Big Players: Emotion in Twitter Communities Tweeting about Global Warming

Dennis J. Drown[1], Roger Villemaire[1], and Serge Robert[2]

[1] Département d'informatique, Université du Québec à Montréal
[2] Département de philosophie, Université du Québec à Montréal

Abstract This work considers exceptionally active users on Twitter, the "big players," and analyzes the correlation between the level of emotion these users express in communications tagged with the hashtag *#globalwarming* and the levels expressed by the Twitter community as a whole. Using an emotion lexicon incorporating four of the eight base human emotions according to Plutchik: anger, fear, sadness, and joy, we identify to what extent a small group of these big players may predict the emotion expressed by their online community in their tweets.

1 Introduction

The significant world-wide rise in temperatures is a major challenge facing humanity today. As assessed by the Intergovernmental Panel on Climate Change (IPCC) [26], human-induced changes in atmospheric composition are leading to a rise in global temperatures that has had and will continue to have a severe impact on the earth's climate. Rising temperatures not only challenge human health and food supplies, they are also a threat to biodiversity. Global warming therefore represents a serious social and political issue, and much has been done in recent years to better understand related public perception and behaviour.

Research indicates that there is a complex relation between environmental concern, informal education, experience, and behaviour [20]. Furthermore, environmental risk perception and policy support is in fact strongly influenced by sentiment (positive/negative reactions), emotional responses, imagery, and values [15]. Emotions are a driving force in modifying behaviour in order to avoid risk in dangerous situations [29], and they may be key to understanding how people reason and respond to information about global warming [17].

This work considers emotion expressed on climate-related issues in the context of the social media site Twitter, [3] evaluating affect using sentiment analysis. This relatively new subfield of natural language processing (NLP) is becoming increasingly popular largely due to the enormous amount of opinion that the public expresses online in today's world [16]. Sentiment analysis is often concerned with polarity, determining if a message is positive or negative. Here we focus on affect using a type of analysis commonly called emotion mining.

3. https://twitter.com/

Of particular interest is the behaviour of Twitter users showing exceptionally high participation. One might, for instance, wonder whether a small number of the "top" high-activity users could be particularly representative of the overall community. Hence, we present affective models to predict emotion expressed in the general community [4] from the emotion expressed by these "big players" in their tweets. We also compare these models with predictive models from the emotion expressed by a random set of users of the same size in order to ascertain whether the top high-activity users have a distinct predictive capability.

We show that, while the big players are not necessarily more representative of the whole community than other groups of the same size, they do indeed show a distinct predictive capability. Moreover, this effect sharply depends on the emotion being considered.

2 Related Work

When seeking to identify individuals filling specific roles in their communities, a common approach is to analyze centrality in graph-based representations of the relationships between users. Bigonha et al. [2] create a model to find influencers based on three elements: sentiment polarity in users' tweets; two types of graph representations ("who follows whom" and "who reacts to whose tweets"); and grade-level readability of the messages. Aleahmad et al. [1] propose an algorithm called OLFinder to identify major topics for a given domain in a set of tweets and determine which users are "opinion leaders" for those topics based on the users' calculated competency in that domain and a popularity score determined by graphs of follower relations. A study by Eliacik and Erdogan [6] seeks to boost graph-based methods using a calculated measure of trust that others extend to a user based on her relations, expertise, and activity in a topic-centred community. In research pertaining to climate change, Cody et al. [5] use sentiment analysis to examine changes in polarity in tweets with respect to climate-related events.

The studies above involve measuring the polarity in tweets, but there are also a number of projects which focus on emotion as we do in the present work. Mitchell et al. [19] look at finding the happiest and saddest states and cities in the United States using a large corpus of tweets tagged with geolocation information. Preoţiuc-Pietro et al. [25] work with regression models for predicting the income of Twitter users based partly on polarity and emotion content in their tweets. Finally, Halse et al. [12] use affect models during a crisis to determine if tweets are trustworthy and contain information that may be useful to first responders.

Contrary to research involving graph-based representations of a given role between certain users and their followers, we focus our attention on users set apart by their unusually high level of activity in order to determine to what extent their behaviour can predict that of the general *#globalwarming* community.

4. We use the word *community* to indicate the users sending tweets or being referred to in tweets with a given hashtag.

3 Methodology

To determine which top high-activity users will be the big players, we look at the correlation between the emotion expressed in the online activity of a tentative set of big players and the emotion in tweets published by "regular players," the rest of the Twitter community talking about global warming.

Tweet Dataset: We made use of the Twitter developer platform API [5] to collect 414,035 tweets from 239,590 users, published between January 1, 2018 and August 31, 2019 and incorporating the hashtag *#globalwarming*. We call the analysis of the tweets from this dataset over a fixed period of time a tracking run. We are interested in communications on climate change generally, but this work specifically tracks the hashtag *#globalwarming*. According to a study by Williams et al. [30], this hashtag is more regularly used by Twitter users from both pro-science and skeptic communities, compared to *#climatechange*, which more often appears in tweets from climate activists than from skeptics and deniers. The dataset contains only tweets in English as flagged by accompanying metadata from the Twitter API.

