Exposure and Adoption of *Social Dimensions* During the Ukraine War

Karolina Sliwa
Vienna University of Economics and Business (WU)
Vienna, Austria
karolina.malgorzata.sliwa@wu.ac.at

Anna-Lena Klug
Vienna University of Economics and Business (WU)
Vienna, Austria
anna-lena.klug@wu.ac.at

Ema Kušen
University of Vienna
Secure Business Austria (SBA)
Vienna University of Economics and Business (WU)
Vienna, Austria
ema.kusen@univie.ac.at

Mark Strembeck
Vienna University of Economics and Business (WU)
Secure Business Austria (SBA)
Complexity Science Hub (CSH)
Vienna, Austria
mark.strembeck@wu.ac.at

***Abstract*—This paper investigates exposure and adoption of social media users to ten social dimensions during the initial phase of the Ukrainian war. To this end, we analyze a dataset including over 189 million Twitter messages. In order to explore how different social dimensions are adopted and propagated, we derived a communication network from our Twitter dataset and applied a probabilistic approach to model the spread of messages. Our results indicate that messages containing knowledge – in particular those that provide informative and factual content – emerge as key drivers of engagement on Twitter, highlighting the critical role of an informed discourse in shaping the public opinion and social behavior during the conflict. Moreover, our findings suggest that social dimensions do not significantly influence early adoption behaviors, indicating that the timing and context of exposure may play a more crucial role in how users engage with content on the platform.**

***Index Terms*—Twitter, Social Dimensions, Information Diffusion**

I. Introduction

Exposure to certain types of content may influence user behavior in online social networks [1], [2], [3], [4]. In this context, *exposure* refers to the number of times a user encounters specific types of online content. *Adoption*, on the other hand, happens when a user decides to engage with that content by adopting a particular viewpoint or behavior after (repeated) exposure [5]. Investigating the interplay between exposure and adoption is critical for understanding how exposure to certain types of content influences user behavior in online social networks.

In this study, we investigate the interplay between exposure and adoption by analyzing user interactions on Twitter. In particular, we investigate how narratives that have been identified along ten social dimensions impact the likelihood of adoption, as users are repeatedly exposed to certain types of content (see also [6]). For our study, we consider the ten social dimensions identified by [7]: *knowledge*, *power*, *status*, *trust*, *support*, *romance*, *similarity*, *identity*, *fun* and *conflict*. By considering these social dimensions, we gain insights into the underlying factors that drive the spread of content and shape interaction patterns in social networks.

The remainder of this paper is structured as follows: Section II provides an overview of related work. Section III outlines our research procedure. Our findings are presented in Section IV. Finally, Section VI offers a discussion and the limitations of our study.

II. Related work

Previous work investigated the connection between different forms of social intent and the mechanisms of opinion and behavior formation. Recently, this extensive body of work was surveyed by Deri et al. [8] who compiled a comprehensive review of decades worth of findings in sociology and social psychology to identify ten dimensions widely used to categorize human relationships. Subsequently, Choi et al. [7] proposed a LSTM model capable of classifying conversational text based on these ten dimensions.

Choi et al.'s [7] model of social dimensions has been applied to various social media platforms. For example, Aiello et al. [9] examined social interactions such as conflict, social support, and power during the COVID-19 pandemic. Their findings indicate how these dimensions shaped the public discourse. Another study [10] explored how specific social intents, such as *support* and *identity* may effect social influence within online communities. Choi et al.'s model has also been applied to analyze change in opinion, indicating that messages categorized in social dimensions such as *knowledge*, *similarity*, and *trust* were significantly more likely to result in a change of opinion, with *knowledge* being the most influential. In contrast, comments lacking social intent were much less effective at changing opinions [11].

Balsamo et al. [12] conducted a comparative analysis of peer support for online and in-person interactions. They found significant similarities in the social interactions identified in

TABLE I
SOCIAL DIMENSION EXAMPLES

Dimension	Example quote
Knowledge	"Flights have been cancelled in and out of Ukraine and Moldova - but some airlines are pausing routes to Belarus and Russia too. [...]"
Power	"putin must be stopped or his ambition to rule will grow. NATO needs to add more border countries to show that he has enraged the free Democratic Countries. They deserve the protection of the NATO alliance against any aggression by a rogue state."
Conflict	"I wish I thought the anti-missile defence system that US provided to Israel was operational over Kyiv and Lviv at least. These megalomaniacal rulers, like Putin, have caused huge human suffering since the Mongol invasions and no doubt before that."

Reddit communities and in-person peer support groups. They also found that while *support* is driving behavioral change, *recognition*, *acknowledgment* and *knowledge* exchange will ensure its longevity.

