



HAL
open science

NeMig - A Bilingual News Collection and Knowledge Graph about Migration

Andreea Iana, Mehwish Alam, Alexander Grote, Katharina Luwig, Philipp Müller, Christof Weinhardt, Heiko Paulheim

► **To cite this version:**

Andreea Iana, Mehwish Alam, Alexander Grote, Katharina Luwig, Philipp Müller, et al.. NeMig - A Bilingual News Collection and Knowledge Graph about Migration. Proceedings of the Workshop on News Recommendation and Analytics co-located with RecSys 2023, Sep 2023, Singapore, Singapore. 10.48550/arXiv.2309.00550 . hal-04197081

HAL Id: hal-04197081

<https://telecom-paris.hal.science/hal-04197081v1>

Submitted on 5 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

NeMig - A Bilingual News Collection and Knowledge Graph about Migration

ANDREEA IANA, University of Mannheim, Germany

MEHWISH ALAM, Télécom Paris, Institut Polytechnique de Paris, France

ALEXANDER GROTE and NEVENA NIKOLAJEVIC, Karlsruhe Institute of Technology, Germany

KATHARINA LUDWIG and PHILIPP MÜLLER, University of Mannheim, Germany

CHRISTOF WEINHARDT, Karlsruhe Institute of Technology, Germany

HEIKO PAULHEIM, University of Mannheim, Germany

News recommendation plays a critical role in shaping the public’s worldviews through the way in which it filters and disseminates information about different topics. Given the crucial impact that media plays in opinion formation, especially for sensitive topics, understanding the effects of personalized recommendation beyond accuracy has become essential in today’s digital society. In this work, we present *NeMig*, a bilingual news collection on the topic of migration, and corresponding rich user data. In comparison to existing news recommendation datasets, which comprise a large variety of monolingual news, *NeMig* covers articles on a single controversial topic, published in both Germany and the US. We annotate the sentiment polarization of the articles and the political leanings of the media outlets, in addition to extracting subtopics and named entities disambiguated through Wikidata. These features can be used to analyze the effects of algorithmic news curation beyond accuracy-based performance, such as recommender biases and the creation of filter bubbles. We construct domain-specific knowledge graphs from the news text and metadata, thus encoding knowledge-level connections between articles. Importantly, while existing datasets include only click behavior, we collect user socio-demographic and political information in addition to explicit click feedback. We demonstrate the utility of *NeMig* through experiments on the tasks of news recommenders benchmarking, analysis of biases in recommenders, and news trends analysis. *NeMig* aims to provide a useful resource for the news recommendation community and to foster interdisciplinary research into the multidimensional effects of algorithmic news curation.

CCS Concepts: • **Information systems** → **Recommender systems**; **Test collections**.

Additional Key Words and Phrases: news corpus, user data, recommender system, personalization, bias analysis, aspect diversity, knowledge graph

ACM Reference Format:

Andreea Iana, Mehwish Alam, Alexander Grote, Nevena Nikolajevic, Katharina Ludwig, Philipp Müller, Christof Weinhardt, and Heiko Paulheim. 2018. *NeMig - A Bilingual News Collection and Knowledge Graph about Migration*. In . ACM, New York, NY, USA, 18 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

The large volume of news published online each day creates an information overload that exceeds the consumptive capacities of readers. At the same time, the digitalization of news consumption has sparked an unprecedented development of personalized recommendation algorithms, used by online news platforms to process the continuously growing quantities of news and offer readers personalized suggestions. However, models which are optimized to maximize

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

Manuscript submitted to ACM

congruity to users’ preferences and past behavior, tend to produce recommendations which are highly similar in content to previously read/clicked ones [40, 67]. By filtering out news deemed irrelevant to the user, news recommenders control how and which information is disseminated to the public [47]. Consequently, algorithmic news curation has a large impact on opinion formation and voting behavior [6, 15], as well as the potential to fuel conflicts and polarization [3, 8, 30, 40, 44], especially when dealing with controversial subjects such as war, climate change, or migration. In addition to the power of recommender systems to shape people’s perception of the world, the fundamental role that media plays in today’s society as a public forum [4, 23], have deemed analyzing and understanding the implications of news curation algorithms necessary for modern democracies.

While plethora of recommendation algorithms have been proposed in recent years, resources available to analyze their effects beyond accuracy performance and in multilingual scenarios are scarce. The existing body of work exhibits two main shortcomings: (1) news datasets for recommendation (i) mostly focus on general or less sensitive topics (e.g., entertainment, fashion, sports) [20, 69] and (ii) leave largely unexplored features which are critical for analyzing the algorithmic creation of “filter bubbles” [47] or underlying recommender biases (e.g., sentiment or political orientation); (2) news collections for media analysis (e.g., fake news detection [24, 62], news narratives analysis [36, 55], news bias detection [16, 38]) lack user information, and thus, cannot easily be used in a recommendation scenario.

In this paper, we introduce *NeMig* a bilingual news collection and knowledge graphs (KGs) about refugees and migration, in German and English, as well as corresponding real-world user data. We collect over 7K, and respectively, 10K articles from German and US news outlets. The articles span a large political spectrum in order to cover various viewpoints on the topic. **1)** Compared to existing resources, *NeMig* (i) targets the same polarizing topic in two languages and (ii) is annotated with the sentiment polarization of articles and political leanings of the media outlets, in addition to subtopics and disambiguated named entities. **2)** We construct corresponding knowledge graphs (KGs) from the news’ text, metadata, and extracted named entities, which we further expand with up to two-hop neighbors from Wikidata of the named entities. We provide the KGs in different variants. **3)** We collect and publish in an anonymized fashion both explicit feedback and socio-demographic data for 3K users for each of the two datasets. Although our user dataset is small compared to large benchmark news datasets, it contains information about the users’ media consumption, political attitudes and interests, personality traits and demographics, which makes the dataset valuable for studying the effects of recommender systems beyond accuracy-based performance, e.g., on political polarization. Furthermore, the user data constitutes a starting point for the generation of synthetic user datasets, which comprise not only click behavior information, but also explicit data about the users’ background and preferences.

We demonstrate the utility of the proposed resource through experiments on different downstream tasks. Firstly, we benchmark several news recommenders on *NeMig*. Secondly, we investigate biases and polarization effects of the benchmarked recommendation models, and evaluate the contribution of various features to the quality of the KG. Lastly, we analyze trends in the news by examining the evolution over time and political orientation of the most frequent entities in our corpora, identifying correlations between the patterns of evolution and worldwide events.

NeMig is available in two versions, for research purposes: i) we release the full news datasets in N-Triples format, under restricted access on Zenodo [25], given that due to copyright policies, the body of the news can only be provided upon request, ii) *NeMig*, with all features with the exception of the news bodies, is freely available, under a CC BY-NC-SA 4.0 license, in tabular format on GitHub.¹ The user data is available in anonymized tabular format on both Zenodo [25] and GitHub.²

¹https://github.com/andreeaiana/nemig_nrs/tree/main/data

²https://github.com/andreeaiana/nemig/tree/main/data/user_data

2 RELATED WORK

News Recommendation Benchmarks. Various datasets have been constructed for developing and benchmarking news recommender systems (NRS). Plista [32] comprises a collection of over 70K German articles gathered from 13 news portals, as well as click data for over 14 million users. Globo [13, 14] is a Portuguese dataset for news recommendations collected from Globo.com, and contains data for over 300K users and 46K news distributed in 1.2 million sessions. The Norwegian dataset Adressa [20] contains news collected from the Adresseavisen’s news portal and is provided in two variants: a light version with over 11K articles and click data for 561K users, and a larger one with more than 48K articles and clicks for 3 million users. In addition to metadata (e.g., categories, authors), the articles are annotated with named entities. MIND [69], the most recent large-scale dataset, is constructed from user click logs of Microsoft News. It covers 1 million users and more than 161K news articles in English. Moreover, the dataset is enriched with named entities linked to Wikidata.

News Knowledge Graphs. Knowledge graphs have been shown to model knowledge-level connections between news that cannot be captured by purely text-based models [26, 60]. In the context of the Newsreader³ project, focusing on the multilingual processing of news articles, the authors generate an event-centric KG from a multilingual news dataset that describes which events took place, where, when, and who was involved in them [52, 58]. News Graph [39] is a graph constructed specifically for news recommendation. It is generated from a news corpus collected from MSN News containing over 621k articles, 594k news entities, as well as user-item interaction logs. The graph contains news content, user behaviors, and news topic entities enriched with neighboring triples from Microsoft Satori, along with three types of collaborative relations for entities, i.e., co-occurring in the same news, clicked by the same user, and clicked by the same user in the same browsing session. The use of semantic KGs for the production, distribution, and consumption of news is summarized in [46].

News Datasets for Media Discourse Analysis. The advent of news platforms as an ubiquitous means of information for Internet users has established online news as a fundamental source for various media discourse analysis tasks. For instance, Horne et al. [24] has created a large dataset of political news collected from mainstream and alternative sources, enriched with content-based and social media engagement features, that can be used for news and engagement characterization, news attribution, and content copying, or exploration of news narratives. Similarly, Lim et al. [38] and Färber et al. [16] collected and annotated articles with different types of bias at the sentence level for analyzing news biases. In a related line of work, researchers have designed datasets for monolingual [55, 62] and multilingual cross-domain [54] fake news detection, or for distinguishing between fake and satire news published online [18, 42].

