EuroVis Workshop on Visual Analytics (2023)
M. El-Assady and M. Angelini (Editors)

Why am I reading this?
Explaining Personalized News Recommender Systems

Sverrir Arnórsson[1], Florian Abeillon[1,2], Ibrahim Al-Hazwani[3,4], Jürgen Bernard[3,4], Hanna Hauptmann[5], and Mennatallah El-Assady[1]

[1]ETH Zürich [2]École polytechnique [3]University of Zürich [4]Digital Society Initiative Zürich [5]Utrecht University

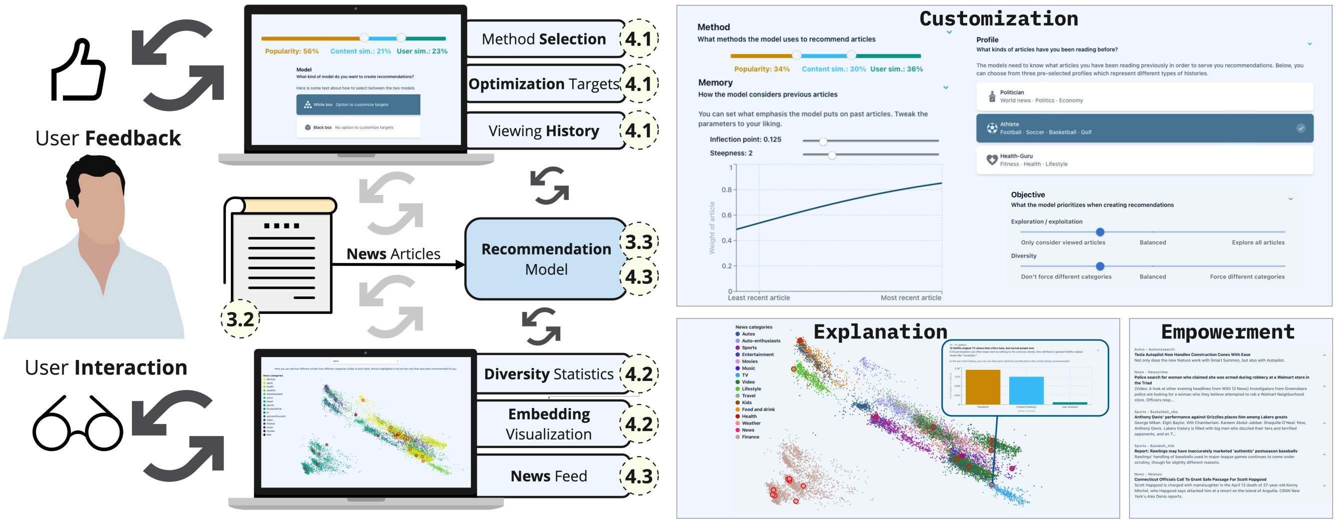

Figure 1: *Our interactive and explainable machine learning approach supports users in understanding, diagnosing, and refining personalized recommender systems. The NewsRecXplain interface allows users to explain and customize news recommendations and contributes to empowering them in contextualizing their filter bubbles. This overview figure depicts the interaction workflow.*

Abstract

Social media and online platforms significantly impact what millions of people get exposed to daily, mainly through recommended content. Hence, recommendation processes have to benefit individuals and society. With this in mind, we present the visual workspace NewsRecXplain, with the goals of (1) explaining and raising awareness about recommender systems, (2) enabling individuals to control and customize news recommendations, and (3) empowering users to contextualize their news recommendations to escape from their filter bubbles. This visual workspace achieves these goals by allowing users to configure their own individualized recommender system, whose news recommendations can then be explained within the workspace by way of embeddings and statistics on content diversity.

1. Introduction

The Internet has made it easier to distribute news articles to more people than ever before. Given this abundance of data content and digital objects, recommender systems (RSs) have become indispensable to keep readers engaged and manage the possible content overflow [KR12]. RSs are software tools and techniques providing suggestions for an user to interact with [RRS15]. While they offer enormous opportunities, they lack transparency and diversity, and provide little or no control in the generation of recommendations [HBMVH19]. Especially in an age where our reliance on digital information has a palpable impact on our lives, it is crucial to make sure such systems are not detrimental to specific individuals nor to society as a whole [IN00, SLL21].