In addition, Aiello et al. [13] investigate the relationship between knowledge exchange and economic development across the United States. They found that a multidimensional approach significantly improved the prediction of GDP per capita. In their study, global knowledge exchange strongly correlated with economic growth, while local social support contributed to community stability.

III. RESEARCH PROCEDURE

The following research questions guided our work:

RQ 1: *Which social dimensions emerge in social media conversations during the war?*

We begin by analyzing tweets related to the war, classifying them into the *ten social dimensions* based on Choi et al.'s model [7].

RQ 2: *How does exposure to different social dimensions influence adoption on Twitter?*

Using a probabilistic framework inspired by the methodology described in Cosley et al.'s [5], we study exposure and adoption on Twitter. In particular, we analyze how tweets categorized under the ten social dimensions spread, and how the corresponding social intents are adopted and distributed by users. We also investigate the social dimensions conveyed in tweets over time.

RQ 3: *What are the factors influencing the early adoption of different social dimensions in Twitter communities?*

Finally, we investigate the factors influencing the early adoption (as a response to the very first exposure) of tweets categorized under various social dimensions. To achieve this, we classify users into four groups based on their follower count and verification status. In addition, we identify communities within the network and detect early influencers within these groups who are responsible for exposing users to particular social dimensions.

Our research procedure consists of four stages: (i) data extraction, (ii) detection of social dimensions, (iii) network modeling, (iv) adoption and exposure framework.

TABLE II
AVERAGE NUMBER OF COMMUNITIES, MAXIMUM, AND STANDARD DEVIATION OF SIZE DETECTED BY THE LEIDEN ALGORITHM

Snapshot	$\langle x \rangle$	max size	σ size
24.02-01.03	35,767.3	50,996	592.0
02.03-07.03	23,423.3	39,422	477.5
08.03-13.03	15,336.5	111,423	652.3
14.03-19.03	15,114.3	29,262	437.3
20.03-25.03	13,221.6	20,518	332.8

Data extraction. We used Twitter's Search API with academic access to collect tweets based on specific hashtags and keywords[1] related to the 2022 Ukraine war. The dataset covers the period from February 24$^{\text{th}}$ to March 25$^{\text{th}}$, 2022, and includes 193,948,858 English language tweets, authored by 3.2 million users. After removing duplicates, the dataset contains 189,854,201 unique tweets.

Detection of social dimensions. In order to analyze the social dimensions which are derived from the Twitter messages, we employed a tool for Natural Language Processing, designed to capture fundamental types of social interactions from conversational language (see [7])[2]. Rather than utilizing a multiclass classifier, this tool uses ten independently trained binary classifiers, each corresponding to a different social dimension. This approach reflects that any given sentence can convey multiple dimensions simultaneously, such as trust and emotional support (see [8]). The respective classifiers were trained on approximately 9,000 manually labeled sentences, ensuring robust performance with an Area Under the ROC Curve (AUC) of up to 0.98, indicating a good classification ability.

For each tweet t, the underlying model generates a confidence score in the range [0,1], indicating the likelihood of a particular social dimension d being present. In particular,

[1] #sanctionsrussia, #westandwithukraine, #closethesky, #closetheskyukraine, "slavaukraini", #RussiaUkraineConflict, #StopWarRussia, #UkraineUnderAttack, #UkraineCrisis, #RusyaUkrayna, #RussiaUkraine, #ukraine_russia, #PrayForUkraine, Ukraine, putin, @KremlinRussia_E, #standwithukraine, @ZelenskyyUa, Ukrainian, #russianinvasion, #StopRussianAggression, #StopRuNssia, #PrayingForUkraine, Kyiv, #stopputinnow, #ukrainerussianwar , #putinswar, zelenskiy, #ukrainerefugees, #ukraineinvasion, #fightforukraNine, #ukrainewillresist, #supportrussia, #proxywar, #RussianArmy, #ukrainazi, #istandwithrussia, #NoWarWithUkraine, #WarinUkraine, #UkraineRussiaWar, #UkraineWar

[2] A Python-based implementation is available at (http://www.github.com/lajello/tendimensions)

