Artificial Intelligence and Social Computing, Vol. 163, 2025, 32–43
https://doi.org/10.54941/ahfe1006033

Convo-Based Attitude Analysis of Twitter Big Data: A Case Study on Ukraine-Russia War Dataset

Ning Sa, Ankita Bhaumik, and Tomek Strzalkowski

Rensselaer Polytechnic Institute, Troy, NY 12180, USA

ABSTRACT

Social media has become a popular platform for studying public perceptions and opinions on important global events like elections, pandemics and international conflicts. Previous studies utilized text mining algorithms to analyse individual messages for references to relevant topics and associated sentiment. Such methods overlook the broader context in which these messages appear and as a result fail to capture often intricate relationships between topics, messages, and their authors. More specifically, these methods do not account for social dynamics among the participants in an online discourse, which typically occurs within a convo (Katsios et al., 2019), a loosely structured cluster of posters interested in a common topic. In this paper, we present a convo-based analysis of a public social media dataset collected over a period of 3 months following the onset of the Ukraine-Russia conflict. In this dataset, we identify the most populous convos, the most influential participants within each, and the topics they discuss. We then demonstrate how the general attitude across these convos shifts over time from a largely pro-Ukraine to an increasingly pro-Russia stance, which we speculate is a result of ongoing influence operations. Our findings provide novel insights into the structure of social media traffic and evolution of attitudes in online populations. This work is a first step towards a more comprehensive framework for social media analysis.

Keywords: Big data, Social media, Attitude detection, Attitude shift, Ukraine-Russia

INTRODUCTION

Social media platforms such as Twitter (now X), and Telegram, are valuable resources for social science research thanks to their large user bases and unedited user generated content (Carr and Hayes, 2015). They have been widely used to study online public opinion and communication patterns during major political/social events and natural disasters. For example, Vargo et al. (2014) analysed how different groups on Twitter processed information during the 2012 U.S. election. Nakazato et al. (2023) focused on the attitudes of healthcare experts on Twitter about the COVID vaccine in 2020. Techniques such as topic modelling and sentiment analysis have been used to rapidly interpret vast amounts of data on social media and public forums. These types of studies aim for broad-brush insights into how people think, feel, and engage during traumatic events and help to uncover general patterns in public discourse and sentiment.

At the same time, social media are readily available platforms for spreading misinformation and disinformation, as well as for running targeted influence campaigns (Sharma et al., 2022). Harmful information has been a significant factor in shaping public opinion, particularly around difficult and controversial topics. Some recent works focused on detection of indicators of deliberate influence campaigns, driven by dedicated actors, often for malicious purposes. For example, Bhaumik et al. (2023; 2024) and Katsios et al. (2024) explored several indicators such as emotions, agendas, and sentiment in social media messages to detect potential campaigns and to understand how these campaigns may affect public opinion. Our work here does not directly detect influence campaigns or measure their effectiveness; instead, we report measurable attitude shifts in the online population that could be caused by ongoing influence campaigns.

Specifically, we focus on topic attitude on Twitter, which is dynamically derived from the attitudes of selected groups of users towards a topic. A user's attitude is derived from the prevailing stance or perspective (e.g., for or against) on a topic: a political issue, a policy, or an event (Hoewe and Peacock, 2020; Mutz et al., 1996). Topic attitude is the aggregated attitude of these selected groups of users. We postulate, that so construed, topic attitude is an important indicator of public opinion about the most salient topics discussed in social media.

Among all the Twitter users engaged in discussing a topic, one crucial subset is the most retweeted users (top influencers) who drive a significant portion of message traffic and whose messages are reinforced by multiple retweets that increase their reach across the platform. Another important group are the users with the highest volume of original tweets (most active users), who may or may not be influential, but they are responsible for generating and propagating most of the social media content, often framing the narratives around the key topics.

We build our analysis framework using convos to identify the most salient topics in a corpus. According to Katsios et al. (2019), a convo refers to an online social phenomenon where people are engaged around a topic or an activity. Examples of convos include a subreddit on Reddit and a repository on GitHub. On Twitter convos can be detected by tracking the hashtags. We follow the methodology introduced by Bhaumik et al. (2024), to identify hashtag convos using groups of frequently co-occurring Twitter hashtags. Bhaumik et al. (2024) focuses on individual convos and extracts each convo's agenda and emotions as indicators of an influence campaign. In this paper, we take a different approach and examine the information conveyed across all convos in the corpus by exploring the convo-user relationship and identifying the overall attitude shift of the convos. We utilize this framework to perform a detailed case study of the Ukraine-Russia conflict. We discover a pronounced shift in public attitude from largely pro-Ukraine at the beginning of the test period to increasingly pro-Russia at the end.

