Exploration of Framing Biases in Polarized Online Content Consumption

Markus Reiter-Haas
supervised by Elisabeth Lex
Graz University of Technology, Institute of Interactive Systems and Data Science
Graz, Austria
reiter-haas@tugraz.at

ABSTRACT

The study of framing bias on the Web is crucial in our digital age, as the framing of information can influence human behavior and decision on critical issues such as health or politics. Traditional frame analysis requires a curated set of frames derived from manual content analysis by domain experts. In this work, we introduce a frame analysis approach based on pretrained Transformer models that let us capture frames in an exploratory manner beyond predefined frames. In our experiments on two public online news and social media datasets, we show that our approach lets us identify underexplored conceptualizations, such as that health-related content is framed in terms of beliefs for conspiracy media, while mainstream media is instead concerned with science. We anticipate our work to be a starting point for further research on exploratory computational framing analysis using pretrained Transformers.

CCS CONCEPTS

• **Information systems** → *Content analysis and feature selection*; **World Wide Web**; Language models.

KEYWORDS

computational frame extraction, content bias, exploratory content analysis, text processing, semantic representations

ACM Reference Format:
Markus Reiter-Haas. 2023. Exploration of Framing Biases in Polarized Online Content Consumption. In *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion), April 30–May 04, 2023, Austin, TX, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3543873.3587534

1 INTRODUCTION

The Web affects society at large, but also reflects the inherent biases of people [4]. Biases have been studied extensively regarding online behavior patterns, e.g., in terms of popularity bias [1, 22, 25] and confirmation bias [20, 21, 43]. Beyond behavioral patterns, biases can also stem from Web content itself. Herein, Draws et al. [16] show that the viewpoint of biased content influences user attitudes, while Rekabsaz et al. [40] highlight the impact of societal biases (e.g., gender) in retrieved content on the representation of particular groups in retrieval results. Similarly, the way content is *framed* can lead to biases and has also been shown to affect human behavior, public opinion and decision-making [45]. Framing corresponds to the selection and saliency of certain aspects in communicating texts [17]. Although research on framing has been thoroughly conducted for media [e.g., 11, 24], framing remains largely unexplored in Web content and its users' consumption patterns, in particular, when content is polarized or negative and thus receiving increased attention [30, 34]. Also, traditional frame analysis techniques frequently require a set of known frames for a topic, which need to be identified manually by domain experts [24].

In this work, we aim to explore framing biases in polarized Web content and their effects on content consumption behavior. We introduce three complementary approaches for exploratory framing analysis based on pretrained Transformers [46]. We subsequently categorize the extracted frames and conduct behavior analysis in openly available online news and social media corpora [28, 48]. We find that polarized health-related news is largely framed in terms of science vs. beliefs. Given that this frame is not part of established prior conceptualizations [e.g., 11], this finding underpins the merits of our approach. We believe that our research will considerably improve the understanding of framing bias and users' consumption patterns on the Web. Besides, we hope that our approach inspires novel debiasing methods to mitigate polarized Web-based retrieval.

2 PROBLEM

Framing of digital media relates to societal effects such as polarization. However, framing biases are difficult to detect and characterize. Moreover, acquiring labeled data on framing is challenging and labor-intensive. This issue becomes even more apparent in non-English settings, where the lack of data is more severe.

In our work, we investigate text representation, such as embeddings, to capture the semantic differences related to framing in text. As our approach is exploratory in nature, it applies to a setting with a low amount of labeled or entirely unlabeled data. Similarly, our approach is not restricted to a single language, as it is based on language models, and hence, can use a multilingual base encoder.

We aim to answer three research questions on framing analysis:

RQ1: How can we extract frames from polarized Web content without prior conceptualization?
RQ2: How can the extracted frames be categorized for specific contexts, e.g., health-related topics?
RQ3: What is the relationship between frames and viewpoint diversity in users' Web content consumption?