To mitigate possible detrimental effects of current RSs, multiple model design alternatives have been researched and proposed [SBCH12, JCL*19]. However, beyond tailoring the recommendation models, it is crucial to give users the agency to determine how such models affect their news consumption. Hence, we argue for a need to explain such models, including their characteristics and parameter effects. This can facilitate bias mitigation and allow for user control, customization, and empowerment. User control in recommendations of news articles has previously been characterized in different ways, e.g., by expressing interest in certain topics [IGÖ15] or by giving feedback on individual recommendations [JJN19].

In this paper, we present an interactive visual analytics approach that allows users to (1) understand news recommendations through

DOI: 10.2312/eurova.20231099

visual *explanations* of the recommended content; (2) steer the recommendation model by configuring a personalized profile and controlling model parameters, i.e., facilitating the *customization* of the RS; and (3) contextualize the configured model through setting it into context to other possible configurations, *empowering* users to take control of their filter bubbles. Rather than only giving users *ex-post* explanations of recommendations, our work tackles the problem of explainability and transparency by raising users' awareness of how recommendation algorithms work and by giving them interactive tools to steer the recommendations they get. We draw on methods used in other systems to design representative RS models, giving users an insight into the behavioral characteristics of RSs.

Overall, we contribute a conceptual description of the problem space of explaining personalized news recommendations through visual analytics; a system implementation of the proposed concepts in the form of *NewsRecXplain*, a visual workspace for explanation customization, and empowerment; lastly, a discussion of usage scenarios and research opportunities based on our proposed approach.

2. Related work

We describe related work from three different points: RSs, explainable RSs, and visualization of RS processes and results.

Recommender Systems–We categorize RSs methods into three main recommendation principles: collaborative filtering, content-based filtering, and popularity-based filtering. Collaborative filtering RSs leverages the ideas that users who agreed in the past will also agree in the future and that users will like similar items as liked in the past [SK09]. Content-based filtering RSs uses item descriptions and profiles of user preferences to recommend similar items that the user has liked in the past [PB07]. In the news domain, content-based filtering is widely used [KJJ18] given the specific domain challenges: recency [DDGR07], diversity [EHWK14], or serendipity [KVW16]. Popularity-based RSs methods follow the principle that item relevance can be defined by the assessment of the frequency of the general item consumption. Accordingly, these methods provide general recommendations to every user depending on the item popularity [JGLA16]. We enable the user to adjust the impact of all three these approaches in their recommendations. In doing so, the user can familiarize themselves with widely used algorithms normally used in news recommendation systems and understand the strengths and weaknesses of the different approaches.

Explainable Recommendations–With the increasing number of parameters needed to consider for generating good recommendations, developers are adopting deep learning techniques to scale the above-mentioned methods. This entailed a decrease of trust in the results and interpretability of the models. The idea of explainable recommendations is to address this lack of trust and interpretability. The possibility to explain why a specific recommendation is prompted to the user is helpful as it can increase the trustworthiness, effectiveness, and persuasiveness of the system [FZ11]. The generation of explanations requires an in-depth characterization of which methods to use [AHSSB22]. In RSs, the users can either explicitly invoke explanations when needed, or be automatically displayed by the system [FZ11]. Moreover, allowing control over the recommendation can lead to a more customized, empowered, and personalized experience [HBMVH19]. Two benefits of explainable and controllable recommendations are to enable personalization and customization [Bea14]. Personalization reflects how the information that the system provides is aligned with the user's individual preferences. In contrast, customization refers to the amount of user involvement in the process, e.g., to personalize recommendations. By both implementing explainable and customizable recommendations, *NewsRecXplain* gives users more trust in their recommendations, empowering them via visualization [WAC*21].

Interactive Visualization for Recommender Systems–Research at the intersection between RSs and interactive visualizations led to exploring how visual interfaces can improve user interaction with recommendation systems. Examples of works that have been developed in this intersection are: *PeerChooser* [OSG*08], *SmallWorlds* [GOB*10], *TasteWeights* [BOH12], and *Relevance Tuner+* [TB21]. *PeerChooser* and *SmallWorlds* visualize the inner logic of the collaborative filtering recommendations. *TasteWeights* and Relevance Tuner+ provide an explanation for a hybrid RS by explaining how the recommendation process works and by enabling the users to elicit their preferences via a user interface. In the news recommendation domain, *NewsViz* [KSZ20] provides the user with a visualization of the category or source of online news by the use of a treemap and enables the users to control their preferences by adjusting the cell size of the treemap. In contrast, *NewsRecXplain* uses a 2D visualization reflecting the internal structures of the embeddings of news articles, accompanied by an encoding of the different new categories. Inspiring works on embedding visualization can be found in [STN*16, KAKC17, LBT*17]. Additionally, *NewsRecXplain* uses statistics visualizations to report the diversity of recommendations as well as the contributions of the three different methods included and combined in our approach (collaborative filtering, content-based filtering, popularity-based filtering).

