

Leveraging Media Frames to Improve Normative Diversity in News Recommendations

Sourabh Dattawad^{1,*}, Agnese Daffara¹ and Tanise Ceron²

¹*Institute for Natural Language Processing, University of Stuttgart, Germany*

²*Department of Computing Sciences, Bocconi University, Italy*

Abstract

Click-based news recommender systems suggest users content that aligns with their existing history, limiting the diversity of articles they encounter. Recent advances in aspect-based diversification – adding features such as sentiments or news categories (e.g. world, politics) – have made progress toward diversifying recommendations in terms of perspectives. However, these approaches often overlook the role of news framing, which shapes how stories are told by emphasizing specific angles or interpretations. In this paper, we treat media frames as a controllable aspect within the recommendation pipeline. By selecting articles based on a diversity of frames, our approach emphasizes varied narrative angles and broadens the interpretive space recommended to users. In addition to introducing frame-based diversification method, our work is the first to assess the impact of a news recommender system that integrates frame diversity using normative diversity metrics: representation, calibration, and activation. Our experiments based on media frame diversification show an improvement in exposure to previously unclicked frames up to 50%. This is important because repeated exposure to the same frames can reinforce existing biases or narrow interpretations, whereas introducing novel frames broadens users' understanding of issues and perspectives. The method also enhances diversification across categorical and sentiment levels, thereby demonstrating that framing acts as a strong control lever for enhancing normative diversity.

Keywords

news recommendation, normative diversity, media frames

1. Introduction

Recent studies have highlighted the importance of diversifying recommendations in news recommendation systems to expose users to a broader range of perspectives [1]. For example, news recommenders with well-designed normative goals ensure exposure to multiple viewpoints, supporting deliberation, and amplifying underrepresented voices [2]. Such diversification is expected to foster democratic participation in public discourse, although empirical evidence remains limited. Recent work highlights that normatively motivated diversification can, under certain conditions, support tolerance and participation, even though its effects are not yet fully established due to the difficulty in evaluation [3]. Existing approaches have made progress in incorporating different perspectives in recommendation models, operationalizing them as "aspects" such as news categories, emotions, and sentiments [4, 5, 6]. However, they have not considered how news stories are narrated or framed. To address this challenge, we propose diversifying perspectives in news recommendations through framing in news coverage. While a previous study by Mulder et al. [7] already explored frame-based news re-ranking, our work is, to the best of our knowledge, the first to apply the full set of media frames from the Media Frames Corpus (MFC) [8] to the core challenge of personalized news recommendation. Besides that, previous works often do not address how encoding different aspects enhances normative diversity, a social-scientific concept reflecting the norms and values of news organizations—such as those emphasized in different

13th International Workshop on News Recommendation and Analytics (INRA), September 22–26, 2025, Prague, Czech Republic

*Corresponding author.

✉ dattawsh@ims.uni-stuttgart.de (S. Dattawad); agnese.daffara@ims.uni-stuttgart.de (A. Daffara);

tanise.ceron@unibocconi.it (T. Ceron)

🌐 <https://sourabhdattawad.github.io/sourabhdattawad/> (S. Dattawad);

<https://scholar.google.com/citations?user=JdOLC-IAAAAJ> (A. Daffara); <https://tceron.github.io> (T. Ceron)

🆔 0000-0002-4642-0743 (S. Dattawad); 0009-0009-6706-0095 (A. Daffara); 0009-0002-4845-2789 (T. Ceron)



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

models of democracy: personal development and autonomy (liberal model), active citizenship (participatory model), access to different viewpoints (deliberative model), access to underrepresented voices (critical model) [2].

In the widespread definition by Entman [9], framing means selecting some parts of reality by making them more salient, with four functions: (i) defining problems, (ii) diagnosing causes, (iii) making judgments, and (iv) suggesting solutions. Framing has received much attention in Media Studies and NLP, as it can shape the dissemination of information and ultimately influence people’s beliefs [10]. For example, it can serve as a powerful tool for agenda-setting and manipulation strategies, because authors can use it to intentionally present some aspects while omitting others [11, 12]. Given its connection to political ideologies [13] and its multifaceted nature, which encompasses different semantic and pragmatic levels [14], framing may prove particularly well-suited for analyzing perspectives in news articles and can effectively serve the purpose of diversifying news recommendations.

We identify two open research questions:

1. Do media frames enhance normative diversity in news recommendations?
2. To what extent can we tune recommender systems with media frames to improve the normative quality of news recommendations?

We implement a frames detection model trained on the MFC [8]. We then integrate these frames as an aspect into the existing MANNeR framework, a news recommendation algorithm that supports aspect-based diversification [15]. The MANNeR framework enables flexible, aspect-aware news recommendation by modularly combining multiple specialized encoders. This design supports both personalization and diversification by allowing targeted control over content dimensions like frames, categories, or sentiment. It is particularly well-suited for our setting due to its efficiency in training and the ease with which it integrates new aspects, eliminating the need for full retraining. To evaluate the system’s ability to reflect diverse frames, we adapt the RADio* metrics, which capture various dimensions of normative diversity [16]. Additionally, we leverage MANNeR’s diversification strategy based on linear score aggregation to fine-tune the amount of frame diversity according to recommendation quality or user needs.

We train and evaluate our models in a multilingual setting using two datasets: the News Portal Recommendations (NPR) dataset in Portuguese [17], and the Ekstra Bladet News Recommendation Dataset (EB-NeRD) in Danish [18]. Our experiments demonstrate that this approach significantly increases exposure to diverse and novel perspectives that users had not previously encountered. Interestingly, we also observe that framing inherently captures both the sentiment and the category of news articles. This enables us to achieve sentiment- and category-based diversification implicitly to an extent, without needing to explicitly encode these features during training. We demonstrate that framing can serve as a powerful lever for steering normative diversity in news recommendation systems, enabling the tuning of aspects such as news category, sentiment, and frames to achieve the desired outcome.

2. Related Work

2.1. Diverse-Aware News Recommenders

The workflow of a news recommender system involves recalling a subset of articles from a large corpus and then ranking them based on user interests. Central to this process is the creation of a user embedding that accurately captures their interests, as well as a news embedding that captures the article’s core semantic meaning. For instance, models like NAML [19] approaches this by using CNNs to encode the news content a user reads and GRUs to learn a sequential representation of their interests. Subsequent research has built upon this, with models like NRMS [20] and HieRec [21] introducing attention mechanisms and hierarchical structures to create even richer user embeddings for the ranking task.

To tackle the impact excessive personalization and of filter bubbles – where users are repeatedly exposed to similar content – diversity has become a central focus in recommender systems. Beyond

algorithmic concerns, user studies have shown that a lack of diversity can lead to frustration [22, 23], and that diversity plays a key role in user acceptance, especially when different models perform similarly in terms of accuracy [24, 25]. Castagnos et al. [26] highlight that diversity increases user confidence in decision-making, while Parapar and Radlinski [27] show it helps elicit user preferences and reduces popularity bias. With respect to diverse news recommendation—where the goal is to recommend content dissimilar to a user’s history in terms of sentiment, emotion, categories, or sources—several approaches have emerged. SentiRec enhances sentiment diversity by explicitly optimizing for a range of sentiment polarities, thus preventing overexposure to emotionally homogeneous content [4]. Emotion-aware models leverage emotional representations of users and articles to recommend affectively novel content, fostering emotional diversity in consumed news [5]. EmoRec incorporates multi-level emotional signals, capturing both coarse and fine-grained affect to provide recommendations that contrast with prior emotional experiences [28]. ProFairRec addresses source-side diversity by promoting provider fairness through debiased embeddings and fairness-aware learning objectives, ensuring exposure to underrepresented news outlets [6]. Graph-based topic nudging exposes users to novel topical areas by expanding and nudging topic representations via knowledge graphs, increasing topic-level dissimilarity from prior consumption [29]. D²RL combines Determinantal Point Processes with actor-critic reinforcement learning to reward both diversity and relevance, promoting dissimilarity in content ranking over time [30]. Finally, MANNeR introduces multi-aspect diversification by simultaneously optimizing for topical variety and sentiment diversity [15].