Limitations of Current News Resources. Existing news datasets have several shortcomings. On the one hand, the large monolingual recommendation benchmarks benefit from sizable user interaction data, but fail to explore news (e.g., sentiment orientation, political leaning) or user features (e.g., demographics, political interests) which can be exploited in the analysis of recommender system effects. On the other hand, the news collections for media discourse analysis contain rich annotations (e.g., news biases, social media features, social context, spatio-temporal information), but are ill-equipped to be seamlessly utilized in training recommendation algorithms due to the lack of user feedback data. In contrast to these resources, NeMig comprises a (i) bilingual news collection, annotated not only with text and metadata information, but also with subtopic, sentiment and political orientation information, and (ii) implicit user feedback,

³<http://www.newsreader-project.eu>

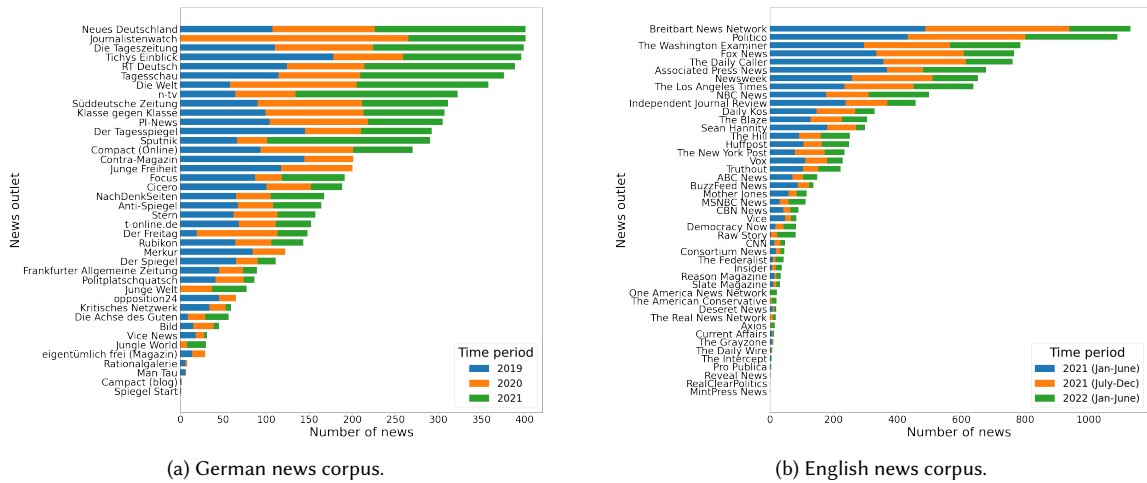


Fig. 1. Distribution of news over time and over outlets.

demographic, and political data. Moreover, we use Wikidata as an open source for entity linking, and not a commercial knowledge base, as done in [39].

3 THE NEMIG NEWS CORPORA

3.1 Data Collection

We construct the news corpora by crawling news articles from 40 German⁴, and respectively 45 US⁵, legacy and alternative media outlets, selected by a team of researchers from the media and communication domain. The news are sampled from the entire political spectrum, such as to cover different viewpoints on the topic. We include German articles published between Jan. 1, 2019 and Dec. 31, 2021⁶, and English ones between Jan. 1, 2021 and July 1, 2022⁷. For each language, we identify relevant articles using 13 keyword stems representative of the migration topic, such as *flüchtl**, *asyl**, or *migration** for the German outlets, and *refugee**, *asylum seeker**, or *migrant** for the US sources. Moreover, only articles written in either German or English (for the German and English corpora, respectively), with a minimum length of 150 words, and containing at least two keywords stems have been included in the two corpora in order to avoid foreign language articles, disclaimers, advertisements, or reader comments. From the resulting raw datasets, we further excluded duplicates, videos, live tickers (articles continuously updated with news headlines), and outliers⁸ (articles which are abnormally long or short). The datasets contain 7,346 (German), and respectively 10,814 (English) filtered articles. Figs. 1a and 1b show the distribution of the articles over time and over media outlets. We observe that the news is unevenly distributed over various publishers, with smaller or niche media outlets being underrepresented. In comparison, with the exception of a few outlets, the articles are relatively evenly distributed over the different time periods.

⁴The implementation of the crawlers for the German news collection is available at <https://github.com/andreeaiana/german-news>

⁵The implementation of the crawlers for the English news collection is available at <https://github.com/andreeaiana/us-news>

⁶The data was collected in Dec. 2020 and Jan. 2021, and updated in Jan. 2022.

⁷The data was collected in Aug. 2022.

⁸Outliers were identified by using two standard deviations from the mean.

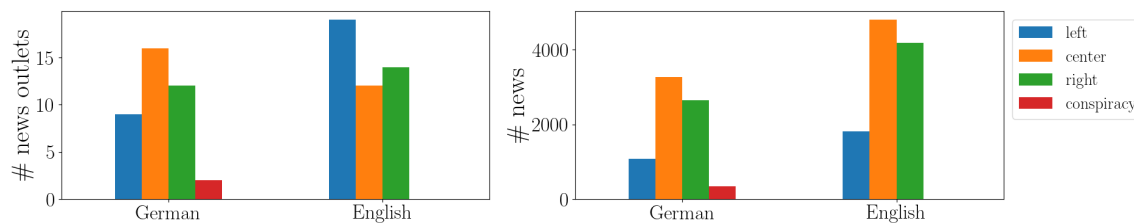


Fig. 2. Distribution of outlets and news over political orientation classes.

Each news article contains a title, a body, and if available, an abstract. Additionally, it contains metadata information specifying its provenance (news outlet, URL), publishing and modification dates, authors (persons or organizations), and keywords provided by the media outlet. Furthermore, we classify the political leaning of the news outlets. Concretely, we categorize the US sources into three political orientation classes, namely *left*, *center*, and *right*, based on the AllSides Media Bias Chart⁹. Since no corresponding classification exists for the German media, the same researcher team classifies the outlets into four political orientation groups: *left*, *center*, *right*, *conspiracy*.¹⁰ We find that the distribution of news outlets over the political orientation classes differs between the German and the English datasets. While the majority of outlets for German are center media, the majority of outlets for English are left. On the article level, however, the distribution for the English dataset is drastically different and similar between the two corpora.

3.2 Data Annotation

Fig. 3 illustrates the corpora annotation and KG construction pipeline.¹¹ Given an input dataset, the majority of steps can be run end-to-end.¹²

Sentiment Analysis. Sentiment polarization information is leveraged by fairness-aware news recommenders to produce sentiment-diverse or agnostic recommendations [67, 68]. Furthermore, such information can be used to investigate whether algorithmic news curation propagates or encourages the sentiment polarization of readers [9, 45, 47] or, more generally, the sentiment bias of different news outlets [17, 35]. We use a multilingual *XLM-R* [11] language model, trained on 198M tweets and fine-tuned for sentiment analysis [5], to classify each article into one of three sentiment classes: positive, neutral, or negative. We perform sentiment classification on the concatenation of the news’ title and abstract.¹³ Table 1 shows the distribution of articles over sentiment classes in the two datasets. For both corpora, only a small number of articles have a positive sentiment. However, while the distribution of articles in the neutral and negative classes is relatively balanced in the German corpus, the English corpus has significantly more negative articles.

Sub-topic Modeling. We observe that news cover discourses about the refugee migration from different geographical areas or related to various political events. In order to extract sub-topics from each dataset, we use the neural topic modeling approach proposed in [19], with a pre-trained English Sentence Transformer for the English news, and multilingual Sentence Transformer [51] for the German ones. We assign labels to each of the resulting topics based on

⁹<https://www.allsides.com/media-bias/media-bias-chart>

¹⁰Disclaimer: We excluded the outlet *Man Tau* from further analysis regarding the political orientation as it cannot be categorized in any of the four groups, but we include the corresponding news in the dataset.

¹¹The implementation of the annotation and KG construction pipeline, and the intermediary files produced by the annotation process are available at <https://github.com/andreeaiana/nemig>

¹²The sub-topic extraction step constitutes the only exception, as it requires setting parameters (e.g., threshold for pruning topics) based on intermediary results. We run this process separately and integrate its output within the pipeline.

¹³Note that we use the concatenation of the title and the first sentences of the news body as input when abstracts are not provided.

