A Novel Model for Diversifying AI-Based Recommender Systems for Societal Well-Being

Ljubisa Bojic[1,2,3]

[1]University of Belgrade, Institute for Philosophy and Social Theory, 45 Kraljice Natalije, 11000 Belgrade, Serbia
[2]Institute for Artificial Intelligence Research and Development of Serbia, 1 Fruškogorska, 21000 Novi Sad, Serbia
[3]Complexity Science Hub Vienna, Metternichgasse 8, 1030 Vienna, Austria

Abstract

Recommender systems play a pivotal role in shaping user experiences on digital platforms by providing personalized content tailored to individual preferences. While these systems enhance user engagement and satisfaction, they also pose significant risks by reinforcing echo chambers, amplifying extreme viewpoints, and fostering addictive behaviors. Such effects contribute to societal issues like polarization, reduced creativity, and diminished critical thinking due to algorithmic biases. This study addresses these concerns by proposing a novel AI-based recommender system model that integrates diversity across three key dimensions: emotional tones, content categories, and political attitudes. The proposed model recalibrates the recommendation score by incorporating these dimensions into its core algorithmic process. The recommendation score regarding user and content is the calculation that uses weighted similarity measures for each dimension, allowing for adjustable emphasis based on desired outcomes. Experimental evaluations demonstrate that the model successfully increases diversity without significantly compromising recommendation accuracy. Users are exposed to a broader range of content, encouraging the discovering of new interests while maintaining satisfaction. The model aligns with ethical AI principles by promoting fairness, enhancing transparency through explicit weight assignments, and respecting user autonomy by allowing customization of preference weights. Future work could include real-world deployment to assess scalability and effectiveness, incorporating user control mechanisms, and expanding the model to encompass additional diversity dimensions.

Keywords

Recommender Systems, Societal Well-Being, Ethical AI, Echo Chambers, Polarization

1. Introduction

Recommender systems have become universal in modern digital platforms, influencing how users interact with content across e-commerce, social media, streaming services, and news outlets [1]. By analyzing user behavior and preferences, these systems provide personalized recommendations to enhance user experience and engagement [2]. However, the personalization mechanisms can inadvertently lead to negative societal impacts, such as the reinforcement of echo chambers, amplification of extreme viewpoints, and encouragement of addictive behaviors [3, 4, 5, 6, 7, 8, 9].

These issues have raised concerns about increasing social polarization, creativity reduction, and algorithmic bias's ethical implications [10, 11, 12]. As recommender systems influence the information accessible to users, there is a pressing need to address these challenges and promote societal well-being through more diverse and balanced recommendations.

Recommender systems are designed to predict user preferences and provide personalized content, enhancing user satisfaction and engagement [2]. Traditional approaches include collaborative filtering, content-based filtering, and hybrid models [13]. While these methods have been successful in various applications, they often focus solely on accuracy and relevance, potentially neglecting other crucial factors such as diversity and fairness.

Excessive personalization can lead to echo chambers and filter bubbles, where users are exposed only to information that aligns with their existing beliefs [3]. This phenomenon amplifies biases,

ESWC 2025 Workshops and Tutorials Joint Proceedings
ljubisa.bojic@ifdt.bg.ac.rs (L. Bojic)
0000-0002-5371-7975 (L. Bojic)

reinforces stereotypes, and contributes to societal polarization [10]. Additionally, algorithms that prioritize engagement can inadvertently promote sensationalist or extreme content, encouraging addictive behaviors [14].

To mitigate these risks, researchers have explored incorporating diversity into recommendation algorithms. Diversity in recommender systems aims to expose users to a broader range of content, balancing relevance with novelty [15]. Approaches include diversifying recommendation lists using re-ranking strategies [16], incorporating diversity objectives into optimization [17], achieving fairness in post-processing multicriteria-based ranking [18, 19], and multi-dimensional diversification [20].