3. Personalized News Recommendation

This section describes the problem design space, specifically, the users and their tasks, and the utilized data and models.

3.1. Users and Tasks

We tailored our approach to tech-savvy readers of online news articles, who are interested in understanding why specific articles were recommended, and how they can achieve other recommendations, measuring the effect of their filter bubble.

To address these needs, we designed *NewsRecXplain* as an interactive visual workspace, as described in Section 4. In concordance with the interactive and explainable AI tasks (understanding, diagnosis, refinement) [SSSEA20], the workspace provides three functionalities to the user:

1. **Explanation**: Explore the recommendation output characteristics to raise the user's awareness about the model's decision-making for their user profile.
2. **Customization**: Choose user profiles and steer the model parameters to control the RSs.
3. **Empowerment**: Contextualize the user's customized models to other configurations, enabling them to control their filter bubble.

3.2. Data

We used the *Microsoft News Dataset* (MiND) [WQC*20], a large-scale dataset for news recommendation research collected from anonymized behavior logs of the Microsoft News website. Overall, 45,463 news articles are provided with protocols of behavior for 29,572 users. The news article dataset includes the title and abstract

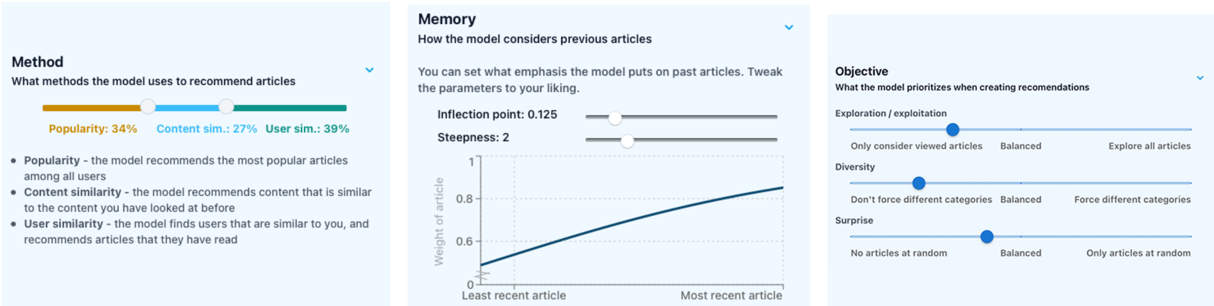

(a) *Methods control enables users to steer the recommendation parameter impact.* **(b)** *Memory function representation of the user's viewing history. Shape can be altered.* **(c)** *Controls of three optimization objectives for customizing the recommendation models.*

Figure 2: *Model customization components, enabling user control of the RSs to customize their recommendations and achieve individualization.*

of each article and its categorization. The user behavior dataset includes a list of every recommendation served to a user, and whether the specific user clicked on the recommendations. This unique combination of articles as well as user history was essential for the purposes of the models used (see Subsection 3.3).

The authors of the dataset extracted any named entity mentioned in the news articles (for example persons and organizations) and matched them with records in Wikidata. Wikidata provides information on the relations between entities, which was used to generate embeddings that encode the connections between entities in a high-dimensional space [BUGD*13]. Closely connected entities are clustered in this embedding space, while entities that don't share any commonalities are further apart, leading to a spatialization of entities reflecting connectivity. We augmented the dataset in two ways to make these embeddings usable for our purposes. First, the news dataset. We built embeddings for every news article, using the centroid of the embeddings of the entities they contain. Then, to group articles within the same category together, we calculated the centroid of the embeddings of news articles within each category and concatenated it to the news articles' original embeddings. Secondly, we augmented the user dataset in order to evaluate user similarity. This was done similarly to the news embeddings by assigning the centroid of the embeddings of the news articles in a user's viewing history as their embedding.

