

A Computational Analysis of Polarization on Indian and Pakistani Social Media

Aman Tyagi^{[0000-0002-6654-0670]*}, Anjalie Field^{[0000-0002-6955-746X]*}, Priyank Lathwal^[0000-0003-3883-3641], Yulia Tsvetkov^[0000-0002-4634-7128], and Kathleen M. Carley^[0000-0002-6356-0238]

Carnegie Mellon University, Pittsburgh PA 15213, USA
 {amant,anjalief,plathwal}@andrew.cmu.edu,
 {ytsvetko,kathleen.carley}@cs.cmu.edu

Abstract. Between February 14, 2019 and March 4, 2019, a terrorist attack in Pulwama, Kashmir followed by retaliatory airstrikes led to rising tensions between India and Pakistan, two nuclear-armed countries. In this work, we examine polarizing messaging on Twitter during these events, particularly focusing on the positions of Indian and Pakistani politicians. We use a label propagation technique focused on hashtag co-occurrences to find polarizing tweets and users. Our analysis reveals that politicians in the ruling political party in India (BJP) used polarized hashtags and called for escalation of conflict more so than politicians from other parties. Our work offers the first analysis of how escalating tensions between India and Pakistan manifest on Twitter and provides a framework for studying polarizing messages.

Keywords: Polarization · Hashtags · Political communication strategies

1 Introduction

While social media platforms foster open communication and have the potential to offer more democratic information systems, they have simultaneously facilitated divisions in society by allowing the spread of polarizing and incendiary content [23,54]. Polarizing content can be beneficial by encouraging pride and solidarity, but it has also become a social cyber-security concern: foreign and domestic actors may employ polarizing social media content to sow divisions in a country, to demean other nations, or to promote political agendas [4,13,15,34].

Using automated methods to analyze social media offers a way to understand the type of content users are exposed to, the positions taken by various users, and the agendas pursued through coordinated messaging across entire platforms. Understanding the dynamics of this information landscape has become critical, because social media can strongly influence public opinion [13]. However, prior computational social science research on polarization has focused primarily on U.S. politics, and much attention has focused on the influence of Russian or

* Equal Contribution

Chinese state actors [4,5,23,30,34,38,52,54]. In contrast, we focus on polarizing social media content in India and Pakistan and how it can contribute to rising tensions between these two nations. Specifically, we examine communication patterns on Twitter following the terrorist attack in the Pulwama district, Jammu and Kashmir, India, on February 14, 2019.

We primarily investigate: *to what extent did entities on social media advocate for or against escalating tensions?*. India and Pakistan are both nuclear-armed countries and have a decades-long history involving multiple armed conflicts. The Pulwama attack in 2019 was followed by an escalation of tensions between these two nations that nearly approached full-fledged war [26,45,46]. Moreover, the relationship between these countries is an important agenda for political parties in both India and Pakistan. India has two primary political parties: the Indian National Congress (INC), which was dominant in the early 21st century, and the Bharatiya Janata Party (BJP), which rose to prominence on a populist and nationalist platform in 2014 and has been in power since [40]. Given this context, we first examine the tweets and communication patterns of general users in order to understand how polarizing the attack was and to what extent users with different viewpoints may have interacted with each other. We then examine the social media messaging of political party members and how it changed over the sequence of events in order to uncover possible political agendas.

Our core methodology uses a network-based label propagation algorithm to quantify the polarity of hashtags along specified dimensions: Pro-India vs. Pro-Pakistan and Pro-Aggression vs. Pro-Peace. We then aggregate the hashtag-level scores into tweet-level and user-level scores, e.g. the polarity of a given user on a given day. Unlike methodology that assumes users’ opinions do not change [20,21,62], focuses on binary stances [12], or requires in-language annotations and feature-crafting [39], our methodology allows us to analyze degrees of polarization in a multilingual corpus and how they change over time.

We begin by providing an overview of the events between February 14 and March 1, 2019 (§2). Next, we describe the Twitter data collection (§3) and discuss methods (§4) and evaluation (§5). Our results (§6) suggest that more members of the BJP propagated a narrative of escalation than members of other political parties. This finding supports anecdotes reported by journalists [61] about these events. Through this research, we develop (1) the first analysis of escalating tensions between India and Pakistan on Twitter, (2) a data-driven investigation of social media messaging following the 2019 Pulwama attack, and (3) a novel and general methodology to examine polarization on multilingual social media.

2 Timeline of Events

We briefly provide background on relevant events, relying primarily on third party newspapers unaffiliated with either nation (The New York Times and BBC News) and noting where official accounts differ.

Feb. 14, 2019 A 22-year old native of Pulwama carried out a suicide attack against a convoy carrying approximately 2,500 security personnel in the Pul-

wama district in Kashmir, India. The attack resulted in the death of more than 40 Indian soldiers. Jaish-e-Mohammad (JeM) a militant group based in Pakistan (the group is formally banned in Pakistan) claimed responsibility [8,58].

Feb. 14-26, 2019 The Indian government responded to the attack with threats of retaliation against Pakistan, even though Pakistani officials denied any role [9]. Diplomatic ties deteriorated, e.g., India revoked Pakistan’s most favored nation status, which had provided trade advantages. Pakistan threatened to retaliate if India pursued military action [10].

Feb. 26, 2019 The Indian Air Force (IAF) conducted a retaliatory airstrike against a JeM training camp inside Pakistan, which the Indian government termed “non-military, preemptive” [57]. According to Indian government officials, the JeM camp targeted by this airstrike was located 70km inside the Line of Control (LoC) – the military line dividing the Indian and Pakistani controlled parts of Jammu and Kashmir. Indian officials reported that the airstrike was “100 percent successful”, went on “exactly as planned”, and killed over 200 terrorists [44,57]. In contrast, Pakistani officials reported that the target of the attacks was located only 5–6km inside the LoC, that the Pakistani air force turned back the Indian fighters, and that the attacks landed in an empty area [22,55].

