

Analyzing Large-Scale Public Campaigns on Twitter

Julia Proskurnia¹, Ruslan Mavlyutov², Roman Prokofyev², Karl Aberer¹, and Philippe Cudré-Mauroux²

¹ École Polytechnique Fédérale de Lausanne, Lausanne—Switzerland
{iuliia.proskurnia, karl.aberer}@epfl.ch

² eXascale Infolab, U. of Fribourg, Fribourg—Switzerland
{firstname.lastname}@unifr.ch

Abstract. Social media has become an important instrument for running various types of public campaigns and mobilizing people. Yet, the dynamics of public campaigns on social networking platforms still remain largely unexplored. In this paper, we present an in-depth analysis of over one hundred large-scale campaigns on social media platforms covering more than 6 years. In particular, we focus on campaigns related to *climate change* on Twitter, which promote online activism to encourage, educate, and motivate people to react to the various issues raised by climate change. We propose a generic framework to identify both the type of a given campaign as well as the various actions undertaken throughout its lifespan: official meetings, physical actions, calls for action, publications on climate related research, etc. We study whether the type of a campaign is correlated to the actions undertaken and how these actions influence the flow of the campaign. Leveraging more than one hundred different campaigns, we build a model capable of accurately predicting the presence of individual actions in tweets. Finally, we explore the influence of active users on the overall campaign flow.

1 Introduction

Social media have become central to our digital lives, as they allow individuals to share news, photos, or opinions, as well as to have online discussions in real-time. One particularly interesting phenomenon is social media marketing, which can be defined as the process of drawing attention to some specific issue or product via social media platforms. Such endeavors often take the form of extensive campaigns, whose aim is to raise the awareness of the public on a particular topic and potentially to engage it into concrete actions.

Social media platforms provide tools to effectively conduct these campaigns; On Twitter, for example, people use so-called *hashtags* to associate their messages to a certain topic. Many campaigns, therefore, have their own hashtags that uniquely identify them. Moreover, many tweets associated with a campaign convey some specific messages to the audience, such as requests for signing a petition, asking for a certain action, attending a demonstration, etc. These messages can be considered as certain *actions*, and their effect on the dynamics of the campaigns remains largely unexplored in the scientific literature. Identifying and categorizing such messages within the context of a campaign would enable us to answer questions such as what drives attention to a particular topic or how to reach a certain target audience. In this work, we propose a number

of categories to classify the actions from the perspective of the goals of the campaigns as well as a methodology to identify them. We build a classifier for the action types based on the tweets content and study the distribution of these action types for different types of campaigns.

In the second part of this work, we analyze the resulting user involvement patterns in order to explore the dynamics of the campaigns. Analyzing such patterns is key to understand how attractive the campaigns are and who are the main contributors to the information dissemination. We perform a comparative analysis of the campaigns and their contents, through which we identify noticeable differences between the various types of campaigns. We observe that campaigns where only a tiny fraction of users create the major part of the content are less likely to attract users on social media. Finally, we cluster the involvement patterns and study their correlations with the campaign types.

This work focuses on campaigns related to climate change and animal welfare. Those two topics recently gained increased attention and have the advantage of gathering a high number of users for relatively long periods of time, thus are well suited for our study. Moreover, these topics are mainly of interest for non-profit and governmental organizations, and this work might help them to better understand the impact of their actions on the audience.

In summary, the main contribution of this paper is a large-scale study on the dynamics of campaigns on social media. This study focuses on the following research questions:

- How to identify and compare various types of public campaigns and their corresponding actions? (Section 4)
- How is the initial goal and contents of a campaign correlated to the user engagement pattern of a campaign? (Section 4.2);
- Is there a relationship between the type of a campaign and the actions undertaken through the campaign lifespan? (Section 5);

The rest of this paper is structured as follows. We start with an overview of related work in the areas of Twitter analytics and social media analysis below in Section 2. Section 3 describes our data collection, aggregation, and cleansing processes. We analyze the collected data in Section 4 by extracting different types of campaigns and clustering them by their user engagement patterns. Section 5 extends our analysis by focusing on various types of tweets and on their distributions in campaigns, as well as building a classifier that is able to predict the type of a tweet. Finally, we discuss our results in Section 6 and draw conclusions in Section 6.

2 Related Work

Social media platforms quickly came to the attention of the research community, since they allow to conduct large-scale studies on various aspects of social network dynamics, such as popularity prediction. Many studies have recently focused on micro-blogging platforms such as Twitter³, which provides an access to a small (compared to the overall data) sample of its data based on keyword queries.

³ <https://twitter.com>

In this work, we study the communication patterns and message type preeminence for various campaigns on climate change. A number of studies have focused on Twitter communication patterns, including studies on hashtag life-cycles [17,19], event detection and their analysis [12,20], food consumption patterns [1], and usage across different languages [10].

Climate change discourse. Climate change issues are receiving increased attention as they lead to a number of global challenges [3,11]. Many studies recently examined how the climate change debate is covered on social media channels [20,22,26]. However, coverage of debates does not reveal how campaigns develop, and how popular they are based on the messages they contain. As users tend to increasingly rely on their social entourage to filter information [9], we examine in this paper how different message types and techniques engage people in different ways throughout the campaign.

Campaign analytics Social media is a very influential tool for widening public awareness on various issues as noted by [25]. Previous work on campaigns on social media mostly focused on political and protest campaigns. [13] used a bispace model based on a Poisson process to capture the propagation of information in both Twitter and non-Twitter environments. Additionally, [8] explores how social networks are used to spread protest information. In our work we focus not on the information dissemination but rather how the campaigns were conducted and what are the main actions that were taken to reach the goal. An in-depth study on the theoretical principles underpinning public communication campaigns is given in [4]. Finally, one of the most recent works on campaign analysis [6] focuses on the behavioural stage sequences of the users during the COP21 and EH2015 forums and proposes a framework to identify a user stage by her tweets. On the contrary, we focus on the campaign actions and the corresponding users' engagement rather than user behavioural stages. Moreover, our analysis is carried out on over a hundred public campaigns.