2.2. Diversity Metrics for Recommender Systems

Diversity in recommender systems can be broadly categorized into two types: descriptive (general-purpose) and normative. Overall, descriptive diversity provides computational tools to increase variety, while normative diversity aligns recommender systems with ethical and democratic responsibilities [31]. The former aims to reduce redundancy in recommendation lists using computational measures of dissimilarity. Ziegler et al. [32] introduce Intra-List Similarity (ILS) and Intra-List Distance as foundational metrics. Other measures include relative diversity, expected intra-list diversity (which incorporates ranking sensitivity), aggregate diversity (across user lists), and the inverse Gini index to capture distributional inequality [33, 34, 35]. However, as Treuillier et al. [36] note, these are often single-number metrics that only partially address broader news recommendation challenges.

Normative diversity extends beyond computational dissimilarity to promote societal goals such as inclusion of voices and opinions, having the objective to strengthen democratic values. Helberger et al. [37] propose a cooperative responsibility framework, emphasizing the shared role of platforms, users, and institutions in realizing public values. Vrijenhoek et al. [31] introduce the DART metrics – calibration, fragmentation, representation, activation, and alternative voices – to capture normative diversity dimensions. While insightful, these lack rank-awareness and uniformity, issues addressed in the RADLo* framework [16], which offers rank-sensitive and consistent metrics for evaluating exposure diversity in recommendation lists. See Section 4 for the precise definitions and computation details of these normative metrics.

2.3. Frames Detection and Application in News Recommendation

Framing categories can be recognized through specific textual and visual devices (e.g. figures, videos) and can be either inducted from the text or established *a priori* [38]. The task of frame detection was first explored in unsupervised settings by treating frames as special types of subtopics [11, 39, 40, 41, 42]. It has since been addressed using supervised or semi-supervised classification techniques, including neural networks [43], fine-tuning of pre-trained models [44], and zero-shot prompting [45, 46].

A widely used inventory of generic frames is the one introduced by Boydston and Gross [40] and adopted by Card et al. [8] in the Media Frame Corpus (MFC), which comprises 15 categories, such as *Politic*, *Morality*, and *Fairness and Equality* (see Table 1 for a complete list with examples). This annotation framework has been applied extensively in subsequent studies [47, 48, 49, 7, 45, 50]. SemEval-2023 Task

3 adopted it for a sub-task on framing detection in online news within a multilingual context—comprising 9 different languages—where all participating teams employed transformer-based models [50]. Although the label set was originally developed from 3 policy issues in the U.S. context, and relying on these coarse-grained categories may lead to some loss of specificity across datasets, the MFC frames have been demonstrated to be generalizable to new contexts and policy issues [46]. We therefore adopt this set in our study.

Mulder et al. [7] are the first to apply framing to diverse recommendation. They propose a re-ranking method for news recommendation lists based on the four functions identified by Entman [9]: (i) Problem Definition, (ii) Causal Attribution, (iii) Moral Evaluation, and (iv) Treatment Recommendation. Specifically, they implement different extraction methods tailored to each constituent: for Problem Definition, they use LDA topic modeling to infer the central issue presented in each article; for Causal Attribution and Moral Evaluation, they apply the IBM Watson Natural Language Processing API to classify news content into topic-related taxonomies and extract sentiment; for Treatment Recommendation, they extract candidate sentences and classify them using the same taxonomy-based classifier. These identified frame functions are then used to measure diversity across articles using different distance metrics, including Kullback-Leibler divergence for LDA-derived topics and a weighted Jaccard index for the taxonomy-based categories. This allows them to re-rank news articles to maximize framing diversity in the recommendation list. While their work focuses on extracting and balancing these four framing functions, our approach shifts attention to the 15 MFC categories, applying them in a topic-generic setting. Rather than focusing on framing constituents, we aim to capture the type of frame conveyed by each article across multiple news categories while leveraging normative diversity.

Frame	Example
Capacity and Resources	<i>Immigration debate: Illegals take jobs from Americans</i>
Crime and Punishment	<i>Two charged in deaths of illegal immigrants in truck</i>
Cultural Identity	<i>Ethnic shift: Immigration—an Irish enclave learns a new language; Mexican immigrants boost a growing Latino population</i>
Economic	<i>Society makes no-interest loans to New York’s immigrants</i>
Fairness and Equality	<i>Strict immigration law unfairly targets Hispanics</i>
External Regulation and Reputation	<i>‘International village’ gets hostile reception</i>
Health and Safety	<i>Colombian drug violence leads to exodus</i>
Legality, Constitutionality and Jurisprudence	<i>House approves bill to abolish INS; The Senate will begin work next week on its own measure dealing with the immigration agency</i>
Morality	<i>County’s undocumented workers say they aren’t here to ‘steal’</i>
Other	<i>U.S. under pressure to carry bigger load</i>
Policy Prescription and Evaluation	<i>President Donald Trump stalls on promise to eliminate J-1 visa program</i>
Political	<i>Following Trump voter fraud allegations, claim that 5.7 million non-citizens voted is wrong</i>
Public Opinion	<i>Immigration: Political evangelicals feel push to take sides</i>
Quality of Life	<i>Big money, cheap labor</i>
Security and Defense	<i>Decision on refugees overdue; U.S. officials must loosen immigration restrictions</i>

Table 1

The 15 frames and their examples extracted from the MFC headlines on migration.

3. Methodology

The process of incorporating media frame awareness into a news recommendation system involves two main stages. In the first stage, we employ a validated classification model to automatically assign a primary media frame to each article. Following Card et al. [8], we define the primary frame as the single, most prominent frame in the text. While articles can contain multiple underlying frames in their title and body, we only focus on the primary one as it has been shown to be the most influential in shaping a reader’s interpretation and serves as a powerful feature for frame-aware recommendation.

Second, we extend the modular MANNr recommendation framework by creating a dedicated module that represents articles based on their frame similarity. This allows us to combine content-based personalization with a new, frame-based dimension, enabling explicit control over the diversity of

recommended news. Lastly, we evaluate the frame-aware news recommenders with normative metrics adapted from RADIo*.

3.1. Media Frames Detection

In our study, we adopt the methodology established by prior work [46] to identify the dominant media frame for each article in our corpus. Specifically, we use a pre-trained XLM-RoBERTa-base model that has been fine-tuned on the MFC dataset [8]. The model assigns a probability distribution over all the 15 frame labels for each article, and we select the one with the highest probability as the primary frame. Table 1 illustrates examples of the frames used in the training.