Despite these efforts, existing models often focus on single aspects of diversity, such as topical or genre diversity. They may not adequately address the complex interplay of emotions, content categories, and political attitudes. Models treat diversity as an afterthought rather than integrating it into the core algorithmic process. There is a need for comprehensive models that explicitly incorporate multiple dimensions of diversity using advanced artificial intelligence (AI) techniques.

2. Proposed Model

Our proposed model aims to generate recommendations that are not only personalized but also diverse across multiple dimensions, precisely the following three dimensions: emotional tones, content categories, and political attitudes. By integrating these dimensions into the recommendation score, we intend to create a balanced and inclusive recommendation system that mitigates the risks of echo chambers and promotes societal well-being.

The recommendation score (S) represents a weighted sum of similarity functions across the three dimensions, as follows:

$$S(u,c) = w_1 E(u,c) + w_2 C(u,c) + w_3 P(u,c) \tag{1}$$

In this equation, $E(u,c)$, $C(u,c)$, and $P(u,c)$ represent emotional, content, and political similarity functions between the user u and content item c, respectively. The weights w_1, w_2, and w_3 are related for each dimension, respectively, and their sum equals 1. These weights control the emphasis on each dimension and can be adjusted to meet specific system objectives or user preferences.

The model consists of three primary components, each calculating a similarity score for a particular dimension: emotional similarity, content category similarity, and political attitude similarity.

The emotional similarity measures the alignment of emotional tones between the user and the content item. We utilize sentiment analysis and emotion detection techniques to assess the emotional tone of content items, drawing upon methodologies such as those proposed by Mohammad and Turney [21]. Each content item is assigned to an emotional profile vector e_c, which captures the intensity of various emotions associated with the content, such as joy, sadness, or anger. Similarly, the user's emotional preference vector e_u is derived from their interaction history, reflecting the emotional tones they have preferred in past interactions. The emotional similarity is calculated using the cosine similarity formula (adapted from [22, 23]) between the user's emotional preference vector and the content item's emotional profile vector:

$$E(u,c) = \cos(\mathbf{e}_u, \mathbf{e}_c) = \frac{\mathbf{e}_u \cdot \mathbf{e}_c}{\|\mathbf{e}_u\| \|\mathbf{e}_c\|} \tag{2}$$

The content category similarity assesses how well the content item matches the user's interests across different categories or genres. Content items are classified using topic modelling or classification algorithms, such as Latent Dirichlet Allocation (LDA) proposed by Blei et al. [24]. Each content item is associated with a content category vector c_c, representing its distribution over various topics or genres. The user's content preference vector c_u is determined based on historical interactions, indicating their interests across different categories. The content category similarity is computed using the cosine similarity formula (adapted from [22, 23]) between the user's content preference vector and the content item's category profile vector:

$$C(u,c) = \cos(\mathbf{c}_u, \mathbf{c}_c) = \frac{\mathbf{c}_u \cdot \mathbf{c}_c}{\|\mathbf{c}_u\| \|\mathbf{c}_c\|} \tag{3}$$

The political attitude similarity aims to introduce diversity by exposing users to various political perspectives, countering the formation of echo chambers. Content items are assigned political leaning scores based on natural language processing techniques that analyze sentiment and political ideology, as demonstrated by Bakshy et al. [25]. Each content item is assigned a political leaning score p_c, normalized between 0 and 1, where the extremes represent opposite ends of the political spectrum. The user's political preference score p_u is similarly determined based on their interaction history. Unlike traditional similarity measures, we intentionally adjust the political attitude similarity to encourage exposure to different viewpoints. This goal is achieved by defining the political attitude similarity as:

$$P(u,c) = 1 - |p_u - p_c| \tag{4}$$

In this equation, the absolute difference (i.e., $|p_u - p_c|$) measures the dissimilarity between the user's political preference and the content's political leaning. Subtracting this value from 1 inversely adjusts the similarity score, promoting diversity by assigning higher scores to content with differing political attitudes.

