KI4DEMOKRATIE: An AI-Based Platform for Monitoring and Fostering Democratic Discourse

Rudy Alexandro Garrido Veliz¹, Till Nikolaus Schaland¹, Simon Bergmoser², Florian Horwege¹, Somya Bansal³, Ritesh Nahar³, Martin Semmann¹, Jörg Forthmann⁴, Seid Muhie Yimam¹

¹Universtät Hamburg, ²NORDAKADEMIE, ³Zirkel Technologies GmbH, ⁴FAKTENKONTOR,

Correspondence: rudy.garrido.veliz@uni-hamburg.de & seid.muhie.yimam@uni-hamburg.de

Abstract

Social media increasingly fuel extremism, especially right-wing extremism, and enable the rapid spread of antidemocratic narratives. Although AI and data science are often leveraged to manipulate political opinion, there is a critical need for tools that support effective monitoring without infringing on freedom of expression. We present KI4DEMOKRATIE, an AI-based platform that assists journalists, researchers, and policymakers in monitoring right-wing discourse that may undermine democratic values. KI4DEMOKRATIE applies machine learning models to a large-scale German online data gathered on a daily basis, providing a comprehensive view of trends in the German digital sphere. Early analysis reveals both the complexity of tracking organized extremist behavior and the promise of our integrated approach, especially during key events.

1 Introduction

Democracy is currently under threat globally and especially across several European nations due to a significant rise in extremist ideologies, particularly from right-wing factions. These groups have used the principle of freedom of speech as a means to spread polarized and often extremist narratives in social networks and beyond, leading to increased societal divisions (Aktas, 2024; Veera Singam and Pantucci, 2024). The presence of right-wing political actors and their influence in digital spaces such as social media is undeniable, amplifying their ideologies (Serrano et al., 2019). This phenomenon is further complicated by the strategic use of artificial intelligence and data science to manipulate public opinion and undermine democratic processes (Heawood, 2018; Khanal et al., 2025; Riedl, 2024).

In response to these challenges, we propose the **KI4DEMOKRATIE** project, a pioneering initiative aimed at leveraging artificial intelligence to counteract the influence of antidemocratic forces. This

project focuses on developing a comprehensive dashboard that helps journalists, politicians, policy makers, and researchers analyze data from social media platforms and online news outlets. The ultimate goal is to preserve democracy and to create safer social media and news portals, fostering the coexistence of diverse viewpoints and mitigating community fragmentation.

The KI4DEMOKRATIE dashboard is composed of three core components: 1) Sentiment and Hate Speech Analysis Module: This component employs advanced AI models to automatically analyze sentiment and hate speech from various data sources. It provides information on the distribution of emotional and hostile content, highlighting prevailing trends in societal polarization (Solovev and Pröllochs, 2022; Yimam et al., 2024). 2) Network Analysis Module: This feature elucidates potential connections between political entities through sophisticated network graphs, providing insight into structural relationships within political ecosystems and the dissemination pathways of ideologies (Leifeld, 2016; Chua, 2024). 3) Topic Extraction and Visualization Module: This component focuses on identifying and visualizing key topics from the collected data, enabling users to derive meaningful insights and enhance their understanding of evolving political narratives (Xu and Xiong, 2020).

The research framework underpinning this project is guided by several key questions: **RQ1**: What are the predominant themes and narratives disseminated by extremist groups on social media? **RQ2**: How effective are AI-driven methodologies in identifying and countering extremist narratives and misinformation? **RQ3**: In what ways can **KI4DEMOKRATIE** contribute to the development of safer online spaces and the promotion of integrated communities? Our contributions are multifold: First, we introduce a novel analytical tool, the **KI4DEMOKRATIE** dashboard, which integrates sentiment analysis, network mapping, and



Figure 1: Complete pipeline of the **KI4DEMOKRATIE** project.

topic modeling to provide a holistic view of political discourse. Second, we demonstrate the application of artificial intelligence in monitoring and countering extremism, thereby offering a blueprint for future initiatives in media oversight and democratic safeguarding. Third, we provide empirical insights into online extremist rhetoric dynamics, contributing to academic knowledge and practical policy-making efforts. These contributions have several implications: They offer a scalable model for monitoring political discourse across various digital platforms, assist stakeholders in developing informed strategies to counteract polarization, and ultimately, serve as a foundation for fostering more constructive and cohesive public dialogue.