Regarding the evaluation of this classifier, on the MFC test set, the fine-tuned model achieved an accuracy of 67% and an F1 score of 68%. Frame-level results indicate high predictive performance on common categories, such as Political (F1: 79.85%) and Capacity and Resources (F1: 70.33%). To assess the model’s robustness in out-of-distribution settings, it is evaluated on *FrameNews-PT*, a dataset of 300 manually annotated Portuguese news articles using the same frame schema [46]. On this dataset, the model achieved a reduced accuracy of 53% and an F1 score of 0.48, though it still outperformed the baseline (accuracy: 23%). Notably, performance was particularly weak on abstract or infrequent frames such as *Morality*, *Fairness and Equality*, and *Quality of Life*, mirroring categories with low inter-annotator agreement (Krippendorff’s $\alpha = 0.78$ overall). These results highlight known challenges in generalizing U.S.-centric frame classifiers across languages and cultural contexts—challenges we account for in our own analysis.

3.2. Integration of Media Frames into Recommendation System

The MANNeR framework introduces a modular approach to news recommendation, enabling the flexible integration of various aspects beyond pure content similarity. It achieves this by training specialized News Encoders (NEs) for different aspects (which the authors call A-module), which are then linearly combined at inference time. Building on this, we extend the MANNeR framework to incorporate media frames as an aspect for recommendation, while disabling other aspects originally supported by the framework, such as topics and sentiments.

The content-based personalization module (CR-Module) remains unchanged. Both the candidate and the clicked news articles are encoded using a dedicated News Encoder (NE). The candidate embedding n_c is compared to each clicked news embedding n_{u_i} using a dot product, and the resulting similarity scores are mean-pooled to compute the overall content relevance score:

$$s_{CR}(n_c, u) = \frac{1}{N} \sum_{i=1}^N n_c \cdot n_{u_i} \quad (1)$$

The CR-Module is trained by fine-tuning the underlying Pretrained Language Model (PLM) using Supervised Contrastive Learning (SCL). For each user, clicked news items are treated as positive samples, while non-clicked items—selected via negative sampling—serve as negative examples. The training objective encourages the model to draw the candidate embedding closer to clicked news and push it away from non-clicked ones, thereby structuring the embedding space according to user relevance.

We introduce a dedicated A-Module for media frames while the original MANNeR framework has one A-module for sentiment and one for topics. Following the MANNeR methodology, the Frame-based A-Module, which we call Frame-Module from now on, is trained to create a specialized news representation space. This is achieved by fine-tuning a separate copy of the initial PLM using a SCL. The objective is to group news articles with the same frame label closer together in the embedding space while simultaneously pushing those with different frames further apart. As with other A-Modules, this module encodes news similarity specifically with respect to frames, independent of user preferences.

At inference time, we leverage the specialized NEs from both the CR-Module and the Frame-Module. For a given candidate news article (n_c) and the user’s click history ($H = \{n_{u_1}, n_{u_2}, \dots, n_{u_N}\}$), we

compute two separate similarity scores. The content similarity (s_{CR}) is calculated as the mean-pooled dot product of the candidate’s and the clicked news’s embeddings from the CR-Module. Similarly, the frame similarity (s_{frame}) is computed using the embeddings from the Frame-Module.

The final ranking score (s_{final}) for a candidate news item (n_c) for a user (u) is a linear aggregation of z-score normalized scores:

$$s_{\text{final}}(n_c, u) = s_{\text{CR}}(n_c, H) + \lambda_{\text{frame}} s_{\text{frame}}(n_c, H) \quad (2)$$

In this equation, the hyperparameter λ_{frame} controls the influence of the media frame aspect on the final recommendation. Setting $\lambda_{\text{frame}} > 0$ promotes personalization towards news with similar frames to what the user has previously engaged with, while $\lambda_{\text{frame}} < 0$ will encourage diversification by recommending news with different frames.

4. Evaluation

Our baseline is the CR-module alone as it does not diversify content based on aspects. Moreover, we evaluate how the frame-aware news recommenders perform at different values of λ , ranging from -1 to 1. Since the underlying scores are z-score normalized prior to the linear combination, their scales are made comparable, which justifies the use of this symmetric range for λ . We evaluate all systems in the following metrics:

Descriptive evaluation We evaluate our news recommendation systems using standard ranking metrics: AUC, MRR, and nDCG@k (with $k \in \{5, 10\}$), which respectively measure the ability to distinguish relevant from non-relevant items, how early relevant items appear in the ranking, and the overall quality of the ranking with greater emphasis on higher-ranked relevant items.

Normative evaluation To assess normative diversity, we adopt the metrics defined in RADio* [16], modifying them to capture the framing dimension. The value for these metrics ranges from 0 to 1. Formally, we adapt the method defined in RADio* and extend it to capture frame diversity. Specifically, we define the diversity of a recommendation list P with respect to a context Q in terms of the distribution of media frames as:

$$D_f^*(P, Q) = \sum_{x \in \text{Frames}} Q^*(x) \cdot f\left(\frac{P^*(x)}{Q^*(x)}\right)$$

Here, $x \in \text{Frames}$ refers to each possible media frame (e.g., economic, morality, legality). Q denotes the list of recommended articles, and P denotes the context list which provides the ground truth, such as a user’s reading history or a reference news corpus. $Q^*(x)$ is the proportion of frame x in the recommendation list, and $P^*(x)$ is the corresponding proportion in the context list. Both the recommendation list Q and the context P can be made rank-aware. The function f denotes a divergence metric, viz. Jensen-Shannon divergence (JSD). Table 2 shows the summary features used for the computation of these metrics. The following measures are used:

- **Calibration** – It measures how well recommendations match a user’s historical preferences. In our setting, it compares the distribution of both news categories and frames in the recommended items with those in the user’s history. *Low scores* indicate strong personalization, while *high scores* suggest more novel or diversified recommendations in terms of news categories and frames.
- **Representation** – It evaluates how closely the distribution of frames in the recommended content mirrors the overall distribution in the dataset. *Low scores* mean recommendations align with the dataset’s frame distribution, while *high scores* indicate a deviation that may highlight over- or under-represented frames.

- **Activation** – This metric measures the emotional intensity of recommended content based on the absolute sentiment score of articles. *Low activation* indicates alignment with the typical tone of the platform, while *high activation* reflects the divergence of recommended content from the typical tone of the platform. This metric helps assess how framing influences emotional engagement. For computing sentiment scores of the articles, we use XLM-T, a multilingual sentiment analysis model [51].

Table 2 shows the summary features used for the computation of these metrics. For representation, we adopt the dataset frame distribution as the contextual baseline, as it reflects the editorial and topical balance of the source corpus. This choice is not without limitations. A uniform distribution could alternatively be used if the objective were to maximize exposure to all frames equally, but such an approach would impose a normative stance that may not be realistic or desirable across all contexts. For example, in some domains, aiming for proportionality with the dataset better captures what users would typically encounter in the news ecosystem, whereas in others, equal representation of frames may be preferable to broaden perspectives. Similarly, while our analysis interprets higher divergence as an indicator of improved alignment or representativeness, we note that this is not a universal truth. The desirability of higher or lower divergence depends on the specific normative goals of the recommendation system (e.g., mirroring existing distributions vs. actively counterbalancing them). We therefore present our results as evidence of controllable diversity effects rather than prescribing a single ‘correct’ target distribution.