3.3. Models and Metrics

To generate the recommendations for *NewsRecXplain*, we present an interpretable and customizable RS.

Recommendation Model The underlying model of the RS includes and combines three types of models with hybrid principles:

- **Content-based filtering:** Recommending news articles most similar to the articles already read by the user. Concretely, news articles are ranked by the cosine similarity of their embeddings to the user's, i.e., the average embedding of their news history.
- **Collaborative filtering:** Recommending news articles read by users similar to the current user. Similar users are identified using cosine similarity between user embeddings. Then, articles popular with the identified users are served for a recommendation.
- **Popularity-based filtering:** Recommending news articles most popular among users. Concretely, news articles are ranked by the number of times they were clicked on by users and by the ratio of the number of clicks to the number of impressions.

Additionally, the weight assigned to each article in the user history can be customized, as demonstrated in Section 4.

Model Characteristics Collaborative and popularity-based filtering models can suffer from echo chambers effects–when the RS only recommend content similar to what users have already viewed before [NLJS20]–and popularity bias–when the more popular an article, the more recommended it gets. To mitigate this, we developed the following optional objectives:

- **Diversity:** Constraining the model to recommend news articles from different categories. Concretely, after each draw, the relative probability of drawing again an article from the same category as the one just drawn decreases by the percentage chosen.
- **Exploration:** Allowing the model to draw news articles with lower scores as well. Concretely, probability weights used to draw news articles are calculated using a logistic function on their scores. A higher exploration score forces a flatter logistic curve, and thus a more uniform probability draw across articles.
- **Surprise:** Requiring some randomness among recommendations. Concretely, the chosen percentage of news articles is drawn at random and placed at random positions in the feed.

4. *NewsRecXplain* Visual Workspace

The *NewsRecXplain* visual workspace consists of three components: model configuration, explanation, and news recommendation. The overall workflow and components are depicted in Figure 1.

The top component, the model configuration, can be seen in Figure 2a. It supports the *customization* task, giving the user control over the RS. Within this component, there are four elements.

Model Configuration The first one allows users to control which news articles are fed into the RS as the user's viewing history. Users can select between three different profiles (sets of user reading histories), whose news categories they draw news articles from, are the following: **Athlete** (football-nfl, basketball-nba, soccer); **Health Guru** (fitness, healthy-living, weight-loss); and **Politician** (news-world, news-politics, finance-news). The remaining three elements facilitate customization of the different attributes described in Subsection 3.3. In the second element (see figure Fig-

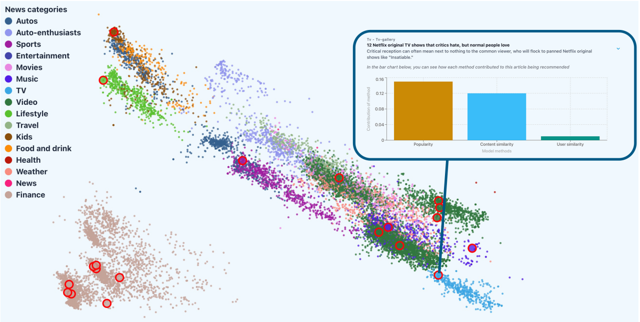

Figure 3: *Scatter plot of the first two principal components of news article embeddings, colored by news categories. Red circles represent the news articles recommended by the model. In the example, the recommendation is wide-spread across the 2D representation of articles, beyond the personal bubble of an individual. One recommended news article is highlighted through tooltip functionality. The bar chart depicts the contribution of the three recommendation methods based on the hybrid principles of content-based, collaborative, and popularity-based filtering.*

ure 2a), the user can tweak the contribution of each method described in Section 5, using a multi-slider. The third element (see Figure 2b) is an interactive visual representation of the memory function of the RS. The memory function is used to weigh older and newer articles differently when selecting articles to recommend. If the memory function is flat, the past articles are weighed equally with the newer ones. If it is left-skewed, the newer articles are weighed more heavily than the older ones. The user can then alter the shape of the function using sliders. In the last element, the user can control the three model characteristics provided by *NewsRecXplain*.

