

Towards Multi-Aspect Diversification of News Recommendations Using Neuro-Symbolic AI for Individual and Societal Benefit

Markus Reiter-Haas^{1,*}, Elisabeth Lex¹

¹ Graz University of Technology, Institute of Human-Centred Computing, Sandgasse 36/III, 8010 Graz, Austria

Abstract

News recommendations are complex, with diversity playing a vital role. So far, existing literature predominantly focuses on specific aspects of news diversity, such as viewpoints. In this paper, we introduce multi-aspect diversification in four distinct recommendation modes and outline the nuanced challenges in diversifying lists, sequences, summaries, and interactions. Our proposed research direction combines symbolic and subsymbolic artificial intelligence, leveraging both knowledge graphs and rule learning. We plan to evaluate our models using user studies to not only capture behavior but also their perceived experience. Our vision to balance news consumption points to other positive effects for users (e.g., increased serendipity) and society (e.g., decreased polarization).

Keywords

Ethical Artificial Intelligence Systems, Balanced News Consumption, Diversifying User Behavior

1. Introduction

News recommender systems (NRS) have become increasingly popular and important for news consumption [1]. An important concern is the news diversity in recommendations [2], as diverse news consumption corresponds to a more balanced information diet [3]. However, compared to other recommendation areas, the news domain has unique challenges, especially temporal considerations due to articles being short-lived and typically consumed at most once per user, which leads to increased sparsity [4]. Besides, NRS also need to deal with limited information about the users, evolving categorical preferences, a broad range of topics being covered, content being framed in certain ways, ideological blind spots in consumption, media biases of the news sources, as well as ethical and regulatory considerations. As news plays a vital role in opinion formation [5] and democracies [6], the diversification in NRS is especially relevant [7]. That is why works like [8, 9] focus directly on diversifying viewpoints. In this paper, we argue that incorporating single aspects is not enough for news diversity and present multi-aspect diversity as an emerging research direction.

The conceptualizations of media and news diversity vary across literature comprising different subdimensions, such as topic and viewpoint diversity [10]. Mattis et al. [11] introduce algorithmic nudges by re-ranking to increasing exposure diversity and presentation nudges via altering the user interface for deliberate changes in consumption diversity. Bernstein et al. [2] outline that the role of news diversity differs from other recommendation domains, with the promotion of diverse content being vital in a democratic society. One particular observation from [2] highlights that current models for NRS typically do not capture the multidimensionality of diversity, which is a driving motivator for the current paper. Moreover, diversification in personalized NRS tends to focus on specific recommendation modes, typically targeting either the (sequential) history of clicked news or the list of recommended news [12]. For instance, Wu et al. [13] trains an end-to-end diversity-aware news recommender based

INRA 2025: 13th International Workshop on News Recommendation and Analytics in Conjunction with ACM RecSys 2025, September 22–26, 2025, Prague, Czech Republic

*Corresponding author.

✉ reiter-haas@tugraz.at (M. Reiter-Haas); elisabeth.lex@tugraz.at (E. Lex)

🌐 <https://iseratho.github.io/> (M. Reiter-Haas); <https://elisabethlex.info/> (E. Lex)

🆔 0000-0001-9852-8206 (M. Reiter-Haas); 0000-0001-5293-2967 (E. Lex)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

on semantic similarity in the recommendation list. Therefore, Wu et al. [12] emphasize diversity-aware NRS concerning both modes, as well as fine-grained diversification that goes beyond content, as a promising research direction. In this vein, Kaya and Bridge [14] considers subprofiles, rather than item features, for intent-aware diversification.

There have been important first strides towards multi-aspect news diversity. Zhou et al. [15] adapt the biology-inspired Simpson’s diversity index for multi-aspect diversity in search results. Li et al. [16] rely on random walks using normative distributions for transparent diversification of news candidates. Vercoutere et al. [17] uses concept graphs for extending the user profile improving diversity and the potential to increase user satisfaction according to their user study. The MANNer model [18] tackles multi-aspect diversity for NRS by incorporating a pairwise similarity matrix. In their formalization, diversity is defined as the normalized entropy over the aspects’ distributions and models aspect via separate labels per aspect in recommendation list.