Feb. 27, 2019 The Pakistan Air Force (PAF) carried out retaliatory airstrikes along the LoC. Indian and Pakistani officials presented different details of the strikes, but both emphasized de-escalation: a Pakistani official reported that the PAF intentionally targeted open spaces, to demonstrate Pakistan’s capabilities without inviting escalation, while an Indian official reported no deaths or civilian casualties [17,18,24]. However, in aerial combat following the strikes, an IAF pilot was captured by the Pakistani Army [7,56].

Mar. 1, 2019 Pakistan returned the IAF pilot to India on March 1 in what Pakistani Prime Minister Imran Khan called “a gesture of peace” [6].

3 Data

We collected tweets related to these events by first identifying a set of relevant hashtags. Our hashtag set is based on hashtags related to #pulwama found on best-hashtags.com.¹ We modified the hashtag set to ensure that it included both hashtags more likely to be used by Pro-India users (e.g., IndiaWantsRevenge) and hashtags more likely to be used by Pro-Pakistan users (e.g., PakistanZindabad). We then collected all tweets using these terms, either as words or as hashtags during the events.²

Our final data set contains 2.5M unique tweets (including retweets) from 567K users that use 67K unique hashtags. All tweets occurred between February

¹ best-hashtags.com uses an algorithm to provide popular hashtags that are similar to the provided seed (#pulwama). Since our analysis, the website has stopped reporting Twitter hashtags.

² We provide further details, including the full list of keywords, data statistics, network densities and evidence that our data set is comprehensive in our project repository: <https://github.com/amantyaag/india-pakistan-polarization>.

14th and March 4th. The data contains a mix of languages including English, Urdu, and Hindi, and many users use multiple languages in the same tweet. While some tweets express neutral opinions, others contain incendiary language, such as: “@PMOIndia @PMOIndia @narendramodi We r eagerly waiting for ur action of revenge...#PulwamaRevenge #IndiaWantsRevenge #PulwamaAttack” and “I feel time has come to give all support to #Balochistan activist. Let us #bleed Pakistan from all fronts. #NeverForget @PMOIndia @narendramodi #IndiaWantsRevenge”.

4 Methodology

We develop a method to assign a polarity score to an aggregate group of tweets, and we analyze how polarities change over time for different groups of users. For instance, given pole A (e.g., Pro-Pakistan) and pole B (e.g., Pro-India), we aggregate all tweets by a given user and assign the user a polarity score between $[a, b]$, where a score close to a indicates the user more likely supports A and a score close to b indicates the user more likely supports B . We could also aggregate only tweets by the user on one day and determine the user’s Pro- A /Pro- B polarity on that day.

In the absence of annotated data, we use a weakly supervised approach. First, for pole A , we hand-select a small seed set of tokens S_A that are strongly associated with A , and we equivalently hand-select S_B . We assign each $s \in S_A$ a polarity score of a , and we assign each $s \in S_B$ a polarity score of b . Then, we use S_A and S_B to infer polarity scores over a larger lexicon of words or hashtags \mathcal{V} , where each $w \in \mathcal{V}$ is assigned a score in $[a, b]$. Finally, we estimate the polarity of an aggregated set of tweets by averaging the inferred polarity scores for all $w \in \mathcal{V}$ used in those tweets.

In order to propagate the hand-annotated labels in S_A and S_B to the larger lexicon \mathcal{V} , we use 3 variants of graph-based label propagation. In each variant, we construct a graph G , whose nodes consist of $w \in \mathcal{V}$ and whose edges and edge weights are defined based on similarity metrics between members of \mathcal{V} . We describe each variant in detail below.

Network-based Hashtag Propagation In the first variant, we define \mathcal{V} to be the set of all hashtags used in our data set. Then, we construct G as a hashtag*hashtag co-occurrence network. Each node in G corresponds to a hashtag. Edges occur between hashtags that co-occur in the same tweet, and edge weights are proportional to how frequently the hashtags co-occur. Then, we use the label propagation algorithm detailed in Algorithm 1 to infer polarity scores for $w \in \mathcal{V}$ from S_A and S_B , where $a = -1$ and $b = 1$. The algorithm uses a greedy approach to assign labels to each node in G . If all nodes connected to a node n have been labeled, then node n is assigned a weighted average of all the adjacent nodes. This step is repeated until the maximum possible number of nodes are labeled. A low value of γ would label nodes neighboring unlabeled nodes, a high value would only label nodes neighboring unlabeled nodes after multiple iterations of the outer loop.

Our algorithm is similar to methods used to infer user-level polarities, in which a small seed of users is hand-annotated and a graph-based algorithm propagates labels to other users by assuming that users who retweet each other share the same views [21,62]. For example, [29] quantify polarity based on a graph structure by assuming that the controversial topics induce clusters of discussions, commonly referred to as echo-chambers. However, we conduct propagation at a hashtag level, by assuming that hashtags that frequently occur in the same tweets indicate similar polarities. Also, our approach does not assume homophily in retweet network nor that user polarities are constant over time. Graph-based approaches have also been used to examine sentiment or for mixed tweet/hashtag/user-level analyses [19,48].

Network-based Word Propagation The second variant is similar to the first; however, instead of restricting \mathcal{V} to be the set of hashtags in the corpus, we define \mathcal{V} to be the set of all tokens, including words and hashtags. We then construct G as a token*token co-occurrence network, and as above, we infer labels using Algorithm 1 and obtain token-level polarity scores in the range $[-1, 1]$. Expanding \mathcal{V} to all tokens instead of just hashtags allows our algorithm to incorporate more information, but also risks introducing noise, as we do not attempt to process nuances in language like negation.

