

Narratives of War: Ukrainian Memetic Warfare on Twitter

YELENA MEJOVA, ISI Foundation, Italy

ARTHUR CAPOZZI, ETH Zürich, Switzerland

CORRADO MONTI, CENTAI, Italy

GIANMARCO DE FRANCISCI MORALES, CENTAI, Italy

The 2022 Russian invasion of Ukraine has seen an intensification in the use of social media by governmental actors in cyber warfare. Wartime communication via memes has been a successful strategy used not only by independent accounts such as @uamemesforces, but also—for the first time in a full-scale interstate war—by official Ukrainian government accounts such as @Ukraine and @DefenceU. We study this prominent example of memetic warfare through the lens of its narratives, and find them to be a key component of success: tweets with a ‘victim’ narrative garner twice as many retweets. However, malevolent narratives focusing on the enemy resonate more than those about heroism or victims with countries providing more assistance to Ukraine. Our findings present a nuanced examination of Ukraine’s influence operations and of the worldwide response to it, thus contributing new insights into the evolution of socio-technical systems in times of war.

CCS Concepts: • **Information systems** → **Social networks**; • **Human-centered computing** → **Empirical studies in HCI**; *Empirical studies in collaborative and social computing*.

Additional Key Words and Phrases: Memes, Propaganda, Social media

ACM Reference Format:

Yelena Mejova, Arthur Capozzi, Corrado Monti, and Gianmarco De Francisci Morales. 2025. Narratives of War: Ukrainian Memetic Warfare on Twitter. *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW139 (April 2025), 28 pages. <https://doi.org/10.1145/3711037>

1 Introduction

On 24 February 2022, Russia began a large-scale military invasion of Ukraine. The United Nations Human Rights Office has confirmed the death of 27 768 civilians up to 8 October 2023.¹ However, their report highlights that the actual figures are believed to be considerably higher, and most estimates suggest more than 300 000 total deaths.

Amid this immense tragedy, the two governments at war—Russia and Ukraine—are also waging an information war on many fronts. Due to the prevalence of the Internet as a primary information source for Ukrainians,² the war has been described as the first social-media war,³ echoing the reputation of the Vietnam war as the “first television war” due to the role this medium had in informing the public about the conflict.

¹<https://ukraine.un.org/en/248799-ukraine-civilian-casualties-8-october-2023>

²<https://www.unian.info/society/media-space-internet-bypasses-tv-as-main-news-source-for-ukrainians-11330729.html>

³<https://www.forbes.com/sites/petersuciu/2022/03/01/is-russias-invasion-of-ukraine-the-first-social-media-war>

Authors’ Contact Information: Yelena Mejova, ISI Foundation, Turin, Italy, yelenamejova@acm.org; Arthur Capozzi, ETH Zürich, Zürich, Switzerland, arthur.capozzi@gess.ethz.ch; Corrado Monti, CENTAI, Turin, Italy, corrado.monti@centai.eu; Gianmarco De Francisci Morales, CENTAI, Turin, Italy, gdfm@acm.org.

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

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 2573-0142/2025/4-ARTCSCW139

<https://doi.org/10.1145/3711037>

Both governments and the common people have been using social media to propagate their views on the conflict and to affect global public opinion. A vivid example is @Ukraine, the official account of the Ukrainian government on Twitter.⁴ This account was started in 2016 by Yarema Dukh while working as a press officer in the Ukrainian presidential office. One of the memes produced by this account, depicting “living next to Russia” as the worst possible headache,⁵ has been shared more than 600 000 times and has reached 55 million views. According to the author, the goal of this meme was “to explain to some large and distant target audiences that Russia is the problem here, not Ukraine, the West, the U.S., NATO, aliens or anyone else”.⁶ This account has now become one of the most visible examples of memetic warfare from a warring government.

An early definition of memetic warfare was given by the NATO Strategic Communications Centre of Excellence in Latvia in a document titled “*It’s time to embrace memetic warfare*” [28]. The document defines it as “the competition over narrative, ideas, and social control in a social-media battlefield” and is viewed as “the digital native’s version of psychological warfare, more commonly known as propaganda”. The authors argue that it needs to be brought into mainstream military thinking by investing in its development. While propaganda has long been a central component of war, memetic warfare introduces a new key aspect: its reach is intrinsically tied to how much its audience is willing to spread it. Governments engaged in psychological warfare used to rely mostly on one-way, broadcast communication media such as posters and radio, and on subsequent word-of-mouth amplification by individuals. However, the two-way nature of social media implies that authors have to persuade their own audience to help spread their message on the platform in order to reach wider audiences. This characteristic creates a particular amalgam of social-media bottom-up dynamics with the top-down approach typical of military objectives and narratives. Different objectives may then require spreading different types of content, which in turn may have varying degrees of success with specific audiences.

Despite the importance of this issue, so far little research effort has been dedicated to the quantitative analysis of the effectiveness of memetic warfare. The goal of the present work is to investigate which audiences it is able to reach and resonate with. These questions are of primary interest not only to the parties involved, but also to global public opinion. In order to prepare, evaluate, and react appropriately, we need to better understand these novel techniques of “soft propaganda” [49], which are likely to become more common in the future.

To concretely and critically evaluate these phenomena, we choose to analyze the main Ukrainian Twitter accounts involved in the Russo-Ukrainian conflict: the two main official Ukrainian government accounts, @Ukraine and @DefenceU, and the most followed Ukrainian meme page, @uamemes-forces. Our analysis focuses on the following research questions:

- RQ1 What are the main characteristics of the messaging used by these popular Ukrainian accounts?
- RQ2 Which attributes of a wartime meme by these popular Ukrainian accounts are good predictors of its virality?

Based on previous literature on memes, we hypothesize that the accounts will employ different visual elements [43]. However, expanding on this literature, we examine the content of memes in terms of the actors involved, their emotional appeal, and the narratives used. Inspired by the narrative psychology and linguistic types analysis [22], we identify four roles within the memes, depending on the moral quality of the actors in the narrative and on their perceived power. In our context, these roles often represent the benevolent moral quality of the author’s side or the malevolent one of the enemy, and we hypothesize that invocation of such qualities will result in

⁴Twitter has been rebranded to “X”. For consistency with previous literature, we prefer using the former name.

⁵<https://twitter.com/Ukraine/status/1468206078940823554>

⁶<https://www.washingtonpost.com/world/2022/01/26/ukraine-russia-memes-social>

greater popularity of the content. Indeed, we find that the narratives we identify are a powerful predictor of the virality of a meme authored by these accounts. Content with a victim narrative—a benevolent, weak actor—has a 109% higher reach w.r.t. content with no narrative, while a villain narrative—depicting the enemy as malevolent and strong—increases reach by 32%.

Finally, since one of the goals of memetic warfare is to influence global public opinion, we look at the success of the content geographically. Considering the importance of the narratives we identify, we ask ourselves the following research question:

RQ3 What is the geographic resonance of each narrative, and how does it relate to each country's actions in the conflict?

As the Web platforms have been employed to rally support for war efforts [75], we hypothesize that it is possible that the broader affinity to a narrative would be reflected in a country's willingness to assist one of the sides. We find that memetic warfare reach reflects important aspects of the offline world: the number of retweets per capita in a given country is highly correlated ($\rho = 0.787$, $p < 10^{-3}$) with the assistance given by that country to Ukraine. In addition, different narratives resonate differently with different countries. Tweets with a villain narrative spread more effectively in countries that are more involved in the conflict—as measured by the share of GDP devoted to economic and military assistance to Ukraine ($\rho = 0.573$, $p < 0.01$). This connection emerges as significant not only when considering our data set as a whole, but even when analyzing each account individually, thus hinting at a more general pattern that could characterize communication about the Ukrainian war on social media.

The present study has important implications for our understanding of the interaction between communication technologies and their use by both governments and non-governmental actors during a conflict. Our research questions (particularly RQ2 and RQ3) focus on analyzing how the global audience perceives these messages. In doing so, we illustrate how social media can contribute to a country's strategic messaging during war, and how such messaging is related to the attention and aid given to the country. As cyber-influence campaigns become more common and sophisticated, we encourage the research community to apply a multi-factorial view of the relevant content, including visual and narrative aspects, to understand its resonance worldwide.

2 Background and Related Work

The social computing community has long been invested in exposing the horrors of war [9, 35], for example by studying how social media exacerbate mass trauma [61]. The community has analyzed how computer-mediated communication impacts politically-sensitive topics, e.g., by revealing the role of demagoguery in civic engagement on Reddit [54], participating in online activism (“slacktivism”) in donation campaigns [42], and investigating the reach of anti-migration ad targeting on Facebook [10]. Further, recent work on social media communication by politicians during crises has shown that their content receives outsized engagement that is not accounted for by their popularity or engaging tweet features [29]—thus motivating further study on the role of official accounts.

The 2022 Russian invasion of Ukraine happened at a time of wide adoption of social media, when the advancement of Internet infrastructure enables real-time, high-resolution media to be shared with a worldwide audience. As such, the actions on the battlefield and the reactions of those inside and outside Ukraine are streamed on Facebook, Twitter, Instagram, and other social media platforms. This constant stream also “proves to be an essential mechanism in getting on-ground real-time reporting of a dangerous event” for mainstream media sources [66]. Several theoretical frameworks can help us frame such presence. First, as a *cyber-influence campaign*, which seeks to promote a narrative alternative to that of the Russian side [36]. The second framework is that of

nation-building or *nation-branding*: two processes that were originally separated in the domains of the government (former) and PR campaign managers (latter), but that have coalesced into a single process. Bolin and Ståhlberg [4] observe that the recent campaigns by Ukraine to promote “Ukrainian capability” are “aimed to raise support for EU membership, arms deliveries”—a kind of soft power dubbed *selfie diplomacy* [47]. The third framework is that of *community building* during a crisis that “facilitates stakeholder support and builds relationship” [15, 23], and targets both international and domestic audiences.

Since the annexation of Crimea in 2014, the social media presence of both Ukraine and Russia has been closely scrutinized by researchers in communications, media studies, and human-computer interaction. The separatist movement in the Donbass region has intensified a battle of narratives: both sides employed five contextual frames, namely historical, geographical, religious, ethnic, and political [45]. According to Makhortykh and Sydorova [46], those on the pro-Russian side were also more likely to make extensive use of photos of children, the authors speculate, in order to “evoke compassion from the potential audience by using sentimental images”. On the Russian social media VKontakte in 2014, pro-Ukrainian users framed the conflict as a limited military action against local insurgents, while pro-Russian separatists framed it as an “all-out war against the Russian population of Eastern Ukraine” [46]. Such divergent frames might have “led to the formation of divergent expectations in Ukraine and Russia concerning the outcome of the war in Donbass” [46].

After the full-scale Russian invasion, several datasets have been made available to the research community. For instance, Chen and Ferrara [13] collected Twitter posts having various Russo-Ukrainian, war-related keywords. Focusing on the accounts participating in these online conversations, Hare and Jones [33] track the use of the Ukrainian flag as a marker of support on Twitter in late February 2022—a form of identity activism [33]. Such accounts are homophilic, and more likely to share U.S. Democrat-leaning messages. Automated accounts, or bots, are often part of a communication operation. Shen et al. [62] estimate the share of bot accounts around the beginning of Russia’s invasion in 2022. Using the tool Botometer [74], they identify around 13.4% of the tweets as likely to be generated by bots. Most of these tweets espoused a pro-Ukrainian position; it is worth noting that by then Russia had suspended access to Twitter for its citizens.