Our contributions in this idea paper are two-fold. First, we introduce the problem of *multi-aspect news diversity in four distinct news recommendation modes*, where we also model similarities between labels of the same set of aspect. Thereby, our work expands the body of literature on diversity in recommendations [19] by thoroughly examining multi-aspect diversification of news, also in light of distinct novel modes like LLM-generated news summaries. Second, we aim to incorporate this expanded problem formalization to directly in the model architecture, as the training objective, and as transparent rules in deployed systems. Herein, we propose a research direction for *diversifying based on both symbolic and subsymbolic AI using knowledge graphs and rule learning* together with user study evaluations to gauge their beneficial impact. While several neuro-symbolic approaches exist [refer to 20], to the best of our knowledge, none of them provide a universal solution to diversify arbitrary combinations of aspects with both individual and societal benefits in mind.

2. Problem: Multi-Aspect News Diversity

Let $D = (d_1, d_2, \dots, d_n)$ be an arbitrary collection of documents. Each document comprises data and metadata, while $dist : D \times D \rightarrow [0, 1]$ is a distance function between pairs of documents (d_i, d_j) ; we denote $dist(d_i, d_j)$ as their distance value. Furthermore, we assume the distance function satisfies $dist(d_i, d_i) = 0$ and is symmetric: $dist(d_i, d_j) = dist(d_j, d_i)$. The diversity of the whole collection can be measured as average distance between all pairs of documents with $div(D) = \frac{1}{|D| \cdot (|D| - 1)} \sum_{\forall d_i, d_j \in D, d_i \neq d_j} dist(d_i, d_j)$. For simplicity, we assume each document d_i consists of only a topic label t_i and frame label f_i as aspects: $d_i = (t_i, f_i), t_i \in T, f_i \in F$, where T and F represent the set of topic and frame labels, respectively. Note that topic and frame labels can be seen as operating at different levels, which can be parameterized with α in the distance function to account for topical and framing distance separately: $dist(d_i, d_j) = \alpha \cdot dist_t(t(d_i), t(d_j)) + (1 - \alpha) \cdot dist_f(f(d_i), f(d_j))$. The functions $t(\cdot)$ and $f(\cdot)$ retrieve the topic and frame of the document, respectively.

For illustration, we define two topics $T = \{\text{Climate}, \text{Immigration}\}$ and four frames $F = \{\text{Health}, \text{Cultural}, \text{Security}, \text{Economy}\}$ and set $\alpha = 0.5$. We assign the values: $dist_t(\text{Climate}, \text{Immigration}) = dist_f(\text{Health}, \text{Security}) = dist_f(\text{Health}, \text{Economic}) =$

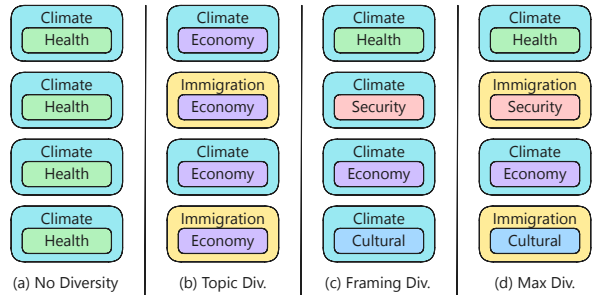


Figure 1: Examples of Multi-Aspect Diversity in Lists. (a) A list of all similar items has the minimum diversity of $div = 0$. (b) A diversified list regarding topic, but no framing diversity ($div = 0.33$). (c) A diversified regarding framing, but no topical diversity ($div = 0.42$). (d) A maximally diversified list regarding both topic and framing ($div = 0.75$).

$dist_f(\text{Cultural}, \text{Security}) = dist_f(\text{Cultural}, \text{Economic}) = 1$ and $dist_f(\text{Health}, \text{Cultural}) = dist_f(\text{Security}, \text{Economic}) = 0.5$. This effectively means the frames Health/Cultural and Security/Economic bear some resemblance (e.g., similarity in ideology), while both topics and other framing pairs are seen as maximally dissimilar to each other.