Metric	Feature	Context	Type	Rank-aware	Desired Value
Calibration (Category & Frame)	Category, Frame	User history	Categorical	Yes	Low: reflects cats, High: diverse cats
Representation (Frame)	Frame class	All articles	Categorical	No	Low: reflect frames, High: diverse frames
Activation (Frame)	Sentiment score	All articles	Continuous (binned)	No	Low: typical tone, High: untypical tone

Table 2

Summary of metrics with features, context, and interpretation. Cats stands for news categories.

4.1. Datasets

We choose to work with the datasets NPR [17] and EB-NeRD [18] because, among the click-history news recommendation datasets, they include the largest amount of hard news (e.g., politics, economics, finance). Datasets with a higher proportion of hard news are more suitable for studying normative diversity, which emphasizes the influence of news on public discourse and civic engagement [17, 52].

NPR It was developed by Globo media, an organization based in Brazil [17]. This dataset includes 1,162,802 randomly sampled users, 148,099 Portuguese news articles, and 1,402,576 impression logs. NPR was specifically designed to support research on normative diversity, making it well-suited for studying the societal impact of hard news content.

EB-NeRD It was compiled from user behavior logs of the Danish newspaper Ekstra Bladet [18]. EB-NeRD features over 1 million users, 37 million impression logs, 251 million interactions, and 125,000 Danish news articles. It emphasizes text-based recommendations for low-resource languages and includes a high proportion of hard news categories such as crime, politics, and economics.

We use the smaller versions of these datasets, *NPR-small* and *EB-NeRD-small*, since the core idea of this research is to analyze the impact of media frames in recommendations. These subsets provide sufficient samples to conduct and validate our experiments while maintaining computational efficiency. Table 5 shows the distribution of articles, users, and impressions across the train, validation, and test splits used for experimentation. Note that the news categories in the NPR dataset are not topical; they primarily consist of abbreviations for Brazilian states. The distribution of news categories across datasets is shown in Table 3. The overall distribution of frames across the two datasets is summarized in Table 4. The observed differences in frame frequencies likely stem from the distinct origins and content

NPR		EB-NeRD	
Category	Prop. (%)	Category	Prop. (%)
sp	14.95	nyheder	20.95
mg	7.37	sport	16.87
rj	6.82	underholdning	16.79
munido	4.63	krimi	14.42
bahia	4.22	side9	7.76
política	4.12	auto	7.07
go	3.83	forbrug	4.58
são paulo	3.58	sex_og_samliv	3.70
economia	3.55	nationen	2.38
pe	2.90	musik	1.87
pr	2.89	ferie	0.76
pa	2.48	play	0.62
ce	2.25	biler	0.60
df	2.07	penge	0.58
to	1.91	opinionen	0.32
rs	1.85	haandvaerkeren	0.19
Other	23.67	Other	4.83

Table 3

Proportion of the top categories in the NPR and EB-NeRD datasets. Note that in the NPR dataset, the two-character categories represent the acronyms of Brazilian states.

Frame	NPR-small (%)	EB-NeRD-small (%)
Economic	8.20	11.65
Capacity and resources	20.08	0.60
Morality	9.70	33.48
Fairness and equality	2.89	17.28
Legality, constitutionality and jurisprudence	14.23	11.23
Policy prescription and evaluation	7.35	5.73
Crime and punishment	1.66	8.27
Security and defense	0.73	2.74
Health and safety	1.39	4.31
Quality of life	27.53	2.58
Cultural identity	4.21	1.19
Public opinion	0.94	0.53
Political	0.95	0.27
External regulation and reputation	0.15	0.14
Other	0.00	0.00

Table 4

Proportion of each frame in the datasets.

of each corpus. Each dataset reflects the unique editorial focus, topical coverage, and national context of its source material. Additionally, variations in writing style across datasets may affect the performance of the frame classification model, potentially contributing to apparent shifts in frame distribution.

5. Experimental Setup


5.1. Training the Recommendation Systems

We adopt training settings similar to those used in MANNer. We use XLM-RoBERTa-base in our experiments on both NPR and EB-NeRD, leveraging its multilingual capabilities to accommodate the


Dataset	#Articles	Users			Impressions		
		Train	Val	Test	Train	Val	Test
NPR-small	65305	57617	9099	17888	64860	9267	18531
EB-NeRD-small	20738	14968	8811	15342	209598	23289	244647

Table 5

Dataset statistics including number of articles, users, and impressions across train/val/test splits.



(a) Frame-shaped embedding space on NPR test dataset.



(b) Frame-shaped embedding space on EB-NeRD test dataset.

Figure 1: t-SNE visualizations of news embeddings after training the frame-based aspect module.

cross-lingual nature of the datasets. The maximum user history length is set to 50 for both datasets. All models are trained using mixed precision and optimized with the Adam optimizer [53]. For the frame-based aspect module, we use a learning rate of $1e-5$ across both datasets. To address class imbalance, we sample 20 instances per class during the aspect module (Frame-Module) training. For the contrastive CR-Module, we adopt a 4:1 negative-to-positive sampling ratio and set the temperature parameter to 0.9 across all datasets. We train the CR-Module and baseline models for 5 epochs on NPR and 20 epochs on EB-NeRD, with a batch size of 8. The Frame-Module is trained for 100 epochs with a batch size of 60. Both modules employ early stopping with a patience of 3 epochs. All experiments are repeated five times with different random seeds, and we report the mean and standard deviation of evaluation metrics across runs. Training is conducted on an NVIDIA RTX A6000 GPU with 48 GB memory. Figure 1 shows the embedding space after training the Frame-Module. The results show that the encoder groups well instances with the same frame closer together in both datasets.

6. Results and Discussion

λ	AUC	MRR	@5	@10	Cal(C)	Cal(F)	Rep(F)	Act
-1.0	66.75 \pm 0.5	40.27 \pm 0.3	45.12 \pm 0.2	54.52 \pm 0.4	87.47\pm0.3	78.45\pm0.2	49.77\pm0.5	58.45 \pm 0.3
-0.4	75.19 \pm 0.3	50.07 \pm 0.3	55.90 \pm 0.4	62.28 \pm 0.3	86.51 \pm 0.3	75.67 \pm 0.2	48.53 \pm 0.3	58.01 \pm 0.3
-0.1	77.74 \pm 0.3	51.29 \pm 0.3	56.94 \pm 0.4	63.10 \pm 0.3	85.64 \pm 0.3	70.15 \pm 0.2	45.51 \pm 0.4	58.46 \pm 0.3
0.0	78.08 \pm 0.3	48.86 \pm 0.3	57.86 \pm 0.3	63.77 \pm 0.3	85.40 \pm 0.2	68.77 \pm 0.2	45.14 \pm 0.4	58.60 \pm 0.3
0.1	78.13\pm0.4	52.73\pm0.3	58.71\pm0.4	64.37\pm0.4	85.39 \pm 0.3	68.40 \pm 0.3	44.89 \pm 0.5	58.61 \pm 0.3
0.4	77.00 \pm 0.3	52.31 \pm 0.3	58.28 \pm 0.3	64.04 \pm 0.3	85.24 \pm 0.2	66.28 \pm 0.3	43.88 \pm 0.4	58.53 \pm 0.3
1.0	73.10 \pm 0.4	48.56 \pm 0.4	56.10 \pm 0.3	60.07 \pm 0.3	84.90 \pm 0.3	63.81 \pm 0.3	43.59 \pm 0.5	58.70\pm0.3