Tweet topic identification In the context of topic identification, recent works focused on classifying and clustering tweets based on their topics [5,16,21]. Those techniques produce different sets of topics for different datasets. [2] performs an extensive evaluation of different tweet topic detection methods, including methods based on the combinations of syntactic and linguistic techniques. In our case, however, such approaches did not result in valid clusters of message types. To the best of our knowledge, this work is the first on tweet action type classification in campaigns. The work of [23] was an important motivation for the definition of further types of tweets, such as official meetings, calls for action and physical actions. Given our objective of comparing campaign agendas, we look into a number of types of campaigns and actions in this paper in order to identify the correlation between the types of campaigns and the different actions.

3 Data Collection and Cleansing

In this section, we first describe the process through which we collected tweets related to the domains of climate change and animal welfare (Section 3.1). We then introduce the strategy we took for identifying campaigns in those domains. We describe the process

of identifying the retweets and duplicated tweets in detail in Appendix B. The resulting dataset, consisting of more than 8.5M tweets, is available online for future research⁴.

3.1 Twitter data collection

We developed a data collection pipeline (see Figure 1) to gather a broad range of Twitter campaigns related to climate change and animal welfare. Those two domains are usually tightly connected [24]. For example, there are multiple articles [15] on the connection between the number of farm animals and the amount of methane released to the atmosphere and thus causing climate change.

First pass We proceeded in two phases in order to identify the campaign. In the first pass, we extracted all available tweets from Topsy⁵ for two very prominent accounts that are related to climate change and animal welfare-related: @AlGore and @GreenPeace (2.77M and 1.33M followers respectively). This first pass resulted in 27K tweets comprising 1250 unique hashtags. To select valid campaign hashtags

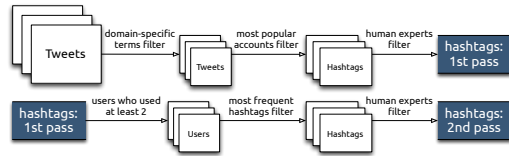


Fig. 1. Data collection pipeline

decided to rely on the annotations made by three authors of this paper. To determine whether a particular hashtag belongs to a given campaign, the authors were asked the following questions: (a) do tweets with the hashtag contain calls for actions, mentions of the campaign or an URL to the campaign website; (b) is there a Twitter account with an identical name and a description that corresponds to a given campaign. A hashtag was considered to be related to a campaign iff it was selected by all of the annotators⁶. This manual annotation produced a set of 52 campaign hashtags.

Second pass To increase the recall of our process, we ran a second pass. We identified further accounts (users) that mentioned at least two campaign tags (out of the first 52) in their messages. In that way, we identified 80 additional accounts for a total of 34K unique hashtags. We filtered out hashtags that appeared in less than 50 tweets, which accounted for 75% of the tweets. Similarly to the first pass, three authors of this paper annotated each resulting hashtag and 56 additional hashtags were identified as relevant. Overall, our process resulted in a dataset of 108 climate and animal welfare-related campaigns⁷, each represented by a distinct Twitter hashtag. The total number of unique tweets (without retweets) in the resulting dataset amounts to 4M.

⁴ <https://github.com/toluoll1/CampaignsDataRelease>

⁵ Topsy (<http://topsy.com/>) is a partner of Twitter delivering search and analytic services and claiming to index all public tweets.

⁶ inter-rater agreement is $\sim 95\%$

⁷ It is worth noticing that many of the hashtags (around 20 each) in our campaign dataset are created using the morphological filters. For example, we collected hashtags that contain words such as save, protect, call, lead, act, 4, forthe, etc. (e.g #savethedolphins, #call4action).

4 Campaign analysis

Given that our research question connects two domains—climate change / animal welfare campaigns and social media content analysis—the framework we propose for campaign analysis is composed of two parts. First, we annotate the campaigns according to their primary goals. Next, we cluster them by examining user engagement patterns and by mining active users for the campaigns (i.e., users who tweet most often for a particular campaign). When organizing our data and defining the annotation process, we turned to dimensions considered in the theory of public communication campaigns [4,23,14]. For each campaign, we consider the major goal of the campaign (increase awareness, mobilize people), user engagement over time (ever-growing, regular, one-day, inactive), as well as user activity.

4.1 Types of campaigns

Following the theoretical analysis of public communication campaigns by [4], we separate the campaigns into two classes based on their primary goals:

- *Mobilization campaigns* refer to the campaigns whose primary goal is to engage and motivate a wide range of partners, allies and individual at the national and local levels towards a particular problem or issue.
- *Awareness campaigns* refer to the campaigns whose primary goal is to raise people’s awareness regarding a particular subject, issue, or situation. As discussed in Section 2, environmental awareness campaigns usually make a large use of mass media, and in particular, of Twitter.

These campaign types represent very different endeavors, which affects both the type of contents used in such campaigns as well as their user involvement pattern over time, which we analyze further.

Three authors of this paper manually annotated the campaigns as either mobilization or awareness campaigns. The category was considered as valid only when all experts agreed on it. This way, 50 awareness and 58 mobilization campaigns were identified. A few sample hashtags are #savesolar, #climateaction for **mobilization** and #cleanair4kids, #worldfoodday for **awareness** campaigns.