Figure 1 visualizes different scenarios for diversity of four documents in a list without considering their order. Note that neither for topic nor framing the maximum diversity of 1 can be reached in this scenario. The number of topics is smaller than the number of items in the list ($|T| < |D|$). While the number of frames is the same as the length of the list, not enough frames are completely dissimilar to each other: $max(\{|F_{dis}|, F_{dis} \subseteq F, \text{ where } \forall f_i, f_j \in F_{dis}, f_i \neq f_j, dist_f(f_i, f_j) = 1\}) < |D|$. Therefore, in this example, only a list of length two (i.e., the minimum number of completely distinct labels per aspect) can achieve the global maximum diversity of 1, where diversity is essentially reduced to item coverage.

The example showcases that even in this almost minimal example, diversifying news is far from trivial. In reality, the complexity of the problem increases tremendously just by scaling the number of aspects (e.g., adding news categories like sport or politics), the number of labels per aspect (e.g., additional topics being covered), and the size of the recommended list. Besides, the distance function is typically not so precisely defined, as several other factors also play a role. For instance, the Health frame in Climate could be semantically different than the Health frame in Immigration, operating at different hierarchical levels. Similarly, users could perceive diversity differently depending on the order, where an interleaved list might be seen as more diverse compared to a neatly divided list (e.g. first half Climate, second half Immigration). Furthermore, considering latent variables, such as derived from the content itself, further exacerbates these problems.

We would like to underscore that for brevity we have discussed diversity as if it were a measure that is completely independent of users, contexts, or other (e.g., ethical) considerations. Of course, this assumption is rather unrealistic, which is why the maximum diversity tends to differ profoundly from the optimal diversity. Diversity must also consider whether the recommendation is relevant for the user in a particular context and whether it adheres to established norms. A user who does not like sports will not benefit from general sport recommendations but might still monitor specific Olympic results, while spreading misinformation is detrimental for society even if considered relevant for a singular user. Therefore, maximizing the diversity cannot be the sole goal of a system. For the purpose of this paper, we nonetheless focus on increasing diversity without addressing these additional constraints.

So far, we have considered diversity in lists, which is a typical mode for recommendations. However, recommendation modes can also have more nuanced characteristics, like a spatial position. Depending on whether the order (or spatial position, for that matter) influences the distance function, the scenarios and resulting diversity differ or stay the same. The extent of these influences also plays a vital role. Recommendations presented as tiles in a grid (e.g., MSN homepage) or as categorical carousels (e.g., Inoreader dashboard) might not fundamentally differ for diversity purposes from a simple list. On the other hand, recommendation modes like endless sequences of news (e.g., on social media platforms like X) or LLM-generated summaries (e.g., on Perplexity AI discover feature) might require fundamentally different approaches.

3. Modes of News Diversification

In this section, we discuss news diversification in four distinct recommendation modes : recommendation lists, sequences, summaries, and interactions (Figure 2). Note that these modes can be combined (e.g., lists of summaries or sequences of lists). For the present paper, we consider them strictly in isolation.

3.1. Diversifying Lists

Diversifying lists closely follows the problem description in Section 2. Variants of this mode use UI elements like the previously mentioned grids, tiles, and carousels, but also cards or different submenus and subpages. A characteristic of list-like recommendations is that they can all be generated at the same time (e.g., at retrieval time or even offline) and then presented to the user at once. While the order