(a) NPR

λ	AUC	MRR	@5	@10	Cal(C)	Cal(F)	Rep(F)	Act
-1.0	56.31 \pm 0.6	33.48 \pm 0.2	39.26 \pm 0.3	47.16 \pm 0.6	74.42 \pm 0.3	65.49 \pm 0.4	71.06 \pm 0.5	60.94\pm0.3
-0.4	59.59 \pm 0.2	35.89 \pm 0.3	42.18 \pm 0.4	49.49 \pm 0.5	74.59\pm0.3	66.92\pm0.3	72.65 \pm 0.4	59.16 \pm 0.2
-0.1	61.04 \pm 0.4	36.08 \pm 0.2	42.31 \pm 0.3	49.61 \pm 0.4	70.78 \pm 0.3	62.87 \pm 0.3	70.70 \pm 0.2	59.40 \pm 0.3
0.0	61.44 \pm 0.2	36.20 \pm 0.3	43.70 \pm 0.3	50.68 \pm 0.5	70.67 \pm 0.3	62.99 \pm 0.3	70.95 \pm 0.4	59.46 \pm 0.2
0.1	61.71 \pm 0.3	38.26\pm0.3	44.75\pm0.3	51.53\pm0.5	70.64 \pm 0.3	63.01 \pm 0.3	70.99 \pm 0.4	59.43 \pm 0.3
0.4	61.72\pm0.2	38.26\pm0.3	44.75\pm0.2	51.53\pm0.4	70.61 \pm 0.3	63.72 \pm 0.3	71.64 \pm 0.5	59.54 \pm 0.2
1.0	60.48 \pm 0.3	37.71 \pm 0.3	44.00 \pm 0.2	50.69 \pm 0.4	71.11 \pm 0.3	65.91 \pm 0.4	73.69\pm0.5	60.34 \pm 0.3

(b) EB-NeRD

Table 6


Performance across λ for (a) NPR and (b) EB-NeRD. @5/@10 = nDCG@5/10; Cal(C/F) = Calibration on category/frame; Rep(F) = Representation on frames; Act = Activation. \pm represents the standard deviation over runs with 5 different random seeds.

Table 6 reports the mean \pm standard deviation of evaluation metrics over 5 independent runs with different random seeds. To assess the effect of the diversification parameter λ , we systematically varied it from -1.0 to 1.0. Figure 2 illustrates the resulting changes in nDCG@10 and the corresponding normative metric scores.


Impact on Personalization and User Satisfaction User satisfaction, measured by nDCG@5 and nDCG@10, peaks at moderate levels of personalization. In NPR, the highest scores occur at $\lambda = 0.1$, with nDCG@5 at 58.71 ± 0.4 and nDCG@10 at 64.37 ± 0.4 . As personalization increases, performance declines—for instance, at $\lambda = 1.0$, nDCG@5 drops to 56.10 ± 0.3 . A similar trend is observed with EB-NeRD, where nDCG@5 and nDCG@10 peak at 44.75 ± 0.3 and 51.53 ± 0.5 for $\lambda = 0.1$ and $\lambda = 0.4$, respectively.

These results indicate that excessive personalization ($\lambda = 1.0$) can introduce redundancy and overly narrow content selection. By relying too heavily on past preferences, recommendation systems risk producing repetitive outputs that miss users’ broader or evolving interests, ultimately reducing content diversity and perceived recommendation quality. This is a tendency also found in Möller and Padó [54].

Impact on Calibration Calibration on frames (Cal(F)) and categories (Cal(C)) is highest when the diversification parameter λ is low and gradually decreases as λ increases. In the NPR dataset, for instance, Cal(F) drops from 78.45 ± 0.2 at $\lambda = -1.0$ to 63.81 ± 0.3 at $\lambda = 1.0$. In the EB-NeRD dataset, on the other hand, the highest scores are reached already at $\lambda = -0.4$ with Cal(C) reaching 74.49 ± 0.3 and Cal(F) equals 66.92 ± 0.3 . This indicates that higher diversification (negative λ values) leads to recommended articles whose frames and categories differ more significantly from the user’s historical consumption. In other words, as diversification increases, the system prioritizes novelty over alignment



(a) Frame aspect diversification for NPR.



(b) Frame aspect diversification for EB-NeRD.

Figure 2: Comparison of diversity-aware scores across EB-NeRD and NPR test datasets.

with past user preferences, resulting in a greater mismatch between recommended content and user history, consequently increasing the number of perspectives the user has access to.

Although categorical information was not explicitly encoded into the frame-based diversification process, it nonetheless influenced categorical diversification. To analyze the impact, we examine the association between news category and frame class. Recognizing that statistical significance is common in large datasets, we calculate Cramér's V to assess the practical significance of the relationship. For the NPR dataset, the analysis yields a Cramér's V of 0.247, indicating a small to moderate association. For the EB-NeRD dataset, the Cramér's V is 0.332, suggesting a moderate association. These results

confirm that while a predictable relationship exists between news category and frame, the effect sizes are not so large as to suggest that frames merely "shadow" news categories. A substantial portion of the variance in framing remains independent of the category. Thus, promoting frame diversity can indirectly enhance category-level diversity as a side effect. This finding highlights a key trade-off: On the one hand, diversifying frames successfully enhances news category diversity, breaking up monolithic recommendation patterns. On the other hand, this side effect can conflict with the goal of providing varied perspectives on a stable news category, as it risks steering users away from their specific subject of interest. Further investigation is needed to understand whether adding news category into another A-module can keep the news categories the same while diversifying frames. This would satisfy the principle of diversifying perspectives on target news categories of interest to users.

Impact on Representation As diversification increases (i.e., lower λ values), the frame representation score ($\text{Rep}(\text{F})$) rises, indicating that the distribution of frames in the recommended content diverges more from that of the candidate pool. This reflects an increase in frame-level diversity, particularly in mitigating the dominance of over- or under-represented frames in the dataset. For example, in the NPR dataset, $\text{Rep}(\text{F})$ reaches its highest value 49.77 ± 0.5 at $\lambda = -1.0$, and gradually declines to 43.59 ± 0.5 at $\lambda = 1.0$. A different, though less pronounced, pattern is observed in EB-NeRD, where $\text{Rep}(\text{F})$ peaks at 73.69 ± 0.5 for $\lambda = 1.0$, but remains consistently high across all λ values due to the dataset’s inherent frame diversity.

Impact on Activation The Activation Score (Act), which measures deviation in sentiment tone, exhibits a non-linear response to the diversification parameter, λ . For the NPR dataset, the score initially decreases from 58.45 ± 0.3 at $\lambda = -1.0$ to a minimum of 58.01 ± 0.3 as λ increases. The trend then reverses, with the Act rising to a peak of 58.70 ± 0.3 at higher λ values. A similar U-shaped trend is observed for the EB-NeRD dataset: the Act is highest at 60.94 ± 0.3 ($\lambda = -1.0$), decreases to 59.16 ± 0.2 at $\lambda = -0.4$, and rises again to 60.34 ± 0.3 as personalization increases ($\lambda = 1.0$). These results suggest that both ends of the λ spectrum, representing strong diversification and personalization, cause the sentiment in recommendations to shift further away from the average tone of the full content set.