4.2 User engagement patterns

In the following, we present an analysis of user engagement in Twitter campaigns. We identify two main axes for analyzing user engagement: the first one focuses on user engagement patterns over time, while the second one analyzes the behavior of the most active users throughout the campaign.

User engagement patterns over time Subsequently, we cluster the campaigns by engagement patterns of their users to detect whether the engagement correlates with the campaign types. In order to do this, we first extract the number of unique daily users for each campaign hashtag and aggregate these numbers with a 30-day time window.

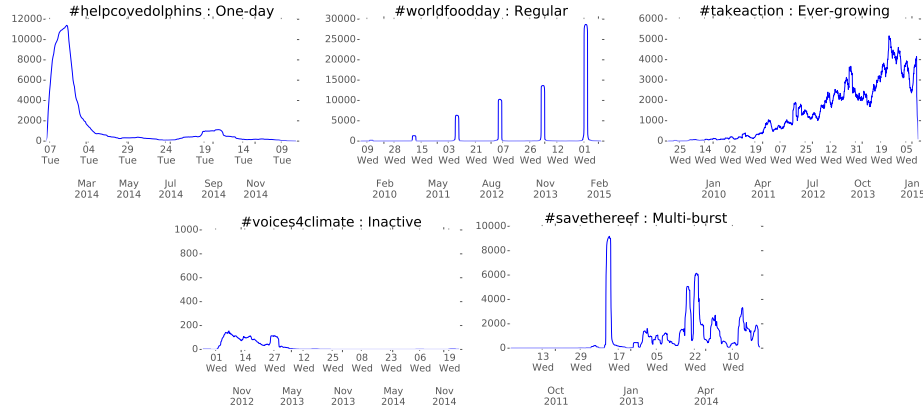


Fig. 2. Different user engagement patterns observed in campaigns.

Then, we compute the similarities between the resulting time series using Dynamic Time Warping [7] and cluster them using K-means by varying the K and chose the setup with the smallest in-cluster distance. This resulted in five major clearly distinguishable clusters. Sample campaigns with the above described types are shown on Figure 2.

From our data collection through this process we have identified several major types of user involvement patterns. Following their overarching distribution, we name them:

- *one-day campaign*, a campaign that is organized over a short period of time to tackle some urgent issue;
- *regular campaign*, a campaign that happens on a regular basis, e.g., annually;
- *ever-growing campaign*, a campaign that gains traction over time;
- *multi-burst campaign*, a campaign that have multiple peaks of activity;
- *inactive campaign*, a campaign that shows a constantly low user engagement throughout its timespan⁸

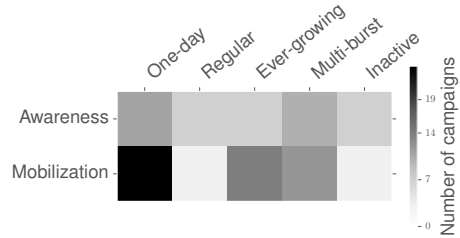


Fig. 3. Distribution of user engagement patterns for the different types of campaigns.

Finally, we compare the representations of two major classes of campaigns with their user involvement patterns. The campaigns are distributed across the aforementioned engagement groups as 36%, 10%, 21% 22%, 11%. As can be observed on Figure 3, most of “regular” and “inactive” campaigns fall in the awareness category, while

⁸ “Inactive” category might be orthogonal to the other ones, however, it gives valuable insights regarding campaigns that have less traction on Twitter

both “one-day” and “ever-growing” campaigns are dominated by the mobilization one. The main reason for the dominance of mobilization campaigns for the “one-day” type is the urgency of their issues and the need for immediate action. On the other hand, “regular” campaigns, that are organized on a periodic basis and pursue long-term goals consist of awareness campaigns mostly. “Ever-growing” campaigns also dominated by mobilization campaigns and focus on global issues, challenges, e.g., #saveanimals, #animalwelfare. “Multi-burst” campaigns are almost equally represented by the both types.

User engagement patterns by volume We observe that in many campaigns, there is a distinct subset of users who are authors of the majority of the campaign tweets. We call such a set of users a *campaign kernel*. A *kernel* identifies users with the most tweets and retweets in the campaign. In order to study the influence of a kernel, we propose the following technique: (1) for each user, compute the total number of original tweets and retweets posted in the campaign; (2) rank all users relatively to the volume of content produced for the campaign; (3) compute the Gini coefficient⁹ based on the normalized per-user impact relative to the volume of messages in the campaign.

Figure 4(a) shows sample distributions of the relative amount of content generated by users participating in campaigns. We observe a clear distinction between campaigns where users are contributing the content almost equally (blue curve) and campaigns where only a tiny fraction of users create the major part of the content (red curve). Table 4(b) shows campaigns with the lowest and the highest Gini coefficient values. High values denote campaigns where the majority of the contents is created by few users only. Values that are close to zero, on the other hand, characterize campaigns where users contribute almost equally. Characteristics of the Gini coefficient play an important role in understanding how an information about campaigns is spread on social media, e.g., if actions of several active people can lead to more participation.

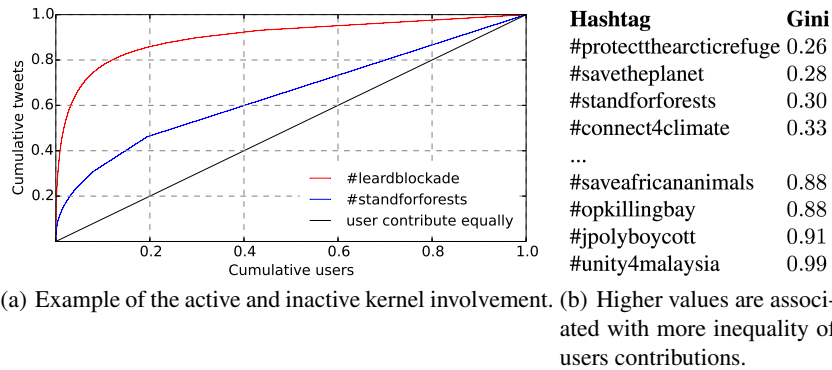



Fig. 4. Campaign kernel contributions presented as gini coefficient.

