

# AI-DRIVEN NEWS RECOMMENDATION SYSTEMS: TECHNICAL CHALLENGES OF PERSONALIZATION WITHOUT AMPLIFYING INFORMATION BUBBLES

Vuk Vujović<sup>1</sup>  
ORCID ID: 0009-0001-0603-8574

**Summary:** This academic paper presents a comprehensive analysis of AI-driven news recommendation systems, with a focus on the central conflict between personalization and informational pluralism. While these systems effectively manage information overload, they risk creating “filter bubbles” and “echo chambers” by optimizing for user engagement, which can potentially isolate users from diverse perspectives. The paper reviews the evolution from classical collaborative filtering to contemporary deep learning models, identifying key technical challenges such as data quality, scalability, and the need for algorithmic diversification. It explores advanced technical solutions for mitigating informational isolation, including diversity-aware model architectures, fairness metrics, and multi-objective optimization frameworks. Specifically, it highlights the potential of reinforcement learning for long-term exposure control and knowledge graphs for contextual expansion. The study concludes that balancing personalization with informational diversity is a crucial interdisciplinary challenge. Achieving this requires sophisticated technical models that move beyond simple engagement metrics to support a more robust and healthy democratic information ecosystem.

**Keywords:** *News recommendation systems, Filter bubbles, Personalization, Algorithmic diversity, Multi-objective optimization*

---

<sup>1</sup> Vuk Vujović, assistant professor, MB University Belgrade, Serbia,  
e-mail: [vuk.vujovic@ppf.edu.rs](mailto:vuk.vujovic@ppf.edu.rs)

# **SISTEMI ZA PREPORUKU VESTI ZASNOVANI NA VEŠTAČKOJ INTELIGENCIJI: TEHNIČKI IZAZOVI PERSONALIZACIJE BEZ POJAČAVANJA INFORMACIONIH BALONA**

**Sažetak:** Ovaj akademski rad pruža sveobuhvatnu analizu sistema za preporuku vesti zasnovanih na veštačkoj inteligenciji, sa posebnim fokusom na centralni sukob između personalizacije i informacionog pluralizma. Iako ovi sistemi efikasno upravljaju preopterećenošću informacijama, istovremeno nose rizik stvaranja „filter balona“ i „odjeka komora“ usled optimizacije za angažman korisnika, što može dovesti do izolacije korisnika od različitih perspektiva. Rad prikazuje evoluciju od klasičnog kolaborativnog filtriranja do savremenih modela dubokog učenja, identifikujući ključne tehničke izazove kao što su kvalitet podataka, skalabilnost i potreba za algoritamskom raznovrsnošću. Istražuju se napredna tehnička rešenja za ublažavanje informacione izolacije, uključujući modele svesne raznovrsnosti, metrike pravičnosti i okvire višeciljne optimizacije. Posebno se ističe potencijal učenja pojačanjem (reinforcement learning) za dugoročno upravljanje izloženošću i upotreba grafova znanja za kontekstualno proširenje. Studija zaključuje da je uravnoteženje personalizacije i informacione raznovrsnosti ključni interdisciplinarni izazov. Njegovo postizanje zahteva sofisticirane tehničke modele koji prevazilaze jednostavne metrike angažovanja korisnika i doprinose stvaranju otpornijeg i zdravijeg demokratskog informacionog ekosistema.

**Ključne reči:** *sistemi za preporuku vesti, filter baloni, personalizacija, algoritamska raznovrsnost, višeciljna optimizacija*

## **1. THEORETICAL FRAMEWORK**

The proliferation of digital media has created an unprecedented challenge for news consumers who now face an overwhelming volume of content daily. Modern artificial intelligence systems for news recommendation have emerged as critical tools for managing this information overload, yet they present a fundamental tension between personalization and pluralism. While these systems promote user engagement, they can lead to the creation of echo chambers or “filter bubbles”, effectively isolating users from diverse perspectives (Bojić, 2024). This research examines how technical recommendation models can optimize personalization while mitigating the reinforcement of filter bubbles and echo effects.