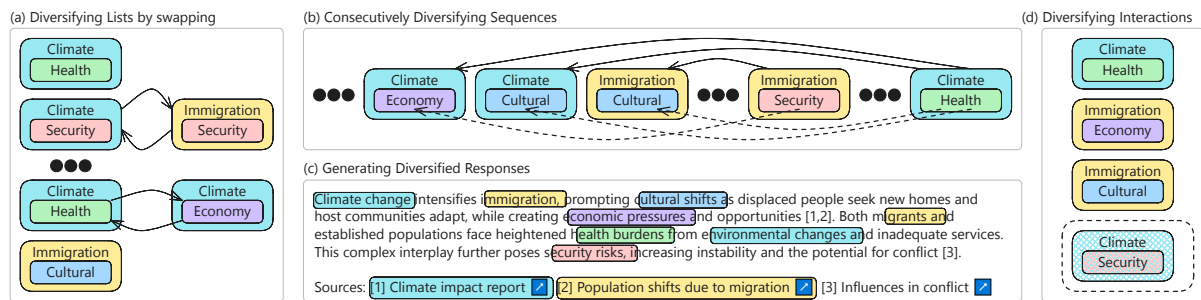


Figure 2: Diversifying news recommendations in different modes. (a) Lists can be diversified by swapping articles with others that have different aspects present. (b) Sequences need to consider items in a specific window, where the focus lies more on the similarity towards more recent items. Arrows denote similarity to previous items. (c) Summaries can be diversified at content or source-level. (d) Considering diversified interactions adds a new layer of complexity to the considerations, i.e., per interaction type such as likes or shares. The suggested interaction (shaded) maximizes both sharing and overall (including likes) diversity.

(or location) is relevant for the user, for measuring diversity it is typically omitted and only evaluated globally on the whole list. Examples for this mode are individual news websites (often regional or specialized ones with limited content), digital versions of print media, or newsletters. A simple way to increase the diversity is by swapping items in the list with dissimilar ones as a post-processing step (Figure 2a). That is, find an item with an overrepresented trait (e.g., label) and replace it with a similar item that has an underrepresented trait in the same aspect instead, which can be repeated. Note that this might not lead to optimal lists and that suitable substitutions might not be available. Furthermore, the presented definition completely dismisses temporal considerations like the presence of a user history. Especially in the news domain, users want to consume novel items rather than old (or repackaged) articles that they have seen previously (i.e., items from their history).

3.2. Diversifying Sequences

Unlike lists, diversifying sequences necessitates temporal considerations. In sequences, recommendations are generated at multiple points and previous items need to be considered. Sequences can generate individual recommendations or sets of recommendations at once. Noteworthy variants are paginated lists or endless scrolling pages. Besides, sequences can often be grouped into distinct sessions, either explicitly (e.g., login and logout) or implicitly (e.g., based on gaps in the consumption). Examples for this mode are large-scale news aggregators, mainstream news websites, RSS feeds/blogs, and social media platforms. When considering the diversity, we can consider all items in a particular window of time, i.e., from a cut-off point in the past until the present, and treat this as a list. When generating new recommendations, consider the history of previous items in the window and the potential candidates, and recommend the ones that maximize the diversity (Figure 2b). Here, the temporal order plays a role. Consider the scenario where two items are equally valid due to a balance in the window. Then recommending the item is more dissimilar to recent items is preferable, as this improves the diversity once the window shifts.

3.3. Diversifying Summaries

When generating summaries, only a single item (i.e., the summary) is recommended. Therefore, the properties of the item itself must be diversified (Figure 2c). This can happen at the source level or content level (or both). At the source level, the articles used to generate the summary should be diverse (similar to list diversity). At the content level, pieces of the content (such as keywords) should have diverse aspects. As summaries with nowadays large language models (LLMs) are generated sequentially, similar principles to diversifying sequences can be applied. Typically, diversity at the source level should lead to diversity at the content level as well. Therefore, a focus lies on diversifying the retrieval

part in retrieval-augmented generation (RAG) to enhance the LLM’s response. Examples for this mode are AI overviews (although typically not displayed for the most novel news), AI search engines, and AI support tools to summarize news pages. A variant of this type is a conversational agent that generates multiple responses. Again, the same principles as in diversifying sequences apply.

3.4. Diversifying Interactions

The three previous modes focus on presenting diversity to the user rather than trying to diversify their actual behavior. Typically, it is assumed that recommender systems suggest relevant items, which the user then consumes (e.g., clicks and then potentially reads). However, different interactions convey different semantics. For instance, a user might like one kind of content but share a different one (Figure 2d). Therefore, another user might observe (and incorrectly assume) a low diversity (e.g., biased) behavior, while a third user might have the completely opposite observation. For that reason, we need to consider the interactions in addition to the interacted items. This can be modeled with a weighted average across different interaction types. Diversified recommendations should consider the type of interaction in the globally interacted items.