Embedding-based Word Propagation (SentProp) In the third variant, we define \mathcal{V} to be the set of all tokens, as in the Network-based Word Propagation approach. Then, we train GloVe embeddings [47] over our entire corpus (limiting vocabulary size to 50K). We then use SentProp [31], a method for inferring domain-specific lexicons to infer labels over \mathcal{V} . In this method, as before, we construct a graph G where each $w \in \mathcal{V}$ is a node. However, rather than relying on raw co-occurrence scores, SentProp uses embedding similarity metrics to define edge weights and a random-walk method to propagate labels. We implement SentProp using the SocialSent package [31], where $a = 0$ and $b = 1$.

Once we have obtained hashtag-level or word-level polarity scores, we infer the polarity of a tweet or a group of tweets (e.g. all tweets by a given user) by averaging the polarity scores inferred by our algorithms for all the hashtags and words used in data subset. This approach is similar to the aggregation conducted in [12], but our label propagation allows for the incorporation of thousands of words and hashtags, rather than relying on only a small hand-annotated set. If the data subset does not contain any of the keywords labelled by our algorithm (e.g. in a hashtag-based approach, the tweet contains no hashtags), we consider it unclassified. In some cases, primarily for evaluation, we convert the polarity scores into a ternary negative/neutral/positive position by using the cut-offs $\{< 0, 0, > 0\}$ for the $[-1, 1]$ scale and $\{< 0.5, 0.5, > 0.5\}$ for the $[0, 1]$ scale.

This methodology allows us to infer the polarity of any group of tweets along any dimensions, provided a small set of seed words or hashtags for each dimension. Thus, we can examine how polarities differed for different groups of

Algorithm 1: Label Propagation Algorithm

Input: Graph G with nodes n and edges e with e_{ij} as the edge weight between $i \in n$ and $j \in n$

initialize $\gamma = 50/100$ and $i=0$;

for each n **do**

define $l = \text{integer}(i/\gamma)$; $i+=1$;

for each n **do**

if n **not labeled** **then**

compute $t = \text{neighbors of } n$;

compute $t_l = \text{labeled neighbors of } n$;

if $|t_l| + l \geq t$ **then**

initialize score, c

for each $t_i \in t$ **do**

$\text{score} += \text{label } t_i * e_{nt_i}$; $c += e_{nt_i}$

update $\text{label } n = \text{score}/c$

users and how they changed over time. The two dimensions we focus on are Pro-India/Pro-Pakistan and Pro-Peace/Pro-Aggression. In practice we found that minor variations in the exact words in the seed set had no noticeable impact on our final results. For the network-based methods, we label Pro-India seeds as +1, Pro-Pakistan seeds as -1, Pro-Peace seeds as +1, and Pro-Aggression seeds as -1. For the embedding-based approach, we label Pro-India seeds as +1 Pro-Pakistan seeds as 0, Pro-Peace seeds as +1, and Pro-Aggression seeds as 0. For all word-based approaches, we limit the vocabulary size to 50K.³

Table 1. Classification results for the 100 most followed Indian and Pakistani Twitter accounts, where Pro-India or Pro-Pakistan are treated as the dominant class, and the nationality of the account owner is treated as a gold label. %Unk denotes accounts that our algorithm was unable to classify and %Incorrect denotes accounts that received polar opposite labels (e.g. Indian accounts classified as Pro-Pakistan)

	Pro-India (84 accounts)				Pro-Pakistan (85 accounts)			
	Prec.	Recall	%Unk.	%Incorrect	Prec.	Recall	%Unk.	%Incorrect
Hashtag	0.91	0.25	0.68	0.07	0.90	0.61	0.36	0.02
Word	0.69	0.69	0.24	0.07	0.83	0.35	0.34	0.31
Sentprop	0.48	0.80	-	0.20	0.43	0.15	-	0.85

³ We provide our manually defined seed sets and label propagation code on https://github.com/amantya/india_pakistan_polarization/.

5 Evaluation

Automated Evaluation We first evaluate our methods by focusing on the Pro-India and Pro-Pakistan dimension and assuming that popular users in India are more likely to post Pro-India content and popular users in Pakistan are more likely to post Pro-Pakistan content. From the [Socialbakers.com](https://socialbakers.com) platform, we identified the 100 most followed Twitter accounts in India and in Pakistan. 16 of the Indian accounts and 15 of the Pakistani accounts do not occur in our data, leaving 84 Indian accounts with 2,199 tweets and 85 Pakistani accounts with 1,456 tweets for evaluation. For each account, we average word and hashtag polarities over all tweets from the account, and binarize the resulting score into a Pro-Pakistan or Pro-India position.

Table 1 reports results. Both of the network-based approaches rely on hashtag or word co-occurrences to propagate labels. Thus, hashtags and words that do not have any co-occurrence links to the original seed list are unable to be labeled. For instance, in the hashtag propagation approach, our method labels 41,700 hashtags out of 67,059 total hashtags in the dataset. Any users who only use unlabeled words or hashtags are therefore unable to be classified by our algorithm, resulting in 88/169 unlabeled accounts for the hashtag approach and 49/169 unlabeled accounts for the word approach (%Unk in Table 1). In contrast, SentProp obtains polarity scores for all accounts, as it relies on embedding similarity and can propagate labels between words, even if they do not ever co-occur.

However, although SentProp labels more accounts, its precision is much lower than the network-based methods. The network-based hashtag propagation approach overall obtains the highest precision and the least explicit errors – lower recall scores occur because of accounts that it leaves unlabeled, rather than because of accounts that it labels incorrectly. Although the word-propagation approach labels more accounts and works well over the Indian accounts, its classification of the Pakistani accounts is close to random. We suspect that our method works well for hashtags, because they tend to be strongly polar and indicative of the overall sentiment of the tweet. A word-based approach likely requires more careful handling of subtle language cues like negation or sarcasm.