The term *meme* was originally defined by the evolutionary biologist Richard Dawkins in 1976 [21] as a replicator analogous to a gene in its ability to transmit information, including cultural artifacts and beliefs. Due to the vague nature of the term’s definition, and the fact that the very nature of memes is a “continuous mutation”, the study of memes, or Memetics, is applied broadly to the study of cultural information transfer [37]. In the age of social media, the concept of “Internet meme” breaks from Dawkin’s analogy to focus on the artifacts—texts, images, or videos—instead of abstract ideas [41]. Indeed, Reese et al. [60] point to the properties of the visual medium, including syntactic implicitness and iconicity, that make it particularly suited for “framing and articulating ideological messages”. As memes are defined by their capability of being modified and reshared, we consider all visual media as potential memes [79]. Their virality potential has been analyzed using a set of visual features—e.g., scale or inclusion of text—by Ling et al. [43]: we use such features as a baseline in our analysis. The usage of memes in political communication have been studied during the 2016 US Presidential election [73], the 2019 Ukrainian election [58], in Brazil [17] and Hong Kong [25], in Germany’s [3] and US’ [18, 30] far right. Military actions in the past decade have been accompanied by meme-supported cultural expression by U.S. troops [64], by anti-Islamic State activists [50], and in the early days of the Russo-Ukrainian conflict around Crimea [72]. Finally, a notable precedent in the usage of memes by state military as a tool of propaganda is that of the Israel Defense Forces [48].

The use of internet memes as a propaganda tool both echoes classic communication strategies, and employs new technologies for a greater impact. Following past tactics, emotional, divisive

messages are used to demonize, discredit, and “other” the opposition [18, 73]. Unlike the past top-down communication, current communication platforms host both official and user-generated content. This content may be designed to look funny, simplistic and amateurish to be perceived as belonging to a grassroots movement instead of a coordinated communication campaign [17]. When the messages are forwarded by their audience, the easily accessible editing technologies allow for their content to be quickly modified and remixed, allowing for a “participatory digital culture” that is diverse and highly personalized [72]. For instance, the hacker collective Anonymous [50], personnel in the US army [64], or individuals bonded by a political view [18] may now produce and propagate politically-salient messages. Moreover, the fractured social Web sphere creates spaces (including highly polarized and radicalized ones) wherein memes originate, develop, propagate, and eventually influence other communities and the mainstream media [77]. Such content is often rife with inside culture references, humor, and attention-grabbing imagery [25, 50, 58], which may then be imitated by state actors [78].

A popular framework for the analysis of memes relies on identifying the overarching *narratives*, whereby each meme constitutes a partial element [22]. Boatwright and Pyle [2] study 163 tweets by the verified @Ukraine and @Kyiv Twitter accounts posted in February 2022, finding pre-invasion promotional content as well as myth-making. Similarly, Yehorova et al. [76] identify different humorous elements in the early wartime tweets. A more structured framework of narrative analysis is that proposed by Stephen Karpman [38] and expanded by de Saint Laurent et al. [22] to include two dimensions: moral quality (benevolent or malevolent) and power (strong or weak). In our study, we employ this framework and the resulting four archetypes: Hero (benevolent, strong); Victim (benevolent, weak); Villain (malevolent, strong); and Fool (malevolent, weak).

3 Methods

3.1 Data collection

To examine Ukraine’s use of memetic engagement during the first 8 months of the 2022 conflict, we begin with a large collection of tweets matching war-related keywords (see Appendix A) spanning from 27 February 2022 to 12 October 2022. As outlined in Section 1, we are particularly interested in the analysis of official, government-run accounts, to study their unprecedented role in a large-scale conflict. Thus, we identify the most popular (in terms of number of followers) non-personal accounts ran by the Ukrainian government:

- (1) **@Ukraine**: self-described as “the official Twitter account of Ukraine”, this verified account is labeled as an “Ukraine government organization” by Twitter. It posts war-related and political updates about Ukraine.
- (2) **@DefenceU**: or “Defense of Ukraine”, is the “Official page of the Ministry of Defense of Ukraine”. Like @Ukraine, this verified account is described by Twitter as a “Ukraine government organization”. Its posts concern battlefield updates and news.

These two accounts have often employed memes to support their message. Our primary research goal is to investigate this activity. Besides these two official accounts, we identify another non-personal account with an important role in terms of memetic warfare:

- (3) **@uamemesforces**: self-described as the “Source of the best Ukrainian memes”, this unverified account states as its goal “to convey the truth about what is happening now in Ukraine with the help of memes”.⁷

This meme page was created in February 2022, when Russia began the invasion of Ukraine. A preliminary analysis of our data confirms that, by the end of our collection period, this account was

⁷<https://www.patreon.com/uamemesforces>

the most followed Twitter page entirely dedicated to creating and spreading pro-Ukrainian memes about the conflict. Scholars have highlighted how @uamemesforces “is followed by academics, journalists from leading publications, diplomats from a host of countries, communications advisers to world leaders and social media managers at multilateral institutions”, and thus how its memes have the potential to “reach global elites”.⁸ For these reasons, we choose to focus our in-depth analysis and labeling efforts on these three accounts, which represent the most prominent examples of memetic warfare in the Ukrainian war.

We retrieve all of the posts by these accounts via Twitter API. We further use the Search API to retrieve the number of retweets of each post at a later date (January 2023), so that this number has stabilized. As explained in Section 2, memes are defined not by any intrinsic characteristic, but by their ability to be reshared. As such, we treat each post with visual media content as a potential meme, irrespective of the format. Section 3.2 illustrates how we incorporate the features that previous research has identified as typically associated with viral memes into our analysis. To ensure preserving content posted by these accounts, we also download the images associated with the posts and the users who retweeted or liked the post. We associate each user with a country in the GeoNames database⁹ via their location field (a free-text field optionally filled by the user). Considering only tweets with media, out of 586 042 users, 44% are successfully mapped to a country. For geographic analysis, we retain only countries with at least 1000 geolocated retweeting users.

Considering all the tweets obtained through the Search API, @Ukraine consistently garners the highest level of engagement despite having only 54 tweets. This account’s tweets achieved an average of 9670 retweets per post, significantly exceeding @DefenceU (750 retweets per post) and @uamemesforces (531 retweets per post). @DefenceU is the account with the lowest number of retweets, however, its audience is more dedicated and engaged. In fact, on average a user retweets @DefenceU approximately 3.60 times, compared to 2.38 for @uamemesforces and 1.18 for @Ukraine. Among the three Twitter accounts analyzed, @Ukraine exhibits the broadest geographical reach, with a significant portion of retweets originating from outside of Ukraine. Specifically, 27% of retweets are from the United States, 0.9% from the United Kingdom, and 0.8% from Brazil, while only 0.013% from Ukraine. @uamemesforces garners the highest engagement from users in Japan (19%) and the United States (17%), while Ukrainian users contribute only 0.05% of retweets. In contrast, 14% of retweets for @DefenceU originate from Ukraine, with the remaining retweets primarily from the U.S. (14%) and Japan (13%).

To explore the relationship between Ukrainian memes and each country’s actions in the conflict, we use the Ukraine Support Tracker data from the Kiel Institute for the World Economy,¹⁰ which “quantifies military, financial, and humanitarian aid transferred by governments to Ukraine since the end of diplomatic relations between Russia and Ukraine on January 24, 2022”.¹¹ The dataset provides a quantitative measure of a country’s aid to Ukraine. Finally, to examine the attitude of each country’s population towards Ukraine, we use the Eurobarometer STD98 survey, conducted in the winter of 2022-2023 in 30 European countries [26]. Specifically, we use use QE2, answer 6: “To what extent do you agree or disagree with: providing financial support to Ukraine”.

3.2 Annotation

To understand the visual and narrative attributes of the captured data, all authors perform a manual annotation of a selection of the tweets (for a total of 1063 tweets, see Table 1). We consider all tweets with media as candidates for the annotation. We remove self-replies (an account replying to

⁸<https://digdipblog.com/2022/08/08/the-power-of-memes-analyzing-war-time-messaging/>

⁹<https://www.geonames.org>

¹⁰Downloaded on 20 June 2023.

¹¹<https://www.ifw-kiel.de/topics/war-against-ukraine/ukraine-support-tracker>

Table 1. Statistics for the dataset of annotated tweets. The averages are computed only on the subset of annotated tweets. The reported number of followers is the last available one.

user	tweets	avg. likes	avg. RT	num. followers
@Ukraine	51	33 065.50	6245.68	2 274 857
@DefenceU	555	5316.44	852.14	1 557 170
@uamemesforces	457	6750.05	637.33	321 297

its own post) which have no retweets of their own. For the two accounts with a large number of tweets we consider a random sample (31.6% for @DefenceU and 39.4% for @uamemesforces): this way, we balance the labeling effort with the number of annotated tweets. Instead, given the small number of tweets with media by @Ukraine, we annotate all the available tweets that pass the filters (94.4% of @Ukraine tweets with media). Summary statistics for the dataset of annotated tweets (after filtering and sampling) are shown in Table 1.

Guided by previous literature, we compose a codebook to capture visual and narrative aspects of the content. First, we adopt the significant features of the content characterization framework proposed by Ling et al. [43], which is used to model the virality of political memes. These include:

- (1) Number of panels: multiple or single;
- (2) Type of image: photo, screenshot, or illustration;
- (3) Scale: close-up, medium shot, long shot, other;
- (4) Type of subject: object, character, scene, creature, text, other;
- (5) Main attribute of the subject (if character): facial expression, posture, poster
- (6) Character emotion (if character): positive, negative, or neutral;
- (7) Contains words: whether the image itself (not tweet caption) contains words.

Second, we employ the narrative framework by Karpman [38] later extended by de Saint Laurent et al. [22]. It defines two key dimensions to characterize a narrative: moral quality (benevolent or malevolent) and power (strong or weak), thus resulting in four character archetypes: heroes (benevolent, strong), victims (benevolent, weak), villains (malevolent, strong), and fools (malevolent, weak). We also label the actors mentioned in these narratives (as an open category, guided by an initial set of main actors in the conflict, such as Zelensky and Putin). We further contextualize the narratives by their emotional appeal [11], in particular humor, pride, fear, outrage, and compassion (the label was originally open, but the labelers coalesced on this set). Finally, we record whether there is a specific intent in terms of a call to action or information sharing. Note that, as much as possible, we label the narratives as intended from the point of view of the posting account (although some biases may creep in, as we discuss in Section 5). Appendix B reports the full codebook.

Using this literature-driven codebook, all authors of this manuscript perform a deductive coding of the selected posts (for more on such process, see Linneberg and Korsgaard [44]). We rely on Google Translate to translate non-English language content. All coders first annotated a selection of 50 posts, discussed insights and disagreements, and came up with standard labels for emotional appeal, intent, and actor (though the labels remained open). During this process, the coders referred to the coding schemes from previous literature [43] to adapt the concepts to the task at hand. In particular, identifying the presence or absence of an actor proved to be a challenging task. For instance, @DefenceU posts often represented people who may be civilians, combatants, or officials, whose roles may be difficult to distinguish (indeed, this label had the lowest agreement). Concerning the narrative, often disagreement arises on whether one or several narratives are present in the content. For instance, when both parties are represented (or even implied), an image can be labeled as focused on the perceived victim, the villain, or both. The remaining posts were then annotated by one author

each, with uncertain cases discussed collectively. Appendix C reports the inter-annotator agreement on a selection of 48 posts (sampled from each account). On average, the agreement (as measured by Krippendorff's alpha) is 0.750 for @Ukraine, 0.604 for @uamemesforces, and 0.667 for @DefenceU (with an average value of around 0.65 for all accounts and tasks). These values are high considering the subjective nature of the tasks and that some features were open and could have multiple labels. The annotated dataset is available at <https://github.com/corradoMonti/ukranian-tweets>.