Filter bubbles, coined by Pariser, are created by personalized recommendation systems that expose users to content similar to their beliefs (Hartman et al., 2025), raising critical questions about the algorithmic design choices that balance engagement with informational diversity. The study addresses the core research question of how technical models can achieve optimal personalization without exacerbating informational isolation. This challenge sits at the intersection of computer science, media studies, and democratic discourse.

## **1.1. Classical recommendation systems architecture**

Traditional news recommendation systems have evolved through several distinct paradigms, each offering different approaches to content filtering and user modeling. Content-based filtering systems analyze the intrinsic features of news articles, including textual content, categories, and metadata, to match user preferences with similar articles. These systems excel in transparency and can provide explanations for their recommendations, but they suffer from a limited scope and tend to reinforce existing preferences.

Collaborative filtering represents the second major approach, leveraging user behavior patterns and community preferences to generate recommendations. This method identifies users with similar consumption patterns and suggests content based on collective behavior. The approach addresses the limitation of content-based systems by discovering unexpected connections between articles and users..

Hybrid models combine both approaches to capitalize on their respective strengths while mitigating individual weaknesses. An example of such a system is the HRS-IU-DL model, a novel hybrid recommendation system that advances the field by integrating multiple sophisticated techniques (Sami et al., 2024). These systems typically employ weighted combinations or switching mechanisms to select the most appropriate recommendation strategy based on available data and user context. The evolution toward hybrid approaches reflects the recognition that single-method systems cannot adequately address the complexity of news consumption patterns.

## **1.2. Contemporary AI Approaches in News Recommendation**

Deep learning has revolutionized news recommendation systems by enabling more sophisticated pattern recognition and user modeling capabilities. For instance, a survey by (Talha et al., 2023) provides a detailed analysis of DL

model-based techniques for building NRS, demonstrating how early systems heavily relied on keyword matching and fundamental semantic analysis. Neural networks, particularly deep neural networks (DNNs), recurrent neural networks (RNNs), and transformer architectures, have demonstrated superior performance in capturing complex user preferences and content relationships. A recent review (Alam & Ahmed, 2024) provides a thorough summary of deep learning in recommendation systems, highlighting how these approaches can simultaneously process vast amounts of textual data and user interaction histories.

Graph neural networks (GNNs) represent a particularly promising direction for news recommendation, as they can model complex relationships between users, articles, topics, and entities. GNNs excel at capturing both local and global patterns in user-item interaction graphs, enabling more accurate predictions of user preferences while maintaining computational efficiency. For example, one study (Alam & Ahmed, 2025) presented a deep learning-based collaborative filtering approach using a Restricted Boltzmann Machine (RBM) on the Movielens-1M dataset, illustrating how machine learning techniques can capture complex user interaction patterns, a principle that is also applied in the news domain

Context-aware recommendation systems incorporate temporal, geographical, and situational factors into the recommendation process. These systems recognize that news consumption patterns vary based on time of day, current events, user location, and device type. The integration of contextual information allows for more dynamic and responsive recommendations that adapt to changing user needs and environmental factors. This approach is particularly relevant for news recommendation, where timeliness and relevance to current events are critical factors.

### **1.3. Information bubbles and algorithmic bias**

The concept of filter bubbles, introduced by Eli Pariser, describes how personalization algorithms can inadvertently create information silos that limit exposure to diverse viewpoints and perspectives. Echo chambers and filter bubbles are potent metaphors that encapsulate widespread public fear that the use of social media may limit the information that users (Kitchens et al., 2020) encounter, potentially undermining democratic discourse and informed decision-making. These phenomena occur when recommendation algorithms prioritize engagement metrics over informational diversity, leading to increasingly narrow content consumption patterns.

Echo chambers represent a related but distinct phenomenon. While filter bubbles are often described as the result of algorithmic filtering creating an individualized information space, echo chambers are characterized by a social dynamic in which individuals actively choose to reinforce their beliefs within a homogenous community. However, the extent of this problem is a subject of academic debate. For instance, a literature review by the Reuters Institute (Arguedas et al., 2022) suggests that echo chambers may be less widespread than commonly assumed, noting that several empirical studies have found no support for the filter bubble hypothesis. This indicates that the relationship between algorithmic recommendation and information diversity is more complex than initially theorized.