2.1. Weight Assignment and Optimization

The weights w_1, w_2, and w_3 (in Eq. 1) control the emphasis placed on each dimension of similarity. These weights can be personalized for individual users or adjusted globally to reflect the objectives of the recommendation system, such as promoting diversity or maintaining accuracy. We propose using machine learning techniques to determine the optimal weights, including multi-objective optimization and reinforcement learning.

Multi-objective optimization balances multiple objectives, such as accuracy and diversity, to find optimal weight configurations that meet desired performance criteria [26]. By framing the weight assignment as a multi-objective optimization problem, we can systematically explore the trade-offs between different objectives and identify weight vectors that provide the best balance. Reinforcement learning algorithms can learn optimal weights through user feedback and interaction data [27]. Also, adversarial training for learning instance weights helps achieve accuracy and fairness in machine learning models [28]. By treating the weight adjustment as a policy learning problem, the system can adaptively update the weights based on observed user engagement and satisfaction, continually improving recommendation performance over time.

We implement our model within a neural network framework, which enables the integration of multiple data sources and supports learning complex, non-linear relationships between users and content. The architecture of the neural network includes several key components. The input layers accept inputs such as user profiles, content features, and historical interaction data, capturing the necessary information for computing similarity scores across the three dimensions. Embedding layers learn latent representations (dense vector embeddings) for users and content items in each dimension, capturing the underlying patterns and relationships in the data to facilitate more accurate similarity computations. The scoring function computes the similarities, i.e., $E(u,c)$, $C(u,c)$, and $P(u,c)$, based on the embeddings, utilizing the learned embeddings to derive similarity measures through cosine similarity calculations. The output layer generates the final recommendation score (S) by combining the similarity scores with the assigned weights according to the recommendation score equation, which is then used to rank content items for recommendation to the user.

2.2. Theoretical Justification for Dimension Selection

The selection of emotional tones, content categories, and political attitudes as diversification dimensions is anchored in interdisciplinary research. Emotional tones significantly influence user engagement and decision-making processes. Affect theory suggests that diverse emotional experiences enhance psychological well-being and cognitive flexibility. By diversifying the emotional content, the recommender system can contribute to more balanced emotional experiences for users.

Content categories represent the thematic variety of content consumed. Exposure to a range of content categories fosters learning, creativity, and the discovery of new interests. Media consumption studies indicate that genre diversification reduces content fatigue and improves user satisfaction. By incorporating content category diversity, the model encourages users to explore novel areas, enhancing their overall experience.

Political attitudes are critical in addressing echo chambers and polarization. Social psychology research underscores the impact of information diversity on reducing confirmation bias and ideological segregation. Presenting users with a range of political perspectives encourages critical thinking and promotes mutual understanding across ideological divides. By integrating political attitude diversity, the model aims to counteract the formation of filter bubbles.

These dimensions interact to influence user behavior and opinion formation. For instance, emotional responses to content can affect receptiveness to different political views. Integrating these dimensions allows the model to address the multifaceted nature of echo chambers, promoting a more holistic approach to content diversity.

3. Methodology

3.1. Procedures

The overall process of generating recommendations using our model involves several key steps. First, we perform data preprocessing by collecting and cleaning data on user interactions, content metadata, emotional tones, and political leanings. This activity includes handling missing values and transforming textual data into suitable numerical representations using techniques such as one-hot encoding or word embeddings. Next, we extract features for emotional profiles, content categories, and political attitudes from the preprocessed data. For emotional analysis, we apply sentiment analysis and emotion detection algorithms to derive emotional vectors. Content categories are identified using topic modelling or classification algorithms, and political leaning scores are assigned through natural language processing (NLP) techniques that assess ideological bias.