WWW '23 Companion, April 30–May 04, 2023, Austin, TX, USA

ACM ISBN 978-1-4503-9419-2/23/04.
https://doi.org/10.1145/3543873.3587534

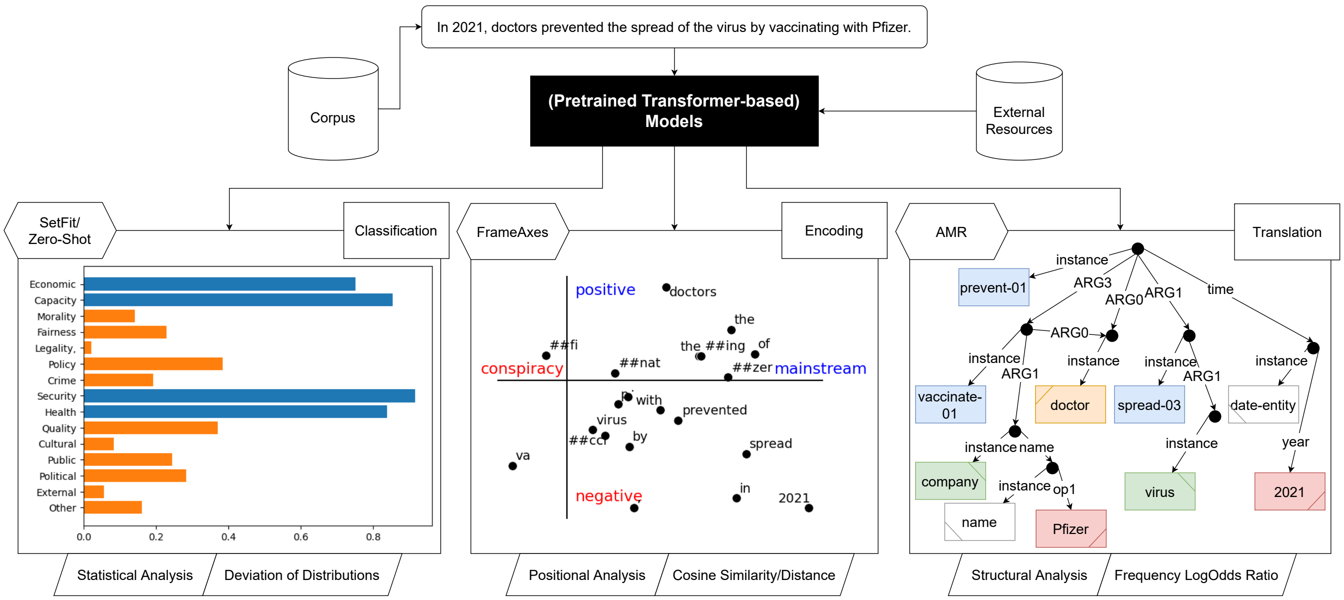

(a) Framing Labels via Label Probabilities

(b) Framing Dimensions via Embedding Space

(c) Framing Structure via Semantic Graph

Figure 1: Overview of the three complementary approaches. Subfigures show the result of a transformation with each approach. In (a), a predefined set of labels is predicted, here in a zero-shot setting, and the label probabilities are plotted (blue for predicted labels with high probability). In (b), the tokens of the sentence are projected onto framing axes in 2-dimensional embedding space, where the axis poles are opposing each other (e.g., positive vs negative). In (c), the text is transformed into a semantic (rooted, directed, and acyclic) graph.

3 STATE OF THE ART

Framing is a fractured paradigm in literature, but essentially deals with the *selection* and *salience* of some aspects of a communicating text [17]. For example, measures against the COVID-19 pandemic can be framed in terms of the *prevention of the spread* or *fight against the virus*, thereby highlighting distinct features of the problem and suggesting opposing solutions. As framing promotes the alteration of the perceived reality and its interpretation [17], it also affects human judgments and choices [45]. Consequently, social scientists have researched the framing of important topics, such as the responses and social movements towards the COVID-19 pandemic [19, 27, 33]. Traditionally, framing analysis involves careful manual analysis of data. More recently, computational methods [e.g., 42, 47] have been suggested to automatically determine the framing of textual content. For example, the *SemEval 2023 Task 3 (Detecting the Category, the Framing, and the Persuasion Techniques in Online News in a Multi-lingual Setup)* [32] [1]. aims to predict the framing of text based on a predefined taxonomy.

Computational frame analysis studies comprise various types of frames, such as war [47], terrorists [14], morality [29], or blame [42]. These studies focus on different conceptualizations of framing and are not necessarily comparable. Moreover, their methodological approaches differ drastically and depend on the preselected frames for the study. Hence, other conceptualizations of framing that would be more characteristic might not be detected.