4. Proposed Research Direction

Although the four listed recommendation modes require tailored solutions, we believe that at the core a common research direction for all of them emerges.

4.1. Algorithmic Advancements

Novel Diversity Metrics. Diversity is often just assessed after the algorithm has been trained. Even so, established diversity metrics consider singular notions of similarity (e.g., regarding topics [21] or viewpoints [8]), which do not explicitly capture nuanced differences such as in the framing of text. Instead, we call for generalized metrics that consider a set of relevant aspects as showcased in Section 2 to enhance existing diversity considerations. These aspects can be modeled as a knowledge graph to consider the interplay of various labels. In the guiding example, the documents can be represented by (hierarchical) knowledge graphs (depicted left in Figure 3), which can be extracted with a topic and framing classifier from the content. Besides considering diversity separately, it can be incorporated with accuracy for a single target metric.

Knowledge-Graph-Enhanced Models. A naive methodological approach is to optimize for these metrics [22, 23]. Furthermore, we can incorporate knowledge graphs into the news production [see 24, 25, for general and news-specific approaches, respectively]. Specific examples of knowledge-graph enhance models comprise considering similarity [26], preference networks [27], concept graphs [17], and including collaborative edges [28]. These knowledge graphs are then (jointly with the other content) fed into the system to (re-)rank the recommendations. However, instead of using them solely as auxiliary information in the input, we argue to exploit it as a regularizer in the loss (i.e., output of the model). Based on this, we can learn an inherent metric that captures the multi-aspect diversity. After bootstrapping it with some initial knowledge graphs, the aspects can automatically be extracted (i.e., derived) from the content in a self-supervised manner. Therefore, our approach combines subsymbolic representation learning with symbolic representations. Concerning the different modes, the application for lists is straightforward and sequential recommendations can use novel architectures like xLSTM [29]. While sequential in nature, summaries would also benefit from retrieval-augmented generation for source diversity. Diversified interactions are more sophisticated and subject to data-driven reasoning like learned rules [e.g., 30].

Incorporating Transparent Rules. Finally, and crucially, we plan to incorporate transparent rules that operate on the knowledge graphs [31], e.g., Figure 3 right depicting examples rules. The rules can be set either (1.) globally, (2.) context-specific, (3.) per recommendation request, or (4.) derived from the user behavior. 1. Globally defined rules are hardcoded in the system, e.g., to avoid spreading potential




Figure 3: Integrating symbolic knowledge into news recommendations. Different aspects can be modeled as a knowledge graph (left) on which rules can operate (right).

misinformation. Global rules are important for recommender platforms to adhere to regulations and internal goals. 2. Context-specific rules can be set to affect a subset of recommendations, e.g., when a local event is relevant for people within a region. Being able to incorporate context-specific rules allows for quick adaptation of the system. 3. Some parameters are set when retrieving the recommendation, e.g., when a filtering operation is applied by the user. This enables interactivity and control for the user. 4. Certain patterns might be predominant, and explicitly learning them could increase system performance. Learning such rules could also lead to increased transparency in the system, as they can be leveraged for explainability.

4.2. Evaluation Design

The algorithms will be evaluated on established news datasets [32, 33, 34]. Unlike accuracy, which can be estimated offline by maximizing using historical behavior data, determining an ideal level of diversity is more challenging. Maximizing diversity could lead to adverse effects (e.g., backfire when exposed to opposite views [35]), which makes determining a ideal level of diversity more challenging. Also, the accepted level of diversity is user-dependent [36]. One offline evaluation approach is to target a specific diversity level, e.g., derived from the data. Yet, aiming for the expected (i.e. mean) diversity per user would not lead to an increase thereof. Consequently, employing the recommender systems in an online scenario is necessary to determine the actual impact of algorithms.