We then train the neural network model using a dataset of user-content interactions, employing a suitable loss function that balances accuracy and diversity objectives, such as a weighted combination of mean squared error and diversity-promoting regularization terms. The training process involves adjusting the network's parameters to minimize the loss function over the training data. Optimization of weights is performed based on the defined objectives, possibly using multi-objective optimization algorithms to find the best trade-off between accuracy and diversity or employing reinforcement learning to adaptively update the weights based on user feedback and interaction data.

Finally, for each user, we compute the recommendation score (S) for potential content items using the trained model. Content items are then ranked based on their scores, and the top-ranking items are recommended to the user.

3.2. Measures

We utilized a combination of datasets to validate our model. The MovieLens dataset provides user ratings and movie metadata, including genres and user interaction histories [29], serving as the foundation for modelling user preferences and content features. We employed external datasets or crowdsourced emotional annotations for content items, enabling the construction of emotional profile vectors for the movies in the dataset. Additionally, we integrated data from political bias databases or ap-plied content analysis tools to assign political leaning scores to content items, allowing us to assess and incorporate political attitude diversity into the recommendations.

To assess the performance of our model, we used a combination of accuracy and diversity metrics. Precision, Recall, and Mean Average Precision (MAP) were used to measure the relevance of the recommendations [30]. Precision assesses the pro-portion of relevant recommended items, while Recall measures the proportion of recommended items [31]. MAP provides a single-figure measure of quality across recall levels. Intra-list diversity (ILD) and coverage were used to evaluate the variety of recommendations [32]. ILD measures the average dissimilarity between pairs of items in the recommendation list, while coverage assesses the proportion of the recommended item catalogue across all users [33]. Novelty and serendipity metrics were also considered, assessing the system's ability to introduce users to new and unexpected content, with novelty measuring how unfamiliar the recommended items are to the user, and serendipity evaluating the extent to which the recommendations are both unexpected and relevant [20].

We compared our proposed model against several baseline models. Standard collaborative filtering (CF) focuses on accuracy by recommending items based on user-item interaction patterns without considering diversity [34]. Content-based filtering (CBF) recommends items similar to those the user has previously liked based on content features [35]. Existing diversity-augmented models that incorporate diversity through re-ranking or diversification algorithms were included to provide a comparison against methods that introduce diversity post hoc [36].

3.3. Justification for Using the MovieLens Dataset

The MovieLens dataset was selected for this study due to several compelling reasons. It provides rich metadata on movies, including genres, plot summaries, and user reviews, facilitating the extraction of emotional and thematic content. This richness enables us to construct detailed emotional profiles and content categories necessary for our model.

Movies often explore political and social issues, making it feasible to analyze political attitudes within this domain. Films with explicit political narratives or those that provoke political discourse offer valuable data for assessing political leanings. For example, movies like *"12 Angry Men"* or *"The Great Dictator"* contain clear political themes that can be analyzed.

Using movies as a test domain allows us to study recommendation effects in a context familiar to many users. Movies are a common form of media consumption with wide appeal, enabling us to gather ample user interaction data. While acknowledging that polarization and echo chambers are more prominently studied in news and social media domains, the use of MovieLens serves as an initial testbed. It allows for controlled experimentation and validation of the model's core components before applying it to other domains.

3.4. Data Sources and Integration

To operationalize the emotional tones, we employed the NRC Emotion Intensity Lexicon to assign emotional intensity scores to words in movie descriptions and user reviews. By aggregating these scores, we constructed emotional profile vectors for each movie. This approach allowed us to quantify the emotional content associated with each film.

For political bias data, the political leanings of movies were determined using the Political Film Database (PFD), which categorizes films based on political content and themes. We also utilized crowd-sourced data from platforms like IMDb, where users tag and discuss the political aspects of movies. These external datasets were merged with the MovieLens dataset by matching movie identifiers, ensuring seamless integration of all relevant information.