Recommender Explanation The center component of the workspace is the recommender explanation, which supports the user task of *explanation*. This component facilitates the exploration of recommendation characteristics, why a certain recommendation was made, and the context for this set of recommendations. The first element of this component is a bar chart showing summary statistics on the number of news articles per category (see Figure 4). This helps the user evaluate how the current model configuration affects the diversity of the recommendation set. The second element is a scatter plot of the first two principal components of news embeddings [AW10] (see Figure 3). As the news embeddings encode similarity, the news articles close to each other in this visualization share common named entities or categories (see Subsection 3.2). News articles that are recommended by the current model configuration are highlighted with red circles. If the set of recommended articles is clustered together, this suggests that a filter bubble of closely connected articles exists. Finally, when the user hovers over a recommended news article, a tooltip view enables the user to gain insights into article details. The bar chart allows assessing how much each of the three recommendation methods contributed to the recommendation of this article.

News Recommendation The bottom component of the *NewsRecXplain* workspace is the news recommendation. This takes the form of a news feed, where the user can browse through the set of recommendations (highlighted points in Figure 3) made based on the model selection. With the news feed, the user can assess the content of the recommendation in detail, gaining a better awareness of the model's decision-making. Together, these three components support the task of *empowerment*, as the user is capable of both analyzing the recommendations and altering the model configuration to gain control of their filter bubble.

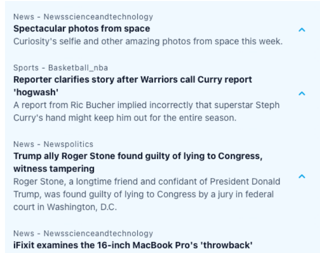

5. Usage Scenarios

We demonstrate the usefulness of our approach in two complementing usage scenarios. These align with the three analysis tasks which *NewsRecXplain* supports: explanation, control, and empowerment.

Raising Awareness through Explanations This usage scenario describes a user who aims to understand why they are receiving a set of recommended news articles. For this purpose, they use the explanatory power of *NewsRecXplain*. At the start, the user receives a set of recommendations, visible in Figure 3. While browsing the news feed, the user encounters an article that they would not expect, given the history they entered in the profile selection Section 4. Curious about why that occurred, the user clicks on the article, revealing the full article accompanied by a bar chart that demonstrates how much each algorithmic method contributed. For the recommended article, the user identifies that the user similarity contributes most to the recommendation. Informed by this finding, the user explores the projection-based view in Figure 3 further. The user identifies the article and sees that it is an outlier within the other recommended articles. Next, the user focuses on the bar chart showing the statistical distribution of articles per news category in Figure 4. The user clearly identifies that sports and finance dominate the collection of news articles. Informed by this finding, the user is inspired to diversify their recommendations as they see that other users likely have more diverse recommendations. The process enhances the user's understanding of the recommendations and their diversity. By analyzing the characteristics of each recommendation, the user can draw parallels with their encounters with other RSs as the contributions displayed in the bar chart are directly linked to the interactive controls provided by *NewsRecXplain*.

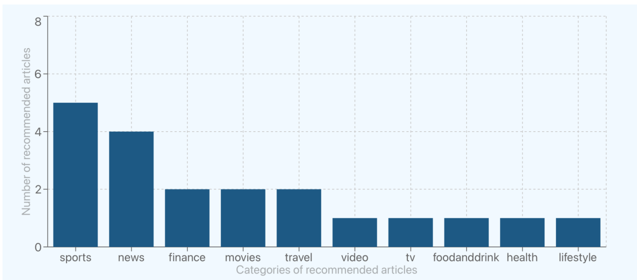

Figure 4: *Statistics showing the number of news articles per category from the sample of recommended articles, to showcase the diversity of news recommendations.*

Developing a Customized Recommender System The second usage scenario describes a process where the user controls the recommendation system, which aligns with their personal goals. The user has a preference for entertainment and sports over other news categories. The user starts by selecting one of the three profiles (see Section 4) that match their interests the most, in this case, the health-guru. This selection generates an initial set of recommendations. Using the similarity-preserving distribution of articles in Figure 3 and the distribution of articles across news categories shown in Figure 4, the user assesses how relevant the recommended news articles are with regard to their preferences. Using this approach, the user evaluates whether the model draws articles from categories they like and if the user also agrees with the range of topics. The user identifies that, while the model draws from categories they like, they would want to also see others, for example, the "news" news category. Following this, the user increases the contribution of the popularity-based filtering in Figure 2c as they have, by analyzing the pie chart in Figure 3, observed that articles from the "news" category tend to be recommended by popularity. The user also increases the contribution of exploration in Figure 2c to expand the number of categories served. Next, the user applies this customization-and-effect analysis process iteratively and analyzes the individual sets of recommendations received until the recommendations align nicely with their goals and preferences. The user can use this exercise to form an opinion on what they want from such a system, leading to more empowerment, as demonstrated in the scenario.

