groups ($G = A_{GC} - A_{GF}$) and operationalized as the pre-post difference $\Delta G = G_{post} - G_{pre}$.

Group assignment: Group assignment followed a two-stage procedure. First, 15% of participants were randomly assigned to the stance-independent baseline condition ($G0$). The remaining participants were assigned using stratified randomization based on their initial ideological self-placement on the gun-legislation feeling thermometer. Participants were divided into three ideological segments: pro-gun freedom ([0-33]), neutral ([34-66]), and pro-gun control ([67-100]). Within each segment, participants were randomly assigned to one of two strategy variants appropriate to their initial stance: pro-gun freedom participants to $G1$ or $G2$, neutral participants to $G3$ or $G4$, and pro-gun control participants to $G5$ or $G6$. As the stance was self reported, it resulted in an imbalanced group distribution: 5 participants were classified as pro-gun freedom, 19 as neutral, and 52 as pro-gun control (after comprehension-based exclusions). The purpose of this stratification was not to maximize polarization change, but to ensure

[1] https://github.com/usergas94/NewsArticlesFromStudies

that participants with different starting positions were exposed to ranking strategies designed relative to their initial stance while retaining a stance-independent baseline for comparison.

Statistical analysis: Analyses were conducted on post-pre change scores. For RQ1 and RQ2, we used one-sample t-tests against zero within each ranking group to test whether ideological and affective change scores differed from no change. To compare ranking strategies, we used Welch two-sample t-tests for pre-specified pairwise contrasts, focusing on comparisons with the random baseline ($G0$) and the stance-mirrored strategy pairs ($G1$ vs. $G2$, $G3$ vs. $G4$, $G5$ vs. $G6$). For RQ3 and RQ4, we computed Pearson correlations between ΔA_{GC} and ΔA_{GF}, and between ΔQ and ΔG, respectively. For RQ5, we examined affective change descriptively across participants grouped by initial stance. Figures report group means with 95% confidence intervals where shown.

Group	n	Strategy	Article Order	User Stance
G0	12	Random (baseline)	A2 → A5 → A4 → A1 → A3	Any
G1	1	Counter-narrative sandwich	A2 → A5 → A3 → A4 → A1	Pro gun freedom [0-33]
G2	4	Balanced alternation	A4 → A2 → A5 → A1 → A3	Pro gun freedom [0-33]
G3	7	Gradient (GF → GC)	A1 → A2 → A3 → A4 → A5	Neutral [34-66]
G4	12	Gradient (GC → GF)	A5 → A4 → A3 → A2 → A1	Neutral [34-66]
G5	29	Counter-narrative sandwich	A4 → A1 → A3 → A2 → A5	Pro gun control [67-100]
G6	23	Balanced alternation	A2 → A4 → A1 → A5 → A3	Pro gun control [67-100]

Table 1
Ranking-strategy groups, article sequences, and stance segments used in the experiment.

4. Results

Analyses use the final sample after excluding participants who failed to meet the reading comprehension threshold ($N = 90$ of 100 recruited).

[RQ1] Influence of Ranking Strategies on Ideological Polarization: To examine whether ranking strategies affected participants' ideological stance, we compared pre-post changes in the feeling thermometer (ΔQ), which ranged from total gun freedom (0) to strict gun control (100).