In our subsequent analysis, we use the network-based hashtag propagation method in order to infer polarities, thus favoring high precision and strong polarization, and choosing not to analyze data where we cannot infer polarity with high-confidence. Additionally, in examining the data set, we found that many of the top-followed accounts in India and Pakistan consisted of celebrities who avoided taking stances on politicized issues, which makes the high number of unclassified accounts in this subset of the data unsurprising.

Manual Evaluation In order to further evaluate our methods, we compare the performance of the network-based hashtag model with a small sample of manually annotated tweets. We randomly sampled 100 users from our data set. For each user, we randomly sampled 1 day on which the user tweeted and aggregated all tweets from that day. Thus, we conduct this evaluation at a per-

Table 2. Inter-annotator agreement and classification accuracy over 100 manually annotated data points

	Krippendorff α	% Agree.	Hashtag Acc.	Soft Hashtag Acc.
India/Pakistan	0.77	88%	74%	89%
Aggression/Peace	0.60	74%	57%	76%

user-per-day level. Two annotators independently annotated each data sample as Pro-India/Pro-Pakistan/Neutral/Can’t Determine and Pro-Peace/Pro-Aggression/Neutral/Can’t Determine. For simplicity, we collapsed Neutral/Can’t Determine and Unclassified into a single “Neutral” label. Notably, the Pro-Peace/Pro-Aggression and Pro-India/Pro-Pakistan dimensions are distinct. For example, users may write tweets that are Pro-Peace and Pro-Pakistan: *“Dont let people pull you into their War, pull them into your Peace... Peace for our World, Peace for our Children, Peace for our Future !! #PakSymbolOfPeace #SayNoToWar”* or that are Pro-Peace and Pro-India: *“Very mature conciliatory speech by #ImranKhan. We now urge him to walk the talk. Please return our #Abhinandan safely back to us. This will go a long way in correcting perceptions and restoring peace. #SayNoToWar”*.

Table 2 reports inter-annotator agreement, which is generally high. Additionally, most disagreements occurred when one annotator labeled Neutral/Can’t Determine and the other did not, meaning polar opposite annotations were rare. If we only count polar opposite labels as disagreements, the percent agreement rises to 94% for both dimensions.

Then, the two annotators discussed any data points for which they initially disagreed and decided on a single gold label for each data point. We compare performance of the network-based hashtag propagation method against these gold annotations in Table 2. In this 3-way classification task, the accuracy of random guessing would be 33%, which our method easily outperforms. In particular, the “Soft” accuracy, in which we only consider the model output to be incorrect if it predicted the polar-opposite label, meaning neutral/unclassified predictions are not considered incorrect, is high for both dimensions.⁴

6 Results and Analysis

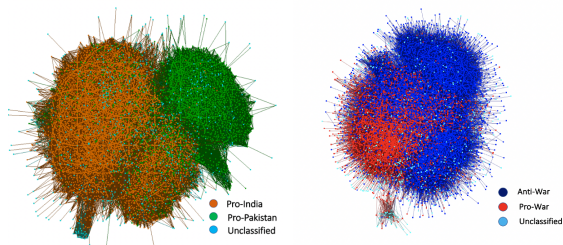
We investigate multiple aspects of our data set, including network structure, polarities of various entities, and changes over time. Based on prior work suggesting that political entities in India and Pakistan may use social media to influence public opinion [2,3,36,51], we pay particular attention to the Twitter accounts of politicians as a method for uncovering political agendas.

What are the overall polarities of our data set? In Table 3, we obtain polarity scores for each user and tweet and then ternarize them into Pro-India/Pro-

⁴ We provide the manual annotations as well as additional metrics on https://github.com/amantyag/india_pakistan_polarization.

Table 3. Overall polarities of users and tweets.

Position	Unique Users	Total Tweets	Position	Unique Users	Total Tweets
Pro-India	125K (23%)	1.16M (46%)	Pro-Aggression	78K (14%)	626K (25%)
Pro-Pakistan	117K (20%)	764K (30%)	Pro-Peace	252K (45%)	1.48M (59%)
Unclassified	325K (57%)	578K (23%)	Unclassified	237K (40%)	351K (16%)

**Fig. 1.** 30-core all communication networks, colored by Pro-India/Pro-Pakistan polarity (left) and Pro-Peace/Pro-Aggression polarity (right). The Pro-India/Pro-Pakistan network displays more homophily than the Pro-Peace/Pro-Aggression network.

Pakistan/Unclassified and Pro-Peace/Pro-Aggression/Unclassified as in §5. At the user level, the classified accounts are approximately balanced between Pro-India and Pro-Pakistan. However, at the tweet level, the classified data contains a high percentage of Pro-India tweets, suggesting Pro-India users tweeted about this issue more prolifically. Further, there is a much higher percentage of Pro-Peace users than Pro-Aggression users. This pattern also holds at the tweet level, where only a small percentage of tweets are unclassified.

What are characteristics of the communication network? Next, we examine the communication network between users, particularly prevalence of echo chambers. Did users with opposite positions interact? Figure 1 shows a 30-core all communication network constructed using ORA-PRO [14]. Accounts are colored based on their Pro-India/Pro-Pakistan polarity (left) and Pro-Peace/Pro-Aggression polarity (right). An edge occurs between two users if one user retweeted, mentioned, or replied to the other and users with ≤ 30 links are not shown. Unsurprisingly, the Pro-India/Pro-Pakistan position is highly segregated, with little interaction between users with different positions. In contrast, the Pro-Peace/Pro-Aggression dimension is more mixed. Although there are some areas of high density for each position, there are interactions between users of different positions, which are potential avenues for users to influence each other’s views.