Our research plan is to determine the impact of the algorithms with user studies, once the offline evaluation shows promising results. We have the following online experiments in mind: using dedicated research platforms like POPROX [37] to recommend diversified lists, employing bots on social media platforms like BlueSky that users can follow to receive personalized sequential recommendations, collaborating with platforms (e.g., the Social Media Accelerator [38]¹) to employ conversational AI agents to discuss news, and building a dedicated dynamic news system for exploring interactivity with rules. For brevity, we only illustrate the design for POPROX². We aim to generate a daily newsletter for a few weeks using a diversified list based on articles from the Associated Press, where we perform the knowledge graph-induced computations for the recommendations while POPROX handles the user recruitment and content delivery. Importantly, POPROX enables users to answer a weekly survey, where we plan to ask general questions about their platform experience (e.g., their enjoyment and engagement), fatigue regarding certain kinds of news (e.g., politics), news habits (e.g., regarding news avoidance). Besides, specific questions on the news items themselves will be asked like whether any surprising news was included (i.e., serendipity) or what feeling a certain news item left the user with (i.e., positivity). We believe that diversified news can favorably influence the subjective perception of users. This is of great societal relevance, as straining the user with repetitive negative news could lead to avoidance and biased consumption patterns that subsequently lead to opinion polarization. We will conduct a randomized controlled trial where we split the user base into groups, with one group being the control, i.e., the standard newsletter, while the rest get the diversified news treatment(s).

¹<https://www.polarizationlab.com/social-media-accelerator>

²See <https://poprox.ai/experimenter> for a platform description

4.3. Research Vision

We believe that diversifying news consumption is great for both societal and individual benefit. First, news recommender systems can indirectly act as bridging algorithms for depolarization. By delivering news that is more moderate to two opposite groups assuming extreme viewpoints, it can lead to higher agreement between them. For instance, by slightly reframing the topic of immigration towards a cultural perspective rather than a security issue (or vice versa). Especially, since it has been shown that news consumers tend to repeatedly consume similarly framed news [39]. Moreover, an audience that has a broader interest diversity, e.g., into a broader spectrum of news, acts as a safeguard for robust democracies [40]. These benefits become more apparent when considering that news consumers might only be exposed to derivatives of news, such as summaries or commentaries, rather than the original sources.

Second, news diversification has benefits for individual users. Related to that is serendipity [41], which is seen as unexpected but useful discoveries. Hence, it requires a broader spectrum of recommendations to include items that the users were unaware of. Therefore, we expect an elevated diversity to not only also increase serendipity but likewise positively influence their perception of the system in general. Besides passive improvements in metrics, users can become more engaged with the system when they have additional options to influence (by enabling/disabling rules to tweak individual preferences) and understand the (more transparent) recommendations. Finally, users could strive to optimize their own behavior. For instance, news recommendations could provide a score on their behavior with rewards for positive actions (like constructively engaging with articles of opposing viewpoints).

5. Conclusion

The problem of considering multiple aspects in news diversity is complex, and existing approaches remain insufficient. We present a novel research direction incorporating symbolic representations to diversify news in four distinct recommendation modes. Particularly in light of the increasing utilization of LLMs for news summarization and active user engagement through a variety of news interactions, it is imperative to address the issue holistically. Our research vision points to a future where people can once again discuss important societal topics based on the same foundational facts while also making personalized news consumption more pleasant.

Acknowledgments

This research was funded in whole or in part by the Austrian Science Fund (FWF) 10.55776/COE12.

Declaration on Generative AI

During the preparation of this work, the author(s) used Gemini, LanguageTool in order to: Grammar and spelling check, Improve writing style, Paraphrase and reword. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

- [1] M. Karimi, D. Jannach, M. Jugovac, News recommender systems—survey and roads ahead, *Information Processing & Management* 54 (2018) 1203–1227.
- [2] A. Bernstein, C. De Vreese, N. Helberger, W. Schulz, K. Zweig, C. Baden, M. A. Beam, M. P. Hauer, L. Heitz, P. Jürgens, et al., Diversity in news recommendations, *arXiv preprint arXiv:2005.09495* abs/2005.09495 (2020).