3.3 Popularity model

Using the data collected and annotated as per the previous sections, we construct models to find whether there are trends in (i) what kind of content becomes more popular, and (ii) whom it reaches. Specifically, we use linear regression to model the log-transformed number of retweets via the content variables outlined above as independent variables. As the three accounts have posted a vastly different number of tweets, we weight the data points (tweets) inverse-proportionally to their representation in the dataset, so that each account weighs equally on the outcomes. We build several versions of the model to determine the benefit of each class of variables: visual [43], actor, emotion, intent, and narrative. A baseline model simply uses the identity of the author (i.e., group fixed effects), the (log-transformed) number of followers of that account at the time of posting the meme (available thanks to the Streaming API v1.1)¹² as a proxy for the potential audience, and the number of days from the invasion to account for changing interest in the topic over time. We then select the best model according to the Bayesian Information Criterion (BIC), a model selection metric that captures both unexplained variation in the dependent variable and the number of explanatory variables. The coefficients of this model show the associations between these variables and the virality of the content.

Next, we group the retweets by the country of the user (those that can be geolocated) to understand the preferences of users who spread these messages. We enrich our characterization of the relationships between each country and Ukraine by considering the amount of financial assistance the country provides to Ukraine (normalized by the country's GDP) or by their opinion on this topic (via the Eurobarometer survey). As mentioned above, we consider these measures as different proxies of support for Ukraine's military efforts. First, we examine the retweet rate (normalized by population). Then, we compare the narrative preferences to the amount of support for Ukraine: are there narratives associated with increased support for Ukraine? We operationalize this preference as the log odds of users from a country retweeting content with a specific narrative. Finally, we examine the geographical distribution of the retweet rate and narrative preferences across Europe by plotting theme on choropleth maps.

4 Results

We now analyze the annotated dataset to answer our research questions. Section 4.1 outlines which narratives, emotions, and actors are more prevalent in each of the three accounts. Section 4.2 examines which factors are more important for the virality of this content. Finally, Section 4.3 shows how the audience breaks down across countries.

4.1 Content analysis

Figure 1 outlines which elements are more prevalent in each of the considered accounts. @Ukraine uses a mix of narratives but completely avoids the fool one: portraying the enemy as incompetent is counter-productive when asking for help, something already known since World War 2 [34]. Unsurprisingly, @DefenceU focuses on the heroic narrative by portraying the Ukrainian military as

¹²<https://developer.x.com/en/docs/x-api/v1>

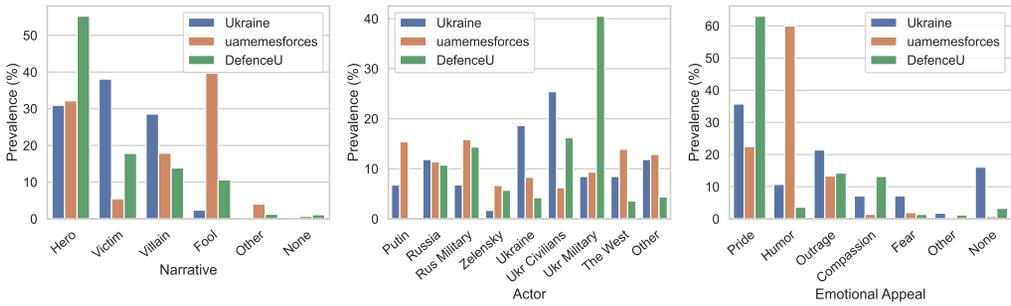


Fig. 1. Distributions of narratives, actors, and emotional appeals in the messages of the three Ukrainian accounts. The accounts show marked differences in their communication strategies. @Ukraine uses a balanced mix, with an emphasis on civilians, but rarely makes fun of their target. Conversely, @uamemesforces favors the fool narrative and uses humor generously. Finally, @DefenceU favors the hero narrative, mostly evoking a sense of pride for Ukraine’s military.

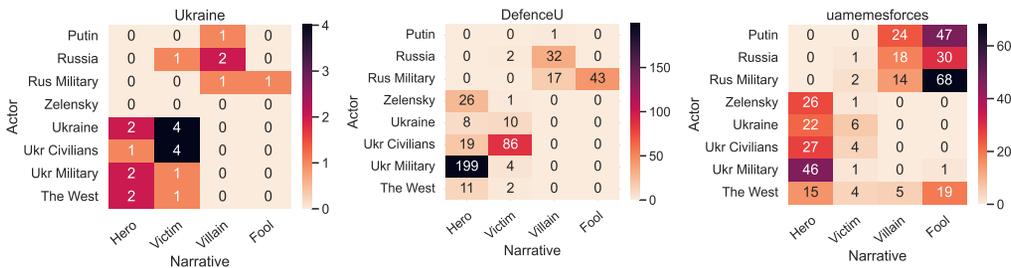


Fig. 2. Distributions of co-mentions of specific actors along with a given narrative for the three Ukrainian accounts. The block structure of the 2d distribution shows that malevolent narratives are used when talking about Russian actors, while benevolent ones when talking about Ukrainian actors. @Ukraine puts an emphasis on the victim narrative and @DefenceU has just two clear narratives that emerge (Ukrainian military as heroes and civilians as victims). @uamemesforces instead focuses most of its messaging on making fun of and Russian military and Putin.

powerful and morally good. Instead, the independent account @uamemesforces focuses on the fool narrative, in which the other side is portrayed as weak and incompetent. At the same time, the account employs a large amount of hero narrative, partially due to numerous memes of Ukrainian farmers resisting Russian soldiers.

The actors mentioned by each account also differ. @Ukraine focuses the most on Ukrainian civilians and Ukraine as a nation, possibly working towards the goal of nation-building, or to keep the global public opinion focused on the main victims of the conflict. @DefenceU puts the spotlight on the Ukrainian military but also dedicates some attention to Ukrainian civilians. Interestingly, Putin is mentioned by @uamemesforces much more than @DefenceU, which almost never mentions him. Similarly, actors associated with the West are mentioned more by @uamemesforces, while @DefenceU tends to ignore them. This result is in line with the more domestic focus of @DefenceU we previously highlighted. The three accounts are also distinct in their emotional appeal: @Ukraine communicates largely outrage and pride, and @DefenceU pride alone, while @uamemesforces focuses on humor. Concerning possible intents (not plotted), @Ukraine posts by far the most calls to action (31%), with many calls to support the war effort, while @DefenceU often posts purely informational material (e.g., war statistics bulletins).

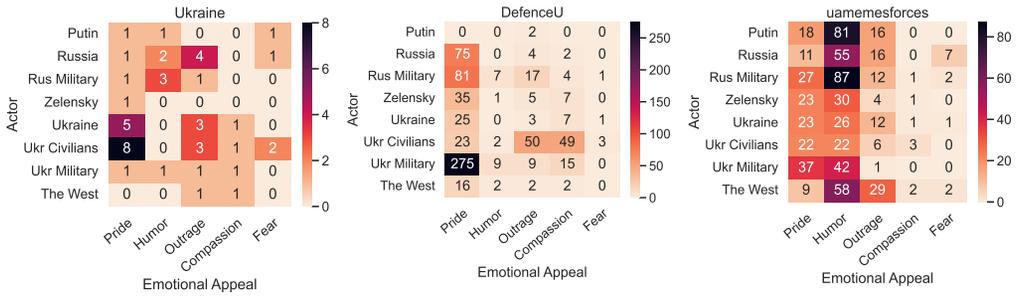


Fig. 3. Distribution of co-mentions of specific actors along with a given emotional appeal. While clear preferences for distinct rhetorical styles are present across the accounts, they do not show consistent associations with specific actors.

Next, we examine the role of the main actors found in the narratives. Figure 2 presents heat maps of the intersection of actors and narratives for each account. Across all three accounts, the dichotomy between the two sides of the conflict is stark: the Russian side is portrayed as either a fool or a villain (negative moral quality), while the Ukrainian side as either a hero or a victim (positive moral quality); the West spans both portrayals. The emphasis by @Ukraine is on Ukrainian civilians and Ukraine as a whole, portrayed in a victim role, whereas less emphasis is put on the actors from the Russian side. The narrative of @DefenceU consists of several well-defined combinations: Ukrainian military as heroes, Russian military as fools, and Ukrainian civilians as victims. It is noteworthy that the Russian military is often portrayed as powerless (fool), while Russia as a country is always portrayed as a villain (powerful), and Putin himself is rarely referred to personally (contrary to the practice of @uamemesforces). Unlike for @Ukraine, @uamemesforces content rarely portrays Ukraine as a victim, even when civilians are involved, and instead attributes them heroic characteristics: recurring examples of this combination are references to a story of Ukrainian farmers who towed away abandoned Russian tanks. The most common combination of actor/narrative is portraying the Russian military as a fool (negative moral quality and powerless), and a similar treatment is reserved for Putin. This powerless portrayal is much more common than the powerful one (in the role of a villain), and is in line with the satirical role of the account. Actors from the West are mostly portrayed as fools or heroes—mostly depending on whether the particular actor is perceived as supporting Ukraine—but in a few cases, they are even portrayed as villains, when they are perceived to be complicit with the enemy. Russian civilians are only briefly mentioned by @uamemesforces (omitted from the plots, 20 tweets in total), often as “conniving” with their government. Thus, we find distinctive narrative styles in the three accounts: @Ukraine focuses on its victimhood, @DefenceU propagates a heroic, powerful view of its own military while portraying its enemy in a foolish, weak perspective, and @uamemesforces emphasizes the weakness (foolishness) of the Russian side and the heroism of Ukrainians.

Finally, Figure 3 shows the emotional appeal of the posts in each account. Since a tweet labeled with a single emotion could refer to several actors, in our analysis we associate such emotion with all the actors present. The emotional focus is very well defined for @DefenceU, where pride is present in an overwhelming majority of the posts. A combination of humor dominates the messaging style of @uamemesforces. Interestingly, Putin is singled out as an object of ridicule only by @uamemesforces. No clear pattern is present for @Ukraine.

Examples. Let us now focus on the narratives, given their explanatory power that we explore in Section 4.2. Here we give a few examples to crystallize the imagery associated with each narrative. These examples are taken from the top 10 most popular tweets per narrative.



Fig. 4. Examples of 'hero' narrative.



Fig. 5. Examples of 'victim' narrative.