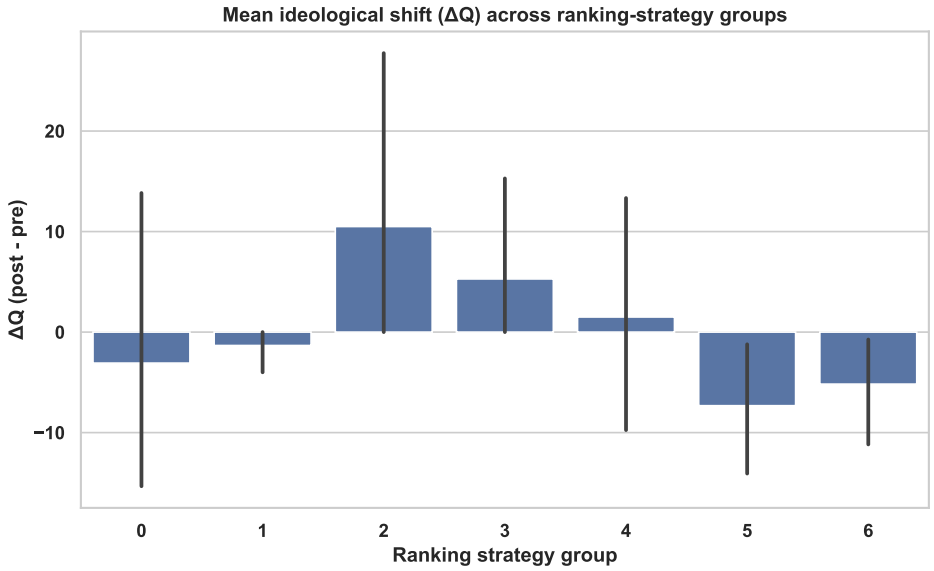

Figure 1: Mean ideological shift (ΔQ) across ranking-strategy groups.

Figure 1 shows that all groups exhibited small mean changes in ideological placement. Group $G2$ displayed the largest positive average change, while groups such as $G5$ and $G6$ showed small negative

shifts. However, all means remained within a relatively narrow range around zero. One-sample t-tests on ΔQ for each group indicated that none of these changes differed significantly from zero ($p > .10$ for all groups), suggesting that no ranking strategy produced a clear ideological shift in this sample. Accordingly, RQ1 is not supported. One interpretation is that article order alone was too weak a manipulation to shift participants' policy self-placement, especially on a topic where many participants may already have stable prior attitudes.

[RQ2] Influence on Affective Polarization: We next examined changes in affective attitudes toward both ideological groups. Affective evaluation indices were constructed as directionally consistent composites. Change scores were defined as $\Delta A_{GC} = A_{GC,\text{post}} - A_{GC,\text{pre}}$ for the group advocating gun control, and $\Delta A_{GF} = A_{GF,\text{post}} - A_{GF,\text{pre}}$ for gun-freedom advocates.

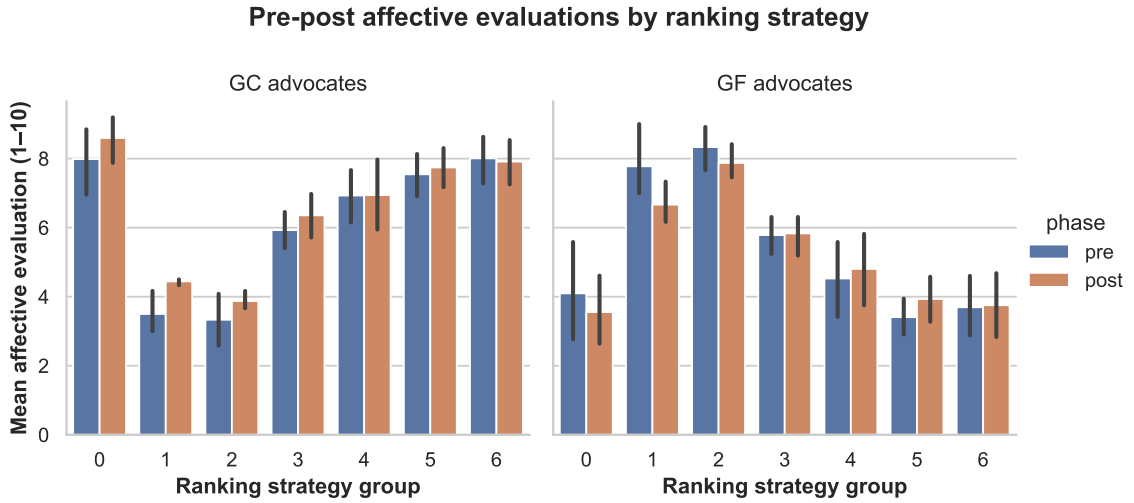

Figure 2: Pre-post affective evaluations of gun-control (A_{GC}) and gun-freedom (A_{GF}) advocates across ranking groups.

As shown in Figure 2, most groups exhibited small positive changes in affective evaluations of gun-control (GC) advocates and near-zero or slightly negative changes in affective evaluations of gun-freedom (GF) advocates. For example, some strategies for initially pro-GC participants resulted in modest increases in A_{GC}, whereas changes in A_{GF} remained close to zero or negative. Patterns for the random baseline $G0$ were comparable in magnitude to those in the other groups. No statistically significant within-group changes or between-group differences were observed for these affective outcomes ($p > .10$). Thus, RQ2 is not supported, which suggests that passive exposure to differently ordered articles does not substantially alter how participants evaluate advocates on either side.