⁹ http://wikipedia.org/wiki/Gini_coefficient

Interestingly, we found a direct correlation between the total number of followers of the kernel users and the total amount of users participating in a campaign. The value of Pearson correlation coefficient for these variables is more than 0.85. This behavior is observed for various kernel sizes (that correspond to various proportions of total tweets they generated). The distribution of Pearson correlation coefficient between the number of kernel followers and the number of users engaged in the campaign for different kernel sizes has the shape of a bell curve , where thresholds for proportions of tweets are displayed on the x axes (50% - 99%), and Pearson correlation is on the y axes (0.70 - 0.858). The maximum correlation is reached for the kernels that produced 75% of the content and on average this corresponds to the 2.5% of the campaign users. Thus, we use this percentage of users further as a kernel of each campaign. Interestingly, we found no clear distinction between awareness and mobilization campaigns with regard to their kernels. However, the activity of kernel users differs with respect to the user engagement patterns described in 4.2. We observe that the majority of content in inactive campaigns is produced by a tiny fraction of users, while regular and ever-growing campaigns accumulate tweets from a much larger subset of users. Similarly, the kernels of inactive campaigns are 10% smaller than one-day campaigns, while ever-growing campaigns have both small kernels and high participation of the involved users.

5 Tweet type identification and classification

This section presents an in-depth analysis of the tweets from our dataset, focusing on the types of actions they contain (Section 5.1), and their correlations with the campaign types (Section 5.2). Generally, public campaigns perform a specific set of actions to reach their goals. While this information is not always publicly available, we propose a method to retrospectively identify campaign's actions based on their presence in Twitter. Such findings can be of great interest to less experienced campaign stakeholders, who could use this knowledge to adjust their agendas according to the most successful practices given their particular type of a campaign. We collected additional information about the tweets through a large-scale crowdsourcing experiment (Section 5.1), in order to collect enough annotations to build a supervised model capable of accurately predicting the type of action contained in a given tweet (Section 5.1).

5.1 Types of Tweets

As discussed in Section 2, the campaigns that are the most effective at influencing users are typically related to either promoting some positive behavior or preventing some negative actions [4]. In our context, prevention campaigns typically focus their attention on negative consequences rather than on positive alternatives. This introduces our first class of protest-related actions: physical actions [23]. Next, awareness campaigns that promote positive behaviors try to actively connect with either informational or instructional resources [14,22]. This motivates the definitions of two further types of actions: publications and calls for actions. Since most campaigns have some sort of supporters or base community, when running a campaign it is important to focus not only on the general public but also on specific stakeholders, e.g., to empower important communities, activate voluntary associations, or collaborate with governmental agencies.

This often prompts the campaigns to organize official meetings, conferences, and debates [14,22,23]. Taking the above information into account, we consider five different classes of Tweets for our study.

- *Calls for action* correspond to tweets that contain a clear message calling for action, including actions to sign a petition, prevent events from happening, etc.
- *Publications* correspond to tweets that contain a reference to publication, news or some information related to the campaign, including videos, articles or background information on the campaign.
- *Official meetings* correspond to tweets that contain information about an official meeting, a conference, a convention or a debate related to the campaign.
- *Physical actions* correspond to tweets that contain information about past, current or upcoming actions organized by an individual, a group of people, or an organization that is related to the campaign. This includes proposals to participate to challenges, contests or to dedicate some time to a specific issue, e.g., cleaning streets or repairing homes.
- *Others*, finally, correspond to tweets that do not belong to the four categories above, such as content that is indirectly related to climate change or animal welfare domains, as well as personal opinions and experiences, or tweets in other languages.

Tweet filtering and annotation Next, we explain how we classified the tweets from our dataset based on the classes introduced above. Since manually annotating our whole dataset is unrealistic, given the high number of tweets involved, we introduce a two-step process, where we first use micro-task crowdsourcing to annotate parts of the dataset and then leverage the resulting annotations in order to build an effective classifier.

The aim of the first step, i.e., crowdsourcing, is to collect as many high-quality annotations as possible pertaining to the types of tweets while limiting the involvement of the crowd. In order to do this, we first design a set of rules to preselect the tweets given our types. Those rules were created using simple regular expressions based on the analysis of a sample of the tweets, and are presented in Table 3 of Appendix A. In total, we created approximately 40 rules for each message type¹⁰. These rules were geared towards high recall based on the message types, rather than high precision. Nonetheless, they allowed us to significantly narrow down the number of tweets that would be presented to the crowd by focusing on subcategories early in the process. The resulting counts of tweets obtained from this process are given in Table 3 of the Appendix A.

We then crowdsource the action type annotation using the CrowdFlower platform¹¹. The three authors of the paper manually labeled 5% of the tweets beforehand to create a set of test questions for the crowd. Crowd workers could only work on our tasks if they correctly answered at least 7 out of 10 test questions. We additionally selected workers from English-speaking countries only and collected three independent judgments for every tweet. Agreement was obtained through the majority voting. We also made sure to identify and block malicious crowdworkers by leveraging a series of unambiguous test questions, following standard recommendations from CrowdFlower.