Figure 4 illustrates examples of 'hero' narrative. This narrative typically depicts the heroism of Ukrainian military efforts, emphasizing their benevolent moral quality and their power to defeat the enemy. In fact, it is often associated with mentions of Ukraine as an actor. A typical example is the leftmost in Figure 4. This meme was posted by @uamemesforces on 28 March 2023; in the same days, Ukrainian officials were claiming victories against the Russian military forces, such as recapturing the city of Irpin. This meme by @uamemesforces supports the same narrative, but uses typical memetic language: the famous moment at the 2022 Oscars when Will Smith slapped Chris Rock is edited to show Zelenskyy fiercely slapping Putin. This tweet obtained more than 120k likes. Other content, typically by @Ukraine, promotes information on ways to support Ukraine, or celebrates the valor of its civilians, such as their victory in the Eurovision Song Contest, and connects it to the "brave freedom defenders" on the frontline (middle example in the figure). @DefenceU focuses on the military: not-so-veiled military threats (e.g., a video of a military toy floating by the Kerch bridge, not shown) and references to an online social-media organization dedicated to countering Russian online propaganda (the "North Atlantic Fellas Organization #NAFO", rightmost image).

Figure 5 shows examples of 'victim' narrative. The most popular are from @Ukraine: for example, the news that the world's largest plane, the Ukrainian Antonov AN-225, had been destroyed by the Russians is used to illustrate the suffering of the war (leftmost). Another example is shown in the

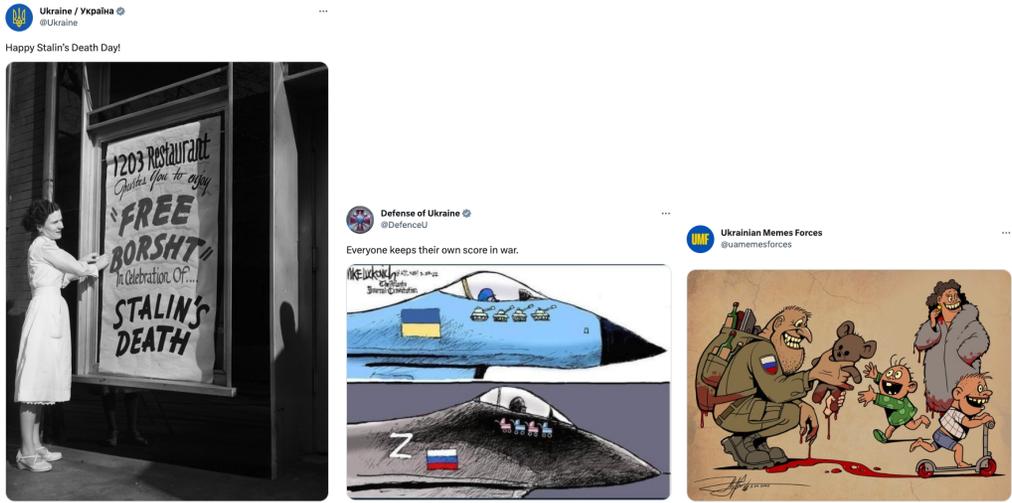


Fig. 6. Examples of ‘villain’ narrative.



Fig. 7. Examples of ‘fool’ narrative.

middle, where @Ukraine uses a photo of civilians whose house was destroyed by Russian bombings in the city of Chuhuiv, near Kharkiv, to ask the global public opinion support for a specific economic sanction against Russia—the closing of seaports. The last example is from @DefenceU, who uses a poster-like design to promote their view that Russia’s actions in Ukraine fall under the definition of genocide—something that the United Nations is currently investigating and that at the time of this writing is inconclusive.¹³

Figure 6 displays examples of ‘villain’ narrative, focusing on the immorality of the enemy and its strength. They often juxtapose Russia with Ukraine and emphasize the civil casualties incurred by

¹³<https://www.reuters.com/world/europe/russia-has-committed-wide-range-war-crimes-ukraine-un-inquiry-finds-2023-03-16>



Fig. 8. Examples of ‘other’ narrative.

Ukraine. Negative historical associations (e.g., with Stalin or Hitler) are often employed (leftmost example). Extreme evilness is attributed to the Russian army: in the middle example, a cartoon shows Russian pilots gloating as they kill Ukrainian babies, a typical atrocity frame employed in propaganda [16]. Even Russian civilians are portrayed as evil: in the rightmost example, they happily play with bloodied toys taken from Ukrainian children.

Figure 7 presents examples of ‘fool’ narrative, which still focuses on the immorality of the enemy, but depicts it as weak and powerless. Many examples focus on Russia’s failed imperial ambitions and portray its soldiers as incompetent and ineffective, as is typical of psychological warfare.

Tweets containing ‘other’ narratives are often associated with humorous posts and mention actors other than Ukraine, bearing an ambivalent attitude. Figure 8 shows several examples, which provide a commentary on the international response to the war, the sanctions, and the role of Hungary in the conflict. While an overarching narrative can be inferred, it does not clearly fit within the narrative framework we employ.

Table 2. Difference in Bayesian Information Criterion score from the baseline model (ΔBIC) and Adjusted R^2 for models with different sets of features. More negative ΔBIC indicates better models. The best model within the group is bolded.

Model (feature groups)	ΔBIC	Adj. R^2
Baseline	0.000	0.557
Visual	-10.080	0.599
Actor	28.257	0.562
Emotion	-21.336	0.578
Intent	3.480	0.560
Narrative	-28.299	0.580
Visual+Actor	7.230	0.608
Visual+Emotion	-18.052	0.613
Visual+Intent	-5.742	0.602
Visual+Narrative	-43.965	0.622
Visual+Narrative+Emotion	-38.423	0.631
Visual+Narrative+Emotion+Intent	-26.376	0.630

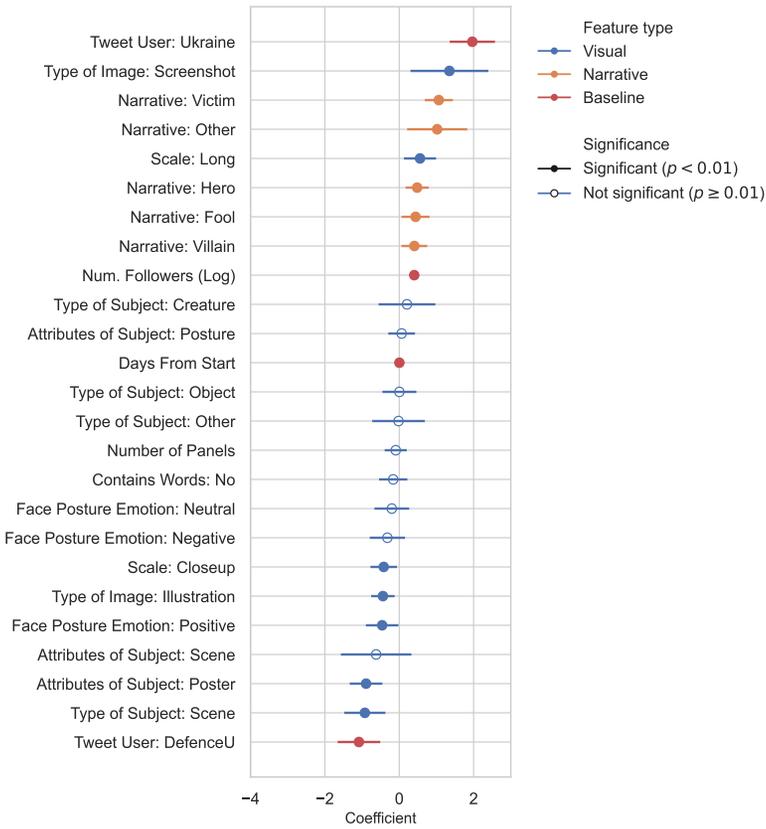


Fig. 9. Coefficients and 95% confidence intervals for the best-performing regression model of the popularity of a meme. The target variable is the \log_2 of the number of retweets: a coefficient of +1 indicates that a feature is associated with double the amount of retweets.

4.2 Modeling popularity

We now move to modeling the popularity of the content shared by these accounts. We use a linear regression with the content variables outlined above as independent variables. Categorical variables are represented via one-hot encoding with “none” as the reference case. The target variable—the number of retweets—is transformed via log base 2, such that a coefficient of 1 indicates the number of retweets doubles, while a -1 that it halves.

First, we investigate which types of features are more effective in predicting the success of a tweet. Table 2 shows the performance of models with different subsets of explanatory variables. All the models include the ‘baseline’ variables: the identity of the author, the number of their followers, and the number of days since the invasion. The other models use the variables grouped as visual, actor, emotion, intent, and narrative. We perform model selection based on the BIC. The best-performing model uses visual and narrative variables, and achieves $-43.9 \Delta\text{BIC}$ compared to the baseline, with an Adjusted R^2 of 0.622. Remarkably, the Narrative model alone (with only 5 features) performs almost as well as the Visual model (state of the art with 16 features) in terms of Adjusted R^2 , and even better in terms of ΔBIC . This result demonstrates the importance of narratives in determining content popularity.

Figure 9 shows the coefficients for the best model and their 95% confidence intervals. As the non-governmental @uamemesforces account is treated as a “base case”, we find @Ukraine to have comparatively more retweets and @DefenceU comparatively fewer. We have also computed a similar model in which, instead of using indicator variables for each account, we use a dummy variable for whether the author of the tweet is a government account (@Ukraine or @DefenceU). In this model, this indicator variable is not significant (the effects of the two government accounts cancel each other out). This result suggests that, at least for the present dataset, the popularity of the message is not necessarily tied to whether the account is governmental, but rather to the identity of the account and to the content itself.

Among the variables with the largest positive coefficients, we find the narrative-related variables, with the addition of screenshot types of images and long-scale images. The fact that all narrative coefficients are positive implies that the very existence of a narrative is associated with greater engagement in terms of retweets. Out of these narratives, the *victim* one is the most resonant: its presence is associated with an increase of 109% in the number of retweets. The other benevolent narrative, *hero*, is associated with an increase of 40%. The two narratives focusing on the malevolence of the enemy are instead less powerful (36% for *fool* and 32% for *villain* respectively), but still positive compared to the absence of a narrative.

Further, less popular content includes posters and illustrations, as well as those of scenes (as opposed to people or objects). However, one visual variable that is positively associated with retweets is whether it is a screenshot: these are often calls for donations and screenshots of news.

Overall, we find narratives to be a powerful predictor of the amount of attention a piece of content receives, while emotional and actor variables have a smaller impact.

4.3 Audience characterization

As summarized by the words of the creator of @Ukraine reported in Section 1, one of the main goals of Ukrainian memetic warfare is to reach “large and distant target audiences”, and thus influence global public opinion. In this section, we study which countries are more receptive to Ukrainian messages, and how this receptiveness relates to narratives and to each country’s actions.

To examine the variability in responses in different countries, we look at retweets and geo-locate their authors to a country (see Section 3). We normalize the number of retweets in a given country by its population size, thus obtaining a measure of retweets per capita. This measure is a proxy for the amount of interest in Ukrainian tweets in a given country from the general population. Note that we choose not to adjust the measure by Twitter or Internet usage within the country: we do not wish to equalize countries where Twitter is more likely to be read with those where it is not. Indeed, Twitter usage is a mediator between the effect of Twitter messages and the support in the general population for sending aid to Ukraine, i.e., the larger the penetration of Twitter, the more chances people in that country might be exposed to Twitter propaganda messages, all other factors being equal. Therefore it would be incorrect to include it as a control variable in this case [55, 56]. Finally, to measure the support of a country to Ukraine in the conflict, we use economic and military assistance normalized by GDP, as explained in Section 3.