Reinforcement learning approaches in recommendation systems can inadvertently amplify these effects through their optimization processes. When systems optimize for immediate engagement metrics such as click-through rates or time spent reading, they may inadvertently reinforce patterns that lead to decreased diversity over time. The mathematical formulation of reward functions in these systems often fails to account for long-term consequences on user information diet and societal outcomes.

## **2. TECHNICAL CHALLENGES IN NEWS RECOMMENDATION**

News recommendation systems face significant challenges in interpreting user feedback, particularly the distinction between explicit and implicit signals (Li et al., 2025). Explicit feedback includes direct user actions such as ratings, shares, and bookmarks, which provide clear preference indicators but are relatively sparse in typical news consumption scenarios. Implicit feedback encompasses clicks, dwell time, scroll behavior, and reading completion rates, which are abundant but potentially misleading.

### **2.1. Data quality and feedback mechanisms**

Click-through rates may reflect curiosity rather than genuine satisfaction, while dwell time can be influenced by factors such as article length, complexity, or multitasking behavior, rather than genuine interest. Engagement bias represents a critical challenge where systems optimize for metrics that may not align with user satisfaction or informational value. Short-term engagement signals, such as

immediate clicks, may conflict with long-term user satisfaction and informational needs. Users might click on sensational headlines but prefer substantive journalism for their overall news consumption experience. This misalignment between optimization targets and user welfare creates a fundamental tension in system design.

The temporal dynamics of news consumption add another layer of complexity to feedback interpretation. Breaking news events can dramatically shift user attention patterns, making historical data less reliable for future predictions. Users may engage heavily with crisis-related content during specific periods but return to different consumption patterns afterward. This temporal instability requires systems to rapidly adapt while avoiding overfitting to temporary trends.

## **2.2. Cold Start and scalability issues**

The cold start problem (Panda & Ray, 2022) manifests in two primary forms: new users and new content. New users lack interaction history, making it difficult to generate personalized recommendations without relying on generic popularity-based suggestions. New articles, particularly breaking news, must be recommended before sufficient user interaction data accumulates, requiring systems to rely on content analysis and early adopter behavior patterns.

Scalability challenges emerge from the need to process millions of articles and user interactions in real-time while maintaining recommendation quality. News platforms often serve millions of users simultaneously, each requiring personalized recommendations updated continuously as new content becomes available. This is compounded by real-time processing requirements that distinguish news from domains like e-commerce, where updates can be less frequent. Breaking news requires immediate incorporation into the system, demanding low-latency processing pipelines and efficient model updating. The computational complexity of advanced recommendation algorithms, particularly deep learning and graph-based approaches, can become prohibitive at scale without careful architectural design. This temporal pressure often forces trade-offs between algorithmic sophistication and system responsiveness.

## **2.3. Diversification and algorithmic balance**

Diversification in news recommendation requires balancing multiple competing objectives: relevance, novelty, diversity, and fairness. One such approach is Sampled Maximal Marginal Relevance (SMMR) (Liakhnovich et al., 2025), which

was proposed to address these trade-offs. Maximal Marginal Relevance (MMR) algorithms attempt to balance relevance with diversity by iteratively selecting items that are both relevant to user preferences and dissimilar to previously selected items. However, implementing MMR in real-time news recommendation systems requires careful tuning of diversity parameters and efficient similarity computation methods.

Determinantal Point Processes (DPPs) provide an alternative mathematical framework for ensuring diversity in recommendations. The application of Determinantal Point Processes (DPPs) at YouTube (Wilhelm et al., 2018) demonstrates how these probabilistic models can be applied to large-scale recommendation systems. DPPs model the probability of selecting sets of items while naturally incorporating both quality and diversity constraints, but they require sophisticated implementation to achieve computational efficiency at scale.

The challenge of explainability in news recommendation systems becomes particularly acute when balancing multiple objectives. Users and regulators increasingly demand transparency in algorithmic decision-making, yet the most effective diversification algorithms often involve complex mathematical formulations that are difficult to explain in intuitive terms. This tension between performance and explainability requires careful consideration of model architecture and user interface design.