[RQ3] Cross-Side Affective Convergence: To test whether participants who increased affective evaluations of one side also increased affective evaluations of the opposing side, we computed the correlation between the two affective change measures. The association was small and not statistically significant. Thus, more positive affective evaluations of one side did not reliably co-occur with more positive affective evaluations of the other side, a pattern that was also observed within individual ranking groups, including the random baseline $G0$. Consequently, RQ3 is not supported.

[RQ4] Relation Between Ideological Moderation and Affective Depolarization: To assess whether ideological moderation was associated with affective depolarization, we computed the affective gap as $G = A_{GC} - A_{GF}$ and its pre-post change (ΔG). We then correlated ideological change (ΔQ) with affective change (ΔG), and found no statistically significant association. Accordingly, RQ4 is not supported.

[RQ5] Moderation by Initial Stance: Participants were assigned to ranking strategies conditional on their initial thermometer score. To examine whether initial stance moderated affective responses, participants were grouped as pro-GF, neutral, or pro-GC based on the pre-exposure feeling thermometer, and changes in affective evaluation indices were inspected across these groups.

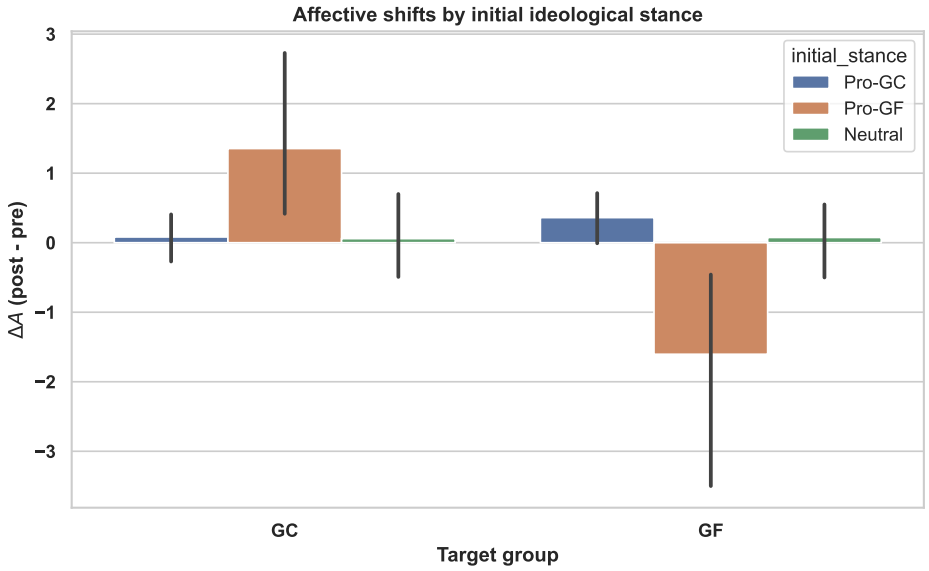

Figure 3: Affective shifts (ΔA_{GC}, ΔA_{GF}) split by initial ideological stance (pro-GF, neutral, pro-GC).

Figure 3 indicates that pro-GC participants tended to show small positive changes in ΔA_{GC} and changes near zero in ΔA_{GF}. Pro-GF participants showed the inverse pattern, with small positive changes in ΔA_{GF} and minimal movement on ΔA_{GC}. Neutral participants exhibited only small changes on either index. The descriptive comparison did not indicate clear stance-based differences in affective change.

[RQ6] Comparing Ranking Strategies to the Baseline: Finally, we tested whether any ranking strategy outperformed the random baseline ($G0$) in reducing polarization across ΔQ, ΔA_{GC}, ΔA_{GF}, and ΔG. All pairwise comparisons with $G0$ were non-significant ($p > .10$). Although $G2$ showed descriptively larger positive changes in ΔQ, these differences were small and accompanied by overlapping uncertainty intervals, suggesting that no strategy clearly outperformed the baseline on any outcome in this sample. Given the small and uneven group sizes, especially in the pro-gun freedom conditions, these comparisons should be read as exploratory.