¹⁰ <https://github.com/toluolll/CampaignsDataRelease>

¹¹ <http://www.crowdflower.com/>

For each type of action, we considered a random sample of 2100 tweets. For more exploration, only half of these tweets is randomly selected from the collection complying with the regular expressions, while the other half is randomly selected from the rest of the campaign tweets. The results obtained through this process were consistent, with an agreement rate of 87.5%. In general, human annotators applied our definitions for the types of actions very strictly. However, this sometimes narrowed the results; For instance, human annotators did not always correctly annotate the tweets related to the attendance of a conference or a meeting when obvious keywords or the acronym of the event were missing, e.g., “conference”. As before, the annotated tweet collection is available online.

Action classification At this stage, we use the results of the crowdsourced annotation campaign as a training set to create an effective type classifier for the tweets. For this task, we consider the following features:

- *Semantic features.* Having a large textual corpus of 10Gb, we trained a Word2Vec model [18] using the implementation from the Gensim library¹² with 200 word vector dimensions. To train the model, we preprocessed each tweet as follows: (a) deleted all punctuation excluding hashtag (#) and handler (@), (b) lowercased the tweets, (c) tokenized the tweets into words. Furthermore, we interpreted each tweet as a bag of word vectors and calculated an averaged vector for every tweet. The main motivation behind the choice of semantic features is their ability to capture the semantic similarity between words and phrases using contextual information [18].
- *Syntactic features.* In addition to the above features, we added manual rules based on the regular expressions from Section 5.1. This resulted in 46, 42, 38, 20 additional features for meetings, actions, calls for action and publications respectively.
- *Contextual features.* Finally, we added a feature whether a particular domain name is contained inside a tweet. We selected the most frequent domain names and used them as binary features for the classifier. The frequency threshold was chosen at one sigma.

In order to predict the type of a tweet, we trained an individual binary classifier for each of our action types. As a classification method, we used a state-of-the-art approach based on Decision Tree Ensembles¹³, which effectively deals with diverse features. Appendix’s A Table 2 shows its precision and recall results for the four types of actions using 10-fold cross-validation. We observe that the physical action type has the lowest precision and recall among all types. We connect this result to the relative subjectivity in the definition of physical actions and to the high linguistic variety of the tweets of this type. The prominence of physical actions is hard to determine in general, since they can encompass anything from territory cleanups and protests to film-making competitions and tweet-a-thons.

¹² <https://github.com/piskvorky/gensim>

¹³ <http://scikit-learn.org/stable/modules/generated/sklearn.ensemble.ExtraTreesClassifier.html>

Further, the introduction of semantic features extracted from the tweet word vectors leads to a loss in precision and to some improvement in recall. This is due to the semantic representation of the tweets, which allows to identify semantically related tweets and words. For example, in the vector space representation produced by the Word2Vec model, the word “debate” is most similar to the words “politics, issue, discuss, policy, conversation”. Overall, due to the very nature of the tweets (i.e, very limited length, use of slang, pictures, videos, or emoticons), recall is relatively low across all the categories.

As expected, we found that manually constructed syntactic rules result in better precision as compared to the Word2Vec features only. This is caused by the fact that the rules are highly representative of the classes they are built for. Additionally, we observed that domain names play a more important role for meetings, calls for actions and publications, which is explained by the presence of conference websites and specialized websites to gather petition signatures.

5.2 Data Analysis

In order to detail content of the campaigns, we ran the tweet type classifiers over all tweets from all campaigns. We relied on the classifiers that were trained on all features from the previous section as they achieved the best F1-scores for all message types.

We applied the models on each campaign to identify the amount of contribution of a particular action to the overall contents of the campaign.

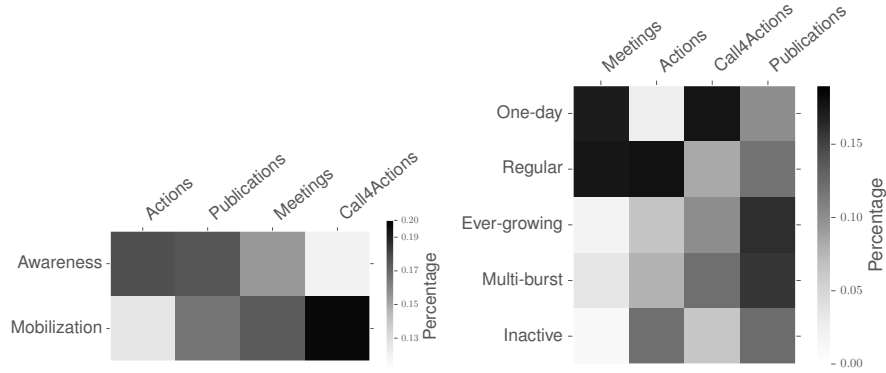
A visual summary of the outcome for the two main classes of campaigns is shown on Figure 5(a). We observe major differences in terms of contents; in particular, we see that mobilization campaigns favor calls for actions that motivate the audience to react on the climate change issues, while, having relatively low physical actions. Interestingly, awareness campaigns encourage more physical actions and publication releases, while mobilization campaigns focus more on calls for actions and official meetings. Mobilization campaigns make a high use of official meetings, probably because they tend to raise more attention from the governments or particular stakeholders. To conclude, we see that mobilization and awareness campaigns get organized in very different ways, thus confirming the initial distinction we make between each other.

Following the analysis given in Section 4, we performed a study on user engagement patterns. As shown on Figure 5(b), “one-day” campaigns¹⁴, focus on call for actions tweets which are mainly duplicated rather than retweeted. On the other hand, “regular” campaigns¹⁵ are mostly represented either by regular meetings or physical actions, e.g., annual conferences, campings, etc. Interestingly, “ever-growing” and “multi-burst” campaigns¹⁶, make larger use of publication and call for actions types, which significantly differs from the awareness campaign strategies in general. This can be explained by the targeted audience and by the issues tackled by those campaigns, such as global poverty, international divestments, dependence on fossil fuels, etc. All of these campaigns share global values and target international audiences around the globe.