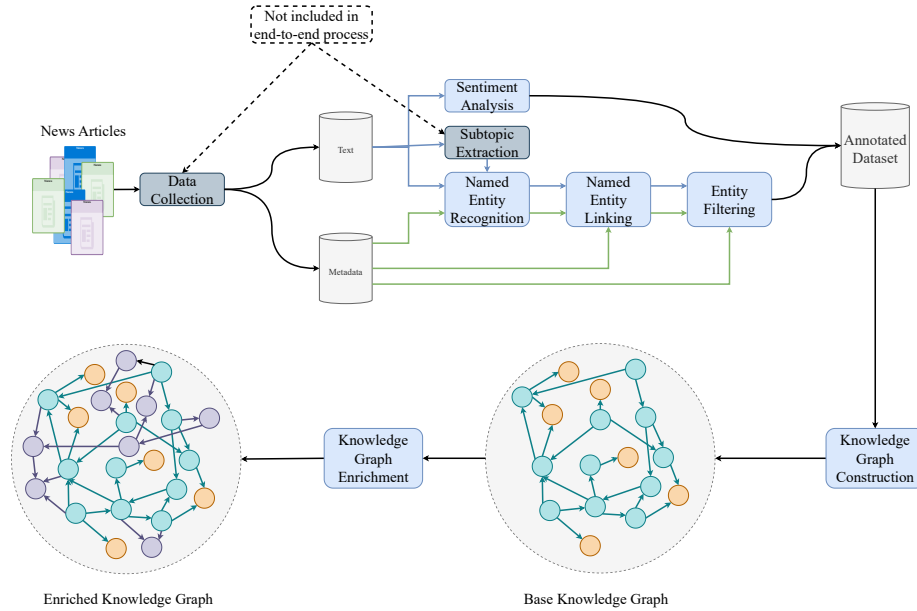


Fig. 3. Corpora annotation and knowledge graph construction pipeline. Arrow colors denote the applicability of steps to: all components (black), text features (blue), and metadata (green). Node colors indicate: custom or linked resources (blue), literals (orange), and Wikidata neighboring nodes (purple). Edge colors denote properties in the original KG (blue), or extracted from Wikidata (purple).

the topmost representative terms of the respective cluster of documents. We group articles that cannot be assigned to any topic appearing in at least 15 news into a separate cluster. We extract 25 sub-topics from the German corpora, and 40 from the English one, respectively. We find, as shown in Fig. 4, that the identified sub-topics indicate migration from (e.g., Libya, Ukraine) or to (e.g., Greece, US) different areas, migration-related political events (e.g., Russian invasion of Ukraine, US army retreating from Afghanistan) or policies (e.g., US immigration and border control laws).

Named Entity Recognition. Events are generally described in news by means of named entities (NEs) that indicate what, when, and where something happened, or who was involved in it [37]. We extract NEs from both the textual content, and the metadata of articles, and classify them into four classes: persons (PER), organizations (ORG), locations (LOC), and miscellaneous (MISC). We use two *XLM-R* [11] language models, one fine-tuned on the German subset¹⁴ and the other on the English subset¹⁵ of the CoNLL03 dataset [53] to perform named entity recognition. Concretely, we input the title, abstract, or body of the news after applying sentence segmentation [49] for text components, whereas for metadata, we use author names, lists of keywords from the news outlets, and those used to describe sub-topics. In the former case, we output only the NEs recognized in these original input. In the latter scenario, the output consists not only in NEs extracted by the model but also of the original inputs, where no entity was extracted. In the KG construction step, we include NEs extracted from metadata, as well as other information (e.g., frequent keywords, unrecognized authors) to model connections between articles.

Entity Linking. Next, we disambiguate the NEs through named entity linking to Wikidata [59] using a multilingual entity linking model [12], which is based on a sequence-to-sequence architecture [41] that generates entity names in

¹⁴<https://huggingface.co/xlm-roberta-large-finetuned-conll03-german>

¹⁵<https://huggingface.co/xlm-roberta-large-finetuned-conll03-english>

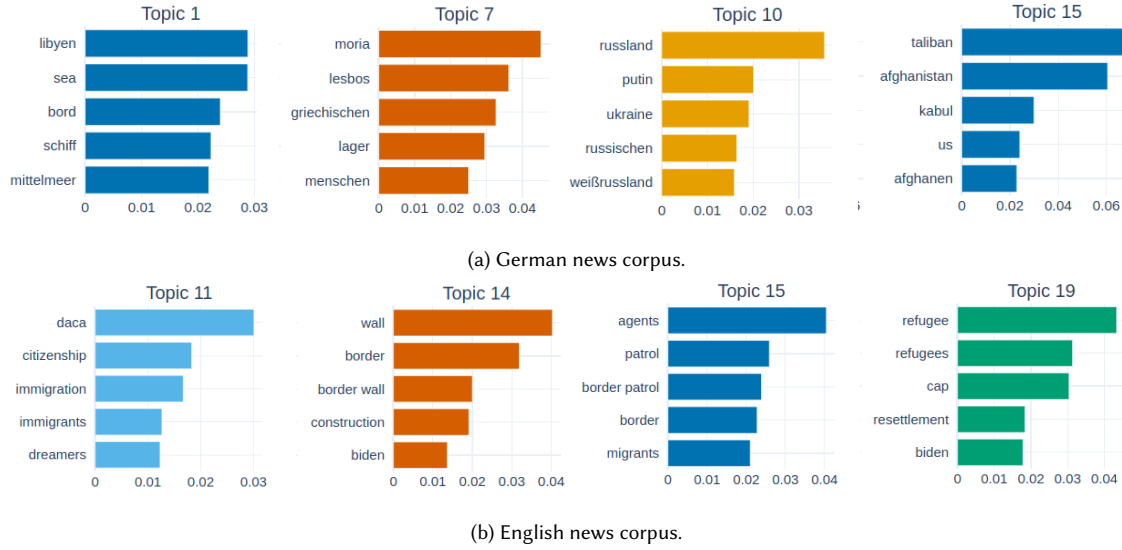


Fig. 4. Examples of words assigned to different sub-topics. The X-axis indicates the c-TF-IDF scores of the words representative of each sampled topic.

Table 1. News corpora statistics.

		German	English
# News	raw	8,647	12,220
	filtered	7,346	10,814
	positive	144	590
	negative	3,476	2,933
	neutral	3,726	7,291
Avg #words	title	7.9	11.8
	abstract	29.4	24.9
	body	898.4	1,013.9

Table 2. Statistics of extracted named entities

	German			English		
	Total	Linked	Not linked	Total	Linked	Not linked
Title	1,075	1,075	0	1,542	1,542	0
Abstract	2,383	2,383	0	2,365	2,365	0
Body	19,683	19,683	0	33,857	33,857	0
Publishers	40	34	6	45	44	1
Authors	490	193	297	1,242	386	856
Keywords	4,481	2,395	2,086	5,175	2,636	2,539

over 100 languages. As input, we use a text sequence containing one named entity annotated with special start and end tokens. The input text sequence is either the title, abstract, or sentence-segmented news body, or for metadata, either the extracted entity from the previous step or the original keyword or author name. For both types of input, the model outputs the linked entity, the corresponding Wikidata QID, the language of the Wikidata entity label, as well as a score. We select Wikidata as the external knowledge base for linkage and graph expansion due to (i) being open source, (ii) its wide coverage, and (iii) its up-to-dateness [22], which is particularly important in the context of news.

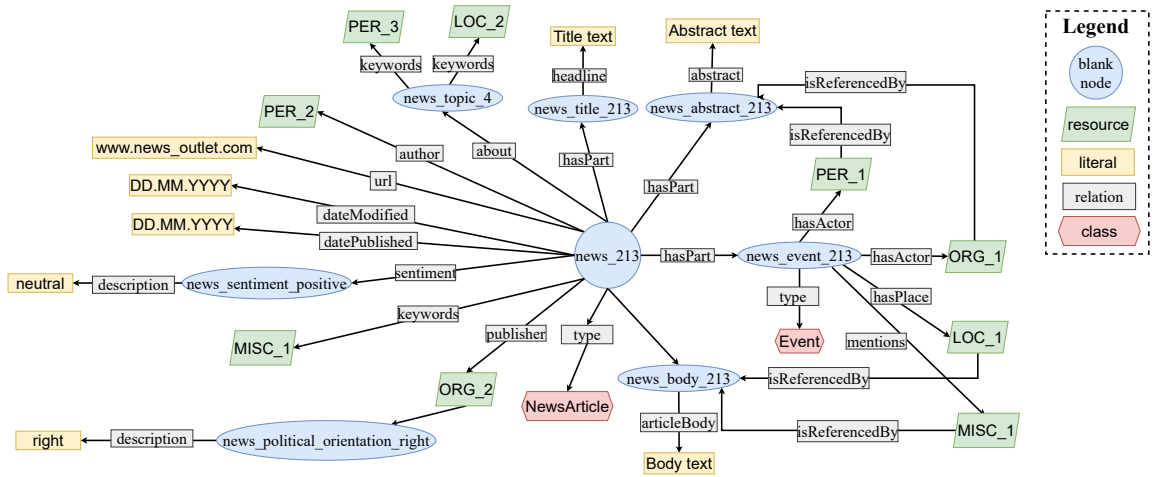


Fig. 5. Example of an article’s representation in NeMigKG. The node colors indicate: (i) blank nodes (blue); (ii) disambiguated or custom resources (green); (iii) literals (yellow); (iv) classes (red). The gray boxes denote relation types.

Entity Filtering. The previous step can yield both entities not identified in Wikidata although they exist (e.g., keyword *academia* corresponds to Wikidata QID *Q1211427*), as well as incorrect links. Therefore, we filter incorrectly extracted or linked entities. For NEs originating from textual components, we (1) threshold on the confidence score of the named entity recognition model; (2) remove entities without a Wikidata page or linked to a Wikimedia disambiguation page; (3) remove entities whose type does not correspond to the entity type in Wikidata based on type-specific properties¹⁶; (4) threshold on the score of the named entity linking model; (5) remove entities whose Wikidata label is in a foreign language and are only observed once in the dataset. For authors and keywords, we use similar filtering pipelines. Table 2 shows the statistics for the named entity recognition and linking steps.