Figure 10 shows these two measures, retweet per capita and assistance, for each country present in our dataset. First, countries bordering Russia have a high retweet rate per capita (above 0.5‰): their attention to Ukrainian tweets is very high. In particular, Baltic countries bordering Russia—Latvia, Lithuania, and Estonia—are outliers both in terms of assistance to Ukraine and of engagement. Finland, also bordering Russia, presents a high retweet rate, albeit with a lower aid statistic. Other European countries—northern (Norway, Denmark), eastern (Czech Republic, Poland), and English-speaking ones (Ireland, United Kingdom)—have high rates of retweets per capita (above 0.3‰).

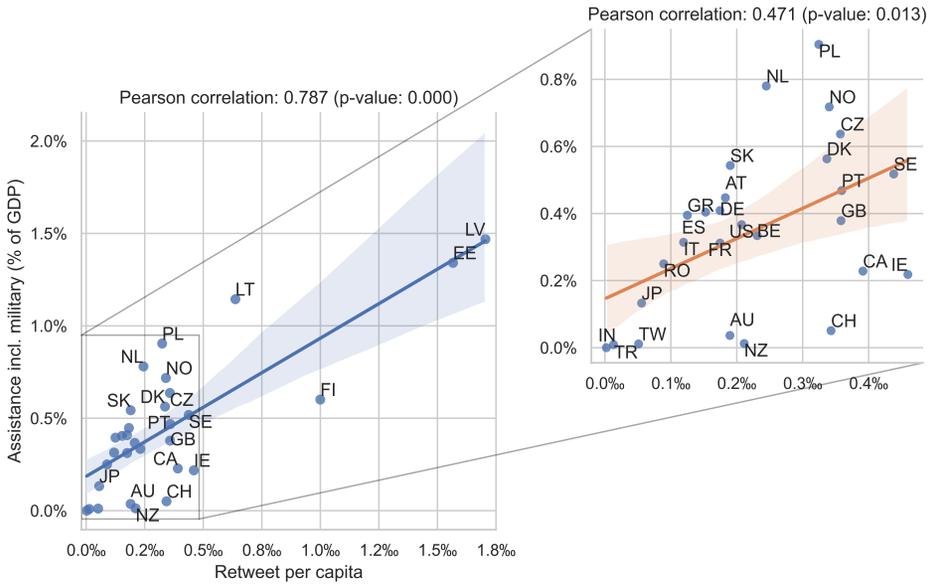


Fig. 10. Number of retweets normalized by population, versus total bilateral and EU assistance to Ukraine, by country. Left: all countries. Right: excluding outliers. Both the regression including outliers (blue) and the one excluding them (orange) are significant at $p < 0.05$. The shaded area represents the 95% CI.

The scatterplot shows a strong correlation between the two variables: with a Pearson coefficient of $\rho = 0.787$, attention to Ukrainian content on Twitter is related to the economic and military assistance that the country provides to Ukraine. While European countries bordering Russia are outliers, the correlation after excluding these countries is still a reasonably strong $\rho = 0.471$ ($p < 0.05$, inset on the right). Thus, the relationship extends beyond the countries bordering Russia. To check for robustness, we also perform this analysis disaggregating by account, and the results are in line with those shown in Figure 10 (coefficients > 0.7 and p-values $< 10^{-4}$). In other words, picking any subset of the accounts would not change this result, which therefore seems more general than our particular choice of accounts. Conversely, the same analysis using the Eurobarometer survey shows no significant correlation between retweets per capita and popular support for financial aid to the Ukrainian military effort: the correlation we observe between the effectiveness of Ukrainian content on Twitter and the action of governments does not translate when considering support from the general population.

Let us now break down the geographical resonance of each narrative. To this end, we compute the odds of users in a country to retweet content with a particular narrative. An important working assumption in this analysis is that geolocation is not biased compared to the narrative that is shared, i.e., there is no reason to believe that a person who geolocates themselves prefers one narrative over another. Figure 11 shows the retweet rate and the per-narrative log odds for each country in Europe. The focus on Europe derives from it being the geographic location of the conflict, and because it is the continent with the highest retweet rate per capita (0.3‰, compared to 0.09‰ in North America, which is second). The leftmost map shows the retweet rate normalized by population, which shows again the outsized attention by the Baltic countries and Finland. The retweets by narrative (rest of the plots) often present a gradient from West to East: countries closer to Russia are more likely to

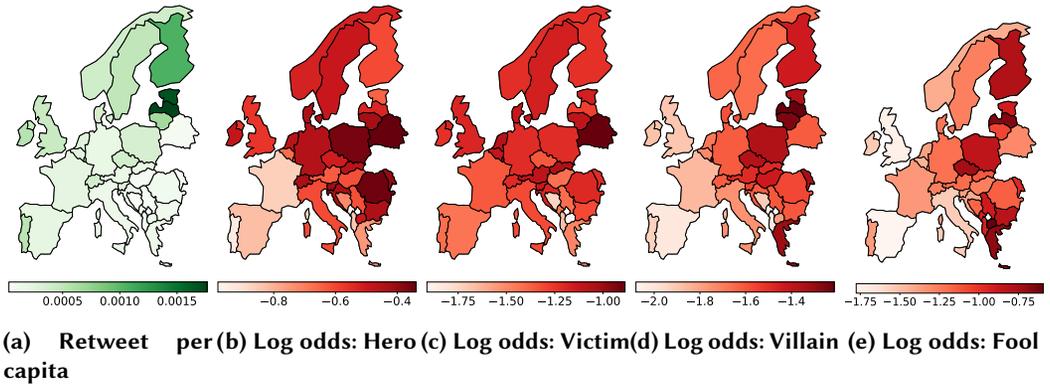


Fig. 11. Geographic distribution of retweets. Left: retweet rate per capita in a country. Rest: log odds of retweeting content with a specific narrative.

retweet posts with a discernible narrative. There are, however, some differences across narratives: the ‘victim’ narrative presents a more uniform geographical spread. In contrast, countries bordering Russia and Ukraine—such as Finland, Poland, and the Baltics—show a stronger resonance with the ‘villain’ narrative focusing on the Russian threat. Baltic states are particularly inclined to amplify the archetypes of villain and fool, interestingly together with Greece (we may speculate that historical parallels with Cyprus’ experience of the contested occupation by Turkey may provoke attention from Greece). The ‘hero’ narrative appeals both to some Western countries—such as the United Kingdom, a strong supporter of Ukraine in the war—as well as those closer to the conflict, including Belarus. Note that Belarus has strong regulations on social media access, thus those who express themselves on social media are likely to be either those in opposition to the government and dedicated enough to use a VPN, or those who are allowed by the government on the site.¹⁴

Finally, we investigate how the receptiveness to different narratives is related to the actions of a country’s government in the conflict. Figure 12 shows the relationship between these log odds and the total assistance to Ukraine. The total aid provided by each country to Ukraine, normalized by GDP, is positively associated with the spread of malevolent narratives (villain) and negatively associated with benevolent ones (hero). We perform the same analysis also disaggregating by account, to evaluate which findings are account-specific and which are robust to our particular choice of accounts to study. In doing so, we wish to consider also the levels of uncertainty of each log odds ratio (i.e., our measure of narrative effectiveness) in each country: a low number of retweets in a country implies that the associated odds ratio is uncertain. Therefore, we use a Bayesian approach through a Monte Carlo simulation with 100k samples, wherein the probability of a retweet is modeled as a Beta distribution with parameters taken from the data (number of retweets with a given narrative vs without that narrative). Table 3 reports the results of this analysis. While the results for some narratives are indeed account-specific, we find that the association between the effectiveness of the *villain* narrative and financial assistance to Ukraine is robust to the choice of the account. This finding seems to indicate that the focus on the evil character of the enemy is the most powerful narrative with governments that provide material assistance. Moreover, we find that these effects are mostly due to military aid, while humanitarian aid has no statistical relation to the resonance of narratives in our dataset (plots omitted for brevity). The Eastern European countries (closest to the conflict) both contribute the most aid (proportional to their GDP) and prefer the

¹⁴https://www.cyberghostvpn.com/en_US/privacyhub/countries-ban-social-media

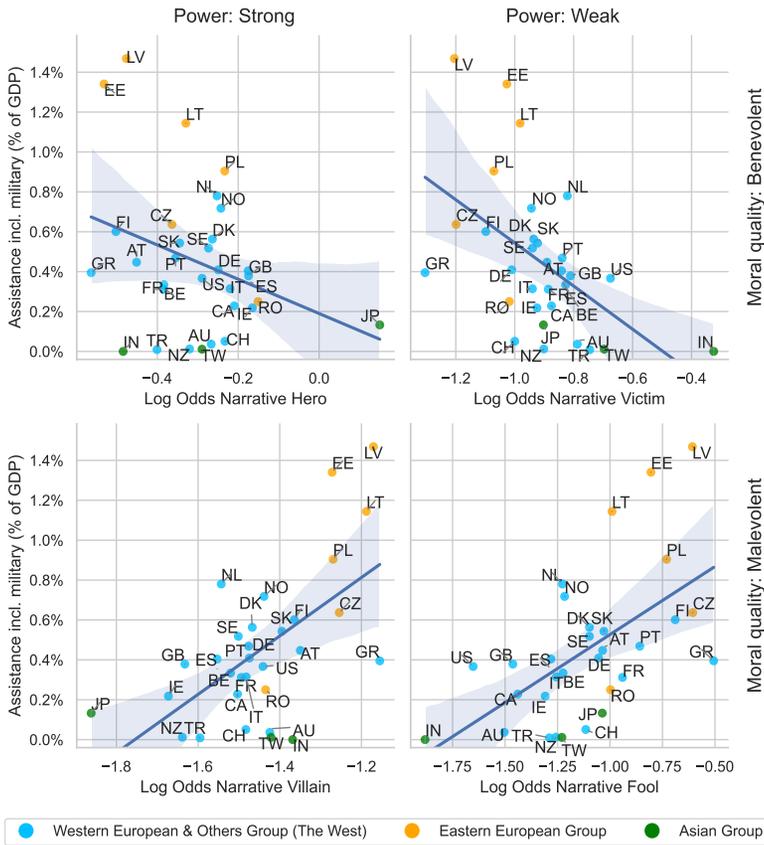


Fig. 12. Log odds of retweeting content with a particular narrative, versus total bilateral and EU assistance to Ukraine, by country. Colors indicate the United Nations Regional Group of each country. The shaded area represents the 95% CI. The columns indicate the power quality of the narrative, while rows indicate their moral quality. The assistance is driven mostly by military aid, whereas non-military aid has no significant correlations with narrative preference. Malevolent narratives that target the malevolent moral quality of the enemy are positively correlated with assistance.

villain narrative. Interestingly, the second most stable relationship is a negative one between the preference for the hero narrative and aid: the countries that contribute the most aid do not show a preference for content portraying Ukraine as a hero. Again, a similar analysis shows no significant results for the Eurobarometer survey.