## **2.4. Multi-objective optimization framework**

Modern news recommendation systems must optimize for multiple, often competing objectives (Hong et al., 2025). These systems must strike a balance between relevance, diversity, and fairness, requiring sophisticated frameworks to handle trade-offs between short-term engagement and long-term user satisfaction. The mathematical formulation of multi-objective optimization typically involves weighted combinations of different loss functions or Pareto optimization approaches. Relevance objectives focus on matching user preferences and predicting engagement, while diversity objectives seek to minimize redundancy and maximize topical coverage. Fairness objectives aim to ensure the equitable representation of diverse viewpoints, sources, and demographic groups in recommendations.

Reinforcement learning approaches offer promising solutions for dynamically managing these trade-offs. As demonstrated in (Mengjie et al., 2025), RL algorithms can learn to balance engagement with polarization reduction over time. These systems can adapt their optimization strategies based on observed long-

term outcomes rather than relying on the fixed weight combinations determined at design time.

### **3. TECHNICAL SOLUTIONS FOR MITIGATING INFORMATION BUBBLES**

Contemporary approaches to diversity-aware news recommendation employ several sophisticated algorithmic strategies designed to expand rather than narrow user information exposure.

#### **3.1. Diversity-aware model architectures**

A recent review on GNN-based recommender systems (Duricic et al., 2023) focuses on these dimensions, moving beyond the conventional accuracy-centric perspective. Graph neural networks hold particular promise for enhancing diversity because they can model complex relationships between users, content, and topics while explicitly optimizing for diverse recommendation sets.

Knowledge graph integration represents another powerful approach for expanding the scope of recommendations beyond immediate user preferences. By incorporating structured knowledge about topics, entities, and their relationships, systems can identify conceptually related but topically distinct content that maintains relevance while promoting exploration. These systems can map user interests to broader knowledge domains and suggest content that expands rather than reinforces existing knowledge boundaries.

Advanced neural architectures now incorporate diversity constraints directly into their loss functions rather than treating diversification as a post-processing step. These end-to-end approaches learn to balance relevance and diversity simultaneously during training, potentially achieving better overall performance than systems that optimize for engagement first and diversify afterward. The mathematical formulation typically involves regularization terms that penalize overly similar recommendation sets or reward coverage of different topical areas.

### **3.2. Fairness metrics and algorithmic accountability**

The integration of fairness considerations into news recommendation systems requires careful definition and measurement of equitable outcomes across different user groups and content categories. A survey on fairness-aware recommender systems (Jin et al., 2023) provides an in-depth taxonomy. It explores the relationship between fairness and accuracy in recommendation systems, highlighting the complex trade-offs involved in achieving both objectives simultaneously.

Fairness metrics in news recommendation often focus on balancing representation across different political viewpoints, demographic groups, and information sources. These metrics might measure the distribution of liberal versus conservative content, local versus national news coverage, or representation of different social groups in recommended articles. Implementation requires careful consideration of how to define and measure these categories without introducing additional biases or oversimplifying complex ideological landscapes.

Algorithmic accountability mechanisms include audit trails, explanation systems, and feedback loops that allow users and oversight bodies to understand and challenge recommendation decisions. These systems must balance transparency with user privacy and competitive concerns while providing meaningful insights into algorithmic behavior. The technical implementation often involves developing interpretable model architectures or post-hoc explanation systems that can articulate why particular recommendations were made.

### **3.3. Reinforcement learning for long-term exposure control**

Reinforcement learning frameworks offer sophisticated approaches to managing long-term user information exposure patterns rather than optimizing for immediate engagement alone. These systems can learn to balance short-term user satisfaction with long-term goals such as maintaining ideological diversity or preventing filter bubble formation. The RL formulation typically treats recommendation as a sequential decision problem where actions (recommendations) have both immediate rewards (engagement) and long-term consequences (exposure patterns).