How polarized were different political entities? We investigate the polarities projected by different political entities: specifically BJP politicians (currently in power in India), INC politicians (largest opposition party), other Indian politicians, and Pakistani politicians. We used the [Socialbakers.com](https://socialbakers.com) platform to obtain the Twitter handles of the 100 most followed politicians in India and

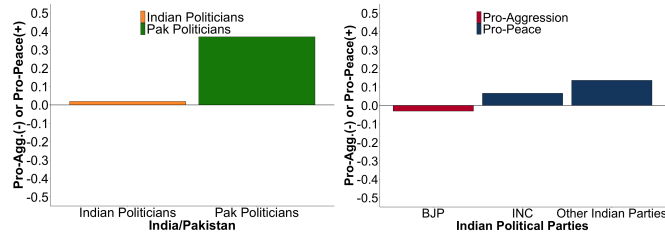


Fig. 2. Aggregate Pro-Peace and Pro-Aggression polarities of the most popular Indian (33/78) and Pakistani (36/66) politicians in our data set (left) and of members of Indian political parties (right).

Pakistan. Our data contained tweets from 66 Pakistani and 78 Indian politicians, and our hashtag model inferred scores for 36 Pakistani and 33 Indian politicians. Figure 2 (left) reports aggregate polarity scores over all tweets from these politicians. Pakistani politicians were predominantly Pro-Peace, while Indian politicians expressed mixed polarities, yielding a near neutral score.

We then examined a broader set of Indian politicians, subdivided by political party based on a list of members running for parliament elections in 2019 [37]. Out of the 1,360 Twitter handles in the list, our data set contained activity from 316 BJP accounts, 281 INC accounts, 204 other Indian party accounts.

Figure 2 (right) shows the overall polarities, aggregated from all tweets by verified members of each party. Strikingly, members of the BJP party are positioned as much more Pro-Aggression than the members of either the INC or other parties, and the party overall obtains a Pro-Aggression polarity score. This score is not dominated by 1-2 strongly polarized members of the party: if we aggregate the polarity scores by individuals instead of by party, 15% of BJP members had net Pro-Aggression scores and 13% had net Pro-Peace scores, in comparison to 10% Pro-Aggression/25% Pro-Peace for INC, and 6% Pro-Aggression/29% Pro-Peace for other parties. The language used by BJP politicians was often openly Pro-Aggression: *#IndiaWantsRevenge We need to give a befitting reply to Pakistan, we will strike back....*

These results support observations made by journalists and community members about the role of the BJP party in these events. BJP is well-known for promoting nationalism, and several journalists have speculated that conflict with Pakistan would increase Prime Minister Modi’s chances of winning the upcoming elections in April and have accused the BJP of war-mongering [28,41,59].

How did polarization change over time? Figure 3 shows how this polarity changed over the two-week period of events: we infer a Pro-Peace/Pro-Aggression polarity score for all tweets posted by members of the specified political subgroup, and we plot the average score across tweets posted each day.

Immediately following the initial attack on 2/14, the tweets from all Indian political party members are inclined towards Pro-Aggression, suggesting initial outrage. However, over the next few days, while tweets from INC and

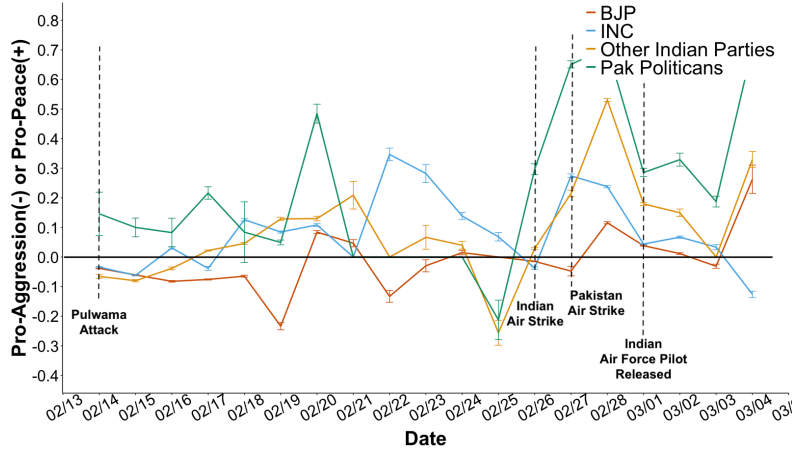


Fig. 3. Daily Pro-Peace/Pro-Aggression positions of political entities. Negative values denote net Pro-Aggression polarity and positive values denote net Pro-Peace. The error bars represent 1 standard deviation.

other Indian political party members switch towards Pro-Peace, tweets from BJP politicians remain consistently Pro-Aggression. There is high volatility between 2/20 and 2/26. However, there was a much lower volume of tweets about the Pulwama incidents during this time period,⁵ and we do not believe these fluctuations are meaningful. The volume of tweets increases once again following the Indian (2/26) and Pakistani (2/27) airstrikes. Tweets by Pakistani politicians generally fall on the Pro-Peace side, but they become more polarized after the Indian airstrike and reach a peak following the Pakistani airstrike. This is consistent with reported quotes by Pakistani officials (§2), saying that the airstrike was designed to avoid escalation. Similarly, tweets by Indian politicians from the INC and other parties become strongly Pro-Peace directly following the Indian airstrike, with polarity increasing after the Pakistani airstrike. In contrast, on the day of the Pakistani airstrike, tweets by BJP politicians remain Pro-Aggression, possibly focusing either on praise for the Indian airstrike or condemnation of the Pakistani airstrike. The polarity of the BJP tweets belatedly switches to Pro-Peace on the following day (2/28), though the strength of the Pro-Peace polarity still remains weaker for BJP tweets than for tweets by other politicians.