4 KNOWLEDGE GRAPH CONSTRUCTION

Wikidata [59] is a large knowledge graph comprising a vast amount of information that is often used to model and integrate common-sense knowledge into recommender systems [21, 26]. Nevertheless, it also contains data which can be irrelevant for domain-specific recommendation tasks. Thus, we use the annotated corpora to build a domain-specific KG, which combines textual content and metadata from news articles with Wikidata triples of the extracted NEs and their *k*-hop neighbors.

4.1 Base graph construction

NeMigKG is a KG in which the nodes denote news content and real-world entities (e.g., authors, locations), while the edges indicate different relation types between them. Fig. 5 illustrates an example of an article in NeMigKG.

Relations. We model relations using several schemas and vocabularies. Specifically, we represent edges between articles and their textual content or metadata with *schema.org*¹⁷ relationships, and those involving extracted NEs with the *Simple Event Model* [56]. Furthermore, we encode information about miscellaneous NEs with the *schema:mentions*

¹⁶We used the following properties to detect the type of an entity in Wikidata: PER is *instance of human*; ORG has *headquarters location*, *inception*, or is *founded by*; LOC has *coordinate location*.

¹⁷<https://schema.org/>

Table 3. Statistics of the German and English NeMigKG.

		German NeMigKG					English NeMigKG				
		Base	Entities	Enriched ($k = 1$)	Enriched ($k = 2$)	Complete	Base	Entities	Enriched ($k = 1$)	Enriched ($k = 2$)	Complete
Triples	Nodes	89,601	59,480	97,021	367,514	397,635	134,775	91,833	138,734	433,965	476,907
	Relations	19	11	768	1,172	1,180	19	11	828	1,195	1,203
	Triples	516,948	458,482	821,529	2,846,849	2,905,315	933,801	847,695	1,354,454	3,599,561	3,685,667
Nodes	% resources	66.38	100	100	100	92.42	68.14	100	100	100	91.00
	% literals	33.62	0	0	0	7.58	31.86	0	0	0	9.00
	% blank	61.28	61.29	37.57	9.92	9.92	58.48	58.48	38.71	12.38	12.38
Resources	% custom (not linked)	4.00	4.00	2.45	0.65	0.65	3.68	3.68	2.43	0.78	0.78
	% Wikidata (linked)	34.71	34.71	59.97	89.43	89.43	37.84	37.84	58.86	86.85	86.85

relation, and map the provenance of NEs to the news text with the *dcterms:isReferencedBy*.¹⁸ Additionally, we create the relations *nemig:sentiment* and *nemig:political_orientation* to represent edges between news and their sentiment labels, and between media outlets and their political leaning, respectively. Lastly, we identify pairs of related news¹⁹, i.e., articles with identical titles, overlapping bodies, but different provenance, which we map with the *schema:isBasedOn* relation. We find that such related news are based on one another, with the most recent published one representing an update, extension, or longer discussion of the original article.

Nodes. NeMigKG contains two kinds of nodes: (i) *literals* (e.g., title, dates), and (ii) *resources* which encode NEs extracted from different parts of an article’s content (e.g. persons, locations) or metadata (i.e., publisher, author). We further distinguish the latter into (i) disambiguated entities from Wikidata (denoted as *Wikidata resources*) and (ii) *custom resources* created from entities that could not be identified or linked to Wikidata, but which still encode meaningful information and provide knowledge-level connections between news (e.g., authors not in Wikidata, or frequent keywords). We represent each article in NeMigKG with a unique identifier. For each title, abstract, body, topic, sentiment, political leaning, or event mapped in the KG, we create additional nodes, as shown in Fig. 5.

Enrichment with external information. We extend NeMigKG with up to 2-hop neighbors from Wikidata of all *Wikidata resources* already included in our graph to inject additional information from external knowledge bases. Firstly, we extract, for all *Wikidata resources*, all triples in Wikidata containing another entity as their tail. We ignore literal neighbors. Secondly, we incorporate the matching triples into the graph and retrieve a new neighbors set for the added tail entities. We repeat this process iteratively until we enrich NeMigKG with all relevant triples two hops away from the original *Wikidata resources*. We post-process the resulting graph by removing all sink entities, i.e., nodes with a degree smaller than two, which cannot be used to derive further connections between the news as they are leaf nodes or not linked to Wikidata.

4.2 Graph types

The majority of KG embedding models used in downstream tasks encode only entities and relations, while ignoring literals. Therefore, we create four variants of NeMigKG for each of the two datasets, as follows: (1) *base* - contains literals and entities from the corresponding news corpus; (2) *entities* - contains only resource nodes from the *base* graph, and no literals; (3) *enriched entities* - contains only resource nodes and k -hop triples from Wikidata; (4) *complete* - represents the *enriched entities* graph with the literals re-added. Table 3 shows the statistics of the different KGs for both languages.

¹⁸<https://www.dublincore.org/specifications/dublin-core/dcmi-terms/terms/isReferencedBy/>

¹⁹Related news are identified using the overlap ratio between pairs of news’ bodies.

Table 4. A sample of questions used during the online user study conducted in the US for collecting user data.

Variable	Question	Answer Scale
Demographics		
Gender	Please indicate your gender.	0 = male, 1 = female, 2 = other, 3 = no answer
Nationality	What is your citizenship?	0 = U.S. citizenship, 1 = U.S. and another non-U.S. citizenship, 2 = Only non-U.S. citizenship, 3 = No Answer
Empathy		
EMPI	How strongly do you agree with the following statement?	When someone else is feeling excited, I tend to get excited too.
EMPI2	It upsets me to see someone being treated disrespectfully.	
Ideological		
IPO1	The U.S. should take in more refugees.	
Ideological		
IPO2	The U.S. has already taken in too many refugees.	
Polarization		
IPO5	Immigrants strive for peaceful cohabitation with U.S.-Americans.	7-point Likert scale: 1=Strongly disagree to 7=Strongly agree
IPO6	Immigrants are hostile toward U.S.-Americans.	
Emotions		
EMO1	Which emotions do you feel towards refugees and immigrants in the USA?	Anger Fear Compassion
EMO2		
EMO9		
EMO13		
Media Usage		
MED1	I receive information about US-American politics via:	TV, radio, newspapers and magazines or their websites (e.g. New York Times, The Wall Street Journal, ...)
MED2		TV networks or their websites (e.g. Fox News, CNN...)
MED4		Facebook (for political information)
Perceived		
PRO1	How strongly do you agree or disagree with the following statement?	Democratic and Republican partisans in the U.S. are increasingly hostile to one another.
Polarization		
PRO3	Democratic and Republican partisans in the U.S. are very polarized.	7-point Likert scale: 1=Strongly disagree to 7=Strongly agree
Affective		
REP	In the following we would like to know about your party identification. Please mark on the scale how warm or cold you feel towards the respective parties.	Republican Democrat
DEM		0 (negative) to 100 (positive)
Political Scale		
POL1	Where on the scale would you place your political point of view?	11-point Likert Scale: 1 = Left, 6 = Center, 11 = Right
POL2	Please indicate how strongly you agree or disagree with the following statements.	I am generally very interested in politics.
Political Topics		
POL3	I regularly inform myself about current political affairs in the U.S.	7-point Likert scale: 1=Strongly disagree to 7=Strongly agree
Participation		
PPA1	Please indicate how likely it is that you will engage in the following activities in the near future.	Participating in an online political discussion on the topic of immigration to the U.S. Participating in a political demonstration on the topic of immigration to the U.S.
Prosocial		
PRO3	Please indicate how strongly you agree or disagree with the following statements.	I am willing to make a one-time donation of ...\$ for refugees in the U.S. How often do you have professional contact (e.g. at work or at school) with immigrants?
Behavior		7-point Likert Scale: 1 = Never to 7 = Very often

5 USER DATA

We collect user data in terms of (i) explicit click feedback, (ii) demographics, and (iii) political information. For each corpus, we conduct one online user study aimed at measuring the political polarization effects of NRS. The participants are recruited through online-access panels and selected using a quote procedure to create a representative sample of German, and respectively, US Internet users aged 18 to 74. Among the German users, 50.2% are male and 49.6% are female, while 44.6% have an Abitur (passed the secondary school final examinations). 46.2% of the US users are male and 52.5% are female, and 33.2% have at least a high school degree.

In the study, each participant is firstly asked to rate 10 articles as *worthy* or *not to be read further* (i.e., explicit feedback). We use the IDs of the news deemed worthy of reading to build the user *Click History*, which is later used to construct the profile of each user. Afterwards, each participant is shown a list of six news (impressions), for which again a binary rating is recorded. This process was repeated five times. We use the information gathered in the second step to construct the *Impressions* log. Similar to [69], we structure the anonymized user behaviors in the format [*impression ID*, *user ID*, *Click History*, *Impression Log*], where the *Impression Log* contains the IDs of the news shown to the user and the label indicating whether the user clicked on them.