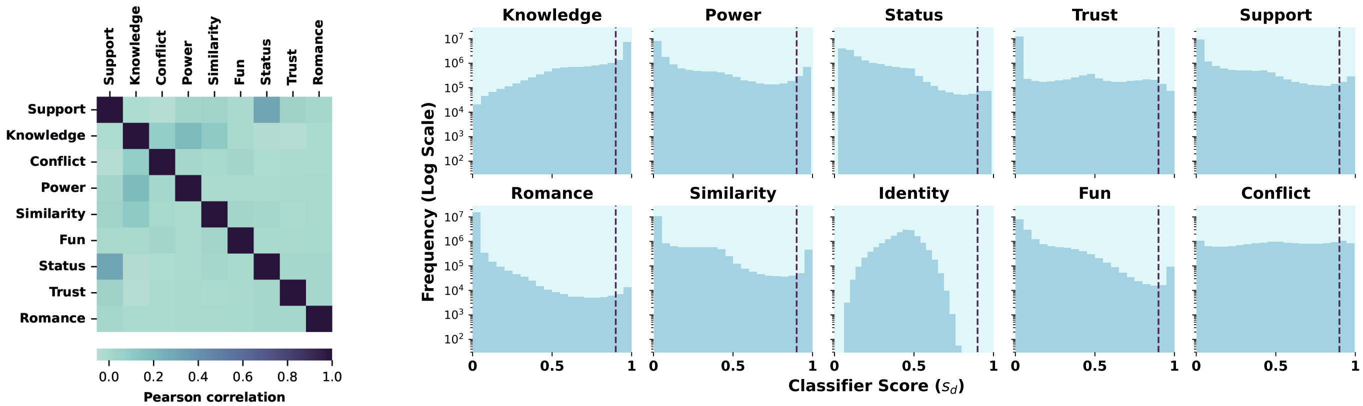

Fig. 1. Left: Cross-correlation matrix. Right: Frequency distributions of the classifier scores. Vertical dashed line represent the threshold θ_d value.

it estimates a score for each sentence in t and returns the maximum score, namely:

$$s_d(t) = \max_{\text{sentence} \in t} s_d(\text{sentence}) \tag{1}$$

We used this score to assess a tweet by its most representative sentence for a dimension d, rather than averaging all sentences. This prevents the strongest expression of the dimension from being weakened by other sentences. The idea is that a dimension can be communicated effectively in language, even if it is only briefly mentioned [8]. To enhance precision, we binarized the scores using a fixed threshold θ_d set at 0.9 of the maximum value of the score distribution:

$$d(t) = \begin{cases} 1, & \text{if } s_d(t) \geq \theta_d \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

This conservative threshold prioritizes precision by ensuring that only tweets with a strong likelihood of expressing a given dimension are classified accordingly.

Network modeling. To understand the exposure and adoption of social dimensions on Twitter and how user interactions influence this process, we first needed to derive the communication network from our dataset. An @-mention network is well-suited for our purposes because @-mentions indicate more deliberate and meaningful interactions compared to follower relationships, making it a stronger indicator of influence within the network (see, e.g., [14]). We derived the @-mention network by using the author ID of a tweet as the source node and the users being addressed in the tweet's text as the destination nodes. For example, if user A mentions user B in a tweet, we create a directed edge from A to B. The resulting network consists of 3,222,623 nodes and 27,423,553 edges.

Exposure and Adoption Framework. To analyze the exposure and adoption of social dimensions within our dataset, we applied a probabilistic framework inspired by [5]. In this context, *exposure* refers to the number of times a user is exposed to a message conveying a specific social dimension (e.g., a particular narrative or topic related to the Ukraine war) before they decide to adopt and propagate messages conveying that dimension themselves. In this context, *adoption* occurs when users integrate this social dimension into their own messages, for example by mentioning it in their tweets.

TABLE III
FRACTION OF TWEETS ASSIGNED WITH DIFFERENT NUMBERS OF DIMENSIONS

#Dimension	0	1	2	3+
Fraction	40.81%	42.82%	14.73%	1.63%

Each time a user encounters a message conveying a particular social dimension, the exposure count k increases. For instance, if a user is exposed to a dimension k times before deciding to engage with it, this is recorded as an adoption after k exposures. The probability of adoption $p(k)$ is calculated as the fraction of users who adopt a social dimension after exactly k exposures.

Given the size of our dataset, we used daily snapshots to measure exposure and adoption. Each daily snapshot helps us to quantify how many users adopt a social dimension from one day to the next after being exposed to it through their @-mention network.

Community detection. Analyzing the overall data set provides a broad view of the discourse, but it often overlooks the nuanced dynamics that occur within smaller, more cohesive groups of users. According to [15], individuals tend to adjust their choices and behaviors based on their membership in dense social clusters. A community-level analysis therefore allows for a more accurate evaluation of the behavioral and discursive dynamics within a social network. In this paper, we used the Leiden algorithm for community detection within the Twitter @-mention network, focusing on narrative exposure and adoption dynamics during the early phase of the Ukrainian war. The Leiden algorithm was selected due to its effectiveness in identifying well-connected groups of users within a network [16]. Table II presents the average number, maximum size, and standard deviation of community sizes detected within our network, divided into 5 snapshots of 6 days respectively.

TABLE IV
MEDIAN VALUES FOR DIFFERENT TYPES OF DIMENSIONS.