¹⁴ #helpcovedolphins, #savebucky, #freethearctic30, etc.

¹⁵ #climatecamp, #climateweek, #worldenvironmentday, etc.

¹⁶ #talkpoverty, #saveanimals, #saveenergy, #actonclimate, #divestment, #fossilfree



(a) Comparison of the distributions of actions for the types of climate change campaigns defined in Section 4.1. Dark indicate greater importance of a particular actions in the campaign type.

(b) Comparison of the distributions of actions for the two main categories of climate change campaigns: awareness and mobilization. Dark indicates a greater importance of a particular actions in the campaign type.

Fig. 5. Action distribution across various campaign types

Duplicate tweets As described in Appendix B, some tweets from our dataset shared the same contents but were not strictly speaking retweets. This is due to some users trying to promote a tweet into a trending topic on Twitter. We decided to compute the proportions of such duplicated messages to see how they are distributed across different campaign types. To select the threshold at which a message should be treated as a duplicate, we considered the distribution of number of similar messages to the total amount of messages with these number as a half-normal distribution. In such way, the tweet was considered to be a duplicate if the number of such tweets exceeded three standard deviations.

Figure 6 illustrates the distribution of duplicate content for the different campaign types. As can be observed, duplicate content is especially significant for the mobilization campaigns, which can be explained by their spontaneous nature and the need to mobilize people in shorter periods of time. Awareness campaigns differ in the sense that they typically operate on longer-terms goals. From a user engagement perspective, both regular and inactive campaigns do not contain much duplicated content, while increasing, multi-burst, one-day campaigns make heavy use of it.

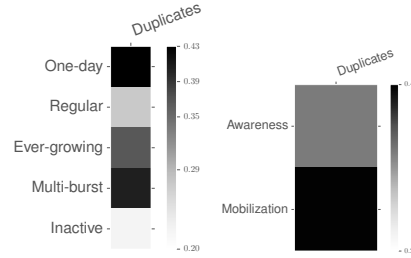


Fig. 6. Comparison of duplicate content between campaigns

Domain usage distribution Users in the climate change community tend to make great use of links to images, facebook pages, youtube videos and petition sites. We explored the general distribution of the top domain names across the campaigns and found that all types of campaigns extensively use visual content (youtube, facebook, photos, etc.). Nevertheless, both ever-growing and regular types of campaigns use such content more parsimoniously comparing to one-day and inactive on average. A similar trend was discovered between awareness and mobilization campaigns respectively. Interestingly, the tendency to overuse visual resources clearly does not affect the popularity of the posted content [27]. Among major domain names whose tweets gained the most retweets, we primarily observe contents related to the campaigns, i.e., the site of the campaigns, news and information about related issues.

6 Discussion

In the following, we take a step back, discuss the results we obtained and also make a series of recommendations in the context of public campaigns on social media. First, we proposed a framework for collecting campaigns and identifying their types. As explained in Section 3, we collected over 100 campaigns that were annotated with types, i.e., awareness and mobilization, as well as clustered by their user engagement patterns (Section 4). This resulted in a large collection of tweets that were partially annotated with action types using crowdsourcing and further generalized based on an annotated corpus using a machine learning classifier. Overall, our tweet action type detection technique showed high precision ($\sim 90\%$) and recall ($\sim 60\%$). This allowed us to automate action identification in tweets and to understand the overall campaigns' contents.

Subsequently, we focused on the analysis of campaigns classified by their initial goal and their user engagement pattern. The goals of awareness and mobilization campaigns differ significantly, and so do their contents. While awareness campaigns often involve physical actions and promote scientific publications, mobilization campaigns make great use of official meetings and calls for actions; For the mobilization campaigns, the more official meetings are organized the more leverage can be obtained from governmental organizations. The analysis of user involvement patterns also showed noticeable differences between campaign types and their agendas. "One-day" campaigns were dominated by calls for actions, while "regular" and "ever-growing" ones contained more physical meetings and publications on climate. This insight represents an important foundation on which specific campaigns studies and their contents can be built.

With the various techniques we leveraged for campaign analysis, we noted major differences in the way users duplicate messages. "One-day" and "ever-growing" campaigns in general contain 20% more duplicated content as compared to the "inactive" campaigns. In the "one-day" campaigns, this phenomenon can be explained by the spontaneous nature of particular tweets and the need to mobilize people in a short period of time. On the other hand, "awareness" campaigns typically operate on a longer basis, therefore their communities do not actively use duplicated tweets, i.e., on average 15% less duplicates than mobilization. This can be explained by the actions required during shorter periods of time for the mobilization campaigns. In a context of user involvement patterns, both "regular" and "inactive" campaigns do not contain as much

duplicated content, while “ever-growing” and “one-day” campaigns make a heavy use of them.

Regarding the effects that drive user engagement, we observe that first-degree neighbors are essential for getting higher numbers of retweets (about half of the retweets from popular tweets originate from direct neighbors), while duplicated content attracts less retweets in general. Finally and most interestingly, the less diverse the main contributors of the campaign, the less likely it is to gain bigger audiences (as shown in Section 4.2).

Overall, this work has a potential to empower governmental and non-profit organizations by facilitating campaign analysis. The analysis of the collected campaigns combined with the analysis of individual tweets provides foundation for many applications, e.g., detecting public campaigns, or identifying means to boost user engagement.