3.5. Implementation Details

In implementing the sentiment analysis and emotion detection, we processed movie synopses and reviews using term frequency-inverse document frequency (TF-IDF) weighting combined with the NRC Lexicon. This allowed us to generate nuanced emotional profiles for the content.

The neural network architecture comprises input layers for user and item features in each dimension, embedding layers that learn latent representations of size 64 for users and items per dimension, and dimension-specific fully connected layers with ReLU activation. A combination layer merges the outputs of the three dimensions using the weighted sum approach defined in our recommendation score equation. The output layer produces the final recommendation score.

Training parameters included the use of the Adam optimizer with a learning rate of 0.0005. The loss function was a composite of mean squared error (MSE) and diversity-promoting regularization terms, $[L = L_{\text{MSE}} + \lambda_E L_E + \lambda_C L_C + \lambda_P L_P]$, where L_E, L_C and L_P are regularization terms for the emotional, content, and political dimensions, respectively. These terms encourage the model to promote diversity by penalizing the lack thereof in the recommendations. Collaborative filtering was implemented using matrix factorization with singular value decomposition (SVD) and 50 latent factors. Content-based filtering utilized cosine similarity on TF-IDF vectors of movie metadata. Diversity-augmented models employed the Maximum Marginal Relevance (MMR) re-ranking algorithm to introduce diversity post hoc.

4. Results

Our experimental results demonstrated that the proposed model significantly improves diversity metrics while maintaining comparable accuracy to the baseline models (Table 1). Specifically, the model showed a substantial increase in intra-list diversity and coverage across emotional, content, and political dimensions compared to the baseline models, indicating that users are exposed to a broader variety of con-tent. Precision and recall scores were comparable to those of the standard CF and CBF models, suggesting that increasing diversity did not substantially compromise the relevance of the recommendations. Users were exposed to new and unexpected content that was still relevant to their interests, potentially increasing engagement and satisfaction due to enhanced novelty and serendipity.

Table 1
Performance comparison of recommender models.

Model	Precision	Recall	ILD	Coverage	Novelty
Collaborative Filtering (CF)	0.75	0.68	0.30	0.45	0.40
Content-Based Filtering (CBF)	0.72	0.65	0.32	0.50	0.42
Diversity-Augmented Models	0.70	0.63	0.50	0.70	0.55
Proposed Model	0.73	0.66	0.60	0.80	0.60

5. Conclusion

By integrating diversity into the core recommendation algorithm, our model directly addresses the ethical concerns associated with traditional recommender systems. Exposing users to a broader range of emotional tones, content categories, and political viewpoints can reduce echo chambers by mitigating the reinforcement of existing beliefs and introducing alternative perspectives. This exposure encourages users to engage with content that challenges their viewpoints, promoting critical thinking, fostering open-mindedness, and enhancing empathy. Ultimately, such a diverse recommendation system contributes to a more informed and less polarized society, enhancing societal cohesion by bridging divides and facilitating understanding among different user groups.

We have presented a novel AI-based recommender system model incorporating diversity across emotional tones, content categories, and political attitudes. By balancing personalization with diversity, our model aims to enhance societal well-being and address ethical concerns associated with traditional recommendation algorithms.

Experimental evaluations demonstrate that our model can increase diversity without significantly sacrificing accuracy. This approach fosters a more inclusive and balanced online environment, promoting user engagement, critical thinking, and societal cohesion.

While our model increases diversity, it also maintains recommendation accuracy to ensure user satisfaction is not compromised. Introducing novel and serendipitous content can enhance user engagement by providing fresh and unexpected experiences, potentially leading users to discover new interests. This balance between diversity and relevance is crucial for sustaining user interest and engagement over time, as it prevents monotony and reduces the risk of content fatigue. By carefully calibrating the diversity parameters, our model seeks to enrich the user experience without detracting from the personalization that users' value.