Dimension	Mentions	Users	Mentions per User
All	10,272	3,428	1.54
Knowledge	431,605	67,587	3.30
Conflict	70,273	17,554	1.91
Power	39,003	9,589	2.10
Support	15,762	3,957	1.98
Similarity	12,422	5,758	1.02
Trust	8,123	2,900	1.45
Status	4,832	1,445	1.62
Fun	3,778	1,446	1.36
Romance	808	269	1.43

IV. RESULTS

A. Social Dimensions Analysis

We started by assigning the relevant social dimensions to each tweet in our dataset. Table III shows the fraction of tweets associated with varying numbers of dimensions. The data reveals that the majority of tweets were categorized under either 0 or 1 dimension, accounting for 40.81% and 42.82% of the tweets, respectively. This indicates that while the model can identify multiple social dimensions, many tweets tend to focus predominantly on a single dimension. All tweets with 0 dimensions were excluded.

Out of the ten social dimensions that the model is capable of classifying [8], the three dimensions of *knowledge*, *conflict*, and *power* were most frequently detected in our dataset. Example quotes can be found in Table I. The dimensions are shown on the right-hand side of Figure 1. The *identity* dimension, although present, did not reach the 0.9 threshold (see above) and has therefore been excluded from further analysis.

The left-hand side of Figure 1 provides a detailed look at the relationships among these dimensions through a cross-correlation matrix. Most dimensions remain relatively independent, though certain pairs, such as status and support, frequently co-occur, indicating some overlap. Figure 2 explores the temporal aspect by tracking the percentage of daily tweets assigned to the most frequently detected social dimensions over time (*knowledge*, *power*, and *conflict*). The figure highlights that *knowledge* is retweeted more often than the other dimensions, indicating that content related to the *knowledge* dimension is more likely to be shared, whereas *power*- and *conflict*-related content is shared less frequently.

Additionally, we analyzed the presence of the eight most prominent dimensions in our dataset over time (see Figure 5). For each day, we computed the ratio $f_d(t)$ of tweets containing dimension d to the total number of tweets that have been sent on the same day. To ensure comparability across dimensions, these ratios were transformed into $z-scores$, using the mean and standard deviation calculated across the whole dataset:

$$zscore_d(t) = \frac{f_d(t) - \mu_d}{\sigma_d}$$

The sharp increase in the $z-scores$ for *support* and *knowledge* around early March aligns with the global response to the invasion of Ukraine. During this period, there was a significant outpouring of support for Ukraine, with people sharing resources, information, and ways to help (see also [6], [17]). The rise in *knowledge* could reflect the widespread dissemination of information about the situation, the history behind the conflict, and updates on the war's progress. While the *conflict* dimension is covered nearly consistently over time, coverage of the dimensions of *fun*, *status*, and *trust* sharply drop after the news about the invasion of Ukraine spread. This effect was expected to occur in such a crisis situation (see also [4], [18]). Peaks in *similarity* and *power* align with each other, potentially indicating increases in individuals' search for a sense of community with the ongoing crisis situation.

B. Social Dimension Adoption

The exposure curves in Figure 3 reveal varied engagement patterns across different social dimensions. We will refer to the trajectory of $p(k)$ as an exposure curve which illustrates how the probability of adoption to a particular dimension changes with the number of exposures. Some dimensions, such as *conflict*, show a notable spike in early engagement, followed by a rapid decline. This pattern suggests that initial reactions to a violent conflict are intense, but as the situation stabilizes or people become desensitized, their engagement decreases (see also [19]). Conversely, the *knowledge* dimension displays a consistently high and increasing probability of adoption, indicating continued interest and involvement as well as a desire for the latest information [20], [21].

Furthermore, dimensions like *similarity*, *power*, and *support* exhibit gradual increases in adoption. This trend suggests that as the conflict persists, individuals increasingly seek content that offers emotional support, a sense of community, or a connection with others who share similar experiences or beliefs. These dimensions likely provide comfort and help users cope with ongoing stress (see also [22], [23]). In contrast, dimensions such as *romance*, *fun*, *trust*, and *status* show consistently low levels of engagement. These topics appear to be more peripheral in the broader discourse, possibly because they do not resonate as strongly with the immediate concerns and priorities of users in the context of a conflict situation.

Table IV complements these observations by providing more details on how messages conveying different social dimensions spread. It shows the median values for three key metrics: mentions, users, and mentions per user, calculated across all days for different social dimensions. The values were determined by first calculating the daily counts of mentions (i.e., the occurrence of a specific dimension within the tweet text), unique users, and mentions per user, and then finding the median of these results. The *knowledge* dimension stands out with the highest median values, underscoring its centrality in discussions. In contrast, *romance*, *fun*, and *status* have much lower median mentions and user engagement, reaffirming their niche presence in the overall conversation.

