

Slava Ukraini: Exploring Identity Activism in Support of Ukraine via the Ukraine Flag Emoji on Twitter

MARGOT A. HARE¹

The Ohio State University, United States

JASON J. JONES

Stony Brook University, United States

Identity Activism is a new phenomenon afforded by the massive popularity of social media. It consists of the prominent display of a social movement symbol within a space reserved for description of the self. The 2022 Russian invasion of Ukraine provides a contemporary (yet unfortunate) opportunity to observe this phenomenon. Here, we introduce and explore this concept in the context of the recent Twitter trend of displaying the Ukraine flag emoji in bios and names to signal support of Ukraine. We explore several questions, including: how has the popularity of this trend changed over time, are users who display the flag more likely to be connected to others who do, and what types of users are and are not participating. We find that Ukraine flag emoji prevalence in both names and bios increased many-fold in late February 2022, with it becoming the 11th most prevalent emoji in bios and the 3rd most prevalent emoji in names during March. We also find evidence that users who display the flag in their bio or name are more likely

Author 1: hare.98@osu.edu

Author 2: jason.j.jones@stonybrook.edu

Date submitted: 2022-10-01

¹ We would like to thank Bailee Wolfe and Laurel Philpott for generously allowing us to use their Twitter account images. We would also like to thank Liam Tucker for his support in the data collection process. This material is based upon work supported by the National Science Foundation under grants IIS-1927227, OAC-1950052 and CCF-2208663.

Copyright © 2023 (Hare, Jones). Licensed under the Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd). Available at: <http://journalqd.org>

to follow and be followed by others who also do so, as compared to users who do not. Finally, we observe that users who share politically left-leaning messages were most likely to display the emoji. Those who share account information from alternative social media sites and non-personal accounts appear least likely. These findings give us insight into how users participate in Identity Activism, what connections exist between participating users, and, in this particular case, what types of users participate.

Keywords: Ukraine, Twitter, emojis, identity activism

Introduction

Identity Activism (IA) is a new phenomenon emerging from the massive popularity of social media. It consists of the prominent display of a social movement symbol within a space reserved for describing oneself. Studying IA is important for two reasons.


First, visibility of movement affiliation is a political tool and an overt goal of organized movements. For example, in a pamphlet titled “Transgender Visibility – A Guide to Being You,” the Human Rights Campaign Foundation strongly encourages individuals to disclose their transgender identity explicitly and publicly. Being visible is associated with “the joy of living an open, honest life and engaging in relationships as a whole and authentic person.” The pamphlet concludes with a message from the President of the Foundation which states, “Progress toward equality is made when we choose to share our lives with others. This simple yet profound step is the greatest political action any of us can take” (Human Rights Campaign Foundation 2014).

Second, IA is a highly observable step in the otherwise obscured process of an individual’s cause affiliation. Social scientists are deeply interested in the stages of individuals’ strengthening or weakening ties to organized social movements (e.g., Klandermans & Oegema 1987; McAdam 1986). The operationalization of IA we present here provides a method to observe individuals’ explicit, public self-affiliation consistently,

persistently, and precisely. The 2022 Russian invasion of Ukraine provides a contemporary (yet unfortunate) opportunity for IA case study. This event has received much attention in the news, triggering popular support among Americans for Ukraine (Jeffery and Breuninger 2022; Lakritz 2022; Morse 2022).

In this paper, we explore this phenomenon in three sections: Prevalence Over Time, Clustering, and Relative Prevalence: Who Does and Does Not Display the Ukraine Flag. First, we find that Ukraine flag emoji prevalence increased many-fold in both names and bios in late February 2022 and present a timeline with daily resolution. We also provide context for typical emoji prevalence levels. Second, we find evidence that users who display the flag in their bio are more likely to follow and be followed by others who also do so, as compared to users who do not. The same pattern is seen in users who display the flag in their names. Next, we find that users who share politically left-leaning messages are most likely to display the emoji, and those who share account information from alternative social media sites and non-personal accounts appear least likely. Finally, we offer a short discussion of bots in our data.

Background

In this work, we define IA as the user’s inclusion of a Ukraine flag emoji in the bio or name field of a Twitter profile. The Ukraine flag emoji depicts the country flag for Ukraine. It consists of equally sized horizontal bands of blue and yellow and is represented in Unicode as U+1F1FA U+1F1E6: ‘’.

What are Twitter Bios?

When someone creates a Twitter profile, they are instructed: “Describe yourself. What makes you special? Don't think too hard, just have fun with it.” They are limited to 160 characters in this description called a “bio” (short for “biography”). This bio appears on their profile page directly below their name and profile picture. Anyone can see this bio, even if the user sets their tweets to protected visibility (i.e., only visible to the user’s followers). The reader may wonder how many bios are left empty. We sampled 100,000

users per year from 2015-2022 and found that approximately 15% of US Twitters users leave their bios blank, while 85% include information. Figure 1 below shows an example of a Twitter user displaying the Ukraine flag emoji in their bio.



Figure 1. Example of Ukraine Flag in Bio

What are Twitter Names?

On Twitter, a user has two types of identifiers: a handle and a username (or simply “name”). A handle is a unique identifier for a user’s account that is chosen when setting up the account. It begins with “@” and cannot contain spaces or emojis. For example, current President Biden’s *handle* is @POTUS. A name is a different identifier for the same account intended to be more flexible than the handle. Users are able to place spaces and emojis into their name. President Biden’s *name* field contains the text “President Biden”. These are two different ways to refer to the same account. Our work focuses on names, not handles, because names can contain emojis. Figure 2 below shows an example of a Twitter user displaying the Ukraine flag emoji in their name.