Successful implementation depends on the state representation and reward function. The state representation must capture not only current user preferences but also historical exposure patterns and diversity metrics, such as the political leaning of recently consumed content, topical coverage over time, or source

diversity. The resulting state space becomes high-dimensional and requires careful design to capture relevant information without becoming computationally intractable. The reward function poses another critical challenge; multi-objective (Wang et al., 2025) optimization approaches can incorporate both engagement and diversity objectives. The reward structure must incentivize both immediate user satisfaction and long-term exposure diversity, often requiring carefully tuned combinations of multiple reward signals.

### **3.4. Knowledge graph enhancement for contextual expansion**

Knowledge graphs offer structured representations of entities, concepts, and relationships, significantly enhancing contextual understanding for recommendation systems (Zhang et al., 2024). By mapping news articles to entities like people, organizations, locations, and events, systems can identify conceptual connections that transcend simple textual similarity. This enables recommendations that maintain topical relevance while introducing users to related but distinct information domains.

Integrating these graphs requires sophisticated algorithms to accurately link entities from news text to their graph representations using modern natural language processing techniques, including named entity recognition and entity linking. Once linked, graph traversal algorithms can systematically explore neighborhoods to identify diverse but related content for recommendation. These algorithms might start from entities in a user's recently consumed articles and explore connected entities to find topically related but distinct content. The challenge lies in balancing relevance (closer graph distances) with diversity (exploring different graph regions) while maintaining computational efficiency.

## **4. COMPARATIVE ANALYSIS OF PLATFORM APPROACHES**

### **4.1. Youtube's diversification strategies**

YouTube's recommendation system has undergone significant evolution in response to concerns about radicalization and filter bubble effects. For example, YouTube's implementation of DPP-based diversification algorithms, detailed in their 2018 paper (Wilhelm et al., 2018), aims to balance engagement with content variety. The system employs multiple stages of filtering and ranking, with diversification occurring both within and across different content categories.

The platform's approach to mitigating problematic recommendation patterns includes both algorithmic modifications and policy interventions. Concerns about the YouTube recommender system and filter bubbles are also reflected in case studies highlighting the potential negative consequences of the (Yesilada & Lewandowsky, 2022) recommendation algorithms, leading to systematic efforts to reduce recommendations of potentially harmful content. These interventions include reducing recommendations of borderline content and increasing promotion of authoritative sources for news and information.

Recent studies have examined the effectiveness of these interventions with mixed results. A substantial body of literature argues that recommendation algorithms contribute to political polarization by creating "filter bubbles" and "rabbit holes" (Liu et al., 2025), yet empirical evidence for significant short-term effects remains limited. The complexity of measuring long-term exposure effects and the challenge of controlled experimentation at scale make definitive evaluation difficult.

### **4.2. Facebook news and information quality controls**

Facebook's approach to news recommendation emphasizes source credibility and information quality alongside personalization factors. The platform has implemented multiple layers of content evaluation, including algorithmic scoring of source reliability, fact-checking partnerships, and user feedback mechanisms for reporting problematic content. These systems aim to strike a balance between user engagement, information quality, and concerns related to democratic discourse.

The platform's News Feed algorithm incorporates signals about source trustworthiness and content authenticity into its ranking mechanisms. This includes down-ranking content from sources that frequently share misinformation and promoting content from sources that meet journalistic standards. The implementation requires maintaining databases of source quality scores and continuously updating these assessments based on new information and user behavior.

Engagement-based optimization remains a central component of Facebook's recommendation system, but the platform has modified its reward functions to incorporate quality and authenticity signals. This includes reducing the weight of engagement metrics for content from unreliable sources and adjusting the time windows over which engagement is measured to account for the difference between immediate viral spread and sustained user value.

### **4.3. Google News and source diversity**

Google News employs a cluster-based approach that groups related stories from multiple sources and presents diverse perspectives on significant events. This architectural choice explicitly promotes source diversity by preventing any single source from dominating coverage of particular topics. The clustering algorithms group articles based on semantic similarity, while maintaining source diversity within each cluster.

The platform's ranking algorithms balance multiple factors including article freshness, source authority, topic relevance, and geographic relevance to the user. In summary, studies in the UK and several other countries, have found that algorithmic selection generally leads to slightly more diverse news (Arguedas et al., 2022) consumption patterns compared to manual selection, suggesting that well-designed algorithms can enhance rather than reduce information diversity.