Before the study’s stimulus phase, we collect information regarding the users’ media usage (*MED*[1-9]), political attitudes (*POL*[1-9]), and empathy levels (*EMP*[1-8]). After the stimulus, we ask participants questions regarding their socio-demographic status (gender, age, qualification, nationality, born in Germany or the US, parents born in Germany or the US, income), emotions (*EMO*[1-14]), levels of ideological (*IPO*[1-12]), affective (i.e. identification with the main political parties), and perceived (*PPO*[1-6]) polarization, of political participation (*PPA*[1-8]), and of prosocial behavior (*PRO*[1-5]). The resulting user datasets contain information about 3,432 German and 3,000 US users, respectively. The German dataset has a sparsity of only 0.02%, meaning that out of the news included in the online user study, less than 1% were never seen by the participants, in either of the two study phases. In contrast, the English dataset is much sparser, with 9.27% of the news never interacted with by any of the users. Table 4 shows an example of questions used during the user data collection process in the US. Please refer to the project page²⁰ or Zenodo [25] for a full set of questions used in both studies, as well as an extensive explanation of each question, and the corresponding answer scale.

²⁰https://github.com/andreeaiana/nemig/tree/main/data/user_data

6 EXPERIMENTS

We analyze the performance of different news recommenders on NeMig on a range of recommendation tasks. Afterwards, we perform news trends analysis using the named entities identified in the corpora.

6.1 Recommendation Models

We benchmark various recommendation models, which differ in their news and user encoders.²¹ The former converts input features (e.g., title, topics, entities) into a news embedding by contextualizing pretrained word embeddings [48]. The later aggregates the embeddings of the clicked news into a user-level representation.

- *NRMS* [65] embeds news titles with a two-layer encoder consisting of multi-head self-attention [57], followed by additive attention [2]; a similar architecture is used to learn user representations.
- *NAML* [63] uses a sequence of convolutional neural network (CNN) [33] and additive attention to learn representations from news titles, and abstracts, and additionally leverages categories through a linear category encoder; its user encoder consists of an additive attention layer.
- *MINS* [61] embeds the titles and abstracts of news with the same encoder as NRMS [65], and categories through a linear embedding layer; it learns user representations with a combination of multi-head self-attention, multi-channel GRU-based recurrent network [10], and additive attention.
- *CAUM* [50] combines the text encoder used by NRMS [65] with an entity embedder composed of attention layers in order to learn news representations; it produces candidate-aware user representations with a candidate-aware self-attention network which models long-range dependencies between clicked news, conditioned on the candidate, combined with a candidate-aware CNN that captures short-term user interests from adjacent clicks, also conditioned on the candidate’s content.
- *DKN* [60] is a knowledge-aware recommender which learns representations from news titles with a knowledge-aware CNN [33] over aligned embeddings of words and entities, and of users with a candidate-aware attention network.
- *TANR* [64] has the same architecture as NAML [63], but does not use a category embedding layer; additionally, it injects information on topical categories, by jointly optimizing the recommender for content personalization and topic classification.
- *SentiDebias* [68], built on the architecture of NRMS [65], addresses the problem of sentiment debiasing using adversarial learning to reduce the model’s sentiment bias (originating from the user data) and generate sentiment-diverse recommendations.

We evaluate the aforementioned NRS not only in terms of the standard content personalization performance, but also w.r.t. aspect-based diversity and personalization of results, meaning that we investigate how diverse or faithful the recommendations are to the user’s consumed news in terms of aspect A_p . Concretely, following [28], we define *aspect-based diversification* as the level of uniformity of an aspect’s distribution among the recommended news. In contrast, we define *aspect-based personalization* as the level of homogeneity between a user’s recommendations and clicked news w.r.t. the distribution of an aspect (e.g., sentiment). We experiment with three aspects: *political leaning*, *topical categories*, and the *sentiment* of news. Note that we use the political leaning of a news outlet as proxy for determining the political leaning of an article originating from that outlet.

²¹The implementation of the news recommenders is available at https://github.com/andreeaiana/nemig_nrs

6.2 Evaluation Metrics

We use the common metrics AUC, MRR, nDCG@5, and nDCG@10 to report content personalization performance. Following [28], we measure aspect-based diversity of recommendations at position k using the normalized entropy of aspects A_p 's distribution in the recommendation list:

$$D_{A_p}@k = - \sum_{j \in A_p} \frac{p(j) \log p(j)}{\log(|A_p|)}, \quad (1)$$

where $|A_p|$ denotes the number of classes of aspect A_p . We evaluate aspect-based personalization²² with the generalized Jaccard similarity [7]:

$$PS_{A_p}@k = \frac{\sum_{j=1}^{|A_p|} \min(\mathcal{R}_j, \mathcal{H}_j)}{\sum_{j=1}^{|A_p|} \max(\mathcal{R}_j, \mathcal{H}_j)}, \quad (2)$$

where R_j and H_j represent the probability of a news with class j of aspect A_p to be contained in the recommendations list \mathcal{R} , and, respectively, in the user history \mathcal{H} . All metrics are bounded to the $[0, 1]$ range.

6.3 Experimental Setup

We use pre-trained 300-dimensional German²³ and English²⁴ GloVe embeddings [48] and 100-dimensional TransD embeddings [31] pretrained on NeMigKG to initialize the word and entity embeddings of the NRS. Note that we use the entity embeddings trained on the variant of NeMigKG enriched with 1-hop Wikidata neighbors. Moreover, we use the extracted sub-topics as topical categories in NAML, MINS, and TANR, and embed them with 100-dimensional vectors. Following [66], we sample four negatives per positive example. We train all models for 10 epochs, using a batch size of 4.²⁵ We optimize with the Adam algorithm [34], with the learning rate set to 1e-5. We set all other model-specific hyperparameters to optimal values reported in the respective papers. We use 70% randomly sampled user interactions for training the NRS, 10% as validation data, and the remaining 20% as a test set. We repeat each experiment five times with different random seeds and report averages and standard deviations. We normalize the scores to the $[0, 100]$ range.

6.4 Results and Discussion

We first discuss the recommendation performance of the aforementioned NRS on NeMig. We then analyze whether they are prone to implicit biases in terms of three aspects: *political leaning*, *topical categories*, and *sentiment orientation*. Lastly, we investigate the influence of different input features on the quality of NeMigKG.

Content Personalization. Table 5 summarizes the results on content personalization for both datasets. We find that, despite the small dataset sizes, the state-of-the-art neural models achieve high predictive performance. We notice a high similarity w.r.t. content personalization performance of NAML and TANR, as well as for NRMS and SentiDebias. Both pairs of models share nearly the same news and user encoder architectures, and differ only in their optimization objectives. These results indicate that secondary optimization goals (i.e., topic classification in the case of TANR, or sentiment debiasing in SentiDebias) have little effect on pure recommendation performance.

Generally, we find that the two knowledge-aware models, DKN and CAUM, outperform all the other recommenders. This suggests that injecting additional external knowledge into the model can help detect additional, knowledge-level

²²Note that perfect aspect-based personalization would imply identical distributions of aspect A_p 's in the recommendations list and user history.

²³<https://www.deepset.ai/german-word-embeddings>

²⁴<https://nlp.stanford.edu/projects/glove/>

²⁵We trained each model on a single NVIDIA RTX 2080Ti GPU.

Table 5. Recommendation performance of different NRS. The best results per column are highlighted in bold, the second best are underlined.

Model	German				English			
	AUC	MRR	nDCG@3	nDCG@6	AUC	MRR	nDCG@3	nDCG@6
NRMS	57.96±0.66	47.83±0.46	44.90±0.65	60.49±0.35	52.40±0.89	41.46±0.59	37.49±1.07	54.21±0.53
NAML	50.49±0.13	47.37±0.84	44.27±1.20	60.13±0.65	50.02±0.20	41.82±0.70	38.14±0.84	54.49±0.62
MINS	57.70±0.53	47.85±0.78	44.96±0.33	60.50±0.71	52.91±0.68	<u>41.96±1.13</u>	38.19±0.41	<u>54.59±0.41</u>
CAUM	<u>58.18±0.85</u>	<u>48.17±1.08</u>	<u>45.27±0.60</u>	<u>60.74±0.55</u>	<u>53.10±0.78</u>	41.83±0.93	38.16±1.42	54.50±0.32
DKN	58.73±0.39	48.20±0.39	45.48±0.43	60.78±0.29	53.69±0.66	42.14±0.66	38.57±0.83	54.74±0.55
TANR	50.62±0.14	47.28±0.44	44.22±0.51	60.06±0.33	50.17±0.10	41.82±0.12	<u>38.26±0.33</u>	54.49±0.17
SentiDebias	56.80±0.30	47.27±0.44	44.09±0.18	60.04±0.44	52.63±0.90	41.63±1.34	37.58±0.13	54.33±0.58

Table 6. Content personalization and aspect diversity (in terms of topical categories, sentiments, and political leaning) performance of different NRS. The best results per column are highlighted in bold, the second best are underlined.