Even though we expanded the campaign coverage by performing several iterations of the data collection, our methodology is focused on English-speaking tweets. The @AlGore and @GreenPeace accounts we used are biased towards the US, and so are the English terms and hashtags that were used for the climate change topic. Therefore, it would be beneficial to further expand the data collection by reiterating over the steps from Section 3.1 to sample more campaign hashtags over other languages and countries. At the same time, a given campaign may leverage multiple hashtags, which can affect the results of the analysis.

7 Conclusions

In this work, we analyzed large-scale social media campaigns related to climate change and animal welfare from various perspectives, including analyses on their primary goals, the types of messages they relay, as well as their user involvement patterns. In the context of climate change and animal welfare, we showed that public campaigns are represented by two main narratives: awareness and mobilization. Our subsequent analysis of user participation revealed that campaigns significantly differ in terms of their user involvement patterns. Finally, we presented a study on the best ways towards increasing user involvement for public campaigns by combining core users, followers, and actions. The high-level patterns that were found in our study lay a solid foundation for future work on specific campaigns and their fine-grained segmentation. As a possible extension, a more fine-grained classification of campaigns and campaign actions could reveal more sophisticated patterns and correlations that appear during the campaign life-span. For example, political or non-profit events might exhibit different user involvement patterns. Similarly, characteristics of geographically focused or global campaigns may differ.

Acknowledgements. The authors would like to thank Alexandra Olteanu for suggestions and feedback. The work was supported by the Sinergia Grant (SNF 147609).

References

1. Abbar, S., Mejova, Y., Weber, I.: You tweet what you eat: Studying food consumption through twitter. arXiv preprint arXiv:1412.4361 (2014)
2. Aiello, L.M., Petkos, G., Martin, C., Corney, D., Papadopoulos, S., Skraba, R., Goker, A., Kompatsiaris, I., Jaimes, A.: Sensing trending topics in twitter. *Trans. Multi.* 15(6), 1268–1282 (Oct 2013), <http://dx.doi.org/10.1109/TMM.2013.2265080>
3. Alley, R.B., Marotzke, J., Nordhaus, W.D., Overpeck, J.T., Peteet, D.M., Pielke, R.A., Pierrehumbert, R.T., Rhines, P.B., Stocker, T.F., Talley, L.D., Wallace, J.M.: Abrupt climate change. *Science (New York, N.Y.)* 299(5615), 2005–10 (Mar 2003), <http://www.sciencemag.org/content/299/5615/2005.short>
4. Atkin, C., Rice, R.: Theory and principles of public communication campaigns. *Public communication campaigns* (2012)
5. Becker, H., Naaman, M., Gravano, L.: Beyond trending topics: Real-world event identification on twitter. In: *ICWSM '11*. pp. 438–441. *ICWSM '11* (2011)
6. Fernandez, M., Piccolo, L.S.G., Maynard, D., Wippoo, M., Meili, C., Alani, H.: Talking climate change via social media: Communication, engagement and behaviour. In: *Proceedings of the 8th ACM Conference on Web Science*. pp. 85–94. *WebSci '16*, ACM, New York, NY, USA (2016), <http://doi.acm.org/10.1145/2908131.2908167>
7. Giorgino, T.: Computing and visualizing dynamic time warping alignments in r: The dtw package. *Journal of Statistical Software* 31(7), 1–24 (8 2009), <http://www.jstatsoft.org/v31/i07>
8. Gonzalez-Bailon, S., Wang, N.: Networked discontent: The anatomy of protest campaigns in social media. Available at SSRN 2268165 (2013)
9. Hermida, A., Fletcher, F., Korell, D., Logan, D.: Share, like, recommend: Decoding the social media news consumer. *Journalism Studies* 13(5-6), 815–824 (2012)
10. Hong, L., Convertino, G., Chi, E.H.: Language matters in twitter: A large scale study. In: *ICWSM '11*. *ICWSM '11* (2011)
11. Hoornweg, D.: Cities and climate change: responding to an urgent agenda. *World Bank Publications* (2011)
12. Jackoway, A., Samet, H., Sankaranarayanan, J.: Identification of live news events using twitter. In: *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks*. pp. 25–32. *ACM* (2011)
13. Jin, F., Khandpur, R.P., Self, N., Dougherty, E., Guo, S., Chen, F., Prakash, B.A., Ramakrishnan, N.: Modeling mass protest adoption in social network communities using geometric brownian motion. In: *KDD '14*. pp. 1660–1669. *KDD '14*, ACM, New York, NY, USA (2014), <http://doi.acm.org/10.1145/2623330.2623376>
14. Kirilenko, A.P., Stepchenkova, S.O.: Public microblogging on climate change: One year of Twitter worldwide. *Global Environmental Change* 26, 171–182 (may 2014), <http://www.sciencedirect.com/science/article/pii/S0959378014000375>
15. Koneswaran, G., Nierenberg, D.: Global farm animal production and global warming: impacting and mitigating climate change. *Environmental Health Perspectives* (2008)
16. Lee, K., Palsetia, D., Narayanan, R., Patwary, M.M.A., Agrawal, A., Choudhary, A.: Twitter trending topic classification. In: *ICDMW '11*. pp. 251–258. *ICDMW '11*, IEEE Computer Society, Washington, DC, USA (2011), <http://dx.doi.org/10.1109/ICDMW.2011.171>
17. Lehmann, J., Gonçalves, B., Ramasco, J.J., Cattuto, C.: Dynamical classes of collective attention in twitter. In: *WWW '12*. pp. 251–260. *WWW '12*, ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2187836.2187871>