7 Discussion and Related Work

The potential that social media platforms have for manipulating public opinion has led to growing interest in information operations and the development of social cyber security as a field of research [15,52]. While we do not claim that social

⁵ Tweet volume is provided in our project repository.

media coverage of the Pulwama incident constituted an information operation, e.g. coordinated efforts to manipulate public opinion and change how people perceive events [52], we do find similarities between our observations and other work in this area. Notably, as described in §2, the Indian and Pakistani governments maintain starkly different accounts about the events that occurred, particularly whether or not the 2/26 airstrikes resulted in 200 casualties. Similarly, Russian and Ukrainian governments circulated conflicting narratives about the cause of the crash of Malaysian Airlines Flight MH17 in 2014, which prompted analyses of information operations about this incident. In a work similar to ours, [30] examine social media coverage of the incident by using a set of hashtags to collect all relevant tweets during a set time frame. Other work has examined the media influence of Chinese and Russian state actors in various domains, including US and UK elections and the Syrian War [27,34,35,50,53]. [4] examine Russian influence in polarizing movements on Twitter, particularly the #BlackLivesMatterMovement, and observe how Russian actors attempted to increase tensions between involved parties. Furthermore, the polarization that we observe in our data align with the “Excite” and “Dismay” strategies, which are tools of public opinion manipulation described in the BEND forms of maneuver [11].

Almost all of these works are focused on U.S. social media, possibly involving Chinese and Russian actors. In general, most work on polarization and public opinion change has focused on U.S. politics [16,21,33], with a few exceptions focusing on Germany [20], Egypt [12,62], and Venezuela [42]. Work on social media in India and Pakistan has focused on healthcare [1], natural disasters [43], self-promotion (e.g. “brand marketing”) primarily in relation to elections [2,3,36], or on election forecasting [32,51], though [60] does argue that the Pakistan Army uses social media to subvert democracy. While these works only focus on intra-country analysis, our work also examines tensions between India and Pakistan. A small selection of work has also looked at the incidents in Pulwama and the implications of rising tensions. [26] and [46] discuss the sociopolitical context and implications of events from a non-computational perspective. [45] additionally conduct a social media analysis, but they use YouTube data and focus on identifying deescalating language. Their timeline of escalation and deescalation is generally consistent with our findings.

Our primary methodology involves using label propagation to infer aggregated polarity scores. In language corpora, label propagation has typically relied on embedding similarity [31,49]. Instead, our approach takes advantage of the short-text nature of Twitter through co-occurrences networks, as well as the strong semantic signals provided by hashtags [25]. Prior methods for analyzing polarization focus on inferring user-level scores [20,21] or require in-language annotations and feature-crafting [39], whereas our method facilitates analyzing how user polarities can change over time in a multilingual corpus.

Conclusions Polarizing language on social media can have long-lasting sociopolitical impacts. Our analysis shows how Twitter users in India and Pakistan used polarizing language during a period of escalating tensions between the two nations, and our methodology offers tools for future work in this area.

Acknowledgements

We thank anonymous reviews and colleagues who provided feedback on this work. The authors would like to acknowledge the support of center for Computational Analysis of Social and Organizational Systems (CASOS), Carnegie Mellon University. This research was also supported in part by Public Interest Technology University Network Grant No. NVF-PITU-Carnegie Mellon University-Subgrant-009246-2019-10-01. The second author of this work is supported by NSF-GRFP under Grant No. DGE1745016. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

1. Abbasi, R.A., Maqbool, O., Mushtaq, M., Aljohani, N.R., Daud, A., Alowibdi, J.S., Shahzad, B.: Saving lives using social media: Analysis of the role of twitter for personal blood donation requests and dissemination. *Telematics and Informatics* **35**(4), 892–912 (2018). <https://doi.org/10.1016/j.tele.2017.01.010>
2. Ahmed, S., Jaidka, K., Cho, J.: The 2014 Indian elections on Twitter: A comparison of campaign strategies of political parties. *Telematics and Informatics* **33**(4), 1071–1087 (2016). <https://doi.org/10.1016/j.tele.2016.03.002>
3. Antil, A., Verma, H.V.: Rahul Gandhi on Twitter: An analysis of brand building through Twitter by the leader of the main opposition party in India. *Global Business Review* **0**(0) (2019). <https://doi.org/10.1177/0972150919833514>
4. Arif, A., Stewart, L.G., Starbird, K.: Acting the part: Examining information operations within# blacklivesmatter discourse. In: *Proc. of CSCW*. p. 20 (2018)
5. Badawy, A., Ferrara, E., Lerman, K.: Analyzing the digital traces of political manipulation: The 2016 Russian interference Twitter campaign. In: *Proc. of ASONAM*. pp. 258–265 (2018)
6. BBC News: Abhinandan: Captured Indian pilot handed back by Pakistan. <https://www.bbc.com/news/world-asia-47412884> (2019)
7. BBC News: India Pakistan: Kashmir fighting sees Indian aircraft downed. <https://www.bbc.com/news/world-asia-47383634> (2019)
8. BBC News: Kashmir attack: Tracing the path that led to Pulwama. <https://www.bbc.com/news/world-asia-india-47302467> (2019)
9. BBC News: Pulwama attack: India will ‘completely isolate’ Pakistan. <https://www.bbc.com/news/world-asia-india-47249133> (2019)
10. BBC News: Pulwama attack: Pakistan warns India against military action. <https://www.bbc.com/news/world-asia-india-47290107> (2019)
11. Beskow, D.M., Carley, K.M.: Social cybersecurity: an emerging national security requirement. *Military Review* **99**(2), 117 (2019)
12. Borge-Holthoefer, J., Magdy, W., Darwish, K., Weber, I.: Content and network dynamics behind egyptian political polarization on twitter. In: *Proc. of CSCW*. p. 700711 (2015). <https://doi.org/10.1145/2675133.2675163>
13. Bradshaw, S., Howard, P.N.: Challenging truth and trust: A global inventory of organized social media manipulation. *The Computational Propaganda Project* (2018)
14. Carley, K.M.: ORA: A toolkit for dynamic network analysis and visualization. *Encyclopedia of Social Network Analysis and Mining* (2017). https://doi.org/10.1007/978-1-4614-7163-9_309-1