Our contributions can be summarized around three research questions: given a Twitter corpus around a major social or political event, 1) how to identify and visualize the most popular topics? 2) how can the attitude towards each topic be studied and visualized? 3) how are these attitudes changing or evolving over time?

We investigate the above RQs using a public Twitter dataset on the Ukraine-Russia War and limit the attitude to 3 categories: pro-Ukraine (pro-U), pro-Russia (pro-R), and None (no stance or irrelevant). In the next section, we will firstly review selected works in the field.

RELATED WORK

There are numerous studies that use social media datasets to understand public opinion around local and global events, such as the COVID-19 pandemic (Nakazato et al., 2023) and the U.S. presidential election (Vargo et al., 2014). To analyze the content of social media data, topic modelling algorithms and clustering algorithms have been used to extract and group key themes from large-scale social media datasets (Chang et al., 2023; Nakazato et al., 2023; Karami et al., 2018). In addition to what topics people talk about, indicators of how they feel and what they may be expected to do are also detected. Such indicators include, but not limited to, sentiment (Yin et al., 2022), emotions (Bhaumik et al., 2023), stance (Mather et al., 2022), and agenda (Katsios et al., 2024).

Our work focuses on topic analysis and prevailing attitudes during the Ukraine-Russia War starting from 2022. This conflict is already drawing significant attention in academic research. Eliguzel (2023) and Chang et al. (2023) combine topic modelling and different clustering methods on Twitter datasets. Maathuis and Kerkhof (2023) track topics and sentiment change by using two months of telegram messages. A major difference between the above studies and our work is that instead of using topic modelling, we identify convos, which in twitter data are hashtag topics with focused groups of messages. Additionally, all previous studies treat the identified topics disjointedly, while we link the topics through shared user attendance in the discussions. Similar to Maathuis and Kerkhof (2023), where the change of sentiments is explored, we detect the change of attitude, which is a multi-dimensional concept (Ajzen, 2001). In this paper, political attitude is simplified to the event specific categories of [pro-Ukraine, pro-Russia, None] and consequently is equivalent to stance.

Biber and Finegan's (1988) work on user stance focuses on the attitude and feelings expressed in texts written by them. Since then, various computational methods have been developed for stance detection across various social media platforms (Al-Ghadir et al., 2021; Glandt et al., 2021). Recently, large language models (LLMs) have shown improvements in capturing the user stance without the need of large annotated datasets or task-specific classifiers (Zhang et al., 2022). Other recent works explore several prompting and fine-tuning strategies to enhance the stance detection performance of LLMs (Cruickshank and Ng, 2024). We rely on these works to design effective prompts for accurate and consistent attitude detection using LLMs.

METHODOLOGY

Identification of Convos

The first step in our analysis is to identify groups of messages. Due to the drawbacks of keyword based clustering or topic modelling algorithms, we use a convo based approach to extract these message groups. To identify hashtag convos in the Twitter dataset, we build a co-occurrence based distance matrix using the top 6000 hashtags in the corpus. Next, we perform dimensionality reduction using UMAP (McInnes et al., 2018), followed by clustering using HDBSCAN (Campello et al., 2013) to obtain a list of frequently co-occurring hashtag groups. Furthermore, we observe that some clusters that contain general hashtags, like #ukraine or #russia, are disproportionately large. For a more granular analysis of these broad topics, we perform an additional clustering step on the clusters containing more than 50,000 tweets or 150 hashtags. This approach has been successfully used by Bhaumik et al. (2024) to analyze agendas and emotions around a set of topics. In this paper, we focus on the overall picture that is revealed by the hashtag convos in a Ukraine-Russia War Twitter dataset.