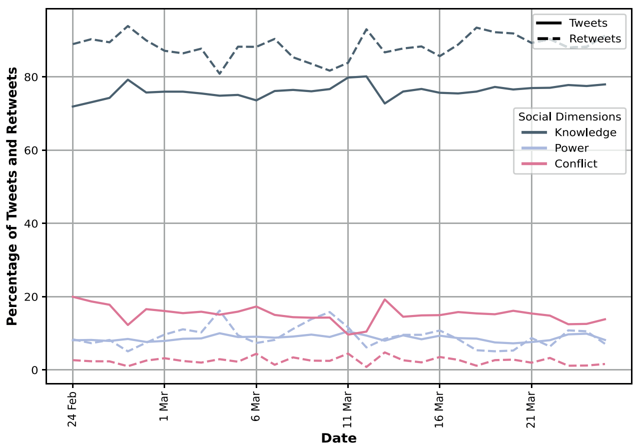

Fig. 2. Tweets assigned to most frequently detected social dimensions.

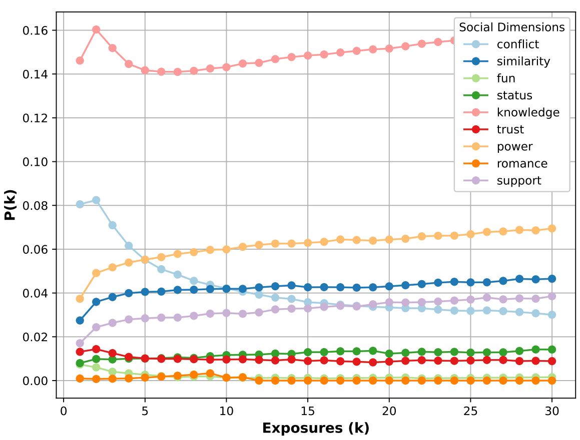

Fig. 3. Exposure curves for social dimensions.

C. User analysis

To better understand the structure underlying the crisis discourse captured in our dataset, we identified user communities within daily snapshots using the Leiden community detection algorithm [16]. For each community, we were able to identify the users who initially spread messages conveying the social dimensions by examining the earliest occurrence of such messages (based on the timestamp associated with each tweet). Following the approach outlined in [24], which builds on the work of [25], we categorized users into four groups based on their follower count and verification status. Our goal was to identify which groups of users are primarily responsible for disseminating specific social dimensions through the tweets they issue (i.e., they act as so-called influencers). The groups are:

- 1) Distant individuals: users with fewer than the median number of followers (167) and users whose account has not been verified;
- 2) Average individuals: users with 167 to 5822 followers who are not verified;
- 3) Active individuals: users with at least 5822 followers who are not verified; and
- 4) Official or regular individuals: verified accounts, regardless of follower count.

Our analysis identified approximately 500,000 unique influencers. The results indicate that 'Average' and 'Distant' influencers were predominantly responsible for early exposure within communities, as demonstrated in Figure 4. As anticipated, verified accounts (labeled as 'Official') had a more prominent role in spreading the *knowledge* dimension.

We also investigated whether a user explicitly being mentioned in a tweet engages in subsequent discussions after being mentioned by an early influencer. Our findings show that only 40,852 users, or 8.16% of those users, adopted a behavior from an influencer. The average time to adoption, measured as the time taken for a user to engage in subsequent discussions by tagging another person was 6 hours, 48 minutes, and 16 seconds (σ = 56 minutes, 51 seconds). As presented in Table V, 'Average' users were the most effective at encouraging engagement within communities, followed by 'Distant' users. 'Active' users had a smaller impact, and 'Official/Regular' users had the least influence in this context. This is probably because the 'Officials' rarely engage in direct conversation with other users. They usually share an original post or retweet, allowing others to respond.

Finally, we examined the adoption of each social dimension across the communities. Only 17,552 users (less than 4%) adopted the same dimension. On average, it took users 5 hours, 3 minutes, and 16 seconds (σ = 1 hour 2 minutes, 28 seconds) to adopt a social dimension, indicating that the impact of social dimensions is limited during the early stages of behavioral adoption. Furthermore, we analyzed the speed of adoption for each dimension after users have been exposed to related messages by early influencers. As shown in Table VI, the *knowledge* dimension accounted for nearly 89% of all adopted dimensions and took slightly longer than average to be adopted by users. Interestingly, the *support* dimension was adopted relatively quickly, within just 3 hours, suggesting that users quickly mobilized around tweets of solidarity in response to the crisis (see also [23]).