Model	German				English			
	nDCG@3	D _{ctg} @3	D _{snt} @3	D _{pol} @3	nDCG@3	D _{ctg} @3	D _{snt} @3	D _{pol} @3
NRMS	44.90±0.65	18.55±0.25	33.75±0.54	30.68±0.32	37.49±1.07	<u>22.19±0.34</u>	<u>32.50±0.27</u>	26.31±0.94
NAML	44.27±1.20	19.31±0.21	33.63±0.18	30.41±0.37	38.14±0.84	23.89±0.65	32.08±0.64	25.66±1.24
MINS	44.96±0.33	18.18±0.44	33.79±0.34	30.10±0.44	38.19±0.41	21.94±1.11	32.34±0.16	25.22±1.11
CAUM	<u>45.27±0.60</u>	18.09±0.33	33.13±0.43	30.42±0.33	38.16±1.42	21.90±0.63	32.09±0.26	25.90±0.63
DKN	45.48±0.43	18.46±0.16	<u>33.79±0.44</u>	<u>30.92±0.59</u>	38.57±0.83	22.04±0.40	32.46±0.66	26.58±0.90
TANR	44.22±0.51	18.52±0.18	33.70±0.36	30.39±0.38	<u>38.26±0.33</u>	21.59±0.19	32.41±0.68	<u>26.47±0.82</u>
SentiDebias	44.09±0.18	<u>18.66±0.30</u>	33.84±0.34	30.99±0.30	37.58±0.13	22.10±1.11	33.26±0.34	25.81±1.11

connections between news, thus resulting in better news representations. The results are partly at odds with the findings of [27], which observed that while CAUM indeed outperforms other models with candidate-agnostic user encoders, DKN underperforms. However, [27] conducted experiments on the MIND dataset [69], which is magnitudes larger than NeMig. This points to the fact that both knowledge-awareness and candidate-awareness user modeling might be particularly beneficial in scenarios where little training data is available. Lastly, we observe that the quality of content personalization is higher for the German than for the English dataset. We hypothesize that this is due to the higher sparsity of the English dataset, in which significantly more news have never been seen by the users.

Aspect-based Diversity and Personalization. Next, we analyze the diversity and personalization of results, using topical categories (ctg), sentiment (snt), and political leaning (pol) as aspects. Concretely, we report the performance for content personalization (nDCG@ k), aspect-based diversity (D_{ctg} for topical categories, D_{snt} for sentiment polarization, and D_{pol} for political leaning) and aspect-based personalization (PS_{ctg} for topical categories, PS_{snt} for sentiment polarization, and PS_{pol} for political leaning). We report results for $k = 3$.

Table 6 summarizes the results on content personalization and aspect diversity. As expected, SentiDebias [68] achieves the highest sentiment diversity on both datasets. NAML [63] outperforms, by a relatively large margin, the other models in terms of category diversity of results. These results are somewhat surprising and at odds with those obtained on other datasets [28], as NAML leverages information on topical categories to customize results to the users’ preferences, and not to address diversity. Furthermore, compared to the first two aspects, we find that all recommenders perform nearly identical w.r.t. degree of political diversification of recommendations. This could be due to the fact that none of the models explicitly target political diversification.

Table 7. Content and aspect personalization (in terms of topical categories, sentiments, and political leaning) performance of different NRS. The best results per column are highlighted in bold, the second best are underlined.

Model	German				English			
	nDCG@3	PS _{ctg} @3	PS _{snt} @3	PS _{pol} @3	nDCG@3	PS _{ctg} @3	PS _{snt} @3	PS _{pol} @3
NRMS	44.90±0.65	20.72±0.45	42.66±0.30	36.55±0.09	37.49±1.07	18.48±0.18	41.38±0.40	40.66±0.22
NAML	44.27±1.20	20.42±0.31	42.59±0.25	36.23±0.32	38.14±0.84	18.24±0.81	41.43±0.19	41.17±0.30
MINS	44.96±0.33	21.10±0.60	42.56±0.34	36.74±0.27	38.19±0.41	<u>18.83±0.41</u>	41.20±0.16	<u>41.15±0.21</u>
CAUM	<u>45.27±0.60</u>	20.85±0.79	42.48±0.43	<u>36.59±0.36</u>	38.16±1.42	<u>18.67±1.45</u>	<u>41.41±0.26</u>	40.95±0.41
DKN	45.48±0.43	<u>20.95±0.35</u>	42.51±0.36	36.08±0.24	38.57±0.83	18.56±0.35	41.25±0.14	40.70±0.15
TANR	44.22±0.51	20.88±0.24	<u>42.64±0.19</u>	36.42±0.34	<u>38.26±0.33</u>	19.10±0.29	41.27±0.24	40.98±0.20
SentiDebias	44.09±0.18	20.50±0.39	42.64±0.34	36.40±0.24	37.58±0.13	18.57±0.33	41.29±0.34	40.84±0.20

Table 8. Content personalization and aspect-based diversity (in terms of topical categories, sentiment, and political leaning) performance of DKN with different variants of NeMigKG.

Model	German					English				
	AUC	nDCG@3	D _{ctg} @3	D _{snt} @3	D _{pol} @3	AUC	nDCG@3	D _{ctg} @3	D _{snt} @3	D _{pol} @3
content entities	58.72±0.45	45.57±0.44	18.44±0.16	33.67±0.38	31.17±0.56	53.67±0.63	38.60±1.01	21.99±0.45	32.33±0.34	26.71±0.75
+ topical categories	58.82±0.23	45.74±0.31	18.43±0.09	33.83±0.39	30.93±0.64	53.67±0.50	38.68±0.72	22.02±0.38	32.20±0.45	26.47±1.12
+ sentiment	58.81±0.26	45.69±0.38	18.48±0.08	33.77±0.44	30.98±0.65	53.70±0.53	38.53±0.73	21.95±0.37	32.17±0.18	26.89±0.72
+ political leaning	58.73±0.42	45.55±0.43	18.45±0.16	33.68±0.47	31.00±0.45	53.73±0.72	38.66±0.79	22.00±0.46	32.28±0.36	26.74±0.52
+ 1-hop neighbors	<u>58.82±0.33</u>	45.74±0.46	<u>18.50±0.09</u>	33.77±0.40	31.20±0.50	53.78±0.44	38.76±0.65	21.95±0.37	32.16±0.25	26.60±0.59
+ 2-hop neighbors	58.73±0.39	45.48±0.43	18.46±0.16	33.79±0.44	30.92±0.59	53.69±0.66	38.57±0.83	22.04±0.40	32.46±0.66	26.58±0.90

Aspect-based diversity is tightly correlated with aspect-based personalization [28]. More specifically, higher levels of diversity come at the cost of personalization, as can be seen in Table 7, which illustrates the results on content and aspect-based personalization. We observe that categorical and political personalization are much more aligned with content personalization performance, then with the sentiment of news, as discussed also in [28]. We find this to be intuitive, as users tend to choose which news to read based on their topics/categories of interest, as well as political preferences, and not on their sentiment polarization.

An earlier version of the German dataset [29] has been used in [1] and [43] to examine sentiment and stance recommender bias, and to investigate polarization and filter bubble creation through online studies, respectively. While in this study we have analyzed political biases of recommenders only from the perspective of users’ click histories, in the future their political information can be further explored in conjunction with the models’ predictions to understand whether any existing biases are reinforced or amplified by the algorithms. Similarly, NeMig can be used in future works to analyze these trends also in multilingual or cross-lingual recommendation scenarios, as it comprises information on the same topic in both German and English.

Knowledge Graph Ablations. The type of data contained in the knowledge graph is paramount for training high quality knowledge graph embeddings, which are further exploited by knowledge-aware recommenders to improve the accuracy of predictions. Thus, we analyze the impact of input features on NeMigKG. Specifically, we pre-train entity embeddings on different versions of NeMigKG using TransD [31], and study the recommendation and aspect-based diversity performance of DKN using these embeddings. We summarize the results in Table 8 in terms of content personalization (nDCG@ k), and aspect-based diversity for topical categories (D_{ctg}), sentiment (D_{snt}), and political leaning (D_{pol}). Adding topical categories, sentiment and political information in NeMigKG increases the diversity of recommendations w.r.t. these aspects, while having minor effects on accuracy-based performance. We find that

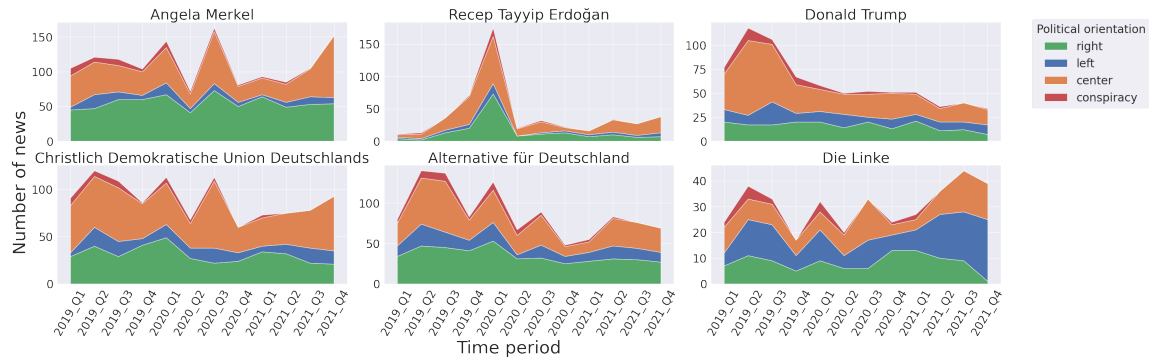


Fig. 6. Evolution over time and political orientation of a sample of the top-20 most popular entities from the German dataset. The first row contains entities of type PER, and the second of type ORG.