Figure 2. Example of Ukraine Flag in Name

Prior Work

Twitter has been a popular research subject ever since its launch (Zimmer and Proferes 2014). However, because 1) most work focuses on the content of tweets and 2) emojis have only recently become a widely available feature, the study of emojis in Twitter bios and names has received limited attention (little to no literature was found investigating

emojis in Twitter names). In *Emoji and Self-Identity in Twitter Bios* (2020), Li et al. found that emojis in bios can reveal what users are more likely to tweet about, who they are more likely to follow, their hobbies, and other identity characteristics. For example, they found that users who display the rainbow emoji ‘’ in their bios are more likely to tweet about life and entertainment topics, while users who display the American flag emoji ‘’ are more likely to tweet about the police and politics. This study sampled 22,000 Twitter users over a period of four months. Our current work focuses more narrowly on one emoji (the Ukraine flag) in both bios and names and follows a larger number of accounts with multiple observations over the first six months of 2022.

Adding something to one’s bio (e.g., hashtag, emoji, motto) to show support for a specific cause is not new. This is often cited as an example of “slacktivism” or “performative activism”, defined by Kutluca and Radke in a 2022 paper as, “easy and costless actions that often do not challenge the status quo and are motivated primarily by the desire to accrue personal benefits”. Other examples may include changing one’s profile picture to black to signal support for the Black Lives Matter movement or signing online petitions. Kutluca and Radke even specifically mention retweeting posts in order to show support for those affected by unjust wars as an example. The literature is divided on whether slacktivism has any benefits or if it only does harm.

Madison and Klang (2020) argue that the ability to “[reach] out, create[e] awareness, [garner] support, and enabl[e] asynchronous political discourse,” using technology is a new tool that should be valued rather than denigrated. Many attack slacktivism by saying that the act is low effort and low cost. Madison and Klang claim it is unreasonable to believe that activism must come at a cost to be meaningful. In fact, the low barriers that allow people to participate in activism online may be a net good by increasing inclusion. However, Katluca and Radke argue if someone is praised for participating in something deemed to be “slacktivism”, it may decrease their likelihood of participating in more significant or costly actions offline. They also add that there is evidence that if a member of a privileged group participates in slacktivism in a way that is perceived as selfish, it may discourage the participation of disadvantaged groups in the

effort. The paper particularly notes black people being discouraged by white “slacktivists”. While displaying the Ukraine flag in one’s name or bio may be an act of slacktivism, it has reached a level of popularity worth exploring.

In other related work, Rogers and Jones (2020) presented the first large-scale study of Twitter bios over time, sampling 20 million users over a three-year period. They observed that Americans added political identifiers to their bios at a greater rate than other categories of social identifiers, such as religious terms. Jones, 2021 defined the method of Longitudinal Profile Sampling (LOPS). We used the LOPS method to construct our 2022 Daily Emoji Prevalence datasets.

Eady, Hjorth, and Dinesen (2021) used similar methods to observe a significant decrease in the appearance of Republican-identifying terms in bios immediately following the January 6th Insurrection. This illustrated one example of self-expressed identity in Twitter bios changing in response to offline world events. Although much Twitter bio research relates to political identity, other studies have investigated racial identity and how to analyze social identity within a bio (Firmansyah and Jones 2019; Pathak, Madani, and Joseph 2021). Our work builds on this prior research and serves as a case study of one particular emoji.

Prevalence Over Time

Definition of Prevalence

We wished to estimate the prevalence of Ukraine flag emoji use within the population of active United States Twitter users. We defined the *prevalence* of a given token (generally speaking, a sequence of characters between whitespace) as the number of unique users per 10,000 that displayed that token. When investigating prevalence in bios, we observed whether the user displayed the token in their bio. When investigating prevalence in names, we observed whether the user displayed the token in their name.

The general formula is shown below. Prevalence is multiplied by 10,000 so that prevalence can be discussed in terms of whole numbers.

$$\text{Prevalence} = 10,000 * \frac{\text{Count of Users Displaying the Token}}{\text{Total User Count}}$$

Twitter Data Collection

Our datasets are similar to the dataset introduced in Jones 2021 and were created with the same methods. We used the Twitter Streaming API to collect a random sample of tweets each day in 2022 from January 1 to June 30. The Sample Stream endpoint (Twitter, Inc., 2023) allows for the collection of random tweets without inputting any specific search terms. Next, we gathered information from the accounts that authored the collected tweets. We then filtered these accounts to include only likely US-based accounts. US-based accounts were identified using the profile location field and selecting texts that indicated a US location (e.g., state names and formal abbreviations). The profile information of these accounts came to comprise our dataset. If multiple tweets from the same account were observed within a day, only one tweet (and corresponding profile) was randomly selected. This process gave us a random sample of more than 170,000 unique accounts per day typically (min = 136,920, median = 176,492, and max = 224,023).