Ali and Hassan [2] provide a comprehensive survey of computational framing extraction methods. The main approaches include various kinds of topic modeling and cluster analysis as unsupervised approaches. In comparison, neural networks, also including pretrained Transformer-based language models, are mainly used in a supervised manner. Other methods include parsing semantic relations, frequency-based models, and semantic axes, i.e., FrameAxis [23]. The current state-of-the-art predominantly investigates a predefined set of frames. We strive to alleviate this limitation by utilizing an exploratory approach based on semantic information embedded within the textual content. We adapt several of the previously mentioned methods, i.e., neural networks with pretrained Transformer-based language models, semantic relations, frequency-based models, and semantic axes. Due to the exploratory manner, our approach enables novel and unexpected new conceptualizations predefined selection of frames or labels. This is unlike the existing *OpenFraming* tool [7], which, although exploratory in nature, still requires a preselection of frames and labeling of data.

Finally, framing theory [12] relates to various other concepts, such as public opinion and values. Especially in media frames [13], narratives are another important aspect to consider. Hence, other computational endeavors like computational narrative understanding [31] benefit from improved frame extraction methods. For brevity, we omit a detailed discussion here and refer to Reiter-Haas et al. [37], where we thoroughly discuss the relationship between narratives and framing.

[1]Challenge Website: https://propaganda.math.unipd.it/semeval2023task3/

4 PROPOSED APPROACH

Our approach is based on three complementary sub-approaches, i.e., (a) predicted label probabilities, (b) embedding space of tokens, and (c) semantic graph of content information. All three approaches leverage pretrained language models based on the Transformer architecture [46], as they generalize well to problems with limited available data [10]. An overview of the proposed approach is provided in Figure 1, where we apply each sub-approach to the same specified example[2]. In the following, we detail each sub-approach separately, before detailing how they complement each other.

Framing Labels (a). Assigning and predicting labels (or their associated probabilities) with a supervised classifier is the traditional approach in computational framing analysis, besides unsupervised topic modeling. However, topic models cannot be applied to single examples and thus are not suitable for frame extraction, but only discovery. In Figure 1a, we predicted the characteristic framing labels as defined by Boydstun et al. [9]. As such labeled data is only sparsely available, a few-shot (e.g., with the recently released SetFit [44]) or zero-shot model (e.g., with BART [26] as used for the example) is required. When aggregating over multiple articles, a statistical analysis can be performed, where the deviation of the label distributions is analyzed.

Framing Dimensions (b). When encoding the textual data into an n-dimensional hyperspace (e.g., with BERT [15]), the resulting embedding space comprises latent dimensions that describe tokens, words, paragraphs, and complete articles. An unsupervised way to measure the semantics is by defining axes [3], which can be applied to framing analysis with FrameAxis [23]. Therein, words are projected onto predefined axes characterized by two opposing poles (e.g., positive/negative and mainstream/conspiracy as shown in Figure 1b). When applied to a collection of documents, we can perform a positional analysis, where we consider the distance between points (e.g., with cosine similarity). Individual points such as words, as well as documents, can easily be aggregated using pooling operations, such as mean pooling.

Framing Structure (c). Text can also be converted to other representations, such as graphs. Besides syntax trees, graphs can also represent the semantics of textual data. We use well-established abstract meaning representations [AMR; 6] by translating text via a BART model [26] to a serialized graph containing the semantic information. This graph-based representation can be used for structural analysis, such as which actor relates to which action. For instance, in Figure 1c, it is trivial to observe that the *doctor*, although being mentioned only once, relates to both the *prevent* and *vaccinate* actions, as shown by the reentrants in the graph. For brevity, we refer to the AMR specification for a detailed description [5]. The semantic structure is also closely related to narration, and thus also the analysis of framing regarding dominant narratives in a corpus. For comparing corpora, we can apply frequency-based approaches, such as the log odds ratio [8], to the graph structure (e.g., individual elements or even sub-graphs).

Theoretically, all three approaches are universally applicable for determining arbitrary conceptualizations of framing. A classifier could learn to detect sophisticated narratives by sentence structure, while mining the graph representation could reveal broad labels. Similarly, the embedding space can generalize to labels and structure within the data points. Nevertheless, the three approaches can be applied concurrently. Framing analysis could be performed on predefined labels, while also considering unsupervised similarities and structural information. Moreover, the three approaches could even be combined into a single framework for frame detection. Ideally, a text can be represented as a semantic graph where each instance is additionally represented by an embedding and labels. For instance, the *doctor* instance can be assigned a *health* label and have an embedding similar (i.e., close to) to the *vaccinate* instance. Hence, all three approaches have particular complementary strengths (e.g., as summarized in Table 1).