2 Related Work

The rapid digital spread of antidemocratic ideologies and the organized mobilization of right-wing groups present an escalating challenge for democratic institutions globally. Foundational studies have systematically delineated the mechanisms of digital populism (Bartlett et al., 2011; Mudde, 2019), while research demonstrates how alternative platforms and digital media ecosystems shape and aid radical right agendas (Conway et al., 2019). Advances in automated hate speech and sentiment detection have improved coverage for diverse languages, including significant contributions to model development and resource creation (Schmidt and Wiegand, 2017; Vidgen et al., 2019; Wiegand et al., 2018a), with benchmark datasets such as GermEval supporting evaluation in the German context (Wiegand et al., 2018b).

Extremist activity is increasingly fragmented across platforms, requiring robust and multifaceted network analysis. Methodologies for tracing actor relationships and information flows are detailed in prior work (Morstatter et al., 2013; Zannettou et al., 2017; Ribeiro et al., 2020), yet truly integrative tools that capture influence across multiple platforms are still rare. Similarly, topic modeling, from LDA (Blei et al., 2003) to BERTopic (Grootendorst, 2022), as discussed in applied communication contexts (Jacobi et al., 2016), enables the extraction of latent themes, but dynamic narrative tracking and combined approaches with network or sentiment analytics remain exceptions. Efforts to counter disinformation, such as AI-supported fact-checking (Nakov et al., 2021) and analyses of computational propaganda (Howard and Kollanyi, 2016), have often been limited by a lack of integrated, practitioner-friendly interfaces.

Addressing these gaps, our **KI4DEMOKRATIE** initiative pioneers an integrated solution that combines sentiment analysis, hate speech detection, dynamic topic modeling, and advanced network analysis in an interactive dashboard. This holistic, cross-methodological approach delivers urgently needed scalable tools for researchers, journalists, and policymakers working to protect and strengthen democratic discourse in the digital age.

3 Approaches

The data is obtained daily from varied sources, based on a distilled set of 500 keywords from human experts. The obtained posts go through text classification for sentiment, hate speech, claim identification, and fact checking. In addition to this, a topic modeling analysis is performed and a network graph is created. These steps are detailed in the following sections.

3.1 Analysis Dashboard

The analytical dashboard offers an interactive visualization of sentiment and hate speech trends derived from online content. Developed using Dash and Plotly, the dashboard would integrate the outputs of two types of text classification, sentiment analysis, and hate speech, using small language models (SLMs). For sentiment analysis, the system uses the TimeLM model (Loureiro et al., 2022), which computes compound sentiment scores for each entry, classified as positive, negative, or neutral. For the hate speech classification, we employ the SLM model, LFTW R4 Target (Vidgen et al., 2021), fine-tuned on hate speech data. Each text input is tokenized and passed through the model, which outputs logits. These logits are converted into binary labels using argmax, identifying whether a text contains hate or normal speech. The sentiment and hate speech models used are stateof-the-art and best-performing in English, and we apply translation to German to utilize them. The dashboard empowers users to analyze extremist language and online narratives by offering dynamic visualizations of sentiment and hate speech trends over time and across platforms, enabling targeted exploration of discourse patterns, shifts, and spikes related to political events or media campaigns.

3.2 Topic Modeling

In order to uncover predominant themes and narratives by online news and right-wing groups on social media, we employ topic modeling to identify underlying semantic structures. The topic extraction and visualization component utilizes BERTopic (Grootendorst, 2022), a robust topic modeling technique that relies on transformerbased embeddings to uncover latent themes in large corpora of text and produces interpretable topics by applying class-based TF-IDF. For the generation of document representations, the pre-trained sentence-transformer, Sentence-BERT (Reimers and Gurevych, 2019), is applied. The subsequent dimensionality reduction step in BERTopic was performed using UMAP (McInnes et al., 2018), followed by HDBSCAN (McInnes et al., 2017) as the clustering algorithm. The selection of parameters aimed to mitigate the identification of spurious, unrelated clusters while avoiding the aggregation of distinct themes into overly broad topics and to prevent the fragmentation of coherent topics into excessively fine-grained clusters. Specifically, we apply BERTopic to perform dynamic topic modeling on a daily basis to model the evolving political narratives and themes. By tracking the frequency and content of identified topics daily, journalists and politicians can pinpoint the predominant themes at any given moment in time, observe their trend, and link these dynamics to external resources and real-world events, see examples in Appendix B.