IA in Bios and Names

Here we analyze IA prevalence in bios and in names separately. Both are text fields in which it is expected users describe themselves. However, bios and names differ in several ways. A user's name is displayed any time they tweet, retweet, comment, follow someone, and more. Its visibility is much higher than the bio, which is only seen when someone visits another user's page or hovers over the name on the web platform. Bios also allow more characters. It is possible these slight differences lead users to

perform IA in ways that diverge across the two venues. For instance, we speculate users might retain IA signifiers in bios longer than names. We are not interested in investigating any casual relationships between name and bio differences in this work, but we separate our descriptive analysis to acknowledge that differences might exist.

2022 Daily Emoji Prevalence in Bios Dataset

From the sample of accounts, for each day separately, we tokenized all bios with the regular expression `\b\s+` and calculated the prevalence of each token that appeared that day. For emoji tokens, we included only single emojis after splitting on word boundaries and whitespace. (Some compound emoji strings and combination emoji-punctuation strings were present in the data but not counted here. However, they were present at a very low prevalence, with an average of 1.3 users per 10,000.) The results for each day were compiled into one dataset containing each token that appeared in the first six months of 2022 (“token”), the day it was observed (“obsDate”), its prevalence, the raw number of users with that token in their bio on its observation date (“numerator”), and the number of total users sampled on its observation date (“denominator”).

However, most tokens are words, which are not the object of interest. Therefore, we created a sub-dataset containing only emoji tokens. This was named the 2022 Daily Emoji Prevalence in Bios dataset. The first several rows are pictured below, sorted by date with emojis in no particular order.

	token	obsDate	prevalence	numerator	demoninator
1	❤️	1/1/2022	101	1592	158343
2	🇺🇦	1/1/2022	2	38	158343
3	🇺🇸	1/1/2022	10	152	158343
4	😬	1/1/2022	5	85	158343
5	🇺🇸	1/1/2022	110	1741	158343
6	❤️	1/1/2022	17	267	158343
7	🇺🇸	1/1/2022	68	1077	158343
8	♠️	1/1/2022	3	40	158343

Table 1. Sample of the 2022 Daily Emoji Prevalence in Bios Dataset

With this dataset, we measured prevalence in bios on daily cross-sections of active (i.e., tweeting) users. This means that changes in prevalence are subject to several factors, including: 1) Users with or without the emoji of interest became more or less likely to tweet this day. 2) Some users added or removed the emoji from their bio and then were observed tweeting.

A different method of creating a dataset could be observing the same group of users daily. This would allow for a longitudinal study in which the same sub-population of people is studied over time. (Eady, Hjorth, and Dinesen (2021) used a daily longitudinal panel of accounts that follow US news sources.) This type of data would enable strong inference as to changes at the individual level. For instance, longitudinal analysis would guarantee that observed decreases were caused by more users removing a token than adding it. Increased prevalence could be measured similarly. However, our cross-sectional approach has a particular benefit: it allows one to see population-wide change in active users.

2022 Daily Emoji Prevalence in Names Dataset

We followed a simpler procedure to find daily prevalence in names because we were not interested in non-emoji tokens appearing in names. For each day, we determined what (if any) emojis appeared in each user's name. For each emoji identified, we recorded

how many users had that emoji in their name on that day (“numerator”) and how many users were sampled on that day (“denominator”). Again, the results for each day were compiled into one dataset containing each emoji that appeared in any name during the first six months of 2022 (“token”), the day it was observed (“obsDate”), its prevalence, the numerator, and the denominator. This dataset was named the 2022 Daily Emoji Prevalence in Names dataset.

2015-2022 Profile Dataset

To better understand emoji use on Twitter generally, we examined a random sample of tweets collected for the 2015-2022 period at annual resolution. Again, we filtered these to only include accounts based in the United States. We further randomly sampled and included only 100,000 accounts per year. From each account we collected information such as the bio, name, follower count, and more. This information comprised the 2015-2022 Profile dataset. Only the bios and names were considered here. We counted the number of bios and names containing emojis and the number containing flag emojis for each year. The resulting data is shown in Figures 3a and 3b. This provides context regarding how common it is for a user to display emojis.




Figure 3. Timeline of Prevalence of Emojis in Bios (A) and Names (B)

Note. The scales of the y-axes in Figures 3a and 3b are different. Figure 3a goes to approximately 2,800. Figure 3b goes to approximately 1,700.

As emojis have increased in popularity generally, emojis in Twitter profile bios and names have steadily become more prevalent. In 2022, roughly 27% of US Twitter users (2,706 users per 10,000) displayed at least one emoji in their bios and 7.1% displayed at least one flag emoji. Likewise, the prevalence of emojis in names has similarly increased through the years, with approximately 16% of users including an emoji in their name and 3% including a flag emoji in 2022.

Results

Figure 4 below shows the January - June 2022 prevalence of the Ukraine flag emoji in names and in bios within active US users at daily resolution, with 95% confidence intervals included. There was a sudden increase in prevalence in bios beginning on February 21, when Russian President Vladimir Putin officially declared two separatist regions in eastern Ukraine to be sovereign (Bloomberg 2022). This increase continued as the Russian invasion ensued, until February 27, after which the prevalence in bios fluctuated around a median of 32 until March 27. (These three specific dates are noted in the figure.) The prevalence then declined slightly, fluctuating around a median of 30 through June 30.