5 Discussion

The results of this study have important implications for our understanding of online communication during war, especially within the theoretical frameworks outlined in Section 2. We observed the international reach of the content produced by Ukrainian sources, particularly in reaching Western countries where Twitter’s penetration is considerable. This reach is particularly noteworthy since exposure to political memes is known to be a mediator for political involvement [1]. Ukrainian war memes that contain clear narratives prove to be particularly popular, and in particular, tweets with a victim narrative receive twice the number of retweets. As social media is increasingly

Table 3. Pearson correlation between the log odds of retweeting content with a particular narrative versus total bilateral and EU assistance to Ukraine (only values significant at $\alpha = 5\%$ are reported), combined for all accounts, and for each account separately, as well as the total number of retweets involved in the computation (last column). These values account for the uncertainty in the estimates of the log odds of each country via a Monte Carlo approach with 100k samples, where the probability of retweet is modeled as a Beta distribution with parameters taken from the data.

	Hero	Victim	Villain	Fool	Total Geo. RTs
Combined	-0.312	-0.503	0.573	0.561	224 225
@Ukraine		-0.344	0.281	0.375	70 312
@DefenceU	-0.400	0.328	0.403		91 627
@uamemesforces	-0.400	0.613	0.430	-0.362	62 286

used to mobilize support [12, 75], narrative emphasis may be an important variable to control in communication campaigns. Interestingly, the overall audience preference for the victim narrative captured here does not translate into action. Instead, countries that contribute more to the Ukrainian war effort resonate more with the villain narrative. We can speculate that countries united by the common fear of a powerful enemy are more willing to make sacrifices to help in the conflict. Indeed, the countries that relate more to depictions of Russia as a grave threat are the ones closer to the conflict, such as the Baltic countries, given their shared history with and physical proximity to Russia. For instance, soon after the invasion, Finland reversed a long-time policy of neutrality and applied for membership to NATO on 18 May 2022, which was granted on 4 April 2023. In this sense, the propagation of the content posted by the examined Ukrainian accounts reflects politically relevant signals. At the same time, we find that while economic assistance in the military conflict is correlated with the spread of Ukrainian content on Twitter, it is not correlated with the popular support as measured by surveys (see Section 4.3). This unexpected result could hint at governments and engaged Twitter users being more in tune with each other than with the general population, and highlights the important role played by social media as a mediator for political activity. Whether a government’s decision to support another country monetarily is preceded by social media activity, especially involving state actor accounts, is a fascinating future work direction. In summary, messages focusing on the Ukrainian victims are more effective in reaching a larger audience; but messages focusing on the danger posed by the enemy resonate more with countries willing to offer material assistance in the war.

Our findings suggest several broader considerations on memetic warfare from the perspective of the general public. First of all, since meme creators are likely to be aware of the engagement created by victim narratives, it is important that audiences are able to respond knowledgeably. As such, it would be helpful to increase public awareness of such mechanisms to improve digital media literacy. Currently, the EU’s Digital Skills [69] include the skill to “critically evaluate the credibility and reliability” of information. We argue that evaluating the *intent* behind the posting of content should also be taken into account. Beyond individual literacy, reported grievances from warring governments should be verified by independent sources, such as supranational bodies and humanitarian organizations. As noted by Nyhan [53], in fact, while the efficacy of post-hoc fact checking is severely limited, it is important to prevent and disrupt the formation of linkages between group identities and false claims. In this regard, particular attention is needed towards misusing imagery or taking it out of context, a common misinformation tactic that adds to the perception of authenticity [27].

Another important observation is that messages hinting at the pure evilness of the enemy—e.g., portraying its soldiers as stealing toys from dead children—are particularly likely to spread in countries that are more involved in the war through economic and military support. While we cannot establish the direction of the causal arrow in this case—i.e., we do not know if involved countries are a more fertile ground for these messages, or if such messages encourage involvement in the war—this finding suggests memes might reflect the broader trend of increasing geopolitical tension. Group identification and unempathetic psychological responses toward the rival out-group have been found to drive media consumption and engagement [70]. Negative affect (which subsumes anger) has been recognized as a powerful mediator between “counter-empathy” and social media engagement [71]. As such, a high level of alert is particularly needed regarding the possibility of unsubstantiated wartime “atrocities” [8] in the countries involved in a conflict, lest they provoke the opposite reactions of either inciting the dehumanization of entire populations, or the desensitization of the public. There are historical precedents for this scenario: during World War I, all participating nations widely employed atrocity propaganda, which significantly contributed to fostering the surge of nationalism during the escalation of the war into a global conflict [16]. At the same time, such usage contributed to a growing discrediting towards reports of atrocities, that led to their diminished propagation during World War II [67].

These narrative elements have been previously well recognized as key in the media war. Other analyses of the Ukrainian war [51] also recognized the frames of good vs. evil and us vs. them as essential in understanding the narratives employed by both sides. The focus on the portrayal of the enemy in a ‘Cold-War’ style was already noted by Boyte [6] in 2017. In the Syrian War, while the evil vs. good framing was also employed by traditional media, Cole [14] noted also how the co-presence of different state-sponsored propaganda actors on social media created a confusing and fragmented information environment.¹⁵ That said, not all usage of narrative-laden messaging campaigns on social media is necessarily problematic. Indeed, social media has been used to build collective identities through narratives, which has been widely studied in the field of human-computer interaction. For instance, Das and Semaan [19] postulate that using a social media platform as a space to address colonial grievances and build a self-concept fosters a “narrative resilience”, which serves as a “collective and reflexive mechanism through which people work to generate resilience”. During the annotation process, we encountered many examples of Laenui [40]’s decolonization phases, including mourning, commitment, and calls to action, and yet other emphases may be present as the conflict develops and hopefully ends (although applying the decolonization lens to the Russo-Ukrainian conflict is controversial). Further, narrative building has been used to engage marginalized communities [57], strengthen their voice [31], and promote conflict reconciliation [65]—tasks that may become important towards the end of the conflict. Recently, social media narratives have been shown to provide Ukrainian refugees with a way to maintain an “emotional, dynamic, and constantly updating” bond with their homeland, and a space for the “negotiation and performance of ethnic and national identities” [39]. Given the importance of narratives found in this study, examining their role in establishing the self-concept of Ukraine as an independent state, and its relationship with Russia, the West, and the rest of the world, is an important future research direction.

Another crucial actor impacting the spread of propaganda is the platform, its algorithms, and its moderation policies. The spreading of memes on social media is largely influenced by platform-specific algorithms, which curate content based on user interactions and engagement patterns. It is well-known that content with negative sentiment attracts more engagement from users on social media [7, 32, 59, 68]. These engagement signals are used by the platform algorithms to decide which

¹⁵<https://www.theatlantic.com/international/archive/2018/03/cnn-effect-syria/554387>

content to promote: this algorithmic amplification ultimately determines what becomes viral [24]. Some content may achieve virality across multiple platforms, creating a cross-platform dynamic in which memes spread through different online ecosystems, influenced by different algorithms and user demographics [63]. This self-reinforcing loop not only accelerates the spread of memes [43], but also shapes the cultural relevance of the memes themselves, as they evolve and adapt to the affordances of each platform [52]. At the same time, states and other institutional actors can mobilize their social, economic, and cultural capital to meddle in the work of algorithms: as such they have the means to further their long-term strategic vision [5]. Shifting our attention to moderation policies, we did not find any evidence of moderation of the content posted by the examined accounts. However, platforms have been known to moderate political speech. It has been revealed that Facebook maintains a blacklist of “organizations with a record of terrorist or violent criminal activity”,¹⁶ whereas Twitter has been known to “shadow ban” users by limiting the propagation of their content.¹⁷ The application of moderation actions, especially when applied to political speech, is often controversial. In this sense, platforms’ choices can shape public discourse: this selective moderation, or lack thereof, is a political choice that warrants critical reflection, especially in the context of memetic warfare. The platform’s behavior is a black box: the algorithms are proprietary to the respective platforms and are not available for direct scrutiny, and the moderation decisions are made by in-house teams, possibly assisted by more algorithms. Barring voluntary greater transparency by the platforms, systematic auditing of the political speech and its reach may be the only way to provide a clearer view of the role platforms play in the spread (or limitation) of propaganda.

Limitations. The insights of this study are limited by its unique setting: the period of the Russian invasion, the selection of the accounts, the peculiar affordances the Twitter platform presents, and the uneven adoption of the platform around the world. Nevertheless, our finding that users from countries that support Ukraine the most also retweet the content from major Ukrainian accounts at a higher rate suggests a synchrony between the supportive actions of Twitter users (which could be considered to be a form of slacktivism [42]) and concrete financial decisions by the governments of their countries. Conversely, the responses to the Eurobarometer survey question on sending help to Ukraine had no significant correlation with the propagation of the Ukrainian tweets, which suggests that users on Twitter are more “in tune” with the decisions of their governments than a representative population sample. It is possible that the retweeters include government entities and actors, as well as automated accounts. Unfortunately, we are unable to apply the standard bot detection techniques [20] as the platform closed down its API in early 2023. However, the effect of such accounts should be limited in a retweet analysis, as each account may retweet a post once at most. Our methodology focuses on the popularity of content posted by official accounts. However, the propagation of memes often takes place organically, outside of the retweet mechanism, potentially spreading to other communities. To study the larger evolution of the memes captured in this study, the content and its sources and derivatives should be tracked across accounts, and, ideally, across platforms. Especially since it has been shown by Zannettou et al. [77] that memes often originate in fringe communities, capturing the evolution of the narratives within such communities (e.g., conspiracy theories) and their interaction with propaganda would be an exciting future research direction. Finally, although we strove to identify the narratives from the points of view of the posters, some subjective labels might have been affected by the opinions of the authors (see Positionality Statement below). We hope that sharing this data will allow for reproducibility and an easier re-examination of the narrative extraction process.

¹⁶<https://theintercept.com/2021/10/12/facebook-secret-blacklist-dangerous>

¹⁷<https://mashable.com/article/twitter-visibility-limited-labels-roll-out>

As the content analyzed here was produced during wartime, special care must be taken in an ethical analysis of the potentially sensitive material it contains, despite being published by verified accounts. For instance, @DefenceU posts a variety of reporting from the battlefields, including active war situations involving the use of artillery, images of combatants, and distressing images of affected civilians. We assume that the posting accounts have obtained sufficient permissions to post images of individuals, and have performed intelligence analysis to make sure no important information is being disclosed. Nonetheless, in this manuscript, we endeavored not to use any specific names or locations. Further, by its very nature, this data contains images of and information on vulnerable populations (including children), some of which are affected by the conflict. Special care must be taken that their images and other contexts are not misused and that any harm is minimized in subsequent analyses.

Positionality Statement

All authors reside in “the West”, and one was born in Russia. Although it is our aim to conduct the analysis as objectively as possible, this positionality is likely to affect some interpretation of the results. We make the data and annotations available for the examination of the research community, and possible alternative interpretations.