Geographic and linguistic diversity receive particular attention in Google News's recommendation algorithms. The system attempts to present both local and international perspectives on events while accounting for language preferences and regional relevance. This requires sophisticated natural language processing capabilities and cultural understanding that extends beyond simple translation.

#### **4.4. Cross-platform engagement-diversity trade-offs**

Comparative analysis across platforms reveals different philosophical approaches to balancing engagement with information diversity. YouTube's emphasis on video engagement, for example, leads to different optimization challenges compared to text-based platforms, as video consumption patterns and user attention mechanisms differ significantly from article reading behavior. This creates significant methodological challenges due to differences in content types, user interfaces, and the availability of data. Despite these variations, a common challenge emerges. As one study (Huang & Yang, 2024) suggests, recommendation algorithms on multi-purpose platforms may inadvertently redirect users away from news through two key pathways: direct competition with more engaging entertainment content and the optimization for metrics that do not align with informational value.

The trade-off between engagement and diversity also varies based on content domain, user demographics, and temporal factors. For instance, breaking news events tend to reduce diversity as user attention concentrates on specific topics, while during routine news cycles, recommendation systems have more flexibility to promote diverse content consumption. Understanding these dynamics is crucial for developing effective diversification strategies.

### **5. DISCUSSION AND IMPLICATIONS**

The challenge of integrating personalization with pluralism in news recommendation systems represents a fundamental tension between individual preferences and collective democratic needs (Shin & Zhou, 2024).

#### **5.1. Theoretical integration of personalization and pluralism**

Theoretical frameworks from both computer science and political communication suggest that optimal solutions require moving beyond simple trade-offs toward more sophisticated models that recognize the interdependence of individual and social outcomes. Personalization systems that successfully maintain pluralism may actually achieve better long-term user satisfaction by preventing preference polarization and maintaining user interest in diverse topics.

Mathematical formulations of this integration challenge require multi-objective optimization approaches that can balance immediate user satisfaction with longer-term exposure diversity goals. The temporal dimension introduces complexity because short-term preferences may diverge from long-term user welfare, and individual preferences may not take into account social externalities. Successful integration requires recommendation systems that can model both individual preference evolution and the collective health of the information ecosystem.

Recent advances in reinforcement learning and multi-agent systems offer promising approaches for modeling these complex interactions. Systems that can simulate the long-term consequences of recommendation strategies on both individual users and information ecosystems may identify solution spaces that achieve both personalization and pluralism goals simultaneously. This requires sophisticated modeling capabilities and extensive computational resources but may represent the most promising path forward.

## **5.2. Practical implementation challenges**

Translating theoretical frameworks into practice poses significant technical and organizational challenges. Real-world implementation requires balancing algorithmic sophistication with computational efficiency, particularly given the real-time constraints of news recommendation. Adaptive algorithms are needed to address the dynamic nature of user preferences and the need for diversity in recommended news articles (Raza & Ding, 2022). On the organizational side, challenges include aligning business incentives with diversity goals, training personnel to manage complex algorithmic systems, and developing evaluation metrics that capture both engagement and diversity outcomes. Furthermore, many news organizations lack the technical expertise and resources to implement sophisticated recommendation systems, which may create competitive disadvantages compared to large technology platforms.

Regulatory and policy considerations add another layer of complexity. Different jurisdictions are developing varying approaches to algorithmic accountability and transparency requirements, creating compliance challenges for global platforms. The technical requirements for providing algorithmic explanations and audit capabilities may conflict with system optimization goals and competitive concerns.

### **5.3. Interdisciplinary research opportunities**

The challenges of designing a news recommendation system require interdisciplinary collaboration among computer scientists, media researchers, political scientists, and ethicists. Computer science provides the technical tools for building sophisticated recommendation systems, but understanding their social impacts requires expertise from communication and political science. Ethical frameworks for evaluating algorithmic systems require philosophical analysis combined with empirical research on social outcomes.

Collaborative research opportunities include developing new evaluation metrics that capture both individual and social outcomes, designing user interface approaches that promote diverse information consumption, and creating algorithmic accountability frameworks that balance transparency with system effectiveness. These research directions require sustained collaboration across disciplinary boundaries and significant investment in interdisciplinary research infrastructure.