15. Carley, K.M., Cervone, G., Agarwal, N., Liu, H.: Social cyber-security. In: Proc. of SBP-BRiMS. pp. 389–394 (2018)
16. Chen, S., Khashabi, D., Yin, W., Callison-Burch, C., Roth, D.: Seeing things from a different angle: Discovering diverse perspectives about claims. In: Proc. of NAACL. pp. 542–557 (2019)
17. CNBC: India and Pakistan say theyve launched airstrikes against each other. Heres what you need to know. <https://www.cnbc.com/2019/02/27/india-pakistan-air-strike-claims-what-you-need-to-know.html> (2019)
18. CNBC: Pakistan says it shot down Indian jets, carried out air strikes in Kashmir. <https://www.cnbc.com/2019/02/27/indian-air-force-plane-crashes-in-kashmir-says-indian-police-official.html> (2019)
19. Coletto, M., Esuli, A., Lucchese, C., Muntean, C.I., Nardini, F.M., Perego, R., Renso, C.: Sentiment-enhanced multidimensional analysis of online social networks: Perception of the Mediterranean refugees crisis. In: Proc. of ASONAM. pp. 1270–1277 (2016)
20. Darius, P., Stephany, F.: Twitter “Hashjacked”: Online polarisation strategies of Germany’s political far-right. In: Proc. SocInfo. pp. 188–201 (2019)
21. Darwish, K.: Quantifying polarization on twitter: The Kavanaugh nomination. In: Proc. SocInfo. pp. 188–201 (2019)
22. Dawn: Indian aircraft violate LoC, scramble back after PAF’s timely response: ISPR. <https://www.dawn.com/news/1466038/indian-aircraft-violate-loc-scramble-back-after-pafs-timely-response-ispr> (2019)
23. Demszky, D., Garg, N., Voigt, R., Zou, J., Shapiro, J., Gentzkow, M., Jurafsky, D.: Analyzing polarization in social media: Method and application to tweets on 21 mass shootings. In: Proc. of NAACL. pp. 2970–3005 (2019). <https://doi.org/10.18653/v1/N19-1304>
24. Economic Times: How Pakistan failed to do a Balakot-type strike on India on February 27. <https://economictimes.indiatimes.com/news/defence/how-pakistan-failed-to-do-a-balakot-type-strike-on-india-on-february-27/articleshow/68592269.cms> (2019)
25. Ferragina, P., Piccinno, F., Santoro, R.: On analyzing hashtags in twitter. In: Proc. of ICWSM (2015)
26. Feyyaz, M.: Contextualizing the Pulwama attack in Kashmir—a perspective from Pakistan. *Perspectives on Terrorism* **13**(2), 69–74 (2019)
27. Field, A., Kliger, D., Wintner, S., Pan, J., Jurafsky, D., Tsvetkov, Y.: Framing and agenda-setting in Russian news: a computational analysis of intricate political strategies. In: Proc. of EMNLP. pp. 3570–3580 (2018)
28. Forbes: India’s fight with Pakistan seen lifting Modi’s election chances. <https://www.forbes.com/sites/kenrapoza/2019/02/27/indias-fight-with-pakistan-seen-lifting-modis-election-chances/#1df2795a397c> (2019)
29. Garimella, K., Morales, G.D.F., Gionis, A., Mathioudakis, M.: Quantifying controversy on social media. *Trans. Soc. Comput.* **1**(1) (2018). <https://doi.org/10.1145/3140565>
30. Golovchenko, Y., Hartmann, M., Adler-Nissen, R.: State, media and civil society in the information warfare over Ukraine: citizen curators of digital disinformation. *International Affairs* **94**(5), 975–994 (2018)
31. Hamilton, W.L., Clark, K., Leskovec, J., Jurafsky, D.: Inducing domain-specific sentiment lexicons from unlabeled corpora. In: Proc. of EMNLP. vol. 2016, p. 595 (2016)