18. Mikolov, T., Corrado, G., Chen, K., Dean, J.: Efficient Estimation of Word Representations in Vector Space. *Proceedings of the International Conference on Learning Representations (ICLR 2013)* pp. 1–12 (2013)
19. Norris, P.: *The handbook of comparative communication research. The ICA Handbook Series* (2012)
20. Olteanu, A., Castillo, C., Diakopoulos, N., Aberer, K.: Comparing events coverage in online news and social media: The case of climate change. In: *ICWSM '15. ICWSM '15* (2015)
21. Quercia, D., Askham, H., Crowcroft, J.: TweetLDA. In: *WebSci '12*. pp. 247–250. ACM Press, New York, New York, USA (Jun 2012), <http://dl.acm.org/citation.cfm?id=2380718.2380750>
22. Schmidt, A., Ivanova, A., Schäfer, M.S.: Media attention for climate change around the world: A comparative analysis of newspaper coverage in 27 countries. *Global Environmental Change* 23(5), 1233–1248 (2013)
23. Segerberg, A., Bennett, W.L.: Social media and the organization of collective action: Using twitter to explore the ecologies of two climate change protests. *The Communication Review* 14(3), 197–215 (2011)
24. Shields, S., Orme-Evans, G.: The impacts of climate change mitigation strategies on animal welfare. *Animals* 5(2), 361–394 (2015)
25. Slovic, P.: Informing and educating the public about risk. In: Slovic, P. (ed.) *The perception of risk*, pp. 182–198. Earthscan, London, England (2000), cited by 0000
26. Williams, H.T., McMurray, J.R., Kurz, T., Lambert, F.H.: Network analysis reveals open forums and echo chambers in social media discussions of climate change. *Global Environmental Change* 32, 126 – 138 (2015), <http://www.sciencedirect.com/science/article/pii/S0959378015000369>
27. Xu, Z., Zhang, Y., Wu, Y., Yang, Q.: Modeling user posting behavior on social media. In: *SIGIR '12*. pp. 545–554. *SIGIR '12*, ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2348283.2348358>

A Action Detection

Examples of tweet actions are shown in Table 1.

Type of action	Sample tweets
Official meeting	Monday Dec 1, U.N. COP climate talks begin Lima Peru @YebSano Just witnessed a sign of hope at the climate talks in #Cancun - ... #UNFCCC #tckctck
Physical action	#WorldEnvironmentDay #treeplanting is taking place around 09:00 at Tsarogaphoka in #Soshanguve We came. We swooped. We're camping!!! #climatecamp
Call for action	The #GreatBarrierReef is not a dump! Protect our World Heritage. #UNESCO #FightfortheReef Take Action: Stand with me and support clean #energy and a safer #climate future! #CleanAir4Kids
Publication	660 million Indians could lose 2.1 billion years as a result of air pollution... #gofossilfree Water Fact: Fact: At 1 drip per second, a faucet can leak 3,000 gallons per year. #savewater

Table 1. Sample tweets for each type of action considered.

Results of the action detection based on Decision Tree Ensembles is shown in Table 2. It shows the precision and recall results for the four types of actions using 10-fold cross-validation.

	Meetings			Actions			Call for actions			Publications		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
All features	0.896	0.616	0.730	0.771	0.577	0.660	0.902	0.664	0.765	0.897	0.528	0.665
Sem	0.723	0.605	0.659	0.707	0.510	0.592	0.751	0.441	0.556	0.857	0.461	0.599
Sem + Cont	0.788	0.531	0.635	0.703	0.518	0.597	0.792	0.439	0.565	0.865	0.487	0.623
Sem + Synt	0.912	0.590	0.717	0.754	0.480	0.587	0.920	0.563	0.699	0.862	0.464	0.603
Synt	0.895	0.375	0.529	0.816	0.276	0.412	0.920	0.472	0.624	0.890	0.134	0.232
Synt + Cont	0.901	0.384	0.538	0.812	0.300	0.438	0.921	0.643	0.758	0.911	0.325	0.479

Table 2. Precision, Recall and F1-score values for classification of different types of actions with different sets of features.

Below we show examples of manually created rules to preselect the tweets given specified action types.

Type of action	Sample rules	N# of tweets
Official meeting	<i>/speakingat(demo the)/</i>	113 989
Physical action	<i>/actionat(the)?park/</i>	154 874
Call for action	<i>/tell(the)?to(keep protect)/</i>	328 603
Publication	<i>/great news/</i>	2 559 063

Table 3. Examples of rules and number of tweets for each type of action.

B Unique tweets identification and retweets count

One of the main issues with the data collected from Topsy is that the tool does not provide information about retweets. Therefore, we had to create heuristics to make sure that we could properly identify all retweeted messages. Taking into account that all tweets returned by the tool are sorted by timestamp, we can easily figure out the origin of all the tweets using a simple regex pattern `((RT|MT) @author tweet_prefix)`. This approach has a number of limitations, however. It does not identify complex retweet structures, such as where a tweet text is cited using quotes. We found that such cases are quite rare on Twitter and amount for $\sim 0.5\%$ of all tweets.

In order to compute the complex retweet cases, we aggregated the tweets with at most 5 characters edit distance. Further, we discarded explicit retweets `((RT|MT) @author)` and exact duplicates. However, certain retweets can be missing when a hashtag does not fit into the message due to the tweet length limit. To solve this problem, we leveraged the Topsy API, by returning and analyzing related tweets for each requested tweet in order to identify all further retweets. Finally, we note that we apply this process recursively—searching for retweets of retweets iteratively—in order to capture potentially complex retweet patterns. When no new retweets can be identified, we identify content that was not retweeted but duplicated. The practice of duplicating tweets gained traction on the platform as it can help promote topics into *Twitter Trends*. We consider a tweet to be a duplicate whenever at least 80% of its contents exactly matches an original tweet excluding punctuation.