The development of standardized evaluation frameworks for news recommendation systems represents a significant research opportunity. Current evaluation approaches often focus on traditional information retrieval metrics that may not capture the social and democratic functions of news consumption. New frameworks should incorporate measures of information diversity, source credibility, and long-term user learning outcomes alongside traditional engagement metrics.

## **CONCLUSION**

The development of AI-driven news recommendation systems that effectively balance personalization with information diversity represents both a technical challenge and a democratic imperative. This research has examined the theoretical foundations, technical approaches, and practical implementations necessary to achieve this balance. The evidence suggests that personalization and pluralism need not be fundamentally incompatible objectives, but achieving their integration requires sophisticated algorithmic approaches and careful attention to system design choices. Technical solutions including multi-objective optimization, reinforcement learning frameworks, and knowledge graph enhancement offer promising paths forward for creating recommendation systems that promote both user engagement and information diversity. The implementation of these approaches requires moving beyond simple engagement-based optimi-

zation toward more nuanced models that account for long-term user welfare and social outcomes. Platform case studies demonstrate that different approaches to this challenge are possible, but none have yet achieved optimal integration of personalization and pluralism goals. The path forward requires continued interdisciplinary collaboration between computer scientists, media researchers, and policy experts to develop both technical solutions and regulatory frameworks that support democratic information ecosystems. Future research should focus on developing standardized evaluation metrics that capture both individual and social outcomes, creating algorithmic accountability mechanisms that balance transparency with effectiveness, and investigating the long-term consequences of different recommendation strategies on information consumption patterns and democratic discourse. The ultimate success of these efforts will determine whether algorithmic systems can support rather than undermine the informational foundations of a democratic society.

## REFERENCES

1. Alam, M. M., Ahmed, M. (2024). Deep Learning-Based Recommendation Systems: Review and Critical Analysis. In: Swaroop, A., Polkowski, Z., Correia, S.D., Virdee, B. (eds) Proceedings of Data Analytics and Management. ICDAM 2023. Lecture Notes in Networks and Systems, vol 785. Springer, Singapore.
2. Alam, M. M., Ahmed, M. (2025). Deep Learning Based Collaborative Filtering Recommendation System. *Procedia Computer Science*, 258, 2362–2371.
3. Arguedas, A.R., Robertson, C.T., Fletcher, R. & Nielsen, R.K. (2022). Echo chambers, filter bubbles, and polarisation: a literature review. Oxford: Reuters Institute, University of Oxford. Retrieved on August 6th, 2022 from [https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2022-01/Echo\\_Chambers\\_Filter\\_Bubbles\\_and\\_Polarisation\\_A\\_Literature\\_Review.pdf](https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2022-01/Echo_Chambers_Filter_Bubbles_and_Polarisation_A_Literature_Review.pdf)
4. Bojić, L. (2024). AI alignment: Assessing the global impact of recommender systems. *Futures*, 160, 103383.
5. Duricic, T., Kowald, D., Lacic, E., & Lex, E. (2023). Beyond-accuracy: a review on diversity, serendipity, and fairness in recommender systems based on graph neural networks. *Frontiers in big data*, 6, 1251072.
6. Gao, C., Zheng, Y., Li, N. et al. (2023). A Survey of Graph Neural Networks for Recommender Systems: Challenges, Methods, and Directions. *ACM Trans. Recomm. Syst.*, 1 (1), Art. 3.