6. Discussion

The use of explanations, control, and empowerment have been extensively discussed in the field of RSs. However, little attention has been paid to how to use visualizations to facilitate both these steps. This section describes the lessons learned from our approach, limitations, and possible future research opportunities.

User Models and Profiles–Explanations are an increasingly important topic in all fields of human-AI interaction. For RSs, these explanations focus on transparency, scrutability, trust, and in some cases, persuasion. For news recommendations specifically, the concept of self-actualization has recently gained interest. This concept supports users in understanding and exploring their user models and changing them towards their preferred ones. Visualization can support this process by providing visual metaphors for user preferences and other user models. In *NewsRecXplain*, the user can choose between three stereotypical user profiles to see the effects this has on the set of recommendations, visualized as a bar chart topic distribution and a projection scatter plot of all news articles with highlights on the recommended ones. Future research could allow users to visualize stereotypical user models to let them choose which would fit their self-actualization goals.

Control and Effect–Many studies [JJN19, JTHV20] have shown that users appreciate the ability to interact with RSs, either for explorative purposes or to control the output of the systems. Such control can be provided at the data level (explicit user profiling), during the recommendation process (model parameters), or at the recommendation itself (critiquing or rating). However, these controls are hard to use without feedback on their actual impact on the system. Visualizations can support the interactive process by providing responsive effect visualizations for each interaction. In *NewsRecXplain*, we use a three-point slider in the control interface to steer the importance of recommendation features. The three recommendation components are fed back to the user in each recommended item as a bar chart. This way, control, and effect are visually connected. Additionally, model parameters and outcome metrics are configurable. Adequate visual metaphors for the effect of these parameters within the actual recommendation are one opportunity for future research.

Personalization Challenges–Using personalized news RSs has many potentially dangerous side effects, such as filter bubbles, echo chambers, discrimination, and popularity bias which are due to missing diversity, novelty, and serendipity. Visualizations can counteract these effects by empowering users to see the global and local context of the recommendation they are currently receiving. In *NewsRecXplain*, we use the distribution of topics as a first and straightforward visual hint into lacking diversity. Additionally, we contextualize the set of recommended items within the whole spectrum of available options. These are the only two visual metaphors that are quick and easy to understand without interfering a lot with the users' reading experience. More advanced visualizations would be needed to include historic reading behavior and comparison among different user clusters.

Model Limitations–*NewsRecXplain* is a first step toward using visualizations in personalized news RSs. Since this is a visualization-focused work, some underlying algorithmic solutions are still limited. The use of simple embedding and similarity metrics should be replaced by state-of-the-art topic modelling techniques and similarity metrics derived from the actual interpretable recommender algorithms. Similarly, the performance of the chosen recommender algorithms was not optimized or evaluated on a benchmark for this work. In the future, it could be improved with the standard evaluation method of the MiND challenge for accuracy and ideally also for diversity, novelty, and fairness metrics.

In Section 2, it was noted that explainable recommendations lead to a more trustworthy and effective experience. Allowing users to control their recommendations also empowers them. Future research should assess this combination for news recommendations and compare different clustering techniques like t-SNE and UMAP to the principal component approach used in this work.

7. Conclusion

We have presented *NewsRecXplain*, a web-based interactive visual workspace for explainable personalized news recommender systems. The personalized recommendation system includes the three hybrid principles of collaborative filtering, content-based filtering, and popularity-based methods, with user controls to enable users to tune the contribution of each. Users can further adapt the recommendations of *NewsRecXplain* towards individual preferences according to three objectives: diversity, exploration, and surprise, each interactively controllable. We support common user requirements for interactive recommender systems such as explanation, customization, and empowerment with different visual metaphors. Three usage scenarios demonstrate how each of these visual features can be utilized. Finally, we discuss future research for visualization components in explainable personalized news recommender systems.