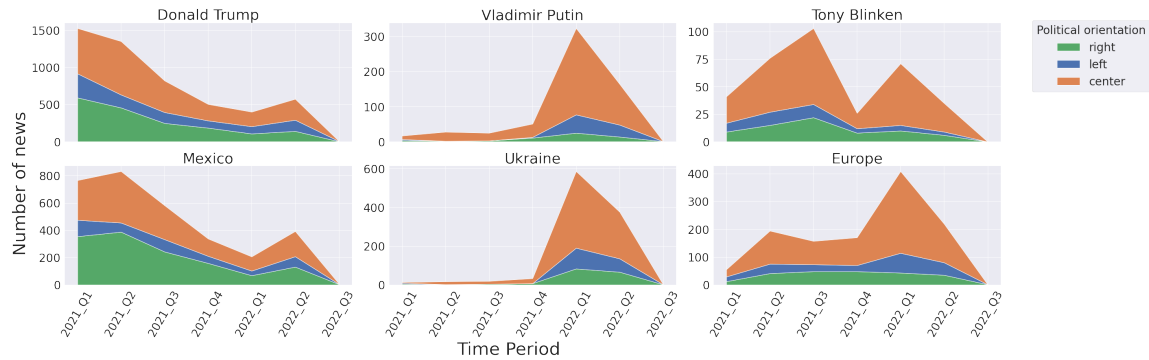


Fig. 7. Evolution over time and political orientation of a sample of the top-20 most popular entities from the English dataset. The first row contains entities of type PER, and the second of type LOC.

extending NeMigKG with k -hop neighbors from Wikidata is beneficial only up to $k = 1$ hops. This indicates that a certain degree of contextualization of named entities with general knowledge is helpful when the data size is small. However, injecting too much external knowledge (i.e., $k = 2$) can dilute the original information and have detrimental effects on the recommender’s downstream performance.

6.5 News Trends Analysis

Lastly, we demonstrate how NeMig can be used to analyze trends and correlations between entities used in the news discourse and events. Concretely, we analyze the evolution of named entities overtime and political groups. Figs. 6-7 show the evolution over time and political orientation of the source outlets of some of the top-20 most frequent entities in our corpora. We find that generally the most frequently sampled entities are covered predominantly in center media, which is also the source of the majority of news in both corpora. Notable exceptions in the German corpus (Fig. 6) are *Angela Merkel*, appearing equally or more often in right-wing media, and *Die Linke*, with a more equal distribution among all political classes, although left-wing coverage dominates the discourse. The evolution of entities over time reveals correlations with major events. In terms of organizations, German news report most frequently about the main political parties, whose evolution coincides with elections on the federal and European level, or about central political

figures in international politics (e.g., the coverage of *Recep Tayyip Erdoğan* appears to be potentially correlated with the Turkish offensive in north-eastern Syria). In contrast to German media, the representation of the top entities from the English corpus in the left US media is generally lower.

7 CONCLUSION

We introduced NeMig, a bilingual news collection and knowledge graphs on the topic of migration, along with user data, including demographic and political information. The news is collected from mainstream and alternative German and US media outlets spanning a wide political scale. We extract sub-topics and named entities linked to Wikidata from the news text and metadata. In contrast to existing datasets, we provide sentiment and political annotations, and use Wikidata for incorporating background knowledge. We hope that this resource will inspire future research on analyzing the multidimensional implications of news curation algorithms, in both monolingual and cross-lingual settings.

ACKNOWLEDGMENTS

This work has been conducted in the ReNewRS project, which is funded by the Baden-Württemberg Stiftung in the Responsible Artificial Intelligence program. The authors would also like to thank Julia Wildgans at the University of Mannheim for her legal advice on data handling.

REFERENCES

- [1] Mehwish Alam, Andreea Iana, Alexander Grote, Katharina Ludwig, Philipp Müller, and Heiko Paulheim. 2022. Towards analyzing the bias of news recommender systems using sentiment and stance detection. In *Companion Proceedings of the Web Conference 2022*. 448–457.
- [2] Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.
- [3] Eytan Bakshy, Solomon Messing, and Lada A Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132.
- [4] Jack M Balkin. 2017. Free speech in the algorithmic society: Big data, private governance, and new school speech regulation. *UCDL rev.* 51 (2017), 1149.
- [5] Francesco Barbieri, Luis Espinosa Anke, and Jose Camacho-Collados. 2022. Xlm-t: Multilingual language models in twitter for sentiment analysis and beyond. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. 258–266.
- [6] Larry M Bartels. 1993. Messages received: The political impact of media exposure. *American political science review* 87, 2 (1993), 267–285.
- [7] Vincenzo Bonnici. 2020. Kullback-Leibler divergence between quantum distributions, and its upper-bound. *arXiv preprint arXiv:2008.05932* (2020).
- [8] Andrei Boutyline and Robb Willer. 2017. The social structure of political echo chambers: Variation in ideological homophily in online networks. *Political psychology* 38, 3 (2017), 551–569.
- [9] Jaeho Cho, Saifuddin Ahmed, Martin Hilbert, Billy Liu, and Jonathan Luu. 2020. Do search algorithms endanger democracy? An experimental investigation of algorithm effects on political polarization. *Journal of Broadcasting & Electronic Media* 64, 2 (2020), 150–172.
- [10] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 1724.
- [11] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 8440–8451.
- [12] Nicola De Cao, Ledell Wu, Kashyap Popat, Mikel Artetxe, Naman Goyal, Mikhail Plekhanov, Luke Zettlemoyer, Nicola Cancedda, Sebastian Riedel, and Fabio Petroni. 2022. Multilingual autoregressive entity linking. *Transactions of the Association for Computational Linguistics* 10 (2022), 274–290.
- [13] Gabriel de Souza Pereira Moreira, Felipe Ferreira, and Adilson Marques da Cunha. 2018. News Session-Based Recommendations using Deep Neural Networks. In *Proceedings of the 3rd Workshop on Deep Learning for Recommender Systems, DLRS@RecSys 2018, Vancouver, BC, Canada, October 6, 2018*, Balázs Hidasi, Alexandros Karatzoglou, Oren Sar Shalom, Bracha Shapira, Domonkos Tikk, Flavian Vasile, and Sander Dieleman (Eds.). ACM, 15–23.
- [14] Gabriel de Souza Pereira Moreira, Dietmar Jannach, and Adilson Marques da Cunha. 2019. Contextual Hybrid Session-Based News Recommendation With Recurrent Neural Networks. *IEEE Access* 7 (2019), 169185–169203.
- [15] Ralf Dewenter, Melissa Linder, and Tobias Thomas. 2019. Can media drive the electorate? The impact of media coverage on voting intentions. *European Journal of Political Economy* 58 (2019), 245–261.