[2]Code for plots available at: https://github.com/Iseratho/web23-phd-symposium

5 METHODOLOGY

To validate our proposed approach, we conduct framing analysis in publicly available online news and social media corpora. For evaluation, we perform a mixed methods-based approach, i.e., quantitative and qualitative. For the quantitative evaluation, we use the limited amount of available labeled data (e.g., from [11]) and consider the coherence of detected frames (i.e., similar to topic coherence [41]). For the qualitative evaluation, we jointly analyze and interpret our results with social scientists. These interpretations are then used to inform possible framing conceptualizations.

In the categorization of the framing concepts, we consider aspects of various granularity. Similar to existing work, we first aim to create broad labels that describe frames, such as a text being politically framed. Moreover, we plan to also consider frame hierarchies (i.e., sub-labels), directionality (i.e., frame bias), and magnitude (i.e., frame intensity). To that end, we consider established theories like the moral foundation theory [18]. For instance, morally framed texts can be strongly framed towards the sub-label harm (i.e., the negative direction of a care/harm axis) while also being mildly framed towards fairness (i.e., the positive direction of a fairness/cheating axis). Finally, we consider how concepts relate to a given text from a structural perspective. As an example, the polarizing topic of vaccination can be assigned opposing sentiments depending on the framing and typically goes along with different actors from a narrative sense (e.g., doctors vs government). Due to the shift from a predictive to an exploratory approach, we expect to find novel conceptualizations (e.g., a belief-oriented framing) while also retaining or expanding upon characteristic labels (e.g., the political orientation) but discarding less pronounced framing concepts (e.g., whether a text reflects a public opinion).

Using the novel categorization, content consumption patterns can then be investigated and novel insights extracted. For instance, we expect a mostly low viewpoint diversity regarding framing, even across different topics. Furthermore, we hypothesize a repeated consumption of content with almost identical framing concepts. This would indicate that repeat consumption patterns hold, as is the case for other domains (e.g., in music consumption [39]), and largely explain pre-existing framing biases.

Regarding data analysis, we deem Web data as a relevant data source to study. Web content is abundantly available and believed to be highly polarized. While individual pieces of text are typically

Criteria	Classifier (a)	Embeddings (b)	Graph (c)
Unsupervised	×[†]	✓	✓
Exploratory	×	~	✓
Narratives	×	×	✓
Challenge	✓	×	×
Dimensions	scalar	n-D	irregular
Data Type	int/float	float	int
Aggregation	trivial	intuitive	challenging

Table 1: Summary of the comparison between the three sub-approaches. The complexity increases from left to right, but similarly increases in exploratory potential.

lacking manual framing annotations, online sources are often associated with certain features, such as political preferences and credibility. Hence, our approach specifically focuses on Web content, such as news websites and social media.

We aim to use recent and large datasets containing textual and log data, such as LOCO [28] for testing the approaches and frame categorization, and MIND [48] for content consumption analysis. LOCO provides online articles regarding a range of topics (including health-related ones), while also containing labels on whether they belong to conspiracy or mainstream media. Hence, we expect a noteworthy difference in framing between the two types of sources. MIND, on the other hand, provides click and impression logs (i.e., consumption data) in addition to content data. Thus, both datasets are suitable for their respective tasks.

6 RESULTS

In our initial paper on framing [38, similar to Figure 1b], we investigate the morality framing of political tweets in the US and Austria with word vectors and FrameAxis [23]. In the study, we find that the framing is coherent with previous findings on US politicians regarding their party's dominant morality. However, followers of Austrian politicians frame their tweets regarding COVID-19 similarly to topic-specific political messages in the public, rather than the usual party-associated dimensions. Opposite to expectations, the left-leaning Social Democratic Party emphasizes authority, while the ruling conservative party focuses on care. For context, at the time of data collection, the conservative party aimed to slow the spread of the virus with public messaging regarding mutual care. Conversely, the leader of the Social Democratic Party, being an epidemiologist herself, repeatedly insisted on listening to doctors and scientists. Hence, just focusing on a predefined conceptualization might lose valuable information regarding the framing of messages that might even lead to counterintuitive pictures.