Although sentiment was not an explicit feature encoded in our model, the diversification of framing strategies resulted in a certain extent of observable diversification in sentiment. To formally investigate this relationship, a one-way analysis of variance (ANOVA) was conducted to determine how much of the variation in sentiment scores could be explained by the different frame classes. The analysis revealed a statistically significant relationship in both datasets ($p < .001$). More importantly, an effect size analysis provided a measure of practical significance. For the NPR dataset, the frame class demonstrated a large effect on sentiment, accounting for 38.3% of the total variance ($\eta^2 = 0.383$). Similarly, for the EB-NeRD dataset, the frame class had a large effect, explaining 20.0% of the variance in sentiment scores ($\eta^2 = 0.200$). These findings show that frame selection has a strong influence on the sentiment of the text in both corpora. This strongly suggests that framing plays a crucial role in shaping the overall emotional tone, or polarity, of the content. Figure 3 illustrates the distribution of average sentiment polarity by media frame for both the NPR and EB-NeRD datasets. Notably, frames such as *Morality*, *Fairness and Equality*, and *Legality* tend to exhibit slightly more positive sentiment polarity compared to frames like *Economic*, *Capacity and Resources*, and *Crime and Punishment*. This indicates that emphasizing certain frames in the recommendation process can lead to a directional shift in the overall sentiment of the recommended content. However, the fact that they are on average negative across frames in both datasets explains why the Act does not vary considerably across setups.

Overall, the cost-benefit of the trade-off between accuracy and other objectives varies significantly between the two datasets. For the EB-NeRD dataset, the relationship is well-balanced; enhancing representation and calibration is achieved with a minimal cost of only 1-2 points to AUC. In contrast, the trade-off is considerably more pronounced for the NPR dataset, where similar enhancements incur a substantial drop of over 11 points in AUC.





Figure 3: Distribution of average sentiment polarity across media frames.

Impact on Frame Novelty Figure 4 illustrates the impact of the diversification parameter λ (-1.0 to 1.0) on the novelty of recommended frames for users of the NPR and EB-NeRD datasets. Three metrics are tracked: *Average Unique Frames per User*, the mean number of distinct frames within each user’s set of recommendations; *KL Divergence* between a user’s viewing history and their recommendations; and *Average Novel Frames per User*, the mean count of recommended frames that the user has not previously encountered in their history. Across both datasets, increasing diversification ($\lambda = -1.0$) leads to a clear rise in the number of novel frames shown to users and greater divergence from their historical preferences. This is evident in the peak values of novel frame exposure: 1.8 in NPR and 0.20 in EB-NeRD at $\lambda = -1.0$. KL divergence follows a similar trend, indicating that recommendations at higher diversification settings are less aligned with users’ prior interactions. A similar trend is observed in number of unique frames remaining in both datasets, showing only a slight decrease as λ increases.


These results suggest that diversification effectively enhances the novelty and distinctiveness of recommended content across frames, although the extent of this effect is shaped by the underlying diversity in the user history within each dataset—more pronounced in NPR and more constrained in EB-NeRD. This difference can be attributed to the inherent frame diversity in users’ reading histories: EB-NeRD users engage with a wider range of frames on average (11.8) compared to NPR users (3.8), leaving less room for further novelty through diversification. Indeed, as seen in the results above, the effect of the diversification in λ is lower in the EB-NeRD dataset.

7. Conclusion

This work addresses the challenge of enhancing normative diversity in news recommendation by incorporating frame-level information in the recommendation system, as aspect of perspectives that is grounded in how news media construct meaning and shape public understanding. For that, we propose a frame-aware diversification approach that builds on the MANNeR framework to implement frame diversity. We evaluate our algorithm on the RADio* metrics to understand the effect of frames in the normative provisions of recommendation systems.



(a) Impact on EB-NeRD test dataset.



(b) Impact on NPR test dataset.

Figure 4: Variation in novel frame exposure, KL divergence, and unique frames as a function of diversification level (λ).

Our experiments demonstrate that this approach effectively exposes users to novel frames that they have not previously encountered, resulting in increased diversity in recommended perspectives. Additionally, we find that framing implicitly captures sentiment and categorical variation, leading to improvements in normative diversity metrics such as calibration – without the need for explicit modeling of these aspects. The increase in the diversity of news categories, a byproduct of frame-aware news recommendation, also enhances the diversity of content.

We also observe a controllable trade-off between personalization and frame diversity. While prioritizing frame diversity improves normative aspects, it can result in a slight reduction in personalization

metrics, such as accuracy and relevance. This trade-off is predictable and tunable, allowing system designers to strike a balanced configuration that maintains user relevance while promoting exposure to diverse frames. The inclusion of two datasets in our study enabled validation across different scenarios. We found that dataset composition can influence diversification outcomes; for example, greater heterogeneity in the EB-NeRD dataset makes differentiation more challenging but also reduces the risk of excessive personalization. These results highlight that the characteristics of the data should be carefully considered when tuning for diversity.

Together, these findings position media frames as a powerful and interpretable control lever for advancing normative diversity in news recommendation systems. Besides opening research directions for the academic community, our work has potential practical implications for media organizations. The adoption of frames on news platforms can support democratic diversification, fostering social and political engagement by the companies, while still allowing diversity tuning to preserve user satisfaction and interest. Moreover, users who are presented with a variety of viewpoints on social media are theoretically stimulated to participate (e.g., commenting, liking, sharing), thereby increasing interaction. To promote transparency and reproducibility for the wider community, we make the code available on GitHub [<https://github.com/sourabhdattawad/framerec>].

8. Future Work

This study employed the MANNeR framework to explore framing-based diversification. Future work can, however, investigate how other diversification methods or frameworks interact with framing. Our experiments were limited to smaller datasets (EB-NeRD-Small and NPR-Small); scaling to full-size versions could improve benchmarking and generalizability.

We also only encode the primary frame of each article, though many articles contain multiple frames. Extending the analysis to capture this complexity could help us understand its impact on diversification and how different framing combinations influence user understanding and engagement.

To address the side effect of frame diversity indirectly enhancing category-level diversity, we could extend the MANNeR framework by adding an additional category module alongside the existing frame module. This would allow us to diversify recommendations by frame while personalizing them by category. In this way, users are exposed to different frames without deviating from their categories of interest.

An important limitation of our approach is error propagation from the frame detection stage into the diversity analysis. Since the frame classifier achieves only moderate out-of-domain performance ($F1 \approx 0.48$ on Portuguese data), misclassifications may distort both calibration and representation scores. As a result, observed differences in divergence metrics should be interpreted with caution. Future work could explore robustness analyses that explicitly quantify the impact of classifier errors on downstream diversity measurements.

While our current approach includes calibration, representation, and activation metrics, future work could expand the RADio* framework with new metrics – namely fragmentation and alternative voices – the former to examine how framing can tackle users’ isolation by reducing exposure to only narrow or isolated viewpoints that are not widely shared among users, and the latter by affecting the visibility of minority perspectives.