32. Kagan, V., Stevens, A., Subrahmanian, V.S.: Using Twitter sentiment to forecast the 2013 Pakistani election and the 2014 Indian election. *IEEE Intelligent Systems* **30**(1), 2–5 (2015)
33. Khosla, S., Chhaya, N., Jindal, S., Saha, O., Srivastava, M.: Do events change opinions on social media? studying the 2016 US Presidential debates. In: *Proc. of SocInfo*. pp. 287–297 (2019)
34. King, G., Pan, J., Roberts, M.E.: How the Chinese government fabricates social media posts for strategic distraction, not engaged argument. *American Political Science Review* **111**(3), 484–501 (2017)
35. Kriel, C., Pavliuc, A.: Reverse engineering Russian Internet Research Agency tactics through network analysis. *Defence Strategic Communication* pp. 199–227 (2019)
36. Kumar, A., Dhamija, S., Dhamija, A.: Political marketing: The horizon of present era politics. *SCMS Journal of Indian Management* **13**(4), 116–125 (2016)
37. Kumaraguru, P., Singh, S., Manu, D., Gupta, K., Sadaria, A., Srikanth, S., Bhatta, H., Garimella, K., Buggana, S., Agarwal, A., Kapoor, A., Gupta, K., Garg, T., Gurjar, O., Saini, S.: Social media to win elections: Analysis of #LokSabhaElections2019 in India. *Precog Technical report* (2019)
38. Le, H., Boynton, G., Shafiq, Z., Srinivasan, P.: A postmortem of suspended Twitter accounts in the 2016 US presidential election. In: *Proc. of ASONAM*. pp. 258–265 (2019)
39. Magdy, W., Darwish, K., Abokhodair, N., Rahimi, A., Baldwin, T.: #isisisnotislam or #deportallmuslims? Predicting unspoken views. In: *Proc. of WebSci*. pp. 95–106 (2016)
40. McDonnell, D., Cabrera, L.: The right-wing populism of India's Bharatiya Janata Party (and why comparativists should care). *Democratization* **26**(3), 484–501 (2019). <https://doi.org/10.1080/13510347.2018.1551885>
41. Mishra, S.: Emerging electoral dynamics after Pulwama tragedy. *Observer Research Foundation* (2019)
42. Morales, A., Borondo, J., Losada, J.C., Benito, R.M.: Measuring political polarization: Twitter shows the two sides of Venezuela. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **25**(3), 033114 (2015)
43. Murthy, D., Longwell, S.A.: Twitter and disasters. *Information, Communication & Society* **16**(6), 837–855 (2013). <https://doi.org/10.1080/1369118X.2012.696123>
44. NDTV: India strikes after Pulwama terror attack, hits biggest Jaish-e-Mohammed camp in Balakot. <https://www.ndtv.com/india-news/india-struck-biggest-training-camp-of-jaish-in-balakot-large-number-of-terrorists-eliminated-governm-1999390> (2019)
45. Palakodety, S., KhudaBukhsh, A.R., Carbonell, J.G.: Hope speech detection: A computational analysis of the voice of peace. In: *Proc. of ECAI* (2020)
46. Pandya, A.: The future of Indo-Pak relations after the Pulwama attack. *Perspectives on Terrorism* **13**(2), 65–68 (2019)
47. Pennington, J., Socher, R., Manning, C.: GloVe: Global vectors for word representation. In: *Proc. of EMNLP*. pp. 1532–1543 (2014)
48. Pollacci, L., Sirbu, A., Giannotti, F., Pedreschi, D., Lucchese, C., Muntean, C.I.: Sentiment spreading: An epidemic model for lexicon-based sentiment analysis on Twitter. In: Esposito, F., Basili, R., Ferilli, S., Lisi, F.A. (eds.) *Proc. of AI*IA*. pp. 114–127 (2017)
49. Rothe, S., Ebert, S., Schütze, H.: Ultradense word embeddings by orthogonal transformation. In: *Proc. of NAACL*. pp. 767–777 (2016)

50. Rozenas, A., Stukal, D.: How autocrats manipulate economic news: Evidence from Russias state-controlled television. *The Journal of Politics* **81**(3) (2019)
51. Singh, P., Kumar, K., Kahlon, K.S., Sawhney, R.S.: Can tweets predict election results? insights from Twitter analytics. In: Luhach, A.K., Jat, D.S., Hawari, K.B.G., Gao, X.Z., Lingras, P. (eds.) *Proc. of ICAICR*. pp. 271–281 (2019)
52. Starbird, K., Arif, A., Wilson, T.: Disinformation as collaborative work: Surfacing the participatory nature of strategic information operations. In: *Proc. of CSCW*. pp. 1–26 (2019)
53. Starbird, K., Arif, A., Wilson, T., Van Koevering, K., Yefimova, K., Scarnecchia, D.: Ecosystem or echo-system? exploring content sharing across alternative media domains. In: *Proc. of ICWSM* (2018)
54. Stewart, L.G., Arif, A., Nied, A.C., Spiro, E.S., Starbird, K.: Drawing the lines of contention: Networked frame contests within #BlackLivesMatter discourse. In: *Proc. of CSCW*. pp. 1–23 (2017)
55. The Express Tribune: No blood. no bodies. no debris. no tragedy. <https://tribune.com.pk/story/1919080/1-no-blood-no-bodies-no-debris-no-tragedy/> (2019)
56. The Hindu: IAF plane shot down, pilot taken captive by Pak. army. <https://www.thehindu.com/news/national/iaf-plane-shot-down-pilot-taken-captive-by-pak-army/article26390980.ece> (2019)
57. The Hindu: India bombs Jaish camp in Pakistans Balakot. <https://www.thehindu.com/news/national/air-strikes-hit-balakot-in-pakistan-initial-assessment-100-hit-sources/article26373318.ece> (2019)
58. The New York Times: Kashmir suffers from the worst attack there in 30 years. <https://www.nytimes.com/2019/02/14/world/asia/pulwama-attack-kashmir.html> (2019)
59. The Week: Imran’s party slams Modi’s ‘warmongering’ for escalating Indo-Pak tension. <https://www.theweek.in/news/india/2019/02/28/imran-s-party-slams-modi-on-escalating-indo-pak-tension.html> (2019)
60. Upadhyay, A.: Decimating democracy in 140 characters or less: Pakistan armys subjugation of state institutions through Twitter. *Strategic Analysis* **43**(2), 101–113 (2019). <https://doi.org/10.1080/09700161.2019.1600823>
61. Washington Post: After Pulwama, the Indian media proves it is the BJP’s propaganda machine. <https://www.washingtonpost.com/opinions/2019/03/04/after-pulwama-indian-media-proves-it-is-bjps-propaganda-machine/> (2019)
62. Weber, I., Garimella, V.R.K., Batayneh, A.: Secular vs. Islamist polarization in Egypt on Twitter. In: *Proc. of ASONAM*. pp. 290–297 (2013). <https://doi.org/10.1145/2492517.2492557>