In our next and most substantial contribution so far [37, i.e., Figure 1c], we demonstrate that semantic graphs are a perfect fit to represent the framing of the narrative information embedded in the textual content. Specifically, we use abstract meaning representations [6] to extract health-related narratives from the LOCO dataset [28] containing articles from both mainstream and conspiracy media. Although the approach requires no predefined conceptualization of framing, we extract differences that support our intuition. Most notably, we find that conspiracy media revolved around belief narratives, whereas mainstream media focused on science instead. Such conceptualization of framing goes beyond the typical analysis of framing classes and dimensions. Moreover, we investigated more specific narratives concerning the interplay of actors and actions. For instance, we find that the *prevent* action is concerned with *the government preventing individuals* for conspiracy media, rather than *the vaccination preventing the virus*, as is the case in mainstream media. This shows that the approach is very suitable for an exploratory framing analysis.

In our recently concluded experiments [35, refer to Figure 1a for an example], we provide our contribution to the SemEval challenge 2023 Task 3 using a SetFit-inspired approach for few-shot predictions [44]. The challenge provides predefined frames on which the performance is measured, but only provides a low amount of labels. Unlike the more exploratory approaches, we deem a classical label prediction approach more applicable for a competitive scenario with predefined labels, i.e., conceptualizations. Hence, our approach adopts contrastive multi-label loss functions for fine-tuning a multilingual base encoder. We achieved the first position on the zero-shot Spanish framing detection subtask.

We summarize our findings regarding the three approaches in Table 1. All three approaches have their merits, but the graph-based approach is best in terms of exploratory potential. Also, all three approaches can be employed in an unsupervised manner, but the classifier can only do so in a zero-shot setting that harms its predictive performance. The graph-based approach is the only one that naively allows the extraction of narratives rather than simpler conceptions. This can be attributed to the irregularity of graphs in comparison to the simpler hyper-dimensional structure of embedding spaces and scalar-valued label probabilities. Regarding the data types, the classifier can be used both for discrete label and continuous label probability predictions. In the embedding space, continuous values are the norm to specify the positions, whereas graph elements and sub-graphs are frequency-based. This also affects the complexity of the aggregation in a corpus, where labels (or their probabilities) are trivial to combine. Embedding spaces, while more complex, are still intuitive to aggregate (e.g., mean embedding). However, the aggregation of sub-graphs or their elements is challenging. Altogether, we highlight that the three approaches are complementary in nature, and considering the information from all three approaches is beneficial.

7 CONCLUSION

In this work, we introduced our efforts towards an exploratory approach for framing analysis, which is a multi-faceted problem. We demonstrated in previous works that the semantic information extracted from pretrained Transformers provides richer representations for comparison between different corpora. There, we also showed that such approaches tie in neatly with the current state-of-the-art, and hence, allow for a more comprehensive analysis.

As currently ongoing research, we aim to consolidate the previously distinct directions of research into a holistic approach and openly available framework for framing analysis (i.e., to conclude *RQ1*). Furthermore, we started working on the categorization of health-related frames using the knowledge of the prior research (i.e.,

[†] A classifier can be used in an unsupervised fashion with zero-shot learning.

RQ2). Afterward, we aim to investigate user behavior from a framing bias perspective that should reveal whether articles containing the same frames are repeatedly consumed (i.e., *RQ3*).

Hence, we pave the way for future work on the long-term dynamics of framing (e.g., to investigate frame adoption), as well as relating framing to other concepts, such as polarization [e.g., 36] and mis-/disinformation. Finally, novel methods should enable debiasing content at data, algorithmic, and presentation-level.

ACKNOWLEDGMENTS

I want to thank my supervisor Elisabeth Lex for her continuous support and shared excitement in my research endeavors.