Acknowledgements

We gratefully acknowledge funding from the German Federal Ministry of Research, Technology and Space (BMFTR) under the grant 01IS23072 for the Software Campus project MULTIVIEW.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT, Grammarly in order to: Grammar and spelling check, 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. Reuver, A. Fokkens, S. Verberne, No NLP task should be an island: Multi-disciplinarity for diversity in news recommender systems, in: H. Toivonen, M. Boggia (Eds.), *Proceedings of the EACL Hackashop on News Media Content Analysis and Automated Report Generation*, Association for Computational Linguistics, Online, 2021, pp. 45–55. URL: <https://aclanthology.org/2021.hackashop-1.7/>.
- [2] N. Helberger, On the democratic role of news recommenders, in: *Algorithms, automation, and news*, Routledge, 2021, pp. 14–33.
- [3] N. Mattis, P. K. Masur, J. Moeller, W. van Atteveldt, It ain't easy: using normatively motivated news diversification to facilitate policy support, tolerance, and political participation, *Information, Communication & Society* 28 (2025) 651–668. URL: <https://doi.org/10.1080/1369118X.2024.2423892>. doi:10.1080/1369118X.2024.2423892. arXiv:<https://doi.org/10.1080/1369118X.2024.2423892>.
- [4] C. Wu, F. Wu, T. Qi, Y. Huang, SentiRec: Sentiment diversity-aware neural news recommendation, in: K.-F. Wong, K. Knight, H. Wu (Eds.), *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, Association for Computational Linguistics, Suzhou, China, 2020, pp. 44–53. URL: <https://aclanthology.org/2020.aacl-main.6/>. doi:10.18653/v1/2020.aacl-main.6.
- [5] N. Babanejad, A. Agrawal, H. Davoudi, A. An, M. Papagelis, Leveraging emotion features in news recommendations, in: *INRA@RecSys*, 2019. URL: <https://api.semanticscholar.org/CorpusID:208614444>.
- [6] T. Qi, F. Wu, C. Wu, P. Sun, L. Wu, X. Wang, Y. Huang, X. Xie, Profairrec: Provider fairness-aware news recommendation, 2022. URL: <https://arxiv.org/abs/2204.04724>. arXiv:2204.04724.
- [7] 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*, 2021, pp. 478–488.
- [8] D. Card, A. E. Boydston, J. H. Gross, P. Resnik, N. A. Smith, The media frames corpus: Annotations of frames across issues, in: *Annual Meeting of the Association for Computational Linguistics*, 2015. URL: <https://api.semanticscholar.org/CorpusID:5636935>.
- [9] R. M. Entman, Framing: Toward Clarification of a Fractured Paradigm, *Journal of Communication* 43 (1993) 51–58. doi:10.1111/j.1460-2466.1993.tb01304.x.
- [10] S. Lecheler, M. Keer, A. R. Schuck, R. Hänggli, The effects of repetitive news framing on political opinions over time, *Communication Monographs* 82 (2015) 339–358.
- [11] O. Tsur, D. Calacci, D. Lazer, A frame of mind: Using statistical models for detection of framing and agenda setting campaigns, in: C. Zong, M. Strube (Eds.), *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Association for Computational Linguistics, Beijing, China, 2015, pp. 1629–1638. URL: <https://aclanthology.org/P15-1157/>. doi:10.3115/v1/P15-1157.
- [12] A. Field, D. Kliger, S. Wintner, J. Pan, D. Jurafsky, Y. Tsvetkov, Framing and agenda-setting in Russian news: a computational analysis of intricate political strategies, in: E. Riloff, D. Chiang, J. Hockenmaier, J. Tsujii (Eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 3570–3580. URL: <https://aclanthology.org/D18-1393/>. doi:10.18653/v1/D18-1393.

- [13] S. Alashri, S. Alzahrani, L. Bustikova, D. Siroky, H. Davulcu, What animates political debates? analyzing ideological perspectives in online debates between opposing parties, in: Proceedings of the ASE/IEEE International Conference on Social Computing (SocialCom-15), Academy of Science and Engineering, Stanford, CA, 2015.
- [14] J. Otmakhova, S. Khanezhari, L. Frermann, Media framing: A typology and survey of computational approaches across disciplines, in: Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2024, pp. 15407–15428.
- [15] A. Iana, G. Glavaš, H. Paulheim, Train once, use flexibly: A modular framework for multi-aspect neural news recommendation, 2024. URL: <https://arxiv.org/abs/2307.16089>. arXiv:2307.16089.
- [16] S. Vrijenhoek, G. Bénédict, M. Gutierrez Granada, D. Odijk, Radio* – an introduction to measuring normative diversity in news recommendations, ACM Trans. Recomm. Syst. 3 (2024). URL: <https://doi.org/10.1145/3636465>. doi:10.1145/3636465.
- [17] J. P. Lucas, J. F. G. da Silva, L. F. de Figueiredo, Npr: a news portal recommendations dataset, in: NORMalize@RecSys, 2023. URL: <https://api.semanticscholar.org/CorpusID:269043679>.
- [18] J. Kruse, K. Lindschow, 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, RecSys Challenge '24, ACM, 2024, p. 1–11. URL: <http://dx.doi.org/10.1145/3687151.3687152>. doi:10.1145/3687151.3687152.
- [19] C. Wu, F. Wu, M. An, J. Huang, Y. Huang, X. Xie, Neural news recommendation with attentive multi-view learning, 2019. URL: <https://arxiv.org/abs/1907.05576>. arXiv:1907.05576.
- [20] C. Wu, F. Wu, S. Ge, T. Qi, Y. Huang, X. Xie, Neural news recommendation with multi-head self-attention, in: K. Inui, J. Jiang, V. Ng, X. Wan (Eds.), Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Association for Computational Linguistics, Hong Kong, China, 2019, pp. 6389–6394. URL: <https://aclanthology.org/D19-1671/>. doi:10.18653/v1/D19-1671.
- [21] T. Qi, F. Wu, C. Wu, P. Yang, Y. Yu, X. Xie, Y. Huang, Hierrec: Hierarchical user interest modeling for personalized news recommendation, 2021. URL: <https://arxiv.org/abs/2106.04408>. arXiv:2106.04408.
- [22] M. Zhang, N. Hurley, Avoiding monotony: improving the diversity of recommendation lists, in: Proceedings of the 2008 ACM Conference on Recommender Systems, RecSys '08, Association for Computing Machinery, New York, NY, USA, 2008, p. 123–130. URL: <https://doi.org/10.1145/1454008.1454030>. doi:10.1145/1454008.1454030.
- [23] N. Lathia, S. Hailes, L. Capra, X. Amatriain, Temporal diversity in recommender systems, in: Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '10, Association for Computing Machinery, New York, NY, USA, 2010, p. 210–217. URL: <https://doi.org/10.1145/1835449.1835486>. doi:10.1145/1835449.1835486.
- [24] M. D. Ekstrand, F. M. Harper, M. C. Willemsen, J. A. Konstan, User perception of differences in recommender algorithms, in: Proceedings of the 8th ACM Conference on Recommender Systems, RecSys '14, Association for Computing Machinery, New York, NY, USA, 2014, p. 161–168. URL: <https://doi.org/10.1145/2645710.2645737>. doi:10.1145/2645710.2645737.
- [25] N. Jones, User Perceived Qualities and Acceptance of Recommender Systems: The Role of Diversity, Ph.D. thesis, EPFL, Lausanne, 2010. URL: <https://infoscience.epfl.ch/handle/20.500.14299/47714>. doi:10.5075/epfl-thesis-4680.
- [26] S. Castagnos, N. Jones, P. Pu, Eye-tracking product recommenders' usage, in: Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10, Association for Computing Machinery, New York, NY, USA, 2010, p. 29–36. URL: <https://doi.org/10.1145/1864708.1864717>. doi:10.1145/1864708.1864717.
- [27] J. Parapar, F. Radlinski, Diverse user preference elicitation with multi-armed bandits, in: Proceedings of the 14th ACM International Conference on Web Search and Data Mining, WSDM '21, Association for Computing Machinery, New York, NY, USA, 2021, p. 130–138. URL: <https://doi.org/10.1145/3437963.3441786>. doi:10.1145/3437963.3441786.