References

- [1] Saifuddin Ahmed and Muhammad Masood. 2024. Breaking Barriers With Memes: How Memes Bridge Political Cynicism to Online Political Participation. *Social Media + Society* 10, 2 (April 2024), 20563051241261277. doi:10.1177/20563051241261277 (Cited on 18)
- [2] Brandon C Boatwright and Andrew S Pyle. 2023. “Don’t Mess with Ukrainian Farmers”: An examination of Ukraine and Kyiv’s official Twitter accounts as crisis communication, public diplomacy, and nation building during Russian invasion. *Public Relations Review* 49, 3 (2023), 102338. (Cited on 5)
- [3] Lisa Bogerts and Maik Fielitz. 2019. Do You Want Meme War? *Post-Digital Cultures of the Far Right: Online Actions and Offline Consequences in Europe and the US* (2019), 137–153. doi:10.14361/9783839446706 (Cited on 4)
- [4] Göran Bolin and Per Ståhlberg. 2023. Nation branding vs. nation building revisited: Ukrainian information management in the face of the Russian invasion. *Place Branding and Public Diplomacy* 19, 2 (2023), 218–222. (Cited on 4)
- [5] Tiziano Bonini and Emiliano Trerè. 2024. *Algorithms of resistance: The everyday fight against platform power*. Mit Press. (Cited on 21)
- [6] Kenneth J Boyte. 2017. An analysis of the social-media technology, tactics, and narratives used to control perception in the propaganda war over Ukraine. *Journal of Information Warfare* 16, 1 (2017), 88–111. (Cited on 20)
- [7] William Brady, Julian Wills, John Jost, Joshua Tucker, and Jay Van Bavel. 2017. Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences* 114 (06 2017), 201618923. doi:10.1073/pnas.1618923114 (Cited on 20)
- [8] David G Bromley, Anson D Shupe Jr, and Joseph C Ventimiglia. 1979. Atrocity tales, the Unification Church, and the social construction of evil. *Journal of Communication* 29, 3 (1979), 42–53. (Cited on 20)
- [9] Vikram Kamath Cannanure, Delvin Varghese, Cuauhtémoc Rivera-Loaiza, Faria Noor, Dipto Das, Pranjali Jain, Meiyin Chang, Marisol Wong-Villacres, Naveena Karusala, Nova Ahmed, Sarina C Till, Bernard Ijesunor Akhigbe, Melissa Densmore, Susan Dray, Christian Sturm, and Neha Kumar. 2023. HCI Across Borders: Towards Global Solidarity. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI EA ’23). Association for Computing Machinery, New York, NY, USA, Article 351, 5 pages. doi:10.1145/3544549.3573806 (Cited on 3)
- [10] Arthur Capozzi, Gianmarco De Francisci Morales, Yelena Mejova, Corrado Monti, André Panisson, and Daniela Paolotti. 2021. Clandestino or Rifugiato? Anti-Immigration Facebook Ad Targeting in Italy (CHI ’21). Association for Computing Machinery, New York, NY, USA, Article 179, 15 pages. doi:10.1145/3411764.3445082 (Cited on 3)
- [11] Viktor Chagas, Fernanda Freire, Daniel Rios, and Dandara Magalhães. 2019. Political memes and the politics of memes: A methodological proposal for content analysis of online political memes. *First Monday* (2019). (Cited on 7)
- [12] Cassandra M Chapman, Morgana Lizzio-Wilson, Zahra Mirnajafi, Barbara M Masser, and Winnifred R Louis. 2022. Rage donations and mobilization: Understanding the effects of advocacy on collective giving responses. *British Journal of Social Psychology* 61, 3 (2022), 882–906. (Cited on 19)

- [13] Emily Chen and Emilio Ferrara. 2023. Tweets in time of conflict: A public dataset tracking the twitter discourse on the war between Ukraine and Russia. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 17. 1006–1013. (Cited on 4)
- [14] Ben Cole. 2022. *Syrian Information and Propaganda War*. Springer. (Cited on 20)
- [15] W Timothy Coombs. 2021. *Ongoing crisis communication: Planning, managing, and responding*. Sage publications. (Cited on 4)
- [16] Nicholas J Cull, David H Culbert, and David Welch. 2003. *Propaganda and mass persuasion: A historical encyclopedia, 1500 to the present*. Bloomsbury Publishing USA. (Cited on 13, 20)
- [17] Guilherme Ghisoni da Silva. 2021. Memes war: The Political Use Of Pictures In Brazil 2019. *Philosophos - Revista de Filosofia* 25, 2 (2021). doi:10.5216/phi.v25i2.64490 (Cited on 4, 5)
- [18] Maxime Dafaure. 2020. The “great meme war.” The alt-right and its multifarious enemies. *Angles. New Perspectives on the Anglophone World* 10 (2020). (Cited on 4, 5)
- [19] Dipto Das and Bryan Semaan. 2022. Collaborative identity decolonization as reclaiming narrative agency: Identity work of Bengali communities on Quora. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–23. (Cited on 20)
- [20] Clayton Allen Davis, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. Botornot: A system to evaluate social bots. In *Proceedings of the 25th international conference companion on world wide web*. 273–274. (Cited on 21)
- [21] Richard Dawkins. 1976. The selfish gene. (Cited on 4)
- [22] Constance de Saint Laurent, Vlad P Glăveanu, and Ioana Literat. 2021. Internet memes as partial stories: Identifying political narratives in coronavirus memes. *Social Media+ Society* 7, 1 (2021), 2056305121988932. (Cited on 2, 5, 7)
- [23] Charmaine Du Plessis. 2018. Social media crisis communication: Enhancing a discourse of renewal through dialogic content. *Public relations review* 44, 5 (2018), 829–838. (Cited on 4)
- [24] Erwan Dujancourt and Marcel Garz. 2023. The effects of algorithmic content selection on user engagement with news on twitter. *The Information Society* 39, 5 (2023), 263–281. (Cited on 21)
- [25] Marta Dynel and Fabio Indio Massimo Poppi. 2021. Caveat emptor: Boycott through digital humour on the wave of the 2019 Hong Kong protests. *Information, Communication & Society* 24, 15 (2021), 2323–2341. (Cited on 4, 5)
- [26] European Commission. 2023. Standard Eurobarometer STD98. https://data.europa.eu/data/datasets/s2872_98_2_std98_eng?locale=en [Accessed 20-06-2023]. (Cited on 6)
- [27] Lisa Fazio. 2020. Out-of-context photos are a powerful low-tech form of misinformation. <https://www.pbs.org/newshour/science/out-of-context-photos-are-a-powerful-low-tech-form-of-misinformation> (Cited on 19)
- [28] Jeff Giese. 2015. It’s time to embrace memetic warfare. *Defence Strategic Communications* 1, 1 (2015), 67–75. (Cited on 2)
- [29] Tristin Glunt and Marina Kogan. 2019. Public Engagement with Official-Source Content in Crisis. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI EA ’19). Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3290607.3313061 (Cited on 3)
- [30] Viveca S Greene. 2019. “Deplorable” satire: Alt-right memes, white genocide tweets, and redpilling normies. *Studies in American Humor* 5, 1 (2019), 31–69. (Cited on 4)
- [31] Brett A Halperin, Gary Hsieh, Erin McElroy, James Pierce, and Daniela K Rosner. 2023. Probing a Community-Based Conversational Storytelling Agent to Document Digital Stories of Housing Insecurity. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18. (Cited on 20)
- [32] Lars Kai Hansen, Adam Arvidsson, Finn Årup Nielsen, Elanor Colleoni, and Michael Etter. 2011. Good friends, bad news-affect and virality in twitter. In *Future Information Technology: 6th International Conference, FutureTech 2011, Loutraki, Greece, June 28-30, 2011, Proceedings, Part II*. Springer, 34–43. (Cited on 20)
- [33] Margot Hare and Jason Jones. 2023. Slava Ukraini: Exploring Identity Activism in Support of Ukraine via the Ukraine Flag Emoji on Twitter. *Journal of Quantitative Description: Digital Media* 3 (2023). (Cited on 4)
- [34] Martin F. Herz. 1949. Some Psychological Lessons From Leaflet Propaganda in World War II. *Public Opinion Quarterly* 13, 3 (1949), 471. doi:10.1086/266096 (Cited on 8)
- [35] Juan Pablo Hourcade and Natasha E. Bullock-Rest. 2011. HCI for Peace: A Call for Constructive Action. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (CHI ’11). Association for Computing Machinery, New York, NY, USA, 443–452. doi:10.1145/1978942.1979005 (Cited on 3)
- [36] Nathan Johnson, Benjamin Turnbull, Thomas Maher, and Martin Reisslein. 2020. Semantically modeling cyber influence campaigns (CICs): Ontology model and case studies. *IEEE Access* 9 (2020), 9365–9382. (Cited on 3)
- [37] Aharon Kantorovich. 2013. An Evolutionary View of Science: Imitation and Memetics. <http://philsci-archive.pitt.edu/9912/> Social Science Information. (Cited on 4)
- [38] Stephen Karpman. 1968. Fairy tales and script drama analysis. *Transactional analysis bulletin* 7, 26 (1968), 39–43. (Cited on 5, 7)

- [39] Ivan Kozachenko. 2021. Transformed by contested digital spaces? Social media and Ukrainian diasporic ‘selves’ in the wake of the conflict with Russia. *Nations and Nationalism* 27, 2 (2021), 377–392. (Cited on 20)
- [40] Poka Laenui. 2000. Processes of decolonization. *Reclaiming Indigenous voice and vision* (2000), 150–160. (Cited on 20)
- [41] Benjamin Lee. 2020. ‘Neo-Nazis Have Stolen Our Memes’: Making Sense of Extreme Memes. *Digital extremism: Readings in violence, radicalisation and extremism in the online space* (2020), 91–108. (Cited on 4)
- [42] Yu-Hao Lee and Gary Hsieh. 2013. Does slacktivism hurt activism? The effects of moral balancing and consistency in online activism. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 811–820. (Cited on 3, 21)
- [43] Chen Ling, Ihab AbuHilal, Jeremy Blackburn, Emiliano De Cristofaro, Savvas Zannettou, and Gianluca Stringhini. 2021. Dissecting the meme magic: Understanding indicators of virality in image memes. *Proceedings of the ACM on human-computer interaction* 5, CSCW1 (2021), 1–24. (Cited on 2, 4, 7, 8, 21)
- [44] Mai Skjott Linneberg and Steffen Korsgaard. 2019. Coding qualitative data: A synthesis guiding the novice. *Qualitative research journal* (2019). (Cited on 7)
- [45] Mykola Makhortykh and Yehor Lyebyedyev. 2015. # SaveDonbassPeople: Twitter, propaganda, and conflict in Eastern Ukraine. *The Communication Review* 18, 4 (2015), 239–270. (Cited on 4)
- [46] Mykola Makhortykh and Maryna Sydorova. 2017. Social media and visual framing of the conflict in Eastern Ukraine. *Media, war & conflict* 10, 3 (2017), 359–381. (Cited on 4)
- [47] Ilan Manor and Elad Segev. 2015. America’s selfie: How the US portrays itself on its social media accounts. In *Digital diplomacy*. Routledge, 89–108. (Cited on 4)
- [48] Alessandra Massa and Giuseppe Anzera. 2022. The platformization of military communication: The digital strategy of the Israel Defense Forces on Twitter. *Media, War & Conflict* (2022), 17506352221101257. (Cited on 4)
- [49] Daniel C. Mattingly and Elaine Yao. 2022. How Soft Propaganda Persuades. *Comparative Political Studies* 55, 9 (Aug. 2022), 1569–1594. doi:10.1177/001041402111047403 (Cited on 2)
- [50] Ally McCrow-Young and Mette Mortensen. 2021. Countering spectacles of fear: Anonymous’ meme ‘war’ against ISIS. *European Journal of Cultural Studies* 24, 4 (2021), 832–849. (Cited on 4, 5)
- [51] Alexandros Moutzouridis et al. 2023. A war beyond the war: the use of propaganda in social media in the 2022 Russo-Ukrainian conflict. *Black Sea and Eastern Mediterranean Review* 1, 1 (2023), 8–55. (Cited on 20)
- [52] Karine Nahon and Jeff Hemsley. 2013. *Going viral*. Polity. (Cited on 21)
- [53] Brendan Nyhan. 2021. Why the backfire effect does not explain the durability of political misperceptions. *Proceedings of the National Academy of Sciences* 118, 15 (2021), e1912440117. (Cited on 19)
- [54] Orestis Papakyriakopoulos, Severin Engelmann, and Amy Winecoff. 2023. Upvotes? Downvotes? No Votes? Understanding the Relationship between Reaction Mechanisms and Political Discourse on Reddit. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI ’23). Association for Computing Machinery, New York, NY, USA, Article 549, 28 pages. doi:10.1145/3544548.3580644 (Cited on 3)
- [55] Judea Pearl. 2012. The causal mediation formula—a guide to the assessment of pathways and mechanisms. *Prevention science* 13 (2012), 426–436. (Cited on 15)
- [56] Judea Pearl. 2013. Linear Models: A Useful “Microscope” for Causal Analysis. *Journal of Causal Inference* 1, 1 (May 2013), 155–170. doi:10.1515/jci-2013-0003 (Cited on 15)
- [57] Lucy Pei, Benedict Salazar Olgado, and Roderic Crooks. 2022. Narrativity, Audience, Legitimacy: Data Practices of Community Organizers. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–6. (Cited on 20)
- [58] Liudmyla Pidkuimukha and Nadiya Kiss. 2020. Battle of narratives: Political memes during the 2019 Ukrainian presidential election. *Cognitive Studies* 20 (2020). (Cited on 4, 5)
- [59] Steve Rathje, Jay Van Bavel, and Sander van der Linden. 2021. Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences* 118 (06 2021), e2024292118. doi:10.1073/pnas.2024292118 (Cited on 20)
- [60] Stephen D Reese, Jr Gandy, and August E Grant. 2001. Prologue—Framing public life: A bridging model for media research. In *Framing public life*. Routledge, 23–48. (Cited on 4)
- [61] Carol F Scott, Gabriela Marcu, Riana Elyse Anderson, Mark W Newman, and Sarita Schoenebeck. 2023. Trauma-Informed Social Media: Towards Solutions for Reducing and Healing Online Harm. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–20. (Cited on 3)
- [62] Fei Shen, Erkun Zhang, Hongzhong Zhang, Wujiong Ren, Quanxin Jia, and Yuan He. 2023. Examining the differences between human and bot social media accounts: A case study of the Russia-Ukraine War. *First Monday* (2023). (Cited on 4)
- [63] Limor Shifman. 2013. *Memes in digital culture*. MIT press. (Cited on 21)
- [64] Lisa Silvestri. 2016. Mortars and memes: Participating in pop culture from a war zone. *Media, War & Conflict* 9, 1 (2016), 27–42. (Cited on 4, 5)