- [16] Michael Färber, Victoria Burkard, Adam Jatowt, and Sora Lim. 2020. A multidimensional dataset based on crowdsourcing for analyzing and detecting news bias. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 3007–3014.
- [17] Rama Rohit Reddy Gangula, Suma Reddy Duggenpudi, and Radhika Mamidi. 2019. Detecting political bias in news articles using headline attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*. 77–84.
- [18] Jennifer Golbeck, Matthew Mauriello, Brooke Auxier, Keval H Bhanushali, Christopher Bonk, Mohamed Amine Bouzaghrane, Cody Buntain, Riya Chanduka, Paul Chekalos, Jennine B Everett, et al. 2018. Fake news vs satire: A dataset and analysis. In *Proceedings of the 10th ACM Conference on Web Science*. 17–21.
- [19] Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794* (2022).
- [20] Jon Atle Gulla, Lemei Zhang, Peng Liu, Özlem Özgöbek, and Xiaomeng Su. 2017. The Adressa Dataset for News Recommendation (*WI '17*). Association for Computing Machinery, New York, NY, USA, 1042–1048.
- [21] Qingyu Guo, Fuzhen Zhuang, Chuan Qin, Hengshu Zhu, Xing Xie, Hui Xiong, and Qing He. 2020. A survey on knowledge graph-based recommender systems. *IEEE Transactions on Knowledge and Data Engineering* 34, 8 (2020), 3549–3568.
- [22] Nicolas Heist, Sven Hertling, Daniel Ringler, and Heiko Paulheim. 2020. Knowledge Graphs on the Web-An Overview.
- [23] Natali Helberger. 2019. On the democratic role of news recommenders. *Digital Journalism* 7, 8 (2019), 993–1012.
- [24] Benjamin Horne, Sara Khedr, and Sibel Adali. 2018. Sampling the news producers: A large news and feature data set for the study of the complex media landscape. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 12.
- [25] Andreea Iana, Mehwish Alam, Alexander Grote, Nevena Nikolajevic, Katharina Ludwig, Philipp Müller, Christof Weinhardt, and Heiko Paulheim. 2022. *NeMig - A Bilingual News Collection and Knowledge Graph about Migration*. <https://doi.org/10.5281/zenodo.7442424>
- [26] Andreea Iana, Mehwish Alam, and Heiko Paulheim. [n. d.]. A survey on knowledge-aware news recommender systems. *Semantic Web Preprint* ([n. d.]), 1–62.
- [27] Andreea Iana, Goran Glavas, and Heiko Paulheim. 2023. Simplifying Content-Based Neural News Recommendation: On User Modeling and Training Objectives. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2384–2388.
- [28] Andreea Iana, Goran Glavas, and Heiko Paulheim. 2023. Train Once, Use Flexibly: A Modular Framework for Multi-Aspect Neural News Recommendation. *arXiv:2307.16089* [cs.IR]
- [29] Andreea Iana, Alexander Grote, Kaharina Ludwig, Mehwish Alam, Philipp Müller, Christof Weinhardt, Harald Sack, and Heiko Paulheim. 2022. *GeNeG: German News Knowledge Graph*. <https://doi.org/10.5281/zenodo.5913171>
- [30] Jacob Jensen, Suresh Naidu, Ethan Kaplan, Laurence Wilse-Samson, David Gergen, Michael Zuckerman, and Arthur Spirling. 2012. Political polarization and the dynamics of political language: Evidence from 130 years of partisan speech [with comments and discussion]. *Brookings Papers on Economic Activity* (2012), 1–81.
- [31] Guoliang Ji, Shizhu He, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Knowledge graph embedding via dynamic mapping matrix. In *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)*. 687–696.
- [32] Benjamin Kille, Frank Hopfgartner, Torben Brodt, and Tobias Heintz. 2013. The Plista Dataset. In *Proceedings of the 2013 International News Recommender Systems Workshop and Challenge (Kowloon, Hong Kong) (NRS '13)*. Association for Computing Machinery, New York, NY, USA, 16–23.
- [33] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Doha, Qatar, 1746–1751.
- [34] Diederik P Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *ICLR* (2014).
- [35] Konstantina Lazaridou and Ralf Krestel. 2016. Identifying political bias in news articles. *Bulletin of the IEEE TCDD* 12 (2016).
- [36] Effi Levi, Guy Mor, Shaul Shenhav, and Tamir Sheaffer. 2020. Compres: A dataset for narrative structure in news. *arXiv preprint arXiv:2007.04874* (2020).
- [37] Lei Li, Ding-Ding Wang, Shun-Zhi Zhu, and Tao Li. 2011. Personalized news recommendation: a review and an experimental investigation. *Journal of computer science and technology* 26, 5 (2011), 754–766.
- [38] Sora Lim, Adam Jatowt, Michael Färber, and Masatoshi Yoshikawa. 2020. Annotating and analyzing biased sentences in news articles using crowdsourcing. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*. 1478–1484.
- [39] Danyang Liu, Ting Bai, Jianxun Lian, Xin Zhao, Guangzhong Sun, Ji-Rong Wen, and Xing Xie. 2019. News Graph: An Enhanced Knowledge Graph for News Recommendation. In *KaRS@ CIKM*. 1–7.
- [40] Ping Liu, Karthik Shivaram, Aron Culotta, Matthew A Shapiro, and Mustafa Bilgic. 2021. The interaction between political typology and filter bubbles in news recommendation algorithms. In *Proceedings of the Web Conference 2021*. 3791–3801.
- [41] Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics* 8 (2020), 726–742.
- [42] Zhan Liu, Shaban Shabani, Nicole Glassey Balet, and Maria Sokhn. 2019. Detection of satiric news on social media: Analysis of the phenomenon with a French dataset. In *2019 28th International Conference on Computer Communication and Networks (ICCCN)*. IEEE, 1–6.
- [43] Katharina Ludwig, Alexander Grote, Andreea Iana, Mehwish Alam, Heiko Paulheim, Harald Sack, Christof Weinhardt, and Philipp Müller. [n. d.]. Divided by the Algorithm? The (Limited) Effects of Content-and Sentiment-Based News Recommendation on Affective, Ideological, and Perceived Polarization. *Social Science Computer Review* ([n. d.]), 08944393221149290.

- [44] Katharina Ludwig and Philipp Müller. 2022. Does social media use promote political mass polarization? A structured literature review. *Questions of Communicative Change and Continuity. In Memory of Wolfram Peiser (2022)*, 118–166.
- [45] Jennifer McCoy, Tahmina Rahman, and Murat Somer. 2018. Polarization and the global crisis of democracy: Common patterns, dynamics, and pernicious consequences for democratic polities. *American Behavioral Scientist* 62, 1 (2018), 16–42.
- [46] Andreas L. Opdahl, Tareq Al-Moslimi, Duc-Tien Dang-Nguyen, Marc Gallofré Ocaña, Bjørnar Tessem, and Csaba Veres. 2022. Semantic Knowledge Graphs for the News: A Review. *ACM Comput. Surv.* (2022). <https://doi.org/10.1145/3543508>
- [47] Eli Pariser. 2011. *The filter bubble: What the Internet is hiding from you*. Penguin UK.
- [48] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [49] Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D Manning. 2020. Stanza: A Python natural language processing toolkit for many human languages. *arXiv preprint arXiv:2003.07082* (2020).
- [50] Tao Qi, Fangzhao Wu, Chuhan Wu, and Yongfeng Huang. 2022. News recommendation with candidate-aware user modeling. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1917–1921.
- [51] Nils Reimers and Iryna Gurevych. 2020. Making Monolingual Sentence Embeddings Multilingual using Knowledge Distillation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- [52] Marco Rospocher, Marieke van Erp, Piek Vossen, Antske Fokkens, Itziar Aldabe, German Rigau, Aitor Soroa, Thomas Ploeger, and Tessel Bogaard. 2016. Building Event-Centric Knowledge Graphs from News. *Journal of Web Semantics* (2016).
- [53] Erik F Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. *arXiv preprint cs/0306050* (2003).
- [54] Gautam Kishore Shahi and Durgesh Nandini. 2020. FakeCovid-A multilingual cross-domain fact check news dataset for COVID-19. *arXiv preprint arXiv:2006.11343* (2020).
- [55] Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2020. Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big data* 8, 3 (2020), 171–188.
- [56] Willem Robert Van Hage, Véronique Malaisé, Roxane Segers, Laura Hollink, and Guus Schreiber. 2011. Design and use of the Simple Event Model (SEM). *Journal of Web Semantics* 9, 2 (2011), 128–136.
- [57] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [58] Piek Vossen, Rodrigo Agerri, Itziar Aldabe, Agata Cybulska, Marieke van Erp, Antske Fokkens, Egoitz Laparra, Anne-Lyse Minard, Alessio Palmero Aprosio, German Rigau, Marco Rospocher, and Roxane Segers. 2016. NewsReader: Using knowledge resources in a cross-lingual reading machine to generate more knowledge from massive streams of news. *Special Issue Knowledge-Based Systems, Elsevier* (2016).
- [59] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Commun. ACM* 57, 10 (2014), 78–85.
- [60] Hongwei Wang, Fuzheng Zhang, Xing Xie, and Minyi Guo. 2018. DKN: Deep knowledge-aware network for news recommendation. In *Proceedings of the 2018 world wide web conference*. 1835–1844.
- [61] Rongyao Wang, Shoujin Wang, Wenpeng Lu, and Xueping Peng. 2022. News recommendation via multi-interest news sequence modelling. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 7942–7946.
- [62] William Yang Wang. 2017. “Liar, Liar Pants on Fire”: A New Benchmark Dataset for Fake News Detection. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 422–426.
- [63] Chuhan Wu, Fangzhao Wu, Mingxiao An, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019. Neural news recommendation with attentive multi-view learning. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 3863–3869.
- [64] Chuhan Wu, Fangzhao Wu, Mingxiao An, Yongfeng Huang, and Xing Xie. 2019. Neural news recommendation with topic-aware news representation. In *Proceedings of the 57th Annual meeting of the association for computational linguistics*. 1154–1159.
- [65] Chuhan Wu, Fangzhao Wu, Suyu Ge, Tao Qi, Yongfeng Huang, and Xing Xie. 2019. Neural news recommendation with multi-head self-attention. 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)*. 6389–6394.
- [66] Chuhan Wu, Fangzhao Wu, and Yongfeng Huang. 2022. Rethinking InfoNCE: How Many Negative Samples Do You Need?. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, Lud De Raedt (Ed.). International Joint Conferences on Artificial Intelligence Organization, 2509–2515.
- [67] Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2020. SentiRec: Sentiment diversity-aware neural news recommendation. In *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*. 44–53. <https://aclanthology.org/2020.aacl-main.6>
- [68] Chuhan Wu, Fangzhao Wu, Tao Qi, Wei-Qiang Zhang, Xing Xie, and Yongfeng Huang. 2022. Removing AI’s sentiment manipulation of personalized news delivery. *Humanities and Social Sciences Communications* 9, 1 (2022), 1–9.
- [69] Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, and Ming Zhou. 2020. MIND: A Large-scale Dataset for News Recommendation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault (Eds.). Association for Computational Linguistics, 3597–3606.