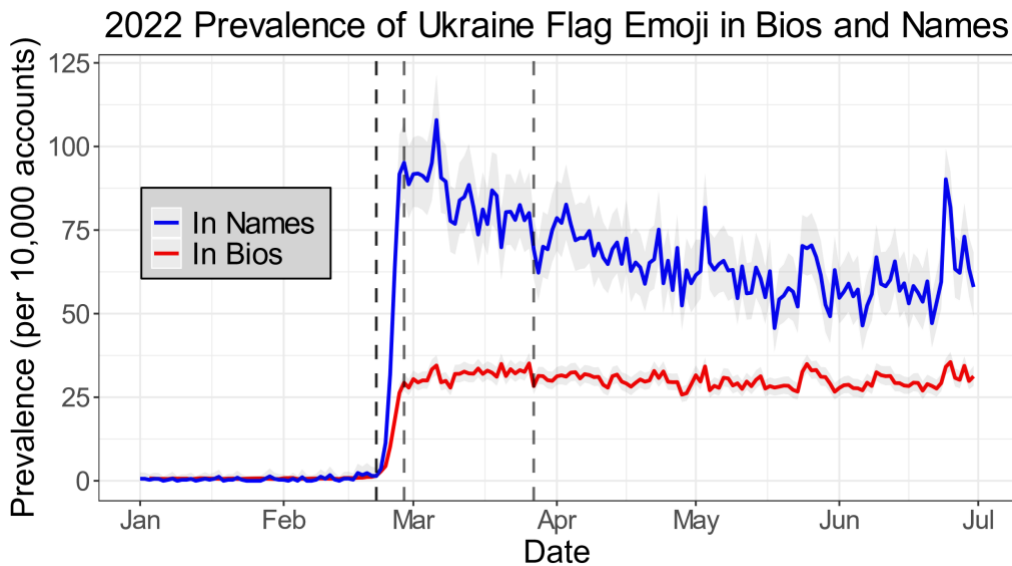



Figure 4. 2022 Prevalence Timeline with 95% Confidence Intervals

The prevalence in names initially increased three times more than the prevalence in bios. Beginning February 17, the prevalence in names increased until reaching 95 on February 27. However, the number did not remain as stable as the prevalence in bios. The prevalence in names began an unsteady decline after its peak at 108 on March 6.

On February 15, before the Ukraine flag emoji increased in prevalence, it was ranked last in prevalence in bios (tied with 483 others) out of all 935 emojis that appeared that day. After its increase, the Ukraine flag emoji jumped to eleventh-ranked in bios. The following figures are purely to give context for what a specific prevalence number means. They are not intended to imply that the other emojis are also examples of identity activism. For added context with another political signifier, the blue wave emoji ‘’ fluctuated between a prevalence of 21 and 35 in bios throughout the first half of 2022 with a median prevalence of 27. This emoji is often used by democrats to show support/hope for an incoming wave of Democratic candidates taking office, taking back the Senate, etc. Figure 5 below shows the eleven emojis with the highest average prevalence in bios in March 2022, the month with the highest average Ukraine flag emoji prevalence. From left to right,

the emojis in Figure 5 are the US Flag, Red Heart, Sparkles, Rainbow Flag, No One Under Eighteen Sign, Blue Heart, Purple Heart, Prohibited Sign, Two Hearts, Rainbow, and Ukraine Flag. Figure 6 shows the emojis with the highest average prevalence in names in March 2022. The Ukraine flag emoji became the second-ranked emoji in names after its increase. From left to right, the emojis in Figure 6 are Sparkles, the Ukraine Flag, the US Flag, Butterfly, Sunflower, Black Heart, Blue Heart, Crown, Rainbow Flag, White Heart, and Purple Heart.

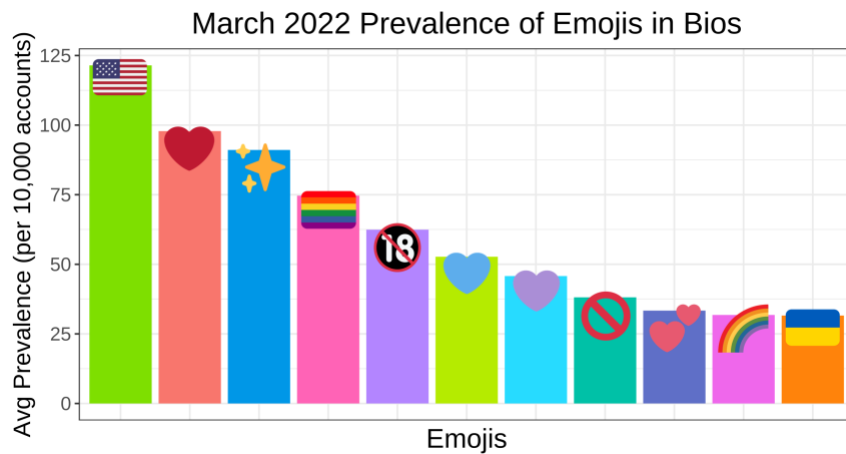


Figure 5. Highest Average Prevalence Emojis in Bios of March 2022




Figure 6. Highest Average Prevalence Emojis in Names of March 2022

For context, the United States flag emoji was the most popular emoji in bios for every day of 2022, indicating that the prevalence level in bios for any single emoji should be expected to not exceed 130 users per 10,000. These results show that Ukraine flag emoji IA is not a rare occurrence. What else, then, can we learn about it?

Clustering

We next asked: are Twitter users who display the flag more likely to be connected to others who also do so? This can also be worded as: is there clustering in the Twitter “friends and followers” network among flag users (FUs, a user who displays the Ukraine flag emoji)?