3.3 Network Graph

Narratives and their influence online extend beyond a single social network. Finding influential actors on multiple social media networks, news articles, blogs, and the Internet in general is not trivial. Some articles find that the heavy reliance of research on social network graphs is based on access to the Twitter/X API (Morstatter et al., 2013). The proposed network extraction method and centrality calculation approach aims to address this issue. The available interaction data provides the author of a post or news snippet and the URL at which this interaction occurs. Both entities can be extracted and transformed into the origins of an edge in a social network graph. The author is extracted directly, and the URL is extracted as a user profile, in which case it is matched to the entity of the author or the organization hosting the interaction. The extraction method then transforms and extracts data from the content of the interaction itself to select what kind of edge in the graph should be attributed to this interaction and what the target node should be.

One way of extracting the target node is by checking for tagged users, for example @mentions, using regular expressions. Another method is by using named entity recognition (NER). The method of extraction determines which of the two labels the edge between the source and the target node has. The entities labeled are the targets of edges labeled intentional, while entities extracted using NER are the targets of the edges labeled inferred. All edges are weighted by the frequency of their occurrence. Some entities, while not very active themselves, are mentioned with a high frequency and are therefore important for the understanding of influential actors in the network. To include this aspect of influence, the method also extracts a third type of edge, passive mutual, which is undirected. It connects nodes that are mentioned within the same content of an interaction by a 3rd party (Räbiger and Spiliopoulou, 2015)

The network graph also includes hashtags and the modeled topics as a separate node label. Hashtags are the targets of the edge labeled *intentional*, while topics extracted using the topic modeling approach are targets of the edge labeled *inferred*. Eigenvector centrality (Newman, 2018) would be used to scale the nodes in a view of the graph to give end-users orientation to identify influential nodes. Nodes should also be filtered based on their frequency of occurrence in the data.

3.4 Fact Checking

In order to assess the factuality of a text, a multistage framework has been proposed (Vlachos and Riedel, 2014; Guo et al., 2022). The stages are (i) claim detection, (ii) evidence retrieval, and (iii) claim verification. We developed a prototype adopting this framework to evaluate its feasibility in connection with a results dashboard. For the stage of claim detection, in which the aim is to identify a check-worthy claim within a text and to extract this claim, one-shot prompting with GPT-3.5 (Brown et al., 2020) was utilized. The model was supplied with a system prompt (see Appendix A), which contained the example, and a user prompt, which contained the post as well as the author, the author party, and the date as context information. The extracted claim was used for the next stage of evidence retrieval. Evidence retrieval aims to find sources in connection to the stated claim from Google News search. For the generation of the search query, one shot prompting with GPT-3.5 was used again. As in the previous setup, the system prompt contained an example, while the user prompt contained the previous reply and context information. The first 3 results of the Google News search were used as knowledge sources and summarized by GPT-3.5. The concatenated summaries were then passed to the next stage of verdict prediction. For the verdict prediction, GPT-3.5 was supplied with the claim, context information and context information. In a zero-shot prompting setup, it was tasked to give a truthfulness score out of five categories (False, Mostly false, Half true, Mostly true, or True) and give a reason of at most 2 sentences, and truthfulness scores were displayed on the dashboard. For each channel, the number of claims in each of the five categories was displayed in a bar chart, giving an overview of how often each author published claims that were supported/not supported by evidence. This planned approach is experimental; among the other considered ones is the one by Özge Sevgili et al. (2024) stands out, for their streamlined approach.

4 Evaluation

To evaluate our approach, 100 random posts across various dates were extracted and given to four different German native speakers. The annotators were asked to categorize the posts regarding sentiment (positive, negative, or neutral) and whether they contain hate speech. Where the four annotators were evenly split, a fifth native speaker acted as a tie-breaker, and the gold label for each post was decided by majority vote. The Kappa score was around 0.43 for both tasks. This shows how challenging even for native speakers to coincide in these matters; and therefore, interpretation of experts, for example, journalists or researchers, is needed to effectively assess the behavior of the right-wing movement. Applying the LFTW R4 Target model for hate speech, the TimeLM model for sentiment analysis, and GPT-40-mini to the gold label dataset shows moderate performance, with GPT-40-mini achieving the best results overall (see Table 1).