Convo Attitude Detection

Once the hashtag convo clusters are identified, we analyze the general attitudes of these convos from two key perspectives: (1) the attitudes of top X influencers (the most retweeted authors), and (2) the attitudes of the X most active users (those with the highest volume of tweets). By comparing the influencers and the active users, we get an idea about potential contrasts in their perspectives. In this study, we use $X = 10$.

We frame the task of user attitude detection as a classification problem where an instruction tuned large language model (LLM) is prompted with a set of at most 25 tweets by a user. The task is to label the overall attitude as one of the 3 classes: pro-Ukraine, pro-Russia, or None. The 'None' label is used when the author has an ambiguous stance, or her tweets are mainly about topics unrelated to the war. We experiment with multiple prompt templates and choose the one that performs the best on a sample set of 120 randomly selected tweets:

You are a stance detection assistant. Your task is to read a set of tweets by an author and output the overall stance of the author as PRO-RUSSIA, PRO-UKRAINE, or NONE.

If the messages by the author strongly support Russia, output PRO-RUSSIA.

If the messages by the author strongly support Ukraine, output PRO-UKRAINE.

If the messages by the author do not support any of them or are not related to the Russia-Ukraine conflict, output NONE.

Output PRO-RUSSIA, PRO-UKRAINE, or NONE for this set of tweets. Output the label only. Tweets:

<input tweets>

With the attitude of the top influencers/active users in a convo, its overall attitude is expressed as {*U*: *m*, *R*: *n*, N: *k*}, where *U*, *R*, N stand for pro-Ukraine, pro-Russia, and None respectively. The values m, *n*, *k* correspond to the number of top influencers/active users of each attitude in that convo.

DATASETS AND IMPLEMENTATION

We perform our analysis on a publicly available Ukraine-Russia War dataset (BwandoWando, 2024). The dataset is collected using a list of general hashtags such as #RussiaUkraine, #Ukraine, and some trending hashtags, like #donbass and #donetsk. In this study, we focus on tweets from May 01 to July 31, 2022.

Table 1: Dataset description before & after pre-processing.

	Before	**After**
Number of original tweets	5,485,469	4,718,213
Number of authors	750,078	592,655

For our initial pre-processing, we remove all retweets, non-English tweets, tweets without hashtags or having less than 3 textual tokens. The final dataset contains 4.7m tweets (Table 1). For the attitude shift analysis, we create snapshots of the data from the first and last 2 weeks of the 3-month period. After preprocessing, there are 1m messages in the first two weeks and 407k messages in the last two weeks.

We use Llama-3.1-8B-Instruct[1] as the primary LLM for attitude detection, but this step could be replicated by other LLMs like Llama-2[2] or ChatGPT[3]. We select this model due to its open-source nature and relatively smaller size, ensuring efficiency and performance comparable to larger models (Dubey et al., 2024).

RESULTS AND DISCUSSIONS

Hashtag Convos

Using our hashtag clustering method, we obtain 184 convos from the input dataset. They encompass a wide range of topics from war-related hashtags like #StandWithUkraine to economic concerns such as #Inflation.

To visualize an overall picture of the convos and their associated authors, we construct a convo-user attendance matrix where each row represents a convo and each column represents an author. We only include authors who have at least 5 tweets and participate in more than one convo. The matrix entries indicate the participation of authors in different convos, providing a clear representation of the interaction patterns across topics. If author *Uj* attends convo *C*i, position (i, *j*) is set to 1, otherwise 0. The final convo-user matrix is composed of 184 convos and over 20k authors. Using this input

[1] https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct

[2] https://huggingface.co/meta-llama/Llama-2-13b-chat

[3] https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/

matrix, we perform dimensionality reduction using UMAP and visualize them on a two-dimensional space. The resulting plot (Fig. 1) illustrates the relationships between the different convos through their participants. The distance between the convos on the plot indicates the user sharing of the convos. The closer the convos are to each other, the more users they share between them.