We explored this question of clustering for both flag-name users (FNUs, a user who displays the Ukraine flag emoji in their name) and flag-bio users (FBUs, a user who displays the Ukraine flag emoji in their bio), and we explored them both in two ways.

Method One: Node-Level Estimation

The first method involved exploring friend and follower lists. On Twitter, there are two ways that a user can be connected to a different user: they can be a friend or a follower (or both). A friend is someone whom the given user follows, and a follower is someone who follows the given user.




Figure 7. Follower and Friend Structure

We randomly selected a roughly equal number of FBUs and non-FBUs whose information was collected between March 1, 2022 – June 30, 2022. This time period was

chosen because it included all dates beginning from the first peak in Ukraine flag emoji prevalence in both names and bios. We did the same for FNUs and non-FNUs. For each selected flag-bio and non-flag-bio user, we counted how many of their friends had the flag in their bio. For each selected flag-name and non-flag-name user, we counted how many of their friends had the flag in their name. The same test was then repeated observing follower lists. The table below shows how many users were sampled for each group.

Table 2. Total Accounts Sampled Across the Four Sample Types

Follower Lists	Accounts Sampled
A) FBUs	7,406
Non-FBUs	7,760
B) FNUs	10,223
Non-FNUs	9,862
Friend Lists	Accounts Sampled
C) FBUs	6,626
Non-FBUs	6,757
D) FNUs	10,183
Non-FNUs	9,916

Note. Results were collected in four ways: A) examining friend lists of both FBUs and non-FBUs, B) examining follower lists of both FBUs and non-FBUs, C) examining friend lists of both FNUs and non-FNUs, and D) examining follower lists of both FNUs and non-FNUs.

The Twitter API returns a maximum of 200 friends/followers at a time for a given user. Because some users have millions of friends/followers, collecting entire lists for every user was not feasible. Instead, we collected a maximum of 1,000 friends/followers for each selected user. If a user had fewer than 1,000, their entire list was collected. For friend lists, entire lists were collected for approximately 69.02% of the selected users. For follower lists, that number is approximately 74.27%. Additionally, the API returns the most recently added friends/followers in reverse chronological order. This means that, for the users with more than 1,000 friends/followers, the results may be skewed by the more recent activity

of the user. However, because entire lists were collected for the majority of users, we do not believe the general results described here will be dominated by recent activity. Also to be noted, friends and followers lists can only be observed for non-protected accounts.

Node-Level Estimation Results

The results for clustering among friends of flag-bio users and flag-name users are shown below in Figure 8. The same pattern of results was observed in followers.




Figure 8. Proportion of Friends with Flag for FBUs and FNUs

Having a friend or follower with the flag is uncommon because of the general low prevalence of the flag, so the proportions are necessarily low. However, in both figures, there is a clear shift of the distribution in the red bars to the right, indicating that flag users (either in the name or in the bio) have a higher proportion of friends who also display the flag, as compared to non-flag users. Results for both friends and followers are shown below in Table 3. A user who engages in IA with the Ukraine flag emoji is 10 times more likely to be connected to another user who does so than a user from the non-IA user comparison group.

Table 3. Percentage of FUs and Non-FUs in Followers and Friends

In Bios	FBUs	Non-FBUs
Followers %	2.80	0.205
Friends %	2.42	0.197
In Names	FNUs	Non-FNUs
Followers %	3.70	0.453
Friends %	2.84	0.369

Note. This table demonstrates the strong clustering of Ukraine flag identity activism. For both Bios and Names, a user who includes the flag (Flag Users) are much more likely to have followers and friends who also do so, as compared to users who do not (Non-Flag Users).

Method Two: Dyad-Level Estimation

We also used a second approach to investigate the question of clustering. Instead of investigating friend and follower lists, we explored connections between individual pairs of users. We selected random pairs of users, again from those observed tweeting March 1, 2022 – June 30, 2022. For clustering between FBUs, we collected three types of pairs: both users are FBUs, neither user is an FBU, one user is an FBU and the other is not. Corresponding pairs were collected for clustering between FNUs. Roughly equal amounts of pairs were collected for each of the three pair types. For every pair, we used the Twitter API to determine if user A was following user B and if user B was following user A.

Dyad-Level Estimation Results

Out of 103,579 pairs of FBUs, 303 pairs were users that follow each other. Out of 105,505 non-FBU pairs, 11 pairs were users that follow each other. Out of 109,551 pairs of FNUs, 112 pairs were users that follow each other. Out of 109,112 non-FNU pairs, 9 pairs were users that follow each other. Again, these numbers are necessarily very small because the probability of two random US Twitter users following each other is very low. Still, the difference is clear. If two users both display the Ukraine flag emoji in the same way (in their bio or in their name), they are more likely to be following each other than two

users who do not display the flag. This is further evidence that flag users are more likely to be connected to other flag users than a non-flag user is.

This clustering could be the result of Ukraine flag emoji use spreading within a community that already existed (i.e., social contagion), or it could be a result of people being more likely to follow someone that they can see agrees with them on something (i.e., homophily). Either way, in this specific case of Identity Activism, we observed evidence of greater connection between those who are activists for the same thing, Ukraine.

Relative Prevalence: Who Does and Does Not Display the Ukraine Flag?