5. Discussion

Across all analyses, we did not find evidence that depolarization-oriented news ranking strategies produced measurable ideological or affective change relative to a random baseline in this short, single-topic exposure study. Changes in ideological self-placement and affective evaluations were small and non-significant, and conditioning ranking strategies on participants' initial stances did not produce detectable improvements in outcomes. Moreover, these findings help explain why ranking-only approaches may have limited depolarization potential. Ideological and affective polarization behaved largely independently in our data. Smaller shifts in self-placement did not coincide with changes in perceptions towards opposing or supporting groups.

From a recommender-systems perspective, the results place empirical limits on ranking-centric approaches to depolarization in short-term exposure settings. Passive exposure to a small number of articles appears insufficient for achieving attitudinal change, affective polarization likely requires intervention beyond ordering alone. This does not show that ranking-based interventions are ineffective in general, but rather that article order alone did not produce detectable effects under the constraints of the present study.

6. Limitations and Future Work

These conclusions should be interpreted within the study's scope. The intervention was brief and non-interactive, focused on a single polarized issue, and deliberately isolated ranking order from selective exposure, social interaction and adaptive feedback. Additionally, because the study focused on gun legislation, a topic that is more politically relevant in the United States of America than elsewhere, participants may have varied in their familiarity with or emotional investment in the issue, which may have reduced the likelihood of observing short-term attitudinal change.

The article stimuli were LLM-generated and internally verified by the authors for intended stance progression, but not validated in a separate participant-based perception study. This increased experimental control, but may have reduced ecological validity, since LLM-generated articles may differ from real partisan journalism in style, emotional intensity and rhetorical extremity. The final sample was also imbalanced across initial stances and ranking groups, with very small cell sizes in some conditions. As a result, the comparisons between specific strategies should be interpreted as exploratory.

Future work should examine longer-term and more interactive interventions, include multiple issue domains, and test whether perceived article stance and ranking order are interpreted by users as intended. Larger and more balanced samples would also make it possible to test fully crossed designs in which each ranking strategy is evaluated within each initial-stance group. Future studies should further compare LLM-generated stimuli with real partisan and mainstream news articles, and examine whether ranking is more effective when combined with explicit reflection prompts, perspective-taking tasks, deliberative interfaces, or social discussion mechanisms.

7. Conclusion

Within these bounds, our findings suggest that ranking order alone should be treated as a limited and uncertain depolarization mechanism. In this study, changing the sequence of viewpoints did not produce detectable ideological or affective depolarization, but this absence of detectable effects should not be interpreted as evidence that ranking-based interventions are ineffective in general. Stronger evidence will require longer, larger, and more ecologically valid studies that combine ranking with explicit affective, reflective, or social intervention mechanisms.

Acknowledgments

This research was funded by the Slovenian Research Agency, grant number N2-0354.

Declaration on Generative AI

During the preparation of this work, the authors used Deepseek, ChatGPT and Deeply tools for grammar and spelling checks, minor language polishing, translating, and improving clarity of the authors' original writing. The authors reviewed and edited the content as needed and take full responsibility for the publication's content.

In addition, OpenAI's GPT-4 was used as part of the experimental design to generate the five news articles used as stimuli in the user study. These stimuli were manually reviewed by the authors for consistency with the intended ideological stance progression before being used in the experiment.

References

[1] S. A. Munson, P. Resnick, Presenting diverse political opinions: how and how much, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10, Association for Computing Machinery, New York, NY, USA, 2010, p. 1457–1466. doi:10.1145/1753326.1753543.

[2] A. M. Guess, N. Malhotra, J. Pan, P. Barberá, H. Allcott, T. Brown, A. Crespo-Tenorio, D. Dimmery, D. Freelon, M. Gentzkow, S. González-Bailón, E. Kennedy, Y. M. Kim, D. Lazer, D. Moehler, B. Nyhan, C. V. Rivera, J. Settle, D. R. Thomas, E. Thorson, R. Tromble, A. Wilkins, M. Wojcieszak, B. Xiong, C. K. de Jonge, A. Franco, W. Mason, N. J. Stroud, J. A. Tucker, How do social media feed algorithms affect attitudes and behavior in an election campaign?, Science 381 (2023) 398–404. doi:10.1126/science.abp9364.