Through analysing the hashtags of the convos, we observe clear thematic clustering in Fig. 1, indicating different focus of different user groups. Convos in the top region focus on the war's progression, with hashtags like #mariupol and #odesa, referencing key Ukrainian cities. Slightly downward, the discussions shift to neighbouring countries such as #belarus and #poland, which are supporters of the two sides of the war. Further down, geopolitical topics like #finland and #iran, emerge. On the left side, we see convos with strong pro-Ukraine attitude, such as #armukrainenow. On the right side, users discuss the war's broader implications with hashtags like #oil and #inflation. Then come the convos covering North American political topics relevant to the war, like #biden and #canada, which are separated from the rest of the countries/regions. Near the bottom, there are convos focusing on fundraising efforts, i.e. #nft, which is used for supporting Ukraine financially (Tolmach et al., 2023). Also, there are convos like #osint (Open-Source Intelligence) which helps in information gathering and analysis (Kotišová and Velden, 2023) and #oprussia which are cyber-attacks toward Russia. Finally, the upper right corner shows convos on entertainment topics, like #metgala and #tennis, whose second top hashtag #wimbledon banned Russian and Belarusian players in 2022[4]. Many of the topics are also detected in previous studies (Eligüzel, 2023; Maathuis and Kerkhof, 2023).

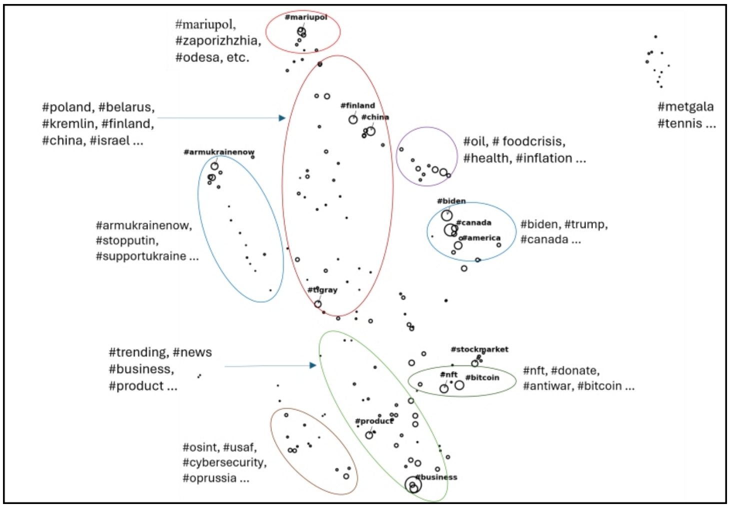

Figure 1: 2-D visualization of convo-author attendance. The size of a point depicts the number of original tweets in the convo.

[4]https://www.washingtonpost.com/sports/2022/06/26/wimbledon-russia-belarus-ban-player-reaction

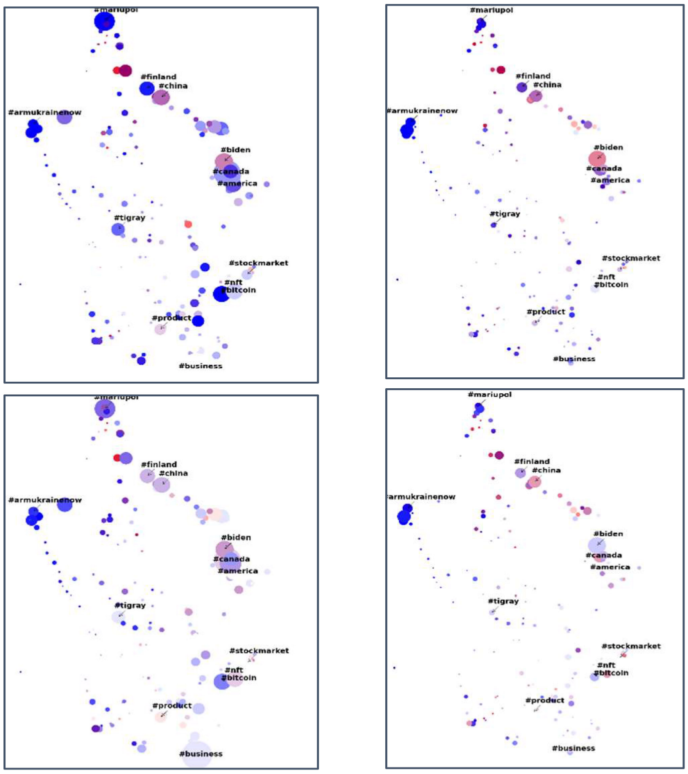

Figure 2: Convo attitudes based on the top influencers (top) and the most active users (bottom), in the first 2 weeks (left) and the last two weeks (right).