7. Guess, A. M. et al. (2023). How do social media feed algorithms affect attitudes and behavior in an election campaign? *Science*, 381, 398–404.
8. Hartmann, D., Wang, S. M., Pohlmann, L. et al. (2025). A systematic review of echo chamber research: comparative analysis of conceptualizations, operationalizations, and varying outcomes. *J Comput Soc Sc*, 8, 52.
9. Hong, Z., Wu, Y., Zhao, Z., Feng, S., Ma, J., Liu, J., & Wei, T. (2025). Multi-Objective Recommendation in the Era of Generative AI: A Survey of Recent Progress and Future Prospects. *ArXiv*, abs/2506.16893.
10. Huang, S., & Yang, T. (2024). Auditing Entertainment Traps on YouTube: How Do Recommendation Algorithms Pull Users Away from News. *Political Communication*, 41(6), 903–920.
11. Jin, D., Wang, L., Zhang, H., et al. (2023). A survey on fairness-aware recommender systems. *Information Fusion*, 100, 101906.
12. Kitchens, B., Johnson, S. L., Gray, P. (2020). Understanding Echo Chambers and Filter Bubbles: The Impact of Social Media on Diversification and Partisan Shifts in News Consumption. *Management Information Systems Quarterly*, 44 (4), 1619–1649.
13. Li, W., Kuo, J.-C., Sheng, M. et al. (2025). Beyond Explicit and Implicit: How Users Provide Feedback to Shape Personalized Recommendation Content. *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, Art. 893, 1–17.
14. Liakhnovich, K., Lashinin, O., Babkin, A. et al. (2025). SMMR: Sampling-Based MMR Reranking for Faster, More Diverse, and Balanced Recommendations and Retrieval. *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2754–2758.
15. Liu, N., Hu, X. E., Savas, Y. et al. (2025). Short-term exposure to filter-bubble recommendation systems has limited polarization effects: Naturalistic experiments on YouTube. *Proc. Natl. Acad. Sci. U.S.A.*, 122 (8).
16. Mengjie, C., Cheng, O., Ofek, E. et al. (2025). Balancing Engagement and Polarization: Multi-Objective Alignment of News Content Using LLMs. 10.48550/arXiv.2504.13444.
17. Panda, D. K., Ray, S. (2022). Approaches and algorithms to mitigate cold start problems in recommender systems: a systematic literature review. *J Intell Inf Syst*, 59, 341–366.
18. Raza, S. & Ding, C. (2022). Relevancy and Diversity in News Recommendations. Retrieved on August 10th, 2025 from <https://ceur-ws.org/Vol-3411/INRA-paper1.pdf>.
19. Sami, A., Adrousy, W. E., Sarhan, S. et al. (2024). A deep learning based hybrid recommendation model for internet users. *Sci Rep*, 14, 29390.
20. Shin, D., & Zhou, S. (2024). A Value and Diversity-Aware News Recommendation Systems: Can Algorithmic Gatekeeping Nudge Readers to View Diverse News? *Journalism & Mass Communication Quarterly*, 0(0).

21. Talha, M. M., Khan, H. U., Iqbal, S. et al. (2023). Deep learning in news recommender systems: A comprehensive survey, challenges and future trends. *Neurocomputing*, 562, 126881.
22. Wang, J., Karatzoglou, A., Arapakis, I. et al. (2025). Beyond Accuracy: Decision Transformers for Reward-Driven Multi-Objective Recommendations. *IEEE Transactions on Knowledge and Data Engineering*, 37 (9), 5004–5016.
23. Wilhelm, M., Ramanathan, A., Bonomo, A. et al. (2018). Practical Diversified Recommendations on YouTube with Determinantal Point Processes. *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, 2165–2173.
24. Yesilada, M., & Lewandowsky, S. (2022). Systematic review: YouTube recommendations and problematic content. *Internet policy review*, 11(1), 1652.
25. Zhang, J.-C., Zain, A. M., Zhou, K.-Q. et al. (2024). A review of recommender systems based on knowledge graph embedding. *Expert Systems with Applications*, 250, 123876.

### **Cite – APA Style:**

Vujović, V. (2025). AI-driven news recommendation systems: technical challenges of personalization without amplifying information bubbles. U: B. Sančanin & A. Penjišević (ur.), *Zbornik radova VII međunarodne naučne konferencije Veštačka inteligencija u kreiranju medijskih sadržaja* (pp. 62-79). Sremski Karlovci: Centar za kulturu, edukaciju i medije; Istočno Sarajevo–Sokolac: Visoka škola za uslužni biznis.

Рад је примљен: 22. септембра 2025.

Received: September 22, 2025

Прихваћен за објављивање: 2. октобра 2025.

Accepted: October 2, 2025