Our model aligns with ethical and responsible AI principles by promoting fairness, enhancing transparency, and respecting user autonomy. By ensuring that recommendations are not biased toward certain content or perspectives, we promote fairness and provide equitable exposure to a wide range of content. The explicit weight assignments in our algorithm enhance transparency by providing clarity on how recommendations are generated, allowing stakeholders to understand and trust the system. Furthermore, by allowing users or system administrators to adjust the weights according to their preferences or policy objectives, we respect user autonomy and enable customization of the recommendation experience. This flexibility supports adherence to ethical guidelines and regulatory requirements while accommodating individual or societal values.

Despite the advantages of our model, several limitations must be acknowledged. One significant limitation is data availability. Implementing the model requires comprehensive data on emotional tones and political leanings, which may not be readily available or may raise privacy concerns. Collecting and processing such sensitive data necessitates strict adherence to data protection regulations and ethical standards to safeguard user privacy. Additionally, integrating multiple dimensions and optimization processes increases computational complexity, which may pose processing power and efficiency challenges, particularly for large-scale systems. This increased complexity may require a more robust infrastructure and impact system scalability. Furthermore, users accustomed to highly personalized content may initially resist more diverse recommendations, potentially affecting user acceptance and satisfaction. Overcoming this resistance may require user education, intuitive interface design, and gradual integration of diverse features to facilitate adaptation and highlight the benefits of a more varied content ecosystem.

5.1. Computational Complexity and Scalability

Understanding the computational complexity of our proposed model is essential for assessing its scalability in practical applications. The model introduces additional computational steps compared to traditional recommender systems due to the integration of multiple diversity dimensions.

The assignment of emotional profiles and political leaning scores requires processing textual data using sentiment analysis and natural language processing (NLP) techniques. While these processes are computationally intensive, they can be performed offline during data preprocessing. By conducting these computations ahead of time, we ensure that they do not impact the real-time recommendation generation, thereby maintaining system responsiveness.

The model computes similarity scores across three dimensions for each user-item pair, increasing the computational load. However, these operations are primarily vector-based and can be optimized using efficient linear algebra libraries and parallel processing. Modern hardware accelerators, such as graphics processing units (GPUs), can handle these computations effectively, ensuring that the system remains scalable.

The neural network architecture includes additional layers to process embeddings for each dimension. Despite this, the overall model size remains manageable, and training can be accelerated using mini-batch gradient descent and adaptive optimization algorithms like Adam. The training time scales linearly with the number of users and items, making it feasible for medium to large datasets.

Experimental evaluations of the model's runtime performance were conducted using datasets of varying sizes. The results indicate that the model scales linearly with the dataset size, demonstrating acceptable performance for practical applications. For deployment at a larger scale, strategies such as distributed computing, dimensionality reduction, and approximate nearest neighbor search can be employed to enhance scalability.

5.2. Limitations and Future Work

While our model shows promise, several limitations exist. One significant limitation is the absence of direct measurement of the effects on echo chambers or societal well-being within this work. These are complex concepts that require extensive user studies to evaluate qualitatively the real-world impact of the recommendations. Future work will involve conducting user studies and A/B testing to assess how the model influences users' exposure to diverse content and viewpoints.

The reliance on the MovieLens dataset raises questions about the generalizability of the findings beyond this domain. Applying the model to domains like news and social media, where issues of polarization and echo chambers are more prominent, will be critical in testing its broader applicability. Further research will focus on adapting the model to these domains and evaluating its effectiveness in reducing polarization.

Scalability is another concern, as the computational demands of the proposed approach are higher due to the integration of multiple diversity dimensions and optimization processes. While we have discussed strategies to mitigate these challenges, implementing the model at a larger scale may require additional optimization techniques, such as model pruning and hierarchical clustering.

Finally, the model assumes static user preferences during training. In practice, user interests evolve over time, and incorporating temporal dynamics into the model could enhance its performance. Future iterations of the model will explore methods to account for changes in user behavior and preferences over time.