- [28] M. Sertkan, J. Neidhardt, Exploring expressed emotions for neural news recommendation, in: Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization, UMAP '22 Adjunct, Association for Computing Machinery, New York, NY, USA, 2022, p. 22–28. URL: <https://doi.org/10.1145/3511047.3536414>. doi:10.1145/3511047.3536414.
- [29] S. Vercoutere, G. Joris, T. De Pessemier, L. Martens, Improving consumption diversity via graph-based topic nudging, *User Modeling and User-Adapted Interaction* 35 (2025). doi:10.1007/s11257-025-09429-1.
- [30] Y. Liu, Y. Zhang, Q. Wu, C. Miao, L. Cui, B. Zhao, Y. Zhao, L. Guan, Diversity-promoting deep reinforcement learning for interactive recommendation, 2019. URL: <https://arxiv.org/abs/1903.07826>. arXiv:1903.07826.
- [31] 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, CHIIR '21, Association for Computing Machinery, New York, NY, USA, 2021, p. 173–183. URL: <https://doi.org/10.1145/3406522.3446019>. doi:10.1145/3406522.3446019.
- [32] 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, WWW '05, Association for Computing Machinery, New York, NY, USA, 2005, p. 22–32. URL: <https://doi.org/10.1145/1060745.1060754>. doi:10.1145/1060745.1060754.
- [33] K. Bradley, Improving recommendation diversity, *Proc. AICS '01* (2001).
- [34] 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, RecSys '11, Association for Computing Machinery, New York, NY, USA, 2011, p. 109–116. URL: <https://doi.org/10.1145/2043932.2043955>. doi:10.1145/2043932.2043955.
- [35] D. Fleder, K. Hosanagar, Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity, Working Papers 07-10, NET Institute, 2007. URL: <https://ideas.repec.org/p/net/wpaper/0710.html>.
- [36] C. Treuillier, S. Castagnos, E. Dufraisse, A. Brun, Being diverse is not enough: Rethinking diversity evaluation to meet challenges of news recommender systems, in: Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization, UMAP '22 Adjunct, Association for Computing Machinery, New York, NY, USA, 2022, p. 222–233. URL: <https://doi.org/10.1145/3511047.3538030>. doi:10.1145/3511047.3538030.
- [37] N. Helberger, J. Pierson, T. P. and, Governing online platforms: From contested to cooperative responsibility, *The Information Society* 34 (2018) 1–14. URL: <https://doi.org/10.1080/01972243.2017.1391913>. doi:10.1080/01972243.2017.1391913. arXiv:<https://doi.org/10.1080/01972243.2017.1391913>.
- [38] C. H. De Vreese, News framing: Theory and typology, *Information design journal+ document design* 13 (2005) 51–62.
- [39] V.-A. Nguyen, J. Boyd-Graber, P. Resnik, K. Miler, Tea party in the house: A hierarchical ideal point topic model and its application to republican legislators in the 112th congress, in: C. Zong, M. Strube (Eds.), Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), Association for Computational Linguistics, Beijing, China, 2015, pp. 1438–1448. URL: <https://aclanthology.org/P15-1139/>. doi:10.3115/v1/P15-1139.
- [40] A. E. Boydston, J. H. Gross, Identifying media frames and frame dynamics within and across policy issues, 2013. URL: <https://api.semanticscholar.org/CorpusID:17157113>.
- [41] M. E. Roberts, B. M. Stewart, D. Tingley, C. Lucas, J. Leder-Luis, S. K. Gadarian, B. Albertson, D. G. Rand, Structural topic models for open-ended survey responses, *American journal of political science* 58 (2014) 1064–1082.
- [42] Y. Ajjour, M. Alshomary, H. Wachsmuth, B. Stein, Modeling frames in argumentation, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019,

pp. 2922–2932.

- [43] N. Naderi, G. Hirst, Classifying frames at the sentence level in news articles, in: R. Mitkov, G. Angelova (Eds.), *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, INCOMA Ltd., Varna, Bulgaria, 2017, pp. 536–542. URL: <https://aclanthology.org/R17-1070/>. doi:10.26615/978-954-452-049-6_070.
- [44] H. Kwak, J. An, Y.-Y. Ahn, A systematic media frame analysis of 1.5 million new york times articles from 2000 to 2017, 2020. URL: <https://arxiv.org/abs/2005.01803>. arXiv:2005.01803.
- [45] F. Gilardi, M. Alizadeh, M. Kubli, Chatgpt outperforms crowd workers for text-annotation tasks, *Proceedings of the National Academy of Sciences* 120 (2023) e2305016120.
- [46] A. Daffara, S. Dattawad, S. Padó, T. Ceron, Generalizability of media frames: Corpus creation and analysis across countries, 2025. URL: <https://arxiv.org/abs/2506.16337>. arXiv:2506.16337.
- [47] S. Khanehzar, A. Turpin, G. Mikołajczak, Modeling political framing across policy issues and contexts, in: *Australasian Language Technology Association Workshop*, 2019. URL: <https://api.semanticscholar.org/CorpusID:209892156>.
- [48] S. Khanehzar, T. Cohn, G. Mikołajczak, A. Turpin, L. Frermann, Framing unpacked: A semi-supervised interpretable multi-view model of media frames, *arXiv preprint arXiv:2104.11030* (2021).
- [49] J. Mendelsohn, C. Budak, D. Jurgens, Modeling framing in immigration discourse on social media, in: K. Toutanova, A. Rumshisky, L. Zettlemoyer, D. Hakkani-Tur, I. Beltagy, S. Bethard, R. Cotterell, T. Chakraborty, Y. Zhou (Eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics, Online, 2021, pp. 2219–2263. URL: <https://aclanthology.org/2021.naacl-main.179/>. doi:10.18653/v1/2021.naacl-main.179.
- [50] J. Piskorski, N. Stefanovitch, G. Da San Martino, P. Nakov, Semeval-2023 task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multi-lingual setup, in: *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, 2023, pp. 2343–2361.
- [51] F. Barbieri, L. Espinosa Anke, J. Camacho-Collados, XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond, in: *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, European Language Resources Association, Marseille, France, 2022, pp. 258–266. URL: <https://aclanthology.org/2022.lrec-1.27>.
- [52] S. Vrijenhoek, Do you mind? reflections on the mind dataset for research on diversity in news recommendations, 2023. URL: <https://arxiv.org/abs/2304.08253>. arXiv:2304.08253.
- [53] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2017. URL: <https://arxiv.org/abs/1412.6980>. arXiv:1412.6980.
- [54] L. Möller, S. Padó, Explaining neural news recommendation with attributions onto reading histories, *ACM Trans. Intell. Syst. Technol.* 16 (2025). URL: <https://doi.org/10.1145/3673233>. doi:10.1145/3673233.