Task	Model	Precision	Recall	F1
Hate speech	LTFTW	0.42	0.85	0.56
	GPT	0.53	0.63	0.58
Sentiment	TimeLM	0.77	0.72	0.74
	GPT	0.86	0.82	0.81

Table 1: Hate speech and sentiment analysis using SLM (LFTW R4 Target and TimeLM respectively) and LLM (GPT 4o-mini) models.

To illustrate the capabilities of the Dynamic Topic Modeling approach described in Section 3.2, we present two example visualizations derived from the dashboard. Specifically, we focused on the period around the Aschaffenburg¹ incident in January 2025. Figure 2 and Figure 3 in Appendix B describe the temporal evolution of selected topics during this timeframe, clearly demonstrating the incident's impact on the social media discourse and the increase in hate speech-related themes. Furthermore, the visualization also demonstrates the subsequent rise in topics related to migration policy, correlating with a proposal introduced by a conservative political party. These examples underline the capability of the models to pinpoint predominant themes and narratives at specific points in time and link their fluctuations to real-world events and news.

5 Conclusion

We have carefully developed a strategy to effectively leverage AI against right-wing antidemocratic movements. The components of the dashboard could support journalists, researchers, and policymakers in identifying and responding to problematic narratives and themes in the digital sphere. In a preliminary qualitative analysis, we have found a cohesive movement of the right-wing discourse when troubling situations arise, for example, the **Aschaffenburg** incident that led to an increase in hate speech. Our work, as it is still in progress, is trying to discover the best ways of using AI to support the resilience of democratic societies.

¹https://www.dw.com/en/

germany-Aschaffenburg-knife-attack-details/ a-71372024

6 Limitations

There are multiple limitations that can be faced when creating such a complex tool. The sheer amount of data and the capabilities for analysis are one of the biggest ones, while also having to mind the resource component.

Another limitation is how hard it is to get direct feedback from the possible users in order to improve it. These users are often in high demand, so it is a challenge to progress at a good pace while satisfying the possible requirements they have.

The accessibility of language models trained in the German language is also a limitation, considering that the ones found are often underdeveloped. The research community must strive for it.

However, the advance of multilingual language models offers fine-tuning opportunities for specific tasks. A possible improvement for the factchecking framework would be fine-tuned models for the specific subtasks. A current work in progress is the fine-tuning of a large language model to identify a claim in a text and extract it in a normalized form, as suggested by Sundriyal et al. (2023).

Ethical Considerations

The main consideration taken is that our work does not assume the political ideology, or all the social expectations that may come with it, of individuals. Rather, it only deals with the texts obtained to get a grasp of the right-wing discourse. We focus exclusively on analyzing publicly shared speech, not individuals, and emphasize transparency and academic rigor in countering anti-democratic actions.

It is also understood how important it is, for a set of tools like the one we intend to create, to keep access to it only for the desired individuals. A curated ingress to the dashboard will be created. Access to this tool is strictly controlled to prevent misuse for anti-democratic purposes.

References

- Murat Aktas. 2024. The rise of populist radical right parties in Europe. *International Sociology*, 39(6):591– 605.
- Jamie Bartlett, Jonathan Birdwell, and Mark Littler. 2011. *The New Face of Digital Populism*. Demos, London.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3:993–1022.

- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. Language models are few-shot learners. In *Proceedings of the* 34th International Conference on Neural Information Processing Systems, NIPS '20, Red Hook, NY, USA. Curran Associates Inc.
- Yi Ting Chua. 2024. "We Want You!" Applying Social Network Analysis to Online Extremist Communities. *Terrorism and Political Violence*, 37(3):409–423.
- Maura Conway, Ryan Scrivens, and Logan Macnair. 2019. Right-Wing Extremists' Persistent Online Presence: History and Contemporary Trends. International Centre for Counter-Terrorism.
- Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794*.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. A Survey on Automated Fact-Checking. *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Jonathan Heawood. 2018. Pseudo-public political speech: Democratic implications of the Cambridge Analytica scandal. *Information Polity*, 23:429–434.
- Philip N. Howard and Bence Kollanyi. 2016. Bots, #StrongerIn, and #Brexit: Computational Propaganda during the UK-EU Referendum. *arXiv preprint arXiv:1606.06356*.
- Carina Jacobi, Wouter van Atteveldt, and Kasper Welbers. 2016. Quantitative analysis of large amounts of journalistic texts using topic modelling. *Digital Journalism*, 4(1):89–106.
- Shaleen Khanal, Hongzhou Zhang, and Araz Taeihagh. 2025. Why and how is the power of Big Tech increasing in the policy process? The case of generative AI. *Policy and Society*, 44:52–69.
- Philip Leifeld. 2016. Discourse network analysis: Policy debates as dynamic networks. pages 301–326.
- Daniel Loureiro, Francesco Barbieri, Leonardo Neves, Luis Espinosa Anke, and Jose Camacho-collados. 2022. TimeLMs: Diachronic Language Models from Twitter. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 251–260, Dublin, Ireland.
- Leland McInnes, John Healy, and Steve Astels. 2017. hdbscan: Hierarchical density based clustering. *The Journal of Open Source Software*, 2(11):205.

- Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. 2018. UMAP: Uniform Manifold Approximation and Projection. *The Journal of Open Source Software*, 3(29):861.
- Fred Morstatter, Jürgen Pfeffer, Huan Liu, and Kathleen M. Carley. 2013. Is the Sample Good Enough? Comparing Data from Twitter's Streaming API with Twitter's Firehose. In *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media (ICWSM 2013)*, pages 400–408, Cambridge, Massachusetts, USA.

Cas Mudde. 2019. The Far Right Today. Polity.

- Preslav Nakov, David Corney, Maram Hasanain, Firoj Alam, Tamer Elsayed, Alberto Barrón-Cedeño, Paolo Papotti, Shaden Shaar, and Giovanni Da San Martino. 2021. Automated Fact-Checking for Assisting Human Fact-Checkers. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4551–4558. Survey Track.
- Mark Newman. 2018. *Networks: An Introduction*, reprinted edition. Oxford University Press, Oxford.
- Stefan Räbiger and Myra Spiliopoulou. 2015. A framework for validating the merit of properties that predict the influence of a twitter user. *Expert Systems with Applications*, 42(5):2824–2834.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. 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), pages 3982–3992, Hong Kong, China.
- Manoel Horta Ribeiro, Raphael Ottoni, Robert West, Virgílio A. F. Almeida, and Wagner Meira. 2020. Auditing radicalization pathways on YouTube. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, page 131–141, New York, NY, USA.
- Martin Riedl. 2024. Political deepfakes and misleading chatbots: understanding the use of genai in recent european elections. *Center for Media Engagement* (*The University of Texas at Austin*).
- Anna Schmidt and Michael Wiegand. 2017. A Survey on Hate Speech Detection using Natural Language Processing. In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, pages 1–10, Valencia, Spain.
- Juan Carlos Medina Serrano, Morteza Shahrezaye, Orestis Papakyriakopoulos, and Simon Hegelich. 2019. The rise of germany's afd: A social media analysis. In *Proceedings of the 10th International Conference on Social Media and Society*, SMSociety '19, page 214–223, New York, NY, USA. Association for Computing Machinery.

- Kirill Solovev and Nicolas Pröllochs. 2022. Hate Speech in the Political Discourse on Social Media: Disparities Across Parties, Gender, and Ethnicity. In *Proceedings of the ACM Web Conference 2022*, WWW '22, page 3656–3661, Virtual Event, Lyon, France.
- Megha Sundriyal, Tanmoy Chakraborty, and Preslav Nakov. 2023. From Chaos to Clarity: Claim Normalization to Empower Fact-Checking. In *Findings* of the Association for Computational Linguistics: EMNLP 2023, pages 6594–6609, Singapore.
- Kalicharan Veera Singam and Raffaello Pantucci. 2024. Extreme Right-Wing in the West. *Counter Terrorist Trends and Analyses*, 16(1):pp. 106–111.
- Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019. Challenges and frontiers in abusive content detection. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 80–93, Florence, Italy.
- Bertie Vidgen, Tristan Thrush, Zeerak Waseem, and Douwe Kiela. 2021. Learning from the Worst: Dynamically Generated Datasets to Improve Online Hate Detection. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1667–1682, Online.
- Andreas Vlachos and Sebastian Riedel. 2014. Fact Checking: Task definition and dataset construction. In Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science, pages 18–22, Baltimore, MD, USA.
- Michael Wiegand, Josef Ruppenhofer, Anna Schmidt, and Christine Greenberg. 2018a. Inducing a Lexicon of Abusive Words — A Feature-Based Approach. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers), pages 1046–1056, New Orleans, LA, USA.
- Michael Wiegand, Melanie Siegel, and Josef Ruppenhofer. 2018b. Overview of the GermEval 2018 Shared Task on the Identification of Offensive Language. In *Proceedings of GermEval 2018, 14th Conference on Natural Language Processing (KONVENS 2018)*, pages 1 – 10, Vienna, Austria.
- Sifan Xu and Ying Xiong. 2020. Setting socially mediated engagement parameters: A topic modeling and text analytic approach to examining polarized discourses on Gillette's campaign. *Public Relations Review*, 46(5):1–12.
- Seid Muhie Yimam, Daryna Dementieva, Tim Fischer, Daniil Moskovskiy Skoltech, Naquee Rizwan, Iit Kharagpur, Punyajoy Saha, Sarthak Roy, Martin Semmann, Alexander Panchenko, Chris Biemann, and Animesh Mukherjee. 2024. Demarked: A Strategy for Enhanced Abusive Speech Moderation through