- [65] Oliviero Stock, Massimo Zancanaro, Chaya Koren, Cesare Rocchi, Zvi Eisikovits, Dina Goren-Bar, Daniel Tomasini, and Patrice Weiss. 2008. A co-located interface for narration to support reconciliation in a conflict: initial results from Jewish and Palestinian youth. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1583–1592. (Cited on 20)
- [66] Peter Suci. 2022. Is Russia's Invasion Of Ukraine The First Social Media War? <https://www.forbes.com/sites/petersuci/2022/03/01/is-russias-invasion-of-ukraine-the-first-social-media-war> (Cited on 3)
- [67] Philip M Taylor. 2013. Munitions of the mind: A history of propaganda. *Munitions of the Mind* (2013), 1–360. (Cited on 20)
- [68] Sho Tsugawa and Hiroyuki Ohsaki. 2015. Negative Messages Spread Rapidly and Widely on Social Media. In *Proceedings of the 2015 ACM on Conference on Online Social Networks* (Palo Alto, California, USA) (COSN '15). Association for Computing Machinery, New York, NY, USA, 151–160. doi:10.1145/2817946.2817962 (Cited on 20)
- [69] Riina Vuorikari, Natalia Jerzak, Zbigniew Karpinski, Artur Pokropek, Jadwiga Tudek, et al. 2022. Measuring Digital Skills across the EU: Digital Skills Indicator 2.0. *Publications Office of the European Union* (2022), 1–21. doi:10.2760/897803 (Cited on 19)
- [70] Robin Wakefield and Kirk Wakefield. 2023. How intergroup counter-empathy drives media consumption and engagement. *Internet Research* (2023). (Cited on 20)
- [71] Julia K Weiss and Elizabeth L Cohen. 2019. Clicking for change: the role of empathy and negative affect on engagement with a charitable social media campaign. *Behaviour & Information Technology* 38, 12 (2019), 1185–1193. (Cited on 20)
- [72] Bradley E Wiggins. 2016. Crimea River: Directionality in memes from the Russia-Ukraine conflict. *International Journal of Communication* 10 (2016), 35. (Cited on 4, 5)
- [73] Heather Suzanne Woods and Leslie A Hahner. 2019. *Make America meme again: The rhetoric of the Alt-Right*. Peter Lang New York. (Cited on 4, 5)
- [74] Kai-Cheng Yang, Emilio Ferrara, and Filippo Menczer. 2022. Botometer 101: Social bot practicum for computational social scientists. *Journal of Computational Social Science* 5, 2 (2022), 1511–1528. (Cited on 4)
- [75] Jinyi Ye, Nikhil Jindal, Francesco Pierri, and Luca Luceri. 2023. Online Networks of Support in Distressed Environments: Solidarity and Mobilization during the Russian Invasion of Ukraine. *arXiv preprint arXiv:2304.04327* (2023). (Cited on 3, 19)
- [76] Olesia Yehorova, Antonina Prokopenko, and Anna Zinchenko. 2023. Towards a typology of humorous wartime tweets: the case of Ukraine 2022. *The European Journal of Humour Research* 11, 1 (2023), 1–26. (Cited on 5)
- [77] Savvas Zannettou, Tristan Caulfield, Jeremy Blackburn, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, and Guillermo Suarez-Tangil. 2018. On the Origins of Memes by Means of Fringe Web Communities. In *Proceedings of the Internet Measurement Conference 2018*. ACM, Boston MA USA, 188–202. doi:10.1145/3278532.3278550 (Cited on 5, 21)
- [78] Savvas Zannettou, Tristan Caulfield, Barry Bradlyn, Emiliano De Cristofaro, Gianluca Stringhini, and Jeremy Blackburn. 2020. Characterizing the use of images in state-sponsored information warfare operations by Russian trolls on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 774–785. (Cited on 5)
- [79] Eline Zenner and Dirk Geeraerts. 2018. One does not simply process memes: Image macros as multimodal constructions. *Cultures and traditions of wordplay and wordplay research* (2018), 167–194. (Cited on 4)

- (5) Attributes of the subject. For images whose subject is one or more characters, we consider whether the image's visual attraction lies with the character's facial expression or with their posture. For the other attributes, we identify five features.
 - Facial expression: only for a character
 - Posture: only for a character
 - Poster: informative large scale image including both textual and graphic elements. There are also posters only with either of these two elements. Posters are generally designed to be displayed at a long distance. It informs or instructs the viewer through text, symbols, graph, or a combination of these.
- (6) Contains words. Whether words are a part of the image, including superimposed on top of an existing image.
- (7) Character face/posture emotion.
 - Positive, Negative, or Neutral: only for a character annotated as facial expression
- (8) Narrative (possible more than one).
 - Hero (benevolent, strong)
 - Victim (benevolent, weak)
 - Villain (malevolent, strong)
 - Fool (malevolent, weak)
 - Other
 - None
- (9) Emotional appeal. [Open label]
 - Humor: satirical, sarcastic, or funny
 - Fear: threats, harms, or calamities
 - Outrage: scandals, corruption, or moral hazards
 - Pride: objects/scenes of adoration or affection
 - Compassion: sympathy and sorrow for another who is stricken by misfortune
 - None
- (10) Intent. [Open label]
 - Informational
 - Call to action
 - Other
- (11) Actor. [Open label]

C Inter-annotator agreement

Table C.1 shows the inter-annotator agreement between the authors, measured by Krippendorff's alpha, for the annotation of the data described in Methods. In aggregate, the labelers are mostly in agreement apart from the open label of Actor (which was later cleaned and aggregated).

D Visual content features

Figure D.2 contains summary statistics of content visual features. The content produced by the three accounts has markedly different features. The @uamemesforces tends to publish content spanning both multiple and single panels, have mostly a medium scale, and have a character present, for which both posture and facial expression may be important, and which vast majority of the time contains words. @Ukraine and @DefenceU accounts are more similar to each other, in the way that most of their content has a single panel, it uses more illustrations/posters, which contain text. @DefenceU is especially more likely to have content without words and with neutral emotive expressions.

Table C.1. Inter-annotator agreement (Krippendorff’s alpha).

Feature	@Ukraine	@uamemesforces	@DefenceU
Number of panels	1.000	1.000	0.825
Type of image	0.754	0.772	1.000
Scale	0.643	0.402	0.675
Type of subject	0.738	0.664	0.846
Attributes of subject	0.597	0.631	0.860
Contains words	0.885	0.848	0.765
Character emotion	0.357	0.582	0.691
Actor (open)	0.329	0.307	0.157
Emotional appeal (open)	0.562	0.195	0.415
Intent (open)	0.627	0.789	0.510
Narrative	0.704	0.590	0.644

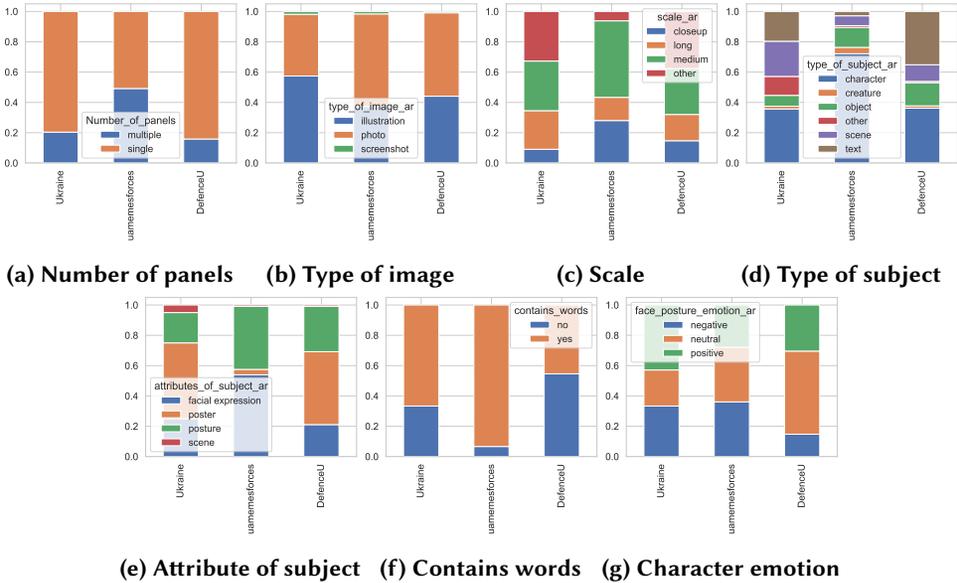


Fig. D.2. Distribution of visual features.

Received January 2024; revised July 2024; accepted October 2024