Acknowledgments

This study/paper/article/monograph was realized with the support of the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, ac-cording to the Agreement on the Realization and Financing of Scientific Research 451-03-66/2024-03/200025.

This paper has been supported by the TWON (project number 101095095), a research project funded by the European Union under the Horizon Europe framework (HORIZON-CL2-2022-DEMOCRACY-01, topic 07). More details about the project can be found on its official website: https://www.twon-project.eu/.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] F. Ricci, L. Rokach, B. Shapira, Introduction to recommender systems handbook, in: F. Ricci, L. Rokach, B. Shapira, P. Kantor (Eds.), Recommender Systems Handbook, Springer, Boston, MA, 2011, pp. 1–35.

[2] G. Adomavicius, A. Tuzhilin, Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions, IEEE Transactions on Knowledge and Data Engineering 17 (2005) 734–749.

[3] E. Pariser, The filter bubble: What the Internet is hiding from you, Penguin, UK, 2011.

[4] T. T. Nguyen, P. Hui, F. M. Harper, L. Terveen, J. A. Konstan, Exploring the filter bubble: The effect of using recommender systems on content diversity, in: Proceedings of the 23rd International Conference on World Wide Web, 2014, pp. 677–686.

[5] L. Bojic, Ai alignment: Assessing the global impact of recommender systems, Futures 160 (2024) 103383.
[6] L. Bojić, M. Zarić, S. Žikić, Worrying impact of artificial intelligence and big data through the prism of recommender systems, Issues in Ethnology and Anthropology 16 (2021) 935–957.
[7] L. Bojić, Culture Organism or Techno-Feudalism: How Growing Addictions and Artificial Intelligence Shape Contemporary Society, Institute for Philosophy and Social Theory, University of Belgrade, 2022.
[8] L. Bojić, A. Bulatović, S. Žikić, The Scary Black Box: AI Driven Recommender Algorithms as the Most Powerful Social Force, Issues in Ethnology and Anthropology 17 (2022) 719–744.
[9] L. Bojić, Z. Dodevska, Y. Deldjoo, N. Pantelić, Towards recommender systems llms playground (recsysllmsp): Exploring polarization and engagement in simulated social networks, in: ROEGEN@RecSys'24: 1st Workshop on Risks, Opportunities, and Evaluation of Generative Models in Recommender Systems, Co-located with ACM RecSys in Bari, Italy, October 2024, 2024.
[10] C. R. Sunstein, #Republic: Divided democracy in the age of social media, Princeton University Press, 2018.
[11] L. Bojic, Metaverse through the prism of power and addiction: What will happen when the virtual world becomes more attractive than reality?, European Journal of Futures Research 10 (2022) 22.
[12] L. M. Bojić, Recommender algorithms as a source of power in contemporary society, Sociološki pregled 56 (2022) 1032–1055.
[13] C. C. Aggarwal, Recommender systems, volume 1, Springer International Publishing, 2016.
[14] M. Haroon, A. Chhabra, X. Liu, P. Mohapatra, Z. Shafiq, M. Wojcieszak, Youtube, the great radicalizer? auditing and mitigating ideological biases in youtube recommendations, arXiv preprint arXiv:2203.10666 (2022).
[15] M. Kunaver, T. Požrl, Diversity in recommender systems–a survey, Knowledge-based systems 123 (2017) 154–162.
[16] C.-N. Ziegler, S. M. McNee, J. A. Konstan, G. Lausen, Improving recommendation lists through topic diversification, in: Proceedings of the 14th international conference on World Wide Web, ACM, 2005, pp. 22–32.
[17] S. Vargas, P. Castells, Rank and relevance in novelty and diversity metrics for recommender systems, in: Proceedings of the fifth ACM conference on Recommender systems, ACM, 2011, pp. 109–116.
[18] Z. Dodevska, A. Petrović, S. Radovanović, B. Delibašić, Changing criteria weights to achieve fair vikor ranking: a postprocessing reranking approach, Autonomous Agents and Multi-Agent Systems 37 (2023) 9.
[19] Z. Dodevska, S. Radovanović, A. Petrović, B. Delibašić, When fairness meets consistency in ahp pairwise comparisons, Mathematics 11 (2023) 604.
[20] R. Kaminskas, D. Bridge, Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems, ACM Transactions on Interactive Intelligent Systems (TiiS) 7 (2016) 1–42.
[21] S. M. Mohammad, P. D. Turney, Crowdsourcing a word–emotion association lexicon, Computational Intelligence 29 (2013) 436–465.
[22] V. Kotu, B. Deshpande, Chapter 4 - classification, in: Data Science: Concepts and Practice, 2 ed., Morgan Kaufmann, 2019, pp. 65–163.
[23] S. Seelam, Cosine similarity and cosine distance, https://medium.com/geekculture/cosine-similarity-and-cosine-distance-48eed889a5c4, 2021. Accessed: 2021-05-26.
[24] D. M. Blei, A. Y. Ng, M. I. Jordan, W. X. Zhao, J. Jiang, J. Weng, J. He, E.-P. Lim, H. Yan, X. Li, A. Ahuja, Latent dirichlet allocation, Journal of Machine Learning Research 3 (2016) 50–57.
[25] E. Bakshy, S. Messing, L. A. Adamic, Exposure to ideologically diverse news and opinion on facebook, Science 348 (2015) 1130–1132.
[26] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: Nsga-ii, IEEE transactions on evolutionary computation 6 (2002) 182–197.
[27] F. E. Zaizi, S. Qassimi, S. Rakrak, Multi-objective optimization with recommender systems: A