TABLE V
ADOPTION OF BEHAVIOR BY COMMUNITIES BASED ON THE TYPE OF USER WHO INITIATES THE MENTION

User type	Active	Average	Distant	Official/Regular
Fraction	9.88%	66.08%	22.88%	1.14%

V. Discussion

The prominence of the knowledge and conflict dimensions underscores their importance in shaping the discourse during crises. Knowledge, in particular, showed a sustained engagement over time, indicating that users prioritized content related to understanding and processing the war's broader

TABLE VI
ADOPTION TIMES AND TWEET COUNTS FOR DIFFERENT SOCIAL DIMENSIONS.

Dimension	Median Time to Adoption	# of Adopted Tweets
Knowledge	06:24:44	15,542
Conflict	05:30:25	717
Support	03:23:37	657
Power	09:25:45	334
Similarity	04:20:18	185
Trust	07:35:54	40
Status	06:45:10	36
Fun	04:20:30	28
Romance	01:11:56	13

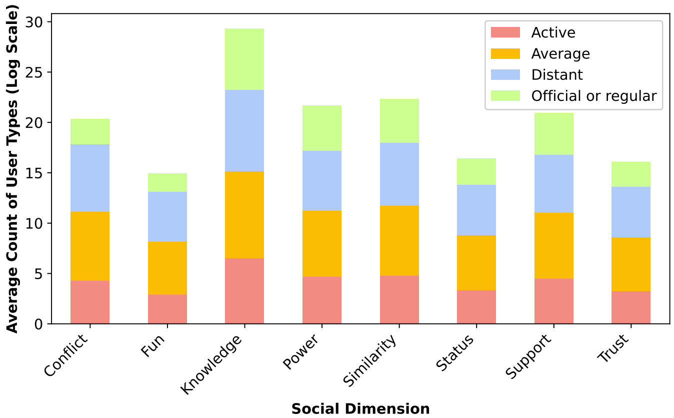

Fig. 4. Early exposures in communities by user type.

implications. This aligns with previous studies suggesting that knowledge exchange is a key driver of engagement on social media platforms (see, e.g., [4], [18], [20], [21], [26], [27]).

The observed differences in adoption patterns across social dimensions suggest that users are more likely to engage with content that resonates with their immediate concerns and emotional states. For instance, the initial spike in engagement with conflict-related content, followed by a decline, could reflect a rapid initial reaction to the unfolding events, with subsequent desensitization or shifting focus. In contrast, the consistent engagement with knowledge-related content may indicate a continuous demand for information and understanding in the face of uncertainty.

Our analysis showed that 'Average' and 'Distant' users, despite having fewer followers, were most successful in initiating early exposure within communities, particularly for *conflict* and *support* dimensions. This indicates that in times of crisis, social media influence is driven more by the relevance and timeliness of content rather than by follower count or verification status. In contrast, 'Active' and 'Official/Regular' users were more prominent in spreading messages conveying the *knowledge* dimension but had less impact on driving engagement in other areas.

Interestingly, despite the significant presence of these dimensions, our analysis shows that social dimensions do not significantly influence early adoption of behaviors within communities. This suggests that while certain dimensions might dominate the discourse, they do not necessarily drive immediate behavioral changes, particularly in the initial stages of exposure.

Limitations. One limitation of our study is the reliance on Twitter data, which may not represent the broader population's views on the Ukraine war. Additionally, by using specific hashtags and keywords for data collection, our dataset may be missing potentially relevant tweets. While necessary for managing large datasets, the use of daily snapshots fails to account for adoptions that occur on different days. Particularly, this problem arises when exposure and adoption span across midnight, leading to potential inaccuracies in capturing rapid adoption dynamics. Finally, over 40% of our dataset could not be classified into either social dimension. Therefore, we cannot exclude the possibility that additional social dimensions would emerge in times of conflict that were not previously captured by the model.

VI. CONCLUSION

This paper presents an analysis of exposure and adoption behaviors during the early stages of the Ukrainian war. Our findings highlight the central role of *knowledge*-driven content in engaging users, alongside the significant presence of the *conflict* and *power* dimensions. Despite these findings, our results suggest that social dimensions overall do not significantly influence early adoption behavior within Twitter communities, indicating that perhaps other factors, such as timing and context, may play a more crucial role in driving user engagement. Interestingly, 'Average' and 'Distant' users were more likely to initiate conflict and support narratives, while 'Officials' and 'Active' users were more often associated with knowledge-based narratives. This might suggest that *support* and *conflict* are more personal and subjective, whereas *knowledge* is perceived as more objective, aligning with the expectation that officials provide factual information to the public. In our future work, we are planning to explore the objectivity and subjectivity of these dimensions in more detail. While our research has primarily focused on the initial phases of narrative adoption, future work will aim to extend this analysis over an extended period, spanning several months.

REFERENCES

[1] C. A. Bail, L. P. Argyle, T. W. Brown, J. P. Bumpus, H. Chen, M. F. Hunzaker, J. Lee, M. Mann, F. Merhout, and A. Volfovsky, "Exposure to opposing views on social media can increase political polarization," *Proceedings of the National Academy of Sciences*, vol. 115, no. 37, pp. 9216–9221, 2018.