Finally, we wanted to investigate which types of users were most and least likely to display the Ukraine flag emoji. To do this, we determined the relative prevalence for each token that appeared alongside the flag from March 1, 2022 – June 30, 2022. We used two types of relative prevalence: relative prevalence among FBUs and relative prevalence among FNUs. The equations are shown below.

$$\text{Relative Prevalence Among FBUs}^* = \frac{\text{Prevalence Among Bios Of FBUs}}{\text{Overall Prevalence}}$$

$$\text{Relative Prevalence Among FNUs}^* = \frac{\text{Prevalence Among Bios of FNUs}}{\text{Overall Prevalence}}$$

$$*\text{If Rev. Prev.} < 1, \text{Relative Prevalence} = \frac{-1.0}{\text{Original Rev. Prev.}}$$

If a token has a high relative prevalence among FBUs, that means the token is more prevalent among flag-bio users than it is overall. Conversely, if a token has a negative relative prevalence among FBUs, that means the token is less prevalent among flag-bio users as compared to overall, which suggests it is less likely to appear in the bio of a flag-bio user. The logic is the same for FNUs.

Results

Figure 9 below shows the tokens with the highest and lowest relative prevalence among FBUs.

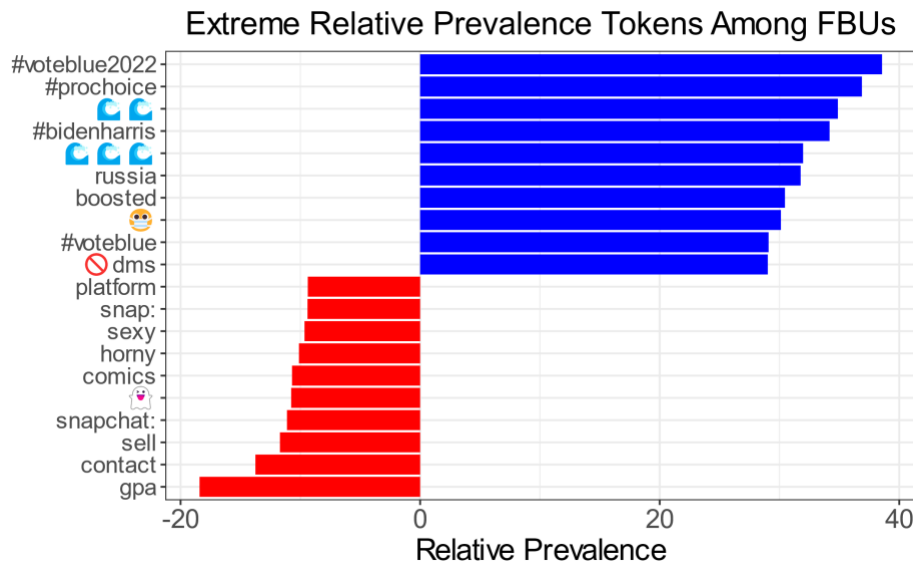


Figure 9. Tokens with the Highest and Lowest Relative Prevalences Among FBUs

As a note, the top six tokens with the highest relative prevalence are not included in Figure 9 because they referred directly to either Ukraine or Russia. In order from highest to lower relative prevalence, the tokens were “ukraine”, “#standwithukraine”, “#istandwithukraine”, “ukrainian”, “putin”, and “stand” (some users say “I stand with Ukraine” instead of using the hashtag). By examining the remaining tokens with the highest and lowest relative prevalence among FBUs, we identified several cooccurring topics that seem more and less likely to be associated with a flag in one’s bio.

Based on the top ten remaining tokens, the major cooccurring topics identified as being most associated with inclusion of the flag in a bio were: pro-democrat messages (“#voteblue2022”), pro-vaccination messages (“boosted”), and generally pro-liberal stances (“#prochoice”), many of which also fall in the “pro-democrat” category. The blue wave emoji ‘🇺🇸’, as mentioned before, is often used by democrats to show support for an

incoming wave of Democratic candidates taking office. The combination of the prohibited emoji “🚫” and “dms” (“🚫dms”) is used by people who do not want to be sent direct messages through Twitter from other users. The connection of this to the Ukraine flag emoji is unclear.

It was more difficult to identify clear cooccurring topics from the bottom ten tokens. This makes sense because there are many more users who do not put the flag into their bio than there are that do. However, there are broad topics we can identify. One of these may be the sharing of account information from alternative social media sites like Snapchat and Instagram. This is based on tokens such as “snapchat:” and the ghost emoji. The token “snapchat:” is commonly used when a user shares their Snapchat username on a different social media site. The ghost emoji is also commonly used for their purpose because the Snapchat logo looks like a ghost. Other tokens like “horny” and “sexy” may be the result of promotional accounts for online pornography. Finally, tokens like “contact”, “sell”, and “consulting” seem to be related to the topic of business. We speculate these are accounts representing companies rather than individuals.

For prevalence among FNUs, the tokens fall into similar categories. Figure 10 below presents some of the highest and lowest ranking relative prevalence tokens among FNUs. Once again, the top five tokens are not presented here because they referred directly to Ukraine or Russia: “istandwithukraine”, “#standwithukraine”, “ukrainian”, “ukraine”, and “putin”. Additionally, only tokens that did not appear in the highest relative prevalence tokens among FBUs are presented here. The tokens that appeared among the highest ranked tokens in both lists were: “#voteblue2022”, “bidenharris”, “prochoice”, “boosted”, and “#voteblue”, and “russia”. The lowest high-ranking token in Figure 10 “#climatechange” is ranked 21 overall.