Snapshots of Convo Attitudes

The authors attitude in a convo is aggregated to obtain the convo's overall attitude, represented as a combination of [pro-Ukraine, pro-Russia, None], and expressed in the format $\{U\text{: } m, R\text{: } n, \text{N: } k\}$. In Fig. 2 (upper-right conner convos not shown), the attitude of a convo is visualized through its color, which combines blue for pro-Ukraine (U), red for pro-Russia (R), and white for None (N). The intensity of each color corresponds to the proportion of top users in the convo holding each attitude. Convos dominated by users with pro-Ukraine or pro-Russia appear as deep blue or deep red, respectively. Similarly, convos with a higher proportion of users holding None attitude appear lighter in color. If all the top users in a convo have None attitude, the convo appears white and is excluded from the visualization. This color-coded approach provides an intuitive overview of the distribution of attitudes across different convos, highlighting different perspectives towards the ongoing event.

Comparing the plots on the left and right side in Fig. 2, we notice that convos in the last two weeks are generally smaller than those in the first two weeks, consistent with the dataset's temporal trends. From the beginning to the end of the three-month period, most convos show a reduction in size, with a few exceptions, such as #russiaisaterroriststate, which increases its prominence.

We observe that the convo colors of Fig. 2-top are darker than those of Fig. 2-bottom, indicating that the attitudes of the top influencers (the most retweeted users) are generally more polarized than those of top active users. There are various factors like emotional content (Brady, 2017) and imageability (Bernhardt, 2023) that affect the retweet count of a message. Using these results, future research could focus on how these factors contribute to the persuasiveness of messages, especially in the context of public attitude shifts.

Convo Attitude Shift

Table 2 compares the dominant attitude distribution for the top influencers and the most active users. We notice that there is only slight change in the number of convos holding different attitudes, except for the pro-R convos based on the most active users.

Table 2: Distribution of convo dominant attitudes based on users.

	Top Influencers			Most Active Users		
	Pro-U	**Pro-R**	**None**	**Pro-U**	**Pro-R**	**None**
First 2 weeks	83	11	88	61	5	116
Last 2 weeks	84	12	86	58	15	109

To further explain the attitude shift, we investigate the number of top users holding different attitudes in each convo, highlighting whether they have increased, decreased, or remained the same. Table 3 summarizes these changes.

Table 3: Number of convos that have an attitude shift in terms of number of top users.

	Top Influencers			Most Active Users		
	Increased	**Same**	**Decreased**	**Increased**	**Same**	**Decreased**
Pro-U	62	36	83	82	40	59
Pro-R	80	63	38	82	61	38
None	59	37	85	43	31	107

While the overall number of convos in each attitude remains relatively similar in Table 2, there is a notable change in the number of top influencers holding different attitude in Table 3. For pro-U attitude, 83 convos show a decrease in top influencers, compared to 62 convos where the number of top influencers increases. A similar pattern is observed for the None attitude, with 85 convos showing a decrease and 59 showing an increase. In contrast, the pro-R attitude shows a notable rise, with 80 convos showing an increase in top influencers, while only 38 convos show a decrease. For the top active users too, we found more increase in pro-R attitude than in pro-U or None attitude. It is also observed in Fig. 2, that compared to the plots on the left, the proportion of red color is higher in the plots on the right.

The result may indicate a significant shift in overall opinion in the corpus and presence of influence operations. That the pro-R convos rising from 5 to

15 in the top active user-based convo attitudes also highlights a major change in the volume of pro-R tweets among the most active users. Our findings align with online news articles that note how the public opinion among Ukrainians is gradually shifting with an increasing number of people willing to concede to stop the effects of the war[5].

From Fig. 2, we observe that most convos shifting attitudes to pro-R are clustered closely together, particularly for the top active user-based attitudes (Fig. 2-bottom). Considering that the proximity of convos on the plot reflects the degree of user overlap, the observed increase in pro-R influencers and convos may be driven by the same group of users. This aligns with findings by Geissler et al. (2023), who find evidence of the role of bots in increasing pro-Russia narratives during the war.

Additionally, the number of top active users with None attitude decreased in over 50% (107) of all the convos (Table 3). Meanwhile, the number of top active users supporting Ukraine or Russia saw an increase in more than 80 convos each. Consequently, the top active authors become more polarized in the 3-month period.