[3] S. Iyengar, G. Sood, Y. Lelkes, Affect, not ideology: A social identity perspective on polarization, Public Opinion Quarterly 76 (2012) 405–431. doi:10.1093/poq/nfs038.

[4] R. K. Garrett, S. D. Gvirsman, B. K. Johnson, Y. Tsfati, R. Neo, A. Dal, Implications of pro- and counterattitudinal information exposure for affective polarization, Human Communication Research 40 (2014) 309–332. doi:10.1111/hcre.12028.

[5] S. Iyengar, Y. Lelkes, M. Levendusky, N. Malhotra, S. J. Westwood, The origins and consequences of affective polarization in the united states, Annu. Rev. Polit. Sci. (Palo Alto) 22 (2019) 129–146.

[6] S. Flaxman, S. Goel, J. M. Rao, Filter bubbles, echo chambers, and online news consumption, Public Opinion Quarterly 80 (2016) 298–320. doi:10.1093/poq/nfw006.

[7] S. Vrijenhoek, G. Bénédict, M. Gutierrez Granada, D. Odijk, M. De Rijke, Radio – rank-aware divergence metrics to measure normative diversity in news recommendations, in: Proceedings of the 16th ACM Conference on Recommender Systems, RecSys '22, Association for Computing Machinery, New York, NY, USA, 2022, p. 208–219. doi:10.1145/3523227.3546780.

[8] L. Heitz, J. A. Lischka, R. Abdullah, L. Laugwitz, H. Meyer, A. Bernstein, Deliberative diversity for news recommendations: Operationalization and experimental user study, in: Proceedings of the 17th ACM Conference on Recommender Systems, RecSys '23, Association for Computing Machinery, New York, NY, USA, 2023, p. 813–819. doi:10.1145/3604915.3608834.

[9] T. Joachims, A. Swaminathan, T. Schnabel, Unbiased learning-to-rank with biased feedback, in: Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM '17, Association for Computing Machinery, New York, NY, USA, 2017, p. 781–789. doi:10.1145/3018661.3018699.

[10] C. Bauer, C. Bagchi, O. A. Hundogan, K. van Es, Where are the values? a systematic literature review on news recommender systems, ACM Trans. Recomm. Syst. 2 (2024). doi:10.1145/3654805.

[11] P. Laban, C.-S. Wu, L. Murakhovs'Ka, X. A. Chen, C. Xiong, Designing and evaluating interfaces that highlight news coverage diversity using discord questions, in: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI '23, Association for Computing Machinery, New York, NY, USA, 2023, pp. 1 – 21. doi:10.1145/3544548.3581569.

[12] C. A. Bail, L. P. Argyle, T. W. Brown, J. P. Bumpus, H. Chen, M. B. F. Hunzaker, J. Lee, M. Mann, F. Merhout, A. Volfovsky, Exposure to opposing views on social media can increase political polarization, Proceedings of the National Academy of Sciences 115 (2018) 9216–9221. doi:10.1073/pnas.1804840115.

[13] A. Akbiyik, J. Bowles, H. Larreguy, S. Liu, Polarization and Exposure to Counter-Attitudinal Media in a Nondemocracy, Technical Report, Working Paper, 2023.

[14] E. Santoro, D. E. Broockman, The promise and pitfalls of cross-partisan conversations for reducing affective polarization: Evidence from randomized experiments, Science Advances 8 (2022) eabn5515. doi:10.1126/sciadv.abn5515.

[15] J. L. Kalla, D. E. Broockman, Voter outreach campaigns can reduce affective polarization among implementing political activists: Evidence from inside three campaigns, American Political Science Review 116 (2022) 1516–1522. doi:10.1017/S0003055422000132.

[16] J. Möller, D. Trilling, N. Helberger, B. van Es, Do not blame it on the algorithm: an empirical assessment of multiple recommender systems and their impact on content diversity, Information, Communication & Society 21 (2018) 959–977. doi:10.1080/1369118X.2018.1444076.

[17] Y. Ge, S. Zhao, H. Zhou, C. Pei, F. Sun, W. Ou, Y. Zhang, Understanding echo chambers in e-commerce recommender systems, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20, Association for Computing Machinery, New York, NY, USA, 2020, p. 2261–2270. doi:10.1145/3397271.3401431.