Figure 10. Tokens with the Highest and Lowest Relative Prevalences Among FNUs

The six tokens shared among both lists again indicate left-leaning and pro-COVID vaccine messaging. The remaining high-ranking tokens in Figure 10 indicate the same thing (“#getvaccinated”, “nevertrump”, “#climatechange”). It was again difficult to identify groups in the bottom tokens, but familiar tokens were seen, such as “snap” (abbreviation for “snapchat”), “ig:” (used to shared Instagram usernames), and “businesses”. The round pushpin emoji “📍” also appears in the lowest ranked tokens. This emoji is often used to share someone’s location.

Overall, this suggests that left-leaning users on Twitter are those more likely to express support for Ukraine using the flag and include that as part of their identity, as opposed to users who uses alternative social media, (presumably) promotional accounts for online pornography, and company/business accounts that may be less likely to do so.

(Lack of) Influence of Bots

To examine the possibility that bot accounts were responsible for the results, we collected Botometer scores for a random sample of accounts. The Botometer API takes in a user's account id and returns a raw score between 0 and 1. A score of 0.5 indicates the classifier is uncertain if the account is a bot or not. First, we collected a random sample of 100 users from each of three groups: non-flag users, flag-bio users, and flag-name users. Most accounts received low Botometer scores. The median score for each group, respectively, was 0.08, 0.12, and 0.07. The mean score for each group, respectively, was 0.23, 0.23, and 0.14. The low FNU scores may indicate that FNUs are less likely to be bot accounts, or it might simply mean there is some correlation between the behaviors FNUs engage in and those Botometer considers non-botlike.

Next, we gathered a random sample of 2,041 non-flag users and 2,009 flag-bio users. We then pared the samples down to include only those accounts with a Botometer score of 0.5 or less. This left us with 1,659 non-flag users and 1,687 flag-bio users. We used these 3,346 accounts to recreate Figure 8a. The results are shown below in Figure 11. We find the same pattern as in Figure 8.

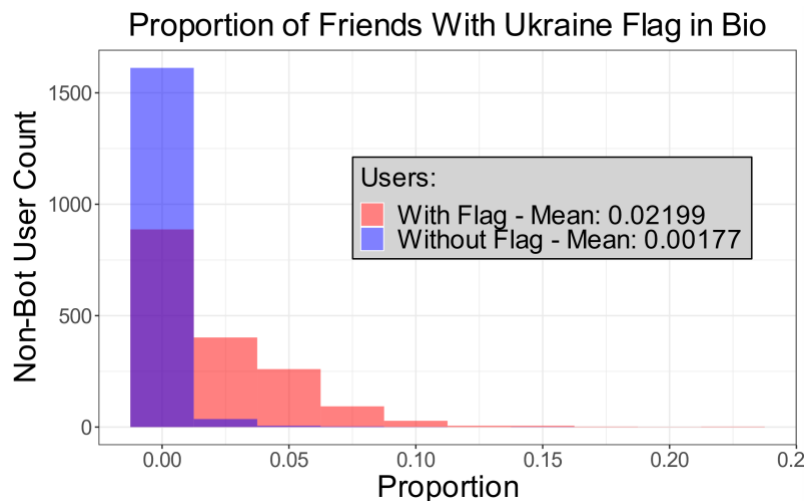


Figure 11: Proportion of Friends with Flag for Low Bot Scoring FBUs

Conclusion and Discussion

Russia's invasion of Ukraine provides a contemporary (and unfortunate) opportunity to observe Identity Activism over time. When the invasion began, Twitter users began displaying the Ukraine flag emoji '🇺🇦' in their names and bios to signal support for Ukraine. The most extreme increase in the prevalence of the Ukraine flag emoji took place during February 21-27, resulting in the flag becoming the 11th most popular emoji in bios and the 2nd most popular emoji in names for the month of March. Although the prevalence in names initially reached levels three times higher than in bios, it has declined substantially since its peak, while the prevalence in bios has not (through June 30). We find that users who choose to display the Ukraine flag emoji in their bio are more likely to follow and be followed by other users who also display the flag in their bio. The same pattern is also found among flag name users. The specific reason for this clustering is unknown, but it reveals a more strongly connected community among Twitter users who are 'identity activists' in support of the same cause as compared to non-activists. Users who choose to display the flag (in either their bio or name) are more likely to include politically left-leaning messages in their bios, such as "#voteblue2022" and "nevertrump". It is more difficult to identify clear groups among users who choose to not display the flag. However, they may be more likely to share usernames from other social media sites or be business accounts.

In the relative prevalence results, particular tokens predict inclusion or exclusion of the Ukraine flag emoji in a bio or name. Using existing tokens as features, a machine learning model could presumably be trained to predict which users will add the flag to their account and which ones will not. Assuming these models' insights could be extracted through explainable artificial intelligence (Miller 2019), automated discovery of what drives IA behavior might replace tedious, laborious single-hypothesis-testing for future study of theories of social movement affiliation.

Our initial investigation focused solely on profile biographies. Exploratory work quickly revealed that users were placing Ukraine flags in their name as well. We consider

both actions Identity Activism by our definition: *display of a social movement symbol within a space reserved for description of the self*. Here, we explored the two categories of users (FBUs and FNUs) separately, but we leave it to future researchers to decide whether to make this distinction.