Counterspeech, Detoxification, and Message Management. *arXiv preprint arXiv:2406.19543v1*.

- Savvas Zannettou, Tristan Caulfield, Emiliano De Cristofaro, Nicolas Kourtelris, Ilias Leontiadis, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. 2017. The Web Centipede: Understanding How Web Communities Influence Each Other Through the Lens of Mainstream and Alternative News Sources. In *Proceedings of the 2017 Internet Measurement Conference (IMC)*, pages 405–417, London, United Kingdom.
- Özge Sevgili, Irina Nikishina, Seid Muhie Yimam, Martin Semmann, and Chris Biemann. 2024. UHH at AVeriTeC: RAG for Fact-Checking with Real-World Claims. pages 55–63, Stroudsburg, PA, USA.

A Fact Checking

The fact-checking prototype used three *one-shot* prompting configurations and one *zero-shot* setup. All prompting was performed with OpenAI's GPT-3.5 model (gpt-3.5-turbo).

Claim Extraction

To extract fact-checkable claims from text, the system was prompted as follows:

System prompt:

You are a fact expert. Your task is to extract the statements from a given text for which the truthfulness can be verified by publicly available knowledge.

Ignore information about personal events, subjective opinions, and statements about upcoming events.

Reformulate the statements into direct statements. Include the date, author, and author's party in your formulation if this is necessary for fact-checking the statement.

Example: Text: "You can see the bad mood among all other parties."; Author: "Christian Lindner from the FDP party"; Date: 1.2.2023; Direct statement: "In February 2023, there is a bad mood among all parties except the FDP."

Return only the direct statements as a JSON object. Name these "statements".

User prompt template:

Extract the statements that can be factually verified with public knowledge from the following text by {author} from the party {author_party} from {date}: {query}.

Query Generation

To rephrase extracted claims into optimized Google search queries, the following prompts were used:

System prompt:

You are a Google search expert. Your task is to create an optimized Google search query from a statement. Example: Input: "The capital of Germany is Berlin." Optimized Google search query: "Capital of Germany"

User prompt template:

Convert the following statement by {author} from {date} into a Google search query: {query}.

Verdict Prediction (Zero-Shot)

The zero-shot setup for predicting truthfulness (verdicts) used the following configuration:

System prompt:

You are a fact-checking expert. Your job is to check the truth of statements that are given to you. You are provided with contextual knowledge for this purpose. Only give your answer in JSON format with the fields "Truthfulness" and "Reason".

Select exactly one of the following categories for the "Truthfulness" field: "False", "Mostly false", "Half true", "Mostly true", or "True". The "Reason" field must not be longer than two sentences.

User prompt template:

Check the following statement by {author} from the party {author_party} from {date}: {claim}. Context: {grounding_context}

B Dynamic Topic Modeling

B.1 Focus Topics



Figure 2: Temporal evolution of a wide range of political and extremist discourse topics on social media from January 20 to February to January 31, 2025, covering the period around the **Aschaffenburg** incident and subsequent proposal by a conservative political party.



Figure 3: Evolution of key political discourse topics on social media from two days before the incident in **Aschaffenburg** (January 22, 2025) until January 31, 2025, coinciding with the joint proposal by a conservative political party. The plot highlights trends in discourse related to migration, border control, asylum policy, and right-wing narratives such as "Remigration".