Evaluation

To evaluate the performance of pre-trained LLMs for the attitude detection task, we conducted a human annotation process. We randomly sampled 5 convos, and selected the top 10 influencers and top 10 active users from each. A maximum of 25 messages from each author were used for annotation. A human annotator reviewed all the messages by each author and assigned a label from: pro-Russia, pro-Ukraine, or None. Overall, 2500 messages were manually examined to evaluate this task. The model's predictions are then compared against these gold labels to achieve a macro F1 score of 0.77 using the instruction-tuned Llama-3.1-8B. Using the predictions of ChatGPT we achieve an F1 score of 0.74. The class wise evaluation scores are listed in Table 4. However, we run our large scale analysis using Llama-3.1 due to its open sourced nature and smaller size.

Table 4: Evaluation of attitude detection using LLMs.

Category	Llama-3.1	Llama-2	ChatGPT
None	0.82	0.75	0.81
Pro-Russia	0.84	0.80	0.91
Pro-Ukraine	0.63	0.47	0.50
Macro F1	0.77	0.67	0.74
Accuracy	0.80	0.69	0.80

Limitations

While our evaluation shows that the LLM agrees well with the human annotator on the user attitude detection task, the evaluation method has certain limitations, including the number of messages, the number of human

[5] https://news.vt.edu/articles/2024/09/ukraine-russia-public-opinion-expert.html

annotators, and potential biases introduced by both the annotator and the LLM. Additionally, we apply the method on a public Twitter dataset, which may introduce platform-specific biases. In future work, we plan to address these limitations by improving the evaluation method and by using datasets from various platforms.

CONCLUSION

In this study we present a framework to analyse and visualize the evolution of public attitudes during significant social or political events, using the Ukraine-Russia conflict as a case study. We use the idea of social convos to identify the largest topics of interest in a Twitter dataset and visualize the overall picture on a two-dimensional author attendance plot. The dominant attitudes of these convos are extracted by prompting a LLM based on the tweets by the top influencers and most active authors in a convo. We demonstrate how the attitudes shift over time across topics. Our analysis reveals several notable trends like increase in polarization and pro-Russia attitudes among these users. This approach provides valuable insights into how public opinions change during global events and presents opportunities for further research into user behaviour and the role of social media in influencing political attitudes. Future work will include examining the convos of major attitude change and the users associated with them to identify potential influential behaviour.

ACKNOWLEDGMENT

This paper is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001121C0186. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA or the U.S. Government.

REFERENCES

Ajzen, I. (2001). 'Nature and operation of attitudes'. *Annual review of psychology*, 52(1), pp. 27–58.

Al-Ghadir, A., Azmi, A., and Hussain, A. (2021). 'A novel approach to stance detection in social media tweets by fusing ranked lists and sentiments'. *Information Fusion* 67, pp. 29–40.

Bernhardt, A., Strzalkowski, T., Sa, N., Bhaumik, A., Katsios, G. (2023). 'Does Imageable Language Make Your Tweets More Persuasive?' *AHFE*. California, USA. 20–24 July.

Bhaumik, A., Bernhardt, A., Katsios, G., Sa, N. and Strzalkowski, T. (2023). 'Adapting emotion detection to analyze influence campaigns on social media'. *ACL-WASSA*.

Bhaumik, A., Sa, N., Katsios, G. and Strzalkowski, T. (2024). 'Social Convos: Capturing Agendas and Emotions on Social Media'. *LREC-COLING 2024*. Torino, Italy.

Biber, D. and Finegan, E. (1988). 'Adverbial stance types in English'. *Discourse processes* 11(1), pp. 1–34.

Brady, W., Wills, J., Jost, J., Tucker, J. and Bavel, J. (2017) 'Emotion shapes the diffusion of moralized content in social networks'. *Proceedings of the National Academy of Sciences*. pp. 7313–7318.

BwandoWando. (2024). Russia Ukraine Conflict Twitter Dataset. Available at https://www.kaggle.com/datasets/bwandowando/ukraine-russian-crisis-twitter-dataset-1-2-m-rows/.

Campello, R., Moulavi, D. and Sander, J. (2013) 'Density-based clustering based on hierarchical density estimates'. *PADKK 2013*. Berlin, Heidelberg.