Using the LOPS method (Jones 2021) to study IA provides multiple benefits. Foremost is the automated production of series of consistent, persistent, precise, high-resolution prevalence estimates. We call these estimates consistent because the operationalization is straightforward (symbol exists or not) and remains the same across time. We call the estimates persistent because they can be continually measured into the future, and precise because it is relatively easy to collect tens of thousands of observations for each estimated proportion. The estimates are available at high spatial resolution (if geolocated) and temporal resolution such as the daily estimates reported here. Because the Twitter platform contains social network and other information, many more possibilities exist than those explored here. For example, one could follow this manifestation (or others) of IA into the future to observe *disaffiliation*. We have observed regularities in the co-occurring tokens of those users currently engaging in Ukraine IA; what tokens might predict those who stop engaging early versus those who stop engaging later? For those who move on from Ukraine IA, where do they take their identity expression next? We encourage all researchers to consider how the concepts demonstrated in the current work could apply to their research questions concerning social movements and identity.

Data and Code

Data and code to reproduce these results are available at <https://osf.io/av7uf/>.

References

- Bloomberg. (2022). A Visual Guide to the Russian Invasion of Ukraine.
<https://www.bloomberg.com/graphics/2022-ukraine-russia-us-nato-conflict/>
- Eady, G., Frederik Hjorth, and Peter Thisted Dinesen. (2021). Do Violent Protests Affect Expressions of Party Identity? Evidence from the Capitol Insurrection.
<https://osf.io/dvgj3/#!>.
- Firmansyah, Firman M., and Jason J. Jones. (2019). Did the Black Panther Movie Make Blacks Blacker? Examining Black Racial Identity on Twitter Before and After the Black Panther Movie Release. *Social Informatics: 11th International Conference, SocInfo 2019, Doha, Qatar, November 18–21, 2019, Proceedings*, pages 66-78.
https://doi.org/10.1007/978-3-030-34971-4_5
- Human Rights Campaign Foundation. (2014). Transgender Visibility: A Guide to Being You. http://assets2.hrc.org/files/assets/resources/trans_guide_april_2014.pdf
- Jeffery, A. and Kevin Breuninger. (2022). Photos Show Global Outrage at Russia's Invasion of Ukraine. *CNBC*. <https://www.cnbc.com/2022/02/24/photos-show-global-outrage-at-russias-decision-to-invade-ukraine.html>
- Jones, Jason Jeffrey. (2021). A Dataset for the Study of Identity at Scale: Annual Prevalence of American Twitter Users with Specified Token in Their Profile Bio 2015–2020. *PLoS ONE* 16(11):e0260185.
<https://doi.org/10.1371/journal.pone.0260185>
- Klandermans, Bert and Dirk Oegema. (1987). Potentials, networks, motivations, and barriers: Steps towards participation in social movements. *American Sociological Review*, 52(4), pages 519-531. <https://doi.org/10.2307/2095297>
- Kutlaca, Maja., & Helena R. M. Radke. (2022). Towards an understanding of performative allyship: Definition, antecedents and consequences. *Social and Personality Psychology Compass*, e12724. <https://doi.org/10.1111/spc3.12724>
- Lakritz, T. (2022). Monuments around the World Are Lighting up in Blue and Yellow in Support of Ukraine. *Insider*. <https://www.insider.com/ukraine-monuments-lit-up-blue-yellow-photos-2022-2>
- Li, Jinhang, Giorgos Longinos, Steven Wilson, and Walid Magdy. (2020). Emoji and

- Self-Identity in Twitter Bios. *Proceedings of the Fourth Workshop on Natural Language Processing and Computational Social Science*, pages 199–211. [10.18653/v1/2020.nlpcss-1.22](https://doi.org/10.18653/v1/2020.nlpcss-1.22)
- Madison, Nora, and Mathias Klang. (2020). The Case for Digital Activism: Refuting the Fallacies of Slacktivism. *Journal of Digital Social Research*, 2(2), Article 2. <https://doi.org/10.33621/jdsr.v2i2.25>
- McAdam, D. (1986). Recruitment to High-Risk Activism: The Case of Freedom Summer. *American Journal of Sociology*, 92(1), 64–90. <https://doi.org/10.1086/228463>
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267, 1-38.
- Morse, J. (2022). People Say They're Using Ukrainian Airbnbs to Send Money to Locals. *Mashable*. <https://mashable.com/article/ukrainian-airbnb-bookings-donations>
- Pathak, Arjunil, Navid Madani, and Kenneth Joseph. (2021). A Method to Analyze Multiple Social Identities in Twitter Bios. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2), pages 1-35. <https://doi.org/10.1145/3479502>
- Rogers, Nick, and Jason J. Jones. (2021). Using Twitter Bios to Measure Changes in Self-Identity: Are Americans Defining Themselves More Politically Over Time?. *Journal of Social Computing*, 2(1), pages 1-13. <https://doi.org/10.23919/JSC.2021.0002>
- Twitter, Inc. (2023). Sample realtime Tweets. <https://developer.twitter.com/en/docs/twitter-api/v1/tweets/sample-realtime/overview>
- Zimmer, Michael, and Nicholas John Proferes. (2014). A Topology of Twitter Research: Disciplines, Methods, and Ethics. *Aslib Journal of Information Management*, 66(3), pages 250-261. <https://doi.org/10.1108/AJIM-09-2013-0083>