- [3] J. Kulshrestha, M. Zafar, L. Noboa, K. Gummadi, S. Ghosh, Characterizing information diets of social media users, in: Proceedings of the international AAAI conference on web and social media, volume 9, AAAI Press, 2015, pp. 218–227.
- [4] S. Raza, C. Ding, News recommender system: a review of recent progress, challenges, and opportunities, *Artificial Intelligence Review* 55 (2022) 1–52.
- [5] R. Kühne, C. Schemer, The emotional effects of news frames on information processing and opinion formation, *Communication Research* 42 (2015) 387–407.
- [6] N. Helberger, On the democratic role of news recommenders, in: Algorithms, automation, and news, Routledge, 2021, pp. 14–33.
- [7] S. Vrijenhoek, M. Kaya, N. Metoui, J. Möller, D. Odijk, N. Helberger, Recommenders with a mission: assessing diversity in news recommendations, in: Proceedings of the 2021 conference on human information interaction and retrieval, ACM, 2021, pp. 173–183.
- [8] T. Draws, N. Tintarev, U. Gadiraju, Assessing viewpoint diversity in search results using ranking fairness metrics, *ACM SIGKDD Explorations Newsletter* 23 (2021) 50–58.
- [9] M. Mulder, O. Inel, J. Oosterman, N. Tintarev, Operationalizing framing to support multiperspective recommendations of opinion pieces, in: Proceedings of the 2021 ACM conference on fairness, accountability, and transparency, ACM, 2021, pp. 478–488.
- [10] F. Loecherbach, J. Moeller, D. Trilling, W. van Atteveldt, The unified framework of media diversity: A systematic literature review, *Digital Journalism* 8 (2020) 605–642.
- [11] N. Mattis, P. Masur, J. Möller, W. van Atteveldt, Nudging towards news diversity: A theoretical framework for facilitating diverse news consumption through recommender design, *New Media & Society* 26 (2024) 3681–3706.
- [12] C. Wu, F. Wu, Y. Huang, X. Xie, Personalized news recommendation: Methods and challenges, *ACM Transactions on Information Systems* 41 (2023) 1–50.
- [13] C. Wu, F. Wu, T. Qi, Y. Huang, End-to-end learnable diversity-aware news recommendation, arXiv preprint arXiv:2204.00539 (2022).
- [14] M. Kaya, D. G. Bridge, Accurate and diverse recommendations using item-based subprofiles., in: FLAIRS, 2018, pp. 462–467.
- [15] J. Zhou, E. Agichtein, S. Kallumadi, Diversifying multi-aspect search results using simpson’s diversity index, in: Proceedings of the 29th ACM International conference on information & knowledge management, 2020, pp. 2345–2348.
- [16] R. Li, L. Heitz, O. Inel, A. Bernstein, D-rdw: Diversity-driven random walks for news recommender systems, arXiv preprint arXiv:2508.13035 (2025).
- [17] S. Vercoutere, G. Joris, T. De Pessemier, L. Martens, Improving selection diversity using hybrid graph-based news recommenders, *User Modeling and User-Adapted Interaction* 34 (2024) 955–993.
- [18] A. Iana, G. Glavaš, H. Paulheim, Train once, use flexibly: A modular framework for multi-aspect neural news recommendation, in: Findings of the Association for Computational Linguistics: EMNLP 2024, 2024, pp. 9555–9571.
- [19] M. Kunaver, T. Požrl, Diversity in recommender systems—a survey, *Knowledge-based systems* 123 (2017) 154–162.
- [20] T. Carraro, et al., Neuro-Symbolic Recommender Systems, Ph.D. thesis, Università degli Studi di Padova, 2025.
- [21] 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.
- [22] X. Li, G. Cong, G. Xiao, Y. Xu, W. Jiang, K. Li, On evaluation metrics for diversity-enhanced recommendations, in: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, ACM, 2024, pp. 1286–1295.
- [23] H. Wu, Y. Zhang, C. Ma, F. Lyu, B. He, B. Mitra, X. Liu, Result diversification in search and recommendation: A survey, *IEEE Transactions on Knowledge and Data Engineering* 36 (2024) 5354–5373.
- [24] A. L. Opdahl, T. Al-Moslmi, D.-T. Dang-Nguyen, M. Gallofré Ocaña, B. Tessem, C. Veres, Semantic