REFERENCES

[1] Himan Abdollahpouri. 2019. Popularity bias in ranking and recommendation. In *2019 AAAI/ACM AIES'19*. 529–530.
[2] Mohammad Ali and Naeemul Hassan. [n. d.]. A Survey of Computational Framing Analysis Approaches. In *EMNLP'22*. Association for Computational Linguistics, 9335—-9348.
[3] Jisun An, Haewoon Kwak, and Yong-Yeol Ahn. 2018. Semaxis: A lightweight framework to characterize domain-specific word semantics beyond sentiment. *ACL'18* (2018).
[4] Ricardo Baeza-Yates. 2018. Bias on the web. *Commun. ACM* 61, 6 (2018), 54–61.
[5] Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2012. Abstract meaning representation (amr) 1.0 specification. In *ACL EMNLP'12*. 1533–1544.
[6] Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. In *LAW VII & ID*. 178–186.
[7] Vibhu Bhatia, Vidya Prasad Akavoor, Sejin Paik, Lei Guo, Mona Jalal, Alyssa Smith, David Assefa Tofu, Edward Edberg Halim, Yimeng Sun, Margrit Betke, Prakash Ishwar, and Derry Tanti Wijaya. 2021. OpenFraming: Open-sourced Tool for Computational Framing Analysis of Multilingual Data. In *ACL EMNLP'21*. 242–250.
[8] J Martin Bland and Douglas G Altman. 2000. The odds ratio. *Bmj* 320, 7247 (2000), 1468.
[9] Amber E Boydstun, Dallas Card, Justin Gross, Paul Resnick, and Noah A Smith. 2014. Tracking the development of media frames within and across policy issues. (2014).
[10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *NeurIPS'20* 33 (2020), 1877–1901.
[11] Dallas Card, Amber Boydstun, Justin H Gross, Philip Resnik, and Noah A Smith. 2015. The media frames corpus: Annotations of frames across issues. In *ACL'20 & IJNLP'20*. 438–444.
[12] Dennis Chong and James N Druckman. 2007. Framing theory. *Annual review of political science* 10, 1 (2007), 103–126.
[13] Paul D'Angelo. 2017. Framing: media frames. *The international encyclopedia of media effects* (2017), 1–10.
[14] Dorottya Demszky, Nikhil Garg, Rob Voigt, James Zou, Matthew Gentzkow, Jesse Shapiro, and Dan Jurafsky. 2019. Analyzing polarization in social media: Method and application to tweets on 21 mass shootings. *NAACL-HLT'19* (2019).
[15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *NAACL-HLT'19* (2019).
[16] Tim Draws, Nava Tintarev, Ujwal Gadiraju, Alessandro Bozzon, and Benjamin Timmermans. 2021. This Is Not What We Ordered: Exploring Why Biased Search Result Rankings Affect User Attitudes on Debated Topics. In *ACM SIGIR'21*. 295–305.
[17] Robert M Entman. 1993. Framing: Towards clarification of a fractured paradigm. *McQuail's reader in mass communication theory* 390 (1993), 397.
[18] Jesse Graham, Jonathan Haidt, Matt Motyl, Peter Meindl, Carol Iskiwitch, and Marlon Mooijman. 2018. Moral foundations theory. *Atlas of moral psychology* 211 (2018).
[19] Kevin Knight, Bianca Badarau, Laura Baranescu, Claire Bonial, Madalina Bardocz, Kira Griffitt, Ulf Hermjakob, Daniel Marcu, Martha Palmer, Tim O'Gorman, et al. 2021. Abstract meaning representation (amr) annotation release 3.0.
[20] Simone Kopeinik, Elisabeth Lex, Dominik Kowald, Dietrich Albert, and Paul Seitlinger. 2019. A Real-Life School Study of Confirmation Bias and Polarisation in Information Behaviour. In *EC-TEL'19*. Springer, 409–422.
[21] Dominik Kowald and Elisabeth Lex. 2018. Studying confirmation bias in hashtag usage on Twitter. *arXiv:1809.03203* (2018).
[22] Dominik Kowald, Markus Schedl, and Elisabeth Lex. 2020. The unfairness of popularity bias in music recommendation: A reproducibility study. In *ECIR'20*. Springer, 35–42.
[23] Haewoon Kwak, Jisun An, Elise Jing, and Yong-Yeol Ahn. 2021. FrameAxis: characterizing microframe bias and intensity with word embedding. *PeerJ Computer Science* 7 (2021), e644.
[24] Sha Lai, Yanru Jiang, Lei Guo, Margrit Betke, Prakash Ishwar, and Derry Tanti Wijaya. 2022. An unsupervised approach to discover media frames. In *Proc. of the LREC 2022 workshop on Natural Language Processing for Political Sciences*. 22–31.