References

[AHSSB22] AL-HAZWANI I., SCHMID J., SACHDEVA M., BERNARD J.: A design space for explainable ranking and ranking models. In *Eurographics Conference on Visualization* (2022), Krone M., Lenti S., Schmidt J., (Eds.), The Eurographics Association. doi:10.2312/evp.20221114. 2

[AW10] Abdi H., Williams L. J.: Principal component analysis. *Wiley interdisciplinary reviews: computational statistics 2*, 4 (2010), 433–459. 4

[Bea14] Beam M. A.: Automating the news: How personalized news recommender system design choices impact news reception. *Communication Research 41*, 8 (2014), 1019–1041. 2

[BOH12] Bostandjiev S., O'Donovan J., Höllerer T.: Tasteweights: a visual interactive hybrid recommender system. In *ACM conference on Recommender systems* (2012), pp. 35–42. URL: https://doi.org/10.1145/2365952.2365964. 2

[BUG*13] Bordes A., Usunier N., Garcia-Duran A., Weston J., Yakhnenko O.: Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems 26* (2013). URL: https://tinyurl.com/3xjtz79u. 3

[DDGR07] Das A. S., Datar M., Garg A., Rajaram S.: Google news personalization: scalable online collaborative filtering. In *International Conference on World Wide Web* (2007), pp. 271–280. doi:10.1145/1242572.1242610. 2

[EHWK14] Ekstrand M. D., Harper F. M., Willemsen M. C., Konstan J. A.: User perception of differences in recommender algorithms. In *ACM Conference on Recommender Systems* (2014), pp. 161–168. doi:10.1145/2645710.2645737. 2

[FZ11] Friedrich G., Zanker M.: A taxonomy for generating explanations in recommender systems. *AI Magazine 32*, 3 (2011), 90–98. doi:10.1609/aimag.v32i3.2365. 2

[GOB*10] Gretarsson B., O'Donovan J., Bostandjiev S., Hall C., Höllerer T.: Smallworlds: visualizing social recommendations. In *Computer Graphics Forum* (2010), vol. 29, Wiley Online Library, pp. 833–842. doi:10.1111/j.1467-8659.2009.01679.x. 2

[HBMVH19] Harambam J., Bountouridis D., Makhortykh M., van Hoboken J.: Designing for the better by taking users into account: A qualitative evaluation of user control mechanisms in (news) recommender systems. In *ACM Conference on Recommender Systems* (2019), pp. 69–77. doi:10.1145/3298689.3347014. 1, 2

[IGÖ15] Ingvaldsen J. E., Gulla J. A., Özgöbek Ö.: User controlled news recommendations. In *IntRS@ RecSys* (2015), pp. 45–48. URL: https://ceur-ws.org/Vol-1438/paper8.pdf. 1

[IN00] Introna L. D., Nissenbaum H.: Shaping the web: Why the politics of search engines matters. *The information society 16*, 3 (2000), 169–185. doi:10.1080/01972240050133634. 1

[JCL*19] Jiang R., Chiappa S., Lattimore T., György A., Kohli P.: Degenerate feedback loops in recommender systems. In *ACM Conference on AI, Ethics, and Society* (2019), pp. 383–390. doi:10.1145/3306618.3314288. 1

[JGLA16] Jonnalagedda N., Gauch S., Labille K., Alfarhood S.: Incorporating popularity in a personalized news recommender system. *PeerJ Computer Science 2* (2016), e63. doi:10.7717/peerj-cs.63. 2

[JJN19] Jannach D., Jugovac M., Nunes I.: Explanations and user control in recommender systems. In *Workshop on Personalization and Recommendation on the Web and Beyond* (2019), pp. 31–31. doi:10.1145/3345002.3349293. 1, 5

[JTHV20] Jin Y., Tintarev N., Htun N. N., Verbert K.: Effects of personal characteristics in control-oriented user interfaces for music recommender systems. *User Modeling and User-Adapted Interaction 30*, 2 (2020), 199–249. doi:10.1007/s11257-019-09247-2. 5

[KAKC17] Kahng M., Andrews P. Y., Kalro A., Chau D. H.: Activis: Visual exploration of industry-scale deep neural network models. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2017), 88–97. doi:10.1109/TVCG.2017.2744718. 2