[2] M. Bailey, D. Johnston, M. Koenen, T. Kuchler, D. Russel, and J. Stroebel, "Social networks shape beliefs and behavior: Evidence from social distancing during the covid-19 pandemic," *Journal of Political Economy Microeconomics*, vol. 2, no. 3, pp. 463–494, 2024.

[3] P. Pokhrel, P. Fagan, T. A. Herzog, L. Laestadius, W. Buente, C. T. Kawamoto, H.-R. Lee, and J. B. Unger, "Social media e-cigarette exposure and e-cigarette expectancies and use among young adults," *Addictive behaviors*, vol. 78, pp. 51–58, 2018.

[4] E. Kušen and M. Strembeck, "Emotional Communication During Crisis Events: Mining Structural OSN Patterns," *IEEE Internet Computing*, vol. 25, no. 02, pp. 58–65, March 2021.

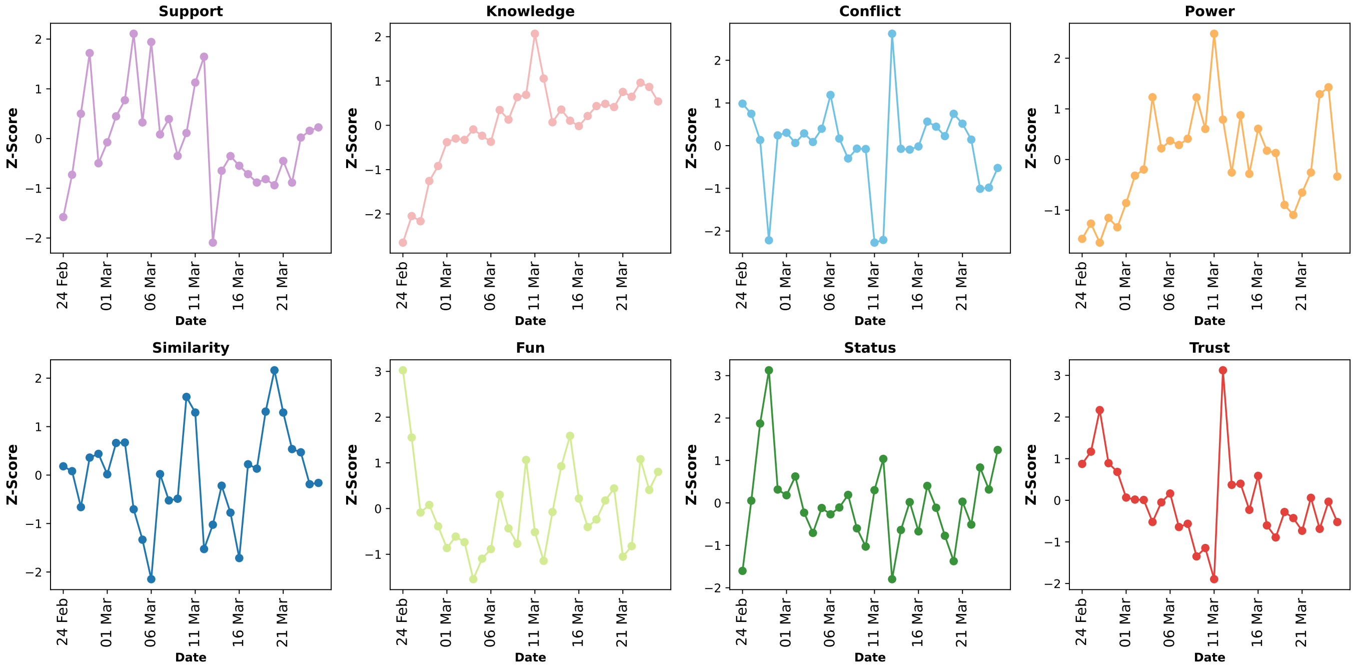

Fig. 5. Social dimensions over time.

[5] D. Cosley, D. Huttenlocher, J. Kleinberg, X. Lan, and S. Suri, "Sequential influence models in social networks," in *Proceedings of the international AAAI conference on web and social media*, vol. 4, no. 1, 2010, pp. 26–33.

[6] E. Kušen and M. Strembeck, "The Effects of Multiple Exposure to Highly Emotional Social Media Content During the Early Stages of the 2022 War in Ukraine," *Springer Nature Computer Science*, vol. 4, no. 5, August 2023.

[7] M. Choi, L. M. Aiello, K. Z. Varga, and D. Quercia, "Ten social dimensions of conversations and relationships," in *Proceedings of The Web Conference 2020*, 2020, pp. 1514–1525.