- knowledge graphs for the news: A review, *ACM Computing Surveys* 55 (2022) 1–38.
- [25] Q. Guo, F. Zhuang, C. Qin, H. Zhu, X. Xie, H. Xiong, Q. He, A survey on knowledge graph-based recommender systems, *IEEE Transactions on Knowledge and Data Engineering* 34 (2020) 3549–3568.
 - [26] S. A. Puthiya Parambath, N. Usunier, Y. Grandvalet, A coverage-based approach to recommendation diversity on similarity graph, in: *Proceedings of the 10th ACM Conference on Recommender Systems*, ACM, 2016, pp. 15–22.
 - [27] R. Xie, Q. Liu, S. Liu, Z. Zhang, P. Cui, B. Zhang, L. Lin, Improving accuracy and diversity in matching of recommendation with diversified preference network, *IEEE Transactions on Big Data* 8 (2021) 955–967.
 - [28] D. Liu, T. Bai, J. Lian, X. Zhao, G. Sun, J.-R. Wen, X. Xie, News graph: An enhanced knowledge graph for news recommendation., *KaRS@ CIKM* 1 (2019).
 - [29] M. Beck, K. Pöppel, M. Spanring, A. Auer, O. Prudnikova, M. Kopp, G. Klambauer, J. Brandstetter, S. Hochreiter, xlstm: Extended long short-term memory, *Advances in Neural Information Processing Systems* 37 (2024) 107547–107603.
 - [30] S. Manoharan, R. Senthilkumar, An intelligent fuzzy rule-based personalized news recommendation using social media mining, *Computational intelligence and neuroscience* 2020 (2020) 3791541.
 - [31] W. Ma, M. Zhang, Y. Cao, W. Jin, C. Wang, Y. Liu, S. Ma, X. Ren, Jointly learning explainable rules for recommendation with knowledge graph, in: *The world wide web conference*, ACM, 2019, pp. 1210–1221.
 - [32] J. Kruse, K. Lindskow, S. Kalloori, M. Polignano, C. Pomo, A. Srivastava, A. Uppal, M. R. Andersen, J. Frellsen, Eb-nerd a large-scale dataset for news recommendation, in: *Proceedings of the Recommender Systems Challenge 2024*, ACM, 2024, pp. 1–11.
 - [33] F. Wu, Y. Qiao, J.-H. Chen, C. Wu, T. Qi, J. Lian, D. Liu, X. Xie, J. Gao, W. Wu, et al., Mind: A large-scale dataset for news recommendation, in: *Proceedings of the 58th annual meeting of the association for computational linguistics*, Association for Computational Linguistics, 2020, pp. 3597–3606.
 - [34] J. A. Gulla, L. Zhang, P. Liu, Ö. Özgöbek, X. Su, The adressa dataset for news recommendation, in: *Proceedings of the international conference on web intelligence*, ACM, 2017, pp. 1042–1048.
 - [35] C. A. Bail, L. P. Argyle, T. W. Brown, J. P. Bumpus, H. Chen, M. F. Hunzaker, J. Lee, M. Mann, F. Merhout, A. Volfovsky, Exposure to opposing views on social media can increase political polarization, *Proceedings of the National Academy of Sciences* 115 (2018) 9216–9221.
 - [36] W. Wu, L. Chen, Y. Zhao, Personalizing recommendation diversity based on user personality, *User Modeling and User-Adapted Interaction* 28 (2018) 237–276.
 - [37] R. Burke, M. Ekstrand, Conducting recommender systems user studies using poprox, in: *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*, ACM, 2025, pp. 1–2.
 - [38] C. A. Bail, D. S. Hillygus, A. Volfovsky, M. Allamong, F. Alqabandi, D. M. Jordan, G. Tierney, C. Tucker, A. Trexler, A. van Loon, Do we need a social media accelerator?, *SocArXiv* doi 10 (2023).
 - [39] M. Reiter-Haas, E. Lex, The framing loop: Do users repeatedly read similar framed news online?, in: *IUI Workshops*, volume 3660, CEUR-WS.org, 2024.
 - [40] J. Bednar, Polarization, diversity, and democratic robustness, *Proceedings of the National Academy of Sciences* 118 (2021) e2113843118.
 - [41] D. Kotkov, S. Wang, J. Veijalainen, A survey of serendipity in recommender systems, *Knowledge-Based Systems* 111 (2016) 180–192.