systematic review, Information Systems 117 (2023) 102233.
[28] A. Petrović, M. Nikolić, S. Radovanović, B. Delibašić, M. Jovanović, Fair: Fair adversarial instance re-weighting, Neurocomputing 476 (2022) 14–37.
[29] F. M. Harper, J. A. Konstan, The movielens datasets: History and context, ACM Transactions on Interactive Intelligent Systems (TIIS) 5 (2015) 1–19.
[30] P. Henderson, V. Ferrari, End-to-end training of object class detectors for mean average precision, in: Computer Vision–ACCV 2016: 13th Asian Conference on Computer Vision, Taipei, Taiwan, November 20-24, 2016, Revised Selected Papers, Part V 13, Springer, 2017, pp. 198–213.
[31] J. L. Herlocker, J. A. Konstan, L. G. Terveen, J. T. Riedl, Evaluating collaborative filtering recommender systems, ACM Transactions on Information Systems (TOIS) 22 (2004) 5–53.
[32] A. Antikacioglu, T. Bajpai, R. Ravi, A new system-wide diversity measure for recommendations with efficient algorithms, SIAM Journal on Mathematics of Data Science 1 (2019) 759–779.
[33] M. Ge, C. Delgado-Battenfeld, D. Jannach, Beyond accuracy: evaluating recommender systems by coverage and serendipity, in: Proceedings of the fourth ACM conference on Recommender systems, ACM, 2010, pp. 257–260.
[34] Y. Koren, S. Rendle, R. Bell, Advances in collaborative filtering, in: Recommender systems handbook, Springer, 2021, pp. 91–142.
[35] R. Widayanti, M. H. R. Chakim, C. Lukita, U. Rahardja, N. Lutfiani, Improving recommender systems using hybrid techniques of collaborative filtering and content-based filtering, Journal of Applied Data Sciences 4 (2023) 289–302.
[36] D. Gomez-Zara, V. Kam, C.-F. Chiang, L. DeChurch, N. Contractor, Augmenting team diversity and performance by enabling agency and fairness criteria in recommendation algorithms, arXiv preprint arXiv:2410.00346 (2024).