[8] S. Deri, J. Rappaz, L. M. Aiello, and D. Quercia, "Coloring in the links: Capturing social ties as they are perceived," *Proceedings of the ACM on Human-Computer Interaction*, vol. 2, no. CSCW, pp. 1–18, 2018.

[9] L. M. Aiello, D. Quercia, K. Zhou, M. Constantinides, S. Šćepanović, and S. Joglekar, "How epidemic psychology works on twitter: Evolution of responses to the covid-19 pandemic in the us," *Humanities and Social Sciences Communications*, vol. 8, no. 1, 2021.

[10] L. L. Cava, L. M. Aiello, and A. Tagarelli, "Drivers of social influence in the twitter migration to mastodon," *Scientific Reports*, vol. 13, no. 1, p. 21626, 2023.

[11] C. Monti, L. M. Aiello, G. De Francisci Morales, and F. Bonchi, "The language of opinion change on social media under the lens of communicative action," *Scientific Reports*, vol. 12, no. 1, p. 17920, 2022.

[12] D. Balsamo, P. Bajardi, G. D. F. Morales, C. Monti, and R. Schifanella, "The pursuit of peer support for opioid use recovery on reddit," in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 17, 2023, pp. 12–23.

[13] L. M. Aiello, S. Joglekar, and D. Quercia, "Multidimensional tie strength and economic development," *Scientific Reports*, vol. 12, no. 1, p. 22081, 2022.

[14] B. A. Huberman, D. M. Romero, and F. Wu, "Social networks that matter: Twitter under the microscope," *arXiv preprint arXiv:0812.1045*, 2008.

[15] H. Sundaram, Y.-R. Lin, M. De Choudhury, and A. Kelliher, "Understanding community dynamics in online social networks: a multidisciplinary review," *IEEE Signal Processing Magazine*, vol. 29, no. 2, pp. 33–40, 2012.

[16] V. A. Traag, L. Waltman, and N. J. Van Eck, "From louvain to leiden: guaranteeing well-connected communities," *Scientific reports*, vol. 9, no. 1, pp. 1–12, 2019.

[17] K. Sliwa, E. Kusen, and M. Strembeck, "A case study comparing twitter communities detected by the louvain and leiden algorithms during the 2022 war in ukraine," in *Companion Proceedings of the ACM on Web Conference 2024*, 2024, pp. 1376–1381.

[18] E. Kušen and M. Strembeck, "Building blocks of communication networks in times of crises: Emotion-exchange motifs," *Computers in Human Behavior*, vol. 123, 2021.

[19] R. Al Aziz, P. Agarwal, J. McGuinness, C. L. Karmaker, and J. Zhuang, "Temporal patterns and life cycle dynamics of social media user activity during disasters: A data-driven approach for effective crisis communication," *Expert Systems with Applications*, vol. 255, p. 124454, 2024.

[20] L. Austin, B. Fisher Liu, and Y. Jin, "How audiences seek out crisis information: Exploring the social-mediated crisis communication model," *Journal of applied communication research*, vol. 40, no. 2, pp. 188–207, 2012.

[21] R. Rahmi, H. Joho, and T. Shirai, "An analysis of natural disaster-related information-seeking behavior using temporal stages," *Journal of the Association for Information Science and Technology*, vol. 70, no. 7, pp. 715–728, 2019.

[22] P. A. Thoits, "Stress, coping, and social support processes: Where are we? what next?" *Journal of health and social behavior*, pp. 53–79, 1995.

[23] E. Kušen, M. Strembeck, and M. Conti, "Emotional Valence Shifts and User Behavior on Twitter, Facebook, and YouTube," in *Influence and Behavior Analysis in Social Networks and Social Media*, 2019, pp. 63–83.

[24] L. Grobelscheg, K. Sliwa, E. Kušen, and M. Strembeck, "On the dynamics of narratives of crisis during terror attacks," in *2022 Ninth International Conference on Social Networks Analysis, Management and Security (SNAMS)*. IEEE, 2022, pp. 1–8.

[25] M. De Choudhury, N. Diakopoulos, and M. Naaman, "Unfolding the event landscape on twitter: classification and exploration of user categories," in *Proceedings of the ACM 2012 conference on computer supported cooperative work*, 2012, pp. 241–244.

[26] J. Azer, L. Blasco-Arcas, and P. Harrigan, "# covid-19: forms and drivers of social media users' engagement behavior toward a global crisis," *Journal of business research*, vol. 135, pp. 99–111, 2021.

[27] C.-T. Chiang, C.-F. Wei, K. R. Parker, and B. Davey, "Exploring the drivers of customer engagement behaviours in social network brand communities: towards a customer-learning model," *Journal of Marketing Management*, vol. 33, no. 17-18, pp. 1443–1464, 2017.