[KJJ18] Karimi M., Jannach D., Jugovac M.: News recommender systems–survey and roads ahead. *Information Processing & Management 54*, 6 (2018), 1203–1227. doi:10.1016/j.ipm.2018.04.008. 2

[KR12] Konstan J. A., Riedl J.: Recommender systems: from algorithms to user experience. *User modeling and user-adapted interaction 22*, 1 (2012), 101–123. 1

[KSZ20] Kunkel J., Schwenger C., Ziegler J.: Newsviz: depicting and controlling preference profiles using interactive treemaps in news recommender systems. In *ACM Conference on User Modeling, Adaptation and Personalization* (2020), pp. 126–135. doi:10.1145/3340631.3394869. 2

[KVW16] Kotkov D., Veijalainen J., Wang S.: Challenges of serendipity in recommender systems. In *International Conference on Web Information Systems and Technologies* (2016), SCITEPRESS. doi:10.5220/0005879802510256. 2

[LBT*17] Liu S., Bremer P.-T., Thiagarajan J. J., Srikumar V., Wang B., Livnat Y., Pascucci V.: Visual exploration of semantic relationships in neural word embeddings. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2017), 553–562. doi:10.1109/TVCG.2017.2745141. 2

[NLJS20] Noordeh E., Levin R., Jiang R., Shadmany H.: Echo chambers in collaborative filtering based recommendation systems. *arXiv preprint arXiv:2011.03890* (2020). 3

[OSG*08] O'Donovan J., Smyth B., Gretarsson B., Bostandjiev S., Höllerer T.: Peerchooser: visual interactive recommendation. In *SIGCHI Conference on Human Factors in Computing Systems* (2008), pp. 1085–1088. doi:10.1145/1357054.1357222. 2

[PB07] Pazzani M. J., Billsus D.: Content-based recommendation systems. In *The adaptive web*. Springer, 2007, pp. 325–341. doi:10.1007/978-3-540-72079-9_10. 2

[RRS15] Ricci F., Rokach L., Shapira B.: Recommender systems: introduction and challenges. In *Recommender systems handbook*. Springer, 2015, pp. 1–34. doi:10.1007/978-1-4899-7637-6_1. 1

[SBCH12] Schwind C., Buder J., Cress U., Hesse F. W.: Preference-inconsistent recommendations: An effective approach for reducing confirmation bias and stimulating divergent thinking? *Computers & Education 58*, 2 (2012), 787–796. doi:10.1016/j.compedu.2011.10.003. 1

[SK09] Su X., Khoshgoftaar T. M.: A survey of collaborative filtering techniques. *Advances in artificial intelligence 2009* (2009). doi:10.1155/2009/421425. 2

[SLL21] Santos F. P., Lelkes Y., Levin S. A.: Link recommendation algorithms and dynamics of polarization in online social networks. *PNAS 118*, 50 (2021), e2102141118. doi:10.1073/pnas.2102141118. 1

[SSSEA20] Spinner T., Schlegel U., Schäfer H., El-Assady M.: explAIner: A visual analytics framework for interactive and explainable machine learning. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (2020), 1064–1074. doi:10.1109/TVCG.2019.2934629. 2

[STN*16] Smilkov D., Thorat N., Nicholson C., Reif E., Viégas F. B., Wattenberg M.: Embedding projector: Interactive visualization and interpretation of embeddings. *arXiv preprint arXiv:1611.05469* (2016). URL: http://arxiv.org/abs/1611.05469. 2

[TB21] Tsai C.-H., Brusilovsky P.: The effects of controllability and explainability in a social recommender system. *User Modeling and User-Adapted Interaction 31*, 3 (2021), 591–627. doi:10.1007/s11257-020-09281-5. 2

[WAC*21] Willett W., Aseniero B. A., Carpendale S., Dragicevic P., Jansen Y., Oehlberg L., Isenberg P.: Perception! immersion! empowerment! superpowers as inspiration for visualization. *IEEE Transactions on Visualization and Computer Graphics 28*, 1 (2021), 22–32. 2

[WQC*20] Wu F., Qiao Y., Chen J.-H., Wu C., Qi T., Lian J., Liu D., Xie X., Gao J., Wu W., et al.: Mind: A large-scale dataset for news recommendation. In *Association for Computational Linguistics* (2020), pp. 3597–3606. doi:10.18653/v1/2020.acl-main.331. 2