[25] Oleg Lesota, Alessandro Melchiorre, Navid Rekabsaz, Stefan Brandl, Dominik Kowald, Elisabeth Lex, and Markus Schedl. 2021. Analyzing item popularity bias of music recommender systems: are different genders equally affected?. In *ACM RecSys'21*. 601–606.
[26] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *ACL'20* (2020).
[27] Richard McNeil-Willson et al. 2020. *Framing in times of crisis: Responses to COVID-19 amongst Far Right movements and organisations*. International Centre for Counter-Terrorism.
[28] Alessandro Miani, Thomas Hills, and Adrian Bangerter. 2021. LOCO: The 88-million-word language of conspiracy corpus. *Behavior research methods* (2021), 1–24.
[29] Negar Mokhberian, Andrés Abeliuk, Patrick Cummings, and Kristina Lerman. 2020. Moral framing and ideological bias of news. In *SocInfo'20*. 206–219.
[30] Maryam Mousavi, Hasan Davulcu, Mohsen Ahmadi, Robert Axelrod, Richard Davis, and Scott Atran. 2022. Effective Messaging on Social Media: What Makes Online Content Go Viral?. In *TheWebConf'22*. 2957–2966.
[31] Andrew Piper, Richard Jean So, and David Bamman. 2021. Narrative Theory for Computational Narrative Understanding. In *EMNLP'21*. 298–311.
[32] Jakub Piskorski, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023. SemEval-2023 Task 3: Detecting the Category, the Framing, and the Persuasion Techniques in Online News in a Multi-lingual Setup. In *SemEval'23*.
[33] Geoffrey Pleyers. 2020. The Pandemic is a battlefield. Social movements in the COVID-19 lockdown. *Journal of Civil Society* 16, 4 (2020), 295–312.
[34] Ashwin Rao, Fred Morstatter, and Kristina Lerman. 2022. Retweets Distort Exposure to Polarized Information. *arXiv:2211.16480* (2022).
[35] Markus Reiter-Haas, Alexander Ertl, Kevin Innerebner, and Elisabeth Lex. 2023. mCPT at SemEval-2023 Task 3: Multilingual Label-Aware Contrastive Pre-Training of Transformers for Few- and Zero-shot Framing Detection. *arXiv:2303.09901* (2023).
[36] Markus Reiter-Haas, Beate Klösch, Markus Hadler, and Elisabeth Lex. 2022. Polarization of Opinions on COVID-19 Measures: Integrating Twitter and Survey Data. *Social Science Computer Review* (2022), 08944393221087662.
[37] Markus Reiter-Haas, Beate Klösch, Markus Hadler, and Elisabeth Lex. 2022. AMR-based Framing Analysis of Health-Related Narratives: Conspiracy versus Mainstream Media. *Manuscript submitted for publication* (2022).
[38] Markus Reiter-Haas, Simone Kopeinik, and Elisabeth Lex. 2021. Studying Moral-based Differences in the Framing of Political Tweets. In *ICWSM'21*. 1085–1089.
[39] Markus Reiter-Haas, Emilia Parada-Cabaleiro, Markus Schedl, Elham Motamedi, Marko Tkalcic, and Elisabeth Lex. 2021. Predicting music relistening behavior using the ACT-R framework. In *ACM RecSys'21*. 702–707.
[40] Navid Rekabsaz, Simone Kopeinik, and Markus Schedl. 2021. Societal biases in retrieved contents: Measurement framework and adversarial mitigation of bert rankers. In *ACM SIGIR'21*. 306–316.
[41] Michael Röder, Andreas Both, and Alexander Hinneburg. 2015. Exploring the space of topic coherence measures. In *ACM WSDM'15*. 399–408.
[42] Chereen Shurafa, Kareem Darwish, and Wajdi Zaghouani. 2020. Political framing: US COVID19 blame game. In *SocInfo'20*. Springer, 333–351.
[43] Masaki Suzuki and Yusuke Yamamoto. 2020. Analysis of relationship between confirmation bias and web search behavior. In *iiWAS'20*. 184–191.
[44] Lewis Tunstall, Nils Reimers, Unso Eun Seo Jo, Luke Bates, Daniel Korat, Moshe Wasserblat, and Oren Pereg. 2022. Efficient Few-Shot Learning Without Prompts. *arXiv:2209.11055* (2022).
[45] Amos Tversky and Daniel Kahneman. 1985. The framing of decisions and the psychology of choice. In *Behavioral decision making*. Springer, 25–41.
[46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *NeurIPS'17* 30 (2017).
[47] Philipp Wicke and Marianna M Bolognesi. 2020. Framing COVID-19: How we conceptualize and discuss the pandemic on Twitter. *PloS one* 15, 9 (2020).
[48] Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, et al. 2020. Mind: A large-scale dataset for news recommendation. In *ACL'20*. 3597–3606.