Carr, C. and Hayes, R. (2015) 'Social media: Defining, developing, and divining'. *Atlantic journal of communication* 23(1), pp. 46–65.

Chang, P., Yu, Y., Sanders, A. and Munasinghe, T. (2023). 'Perceiving the Ukraine-Russia conflict: Topic modeling and clustering on twitter data'. *IEEE BigDataService*.

Cruickshank, I. and Ng, L. 2024. 'Prompting and fine-tuning open-sourced large language models for stance classification'. Preprint.

Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A. (2024). 'The llama 3 herd of models'. *arXiv preprint* arXiv:2407.21783.

Eligüzel, İ. (2023). 'Russia-Ukraine Conflict: A Text Mining Approach through Twitter'. *Bitlis Eren Üniversitesi Fen Bilimleri Dergisi* 12(1), pp. 272–291.

Geissler, D., Bär, D., Pröllochs, N. and Feuerriegel, S. (2023). 'Russian propaganda on social media during the 2022 invasion of Ukraine'. *EPJ Data Science* 12(1) p. 35.

Glandt, K., Khanal, S., Li, Y., Caragea, D. and Caragea, C. (2021) Stance detection in COVID-19 tweets. *ACL-IJCNLP 2021*, Bangkok, Thailand, August 1–6.

Hoewe J. and Peacock, C. (2020). 'The power of media in shaping political attitudes'. *Current Opinion in Behavioral Sciences* 34, pp. 19–24.

Karami, A., Bennett, L. and He, X. (2018). 'Mining public opinion about economic issues: Twitter and the us presidential election'. *IJSDS* 9(1), pp. 18–28.

Katsios, G., Sa, N., Bhaumik, A. and Strzalkowski, T. (2024). 'Uncovering Agendas: A Novel French & English Dataset for Agenda Detection on Social Media'. *LREC-COLING 2024*. Torino, Italy.

Katsios, G., Sa, N. and Strzalkowski, T. (2019). 'Social convos: A new approach to modeling information diffusion in social media'. *AHFE 2019*. Washington USA.

Kotišová, J. and Velden, L. (2023). 'The affective epistemology of digital journalism: emotions as knowledge among on-the-ground and OSINT media practitioners covering the russo-Ukrainian war'. *Digital Journalism* (2023), pp. 1–20.

Maathuis, C. and Kerkhof, I. (2023). 'The first two months in the war in Ukraine through topic modeling and sentiment analysis'. *Regional Science Policy & Practice*, 15(1).

Mather, B., Dorr, B., Dalton, A., Beaumont, W., Rambow, O. and Schmer-Galunder S. (2022) 'From Stance to Concern: Adaptation of Propositional Analysis to New Tasks and Domains'. *Findings of the ACL 2022*, Dublin, Ireland.

McInnes, L., Healy, J., and Melville, J. (2018). 'UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction'. *ArXiv e-prints*.

Mutz, D., Sniderman, P. and Brody, R. (1996). 'Political persuasion and attitude change'. University of Michigan Press.

Nakazato, T., Shibuya, Y., and Takagi, S. (2023) 'Characterizing the Behavior of Healthcare Experts Towards COVID-19 Vaccine on Twitter'. *WI-IAT*. Venice, Italy.

Sharma, K., Zhang, Y. and Liu, Y. (2022). 'COVID-19 vaccine misinformation campaigns and social media narratives'. *ICWSM 2022*. Atlanta, USA.

Tolmach, M, Volynets, V., Trach, Y, Chaikovska, O, Khrushch, S., Kotsiubivska, K., Danieliene, R. and Danielius, P. (2023). 'NFT and Digital Art: Ukrainian Experience of Using Cryptoart'. *ICICT 2023*. London UK.

Vargo, C, Guo, L., McCombs, M. and Shaw, D. (2014). 'Network issue agendas on Twitter during the 2012 US presidential election'. *Journal of communication* 64(2) pp. 296–316.

Yin, H., Song, X., Yang, S. and Li, J. (2022). 'Sentiment analysis and topic modeling for COVID-19 vaccine discussions'. *World Wide Web* 25(3), pp. 1067–1083.

Zhang, B., Ding, D, Jing, L., Dai, G. and Yin, N. (2022). 'How would stance detection techniques evolve after the launch of chatgpt?' *arXiv preprint*.