Emotion Lexicon: We analyze tweets using four base emotions from Robert Plutchik's model [23]: *anger*, *fear*, *sadness*, and *joy*. We treat the tweet text as a bag of words to be checked against the Affect Intensity Lexicon from the National Research Council of Canada (NRC). We chose this lexicon as we see potential in the method known as Best-Worst Scaling, which the NRC used to create it [21]. [6] The lexicon is available via the Affective Tweets plugin [7] for the machine learning platform Weka [10]. Our algorithm uses the plugin to standardize usernames and web URLs and to calculate floating-point values in a tweet's emotion vector, summing intensity levels for each of the four emotions across all words in the text that are contained in the lexicon. When analyzing a tracking run, the first step is to determine this emotion vector for each tweet tagged with *#globalwarming*.

NLP Tools: The Affective Tweets library [21] provides an NLP tokenizer specifically designed for tweets [11] to delimit words, user names, and web links in a user's text. The library also provides other NLP preprocessing tools, allowing the modeller to specify a stopword list to remove common words devoid of analytical value and a stemming algorithm to reduce words to their grammatical roots. When creating our models, we experimented with the English stopword list from Apache's information retrieval package Lucene [8] as well as the Snowball Porter stemmer [24] for the English language. [9] However, our results did not improve substantially, and therefore in this paper we present results obtained without a stop list and without stemming.

The lack of improvement should not be too surprising considering that the utility of traditional NLP tools may suffer when analyzing human language on

5. https://developer.twitter.com/
6. Plutchik's model also includes *anticipation*, *trust*, *surprise*, and *disgust*. The NRC is in the process of expanding the lexicon to include all eight emotions.
7. https://github.com/felipebravom/AffectiveTweets
8. https://lucene.apache.org/
9. http://snowball.tartarus.org/

social media. Users are non-professional writers, tending to express ideas with little thought towards clear content. Tweets contain frequent abbreviations, slang, and (often intentional) errors in grammar and spelling. Twitter is particularly problematic due to the small size of the texts [8].

Communication Categories: There are various ways to demonstrate high activity. We rank users in terms of tweet count for the following categories:

1. **Original Tweeters** (*oter*): Users publishing personally-authored messages.
2. **Retweeters** (*rter*): Users retweeting (resending tweets written by another user) using Twitter's popular "*RT @author* ..." syntax. [10]
3. **Retweeted Authors** (*rted*): Users whose tweets are retweeted by others.
4. **Mentioned Authors** (*tmed*): [11] Users specifically mentioned in the tweets of others via Twitter's "*@user*" notation. This syntax usually serves to address specific users or attract their attention [14].

A single tweet may be accounted for multiple times. For example, if user U_1 publishes tweet T_1, which he has personally authored and which mentions user U_2, then T_1 contributes to U_1's participation in the *oter* category and also to U_2's participation in *tmed*. Similarly, if user U_2 sends a retweet R_2 originally published by user U_1 that mentions user U_3, then R_2 contributes to U_2's participation in *rter*, to U_1's participation in *rted*, and to U_3's participation in *tmed*.

Each category represents a distinct type of participation. For instance, in the last example the original tweet from U_1 would not be considered at all if it was sent before the tracking run. Nevertheless, if others continue retweeting it often enough, then U_1 may become a *rted* big player. Likewise, when considering the category of mentioned authors, one should keep in mind that these users may not have actively participated in the tracking run. For example, the top-mentioned account in our collected *#globalwarming* tweets is *realDonaldTrump*; however, this famous user authored no tweets with this hashtag during the tracking run.

We may also think of *oter* and *rter* as categories of active participation, similarly considering *rted* and *tmed* as representing a passive form of participation. Yet, while the concept of passive big players may be useful, it is not valid in every sense. Twitter users may actively work to be retweeted [3], and users often mention each other in tweets when establishing communication threads [14].

"Top N" Big Players: For each of the four communication categories described above, we identify a set number of users who rank highest with respect to the type of activity that the category represents: the Top N. Note that we are essentially considering four types of big players, and it is possible for a specific user to belong to more than one big player group.

For this study we evaluate a series of big player groups with sizes ranging from 5 to 25. Our reasoning is that N should be small enough that the total number of big players will not be overwhelming to a researcher who must, for example, examine users' account profiles. Ideally, we can identify a size N that

---

10. The syntax for retweeting is not standardized, and alternatives do exist. The Twitter API sends metadata that identifies retweets and indicates the retweeted author.
11. Note that the *tmed* code is backwards: "*m*ention*ed* in *t*weet."

demonstrates a significant correlation between emotion levels in tweets from the big players and emotion levels in the general community.

We define the regular players as users who are not in any of the four big player categories and who have published at least one original tweet with the hashtag *#globalwarming* during the tracking run. We are endeavouring to predict the emotion levels expressed by the regular players in these original tweets. It stands to reason that tweets authored and published by the general *#globalwarming* community are likely a meaningful reflection of what that community is feeling.

Machine Learning Models: We conducted experiments with the following models: linear regression, Gaussian processes, decision lists using separate-and-conquer (M5Rules), random forests, and support-vector machines with first and second degree polynomial kernels (SMOreg using PolyKernel with exponents set to 1 and 2). For each of these we used the implementations in Weka [31] with the default settings. Linear regression set itself apart by consistently showing better results. It also has the advantage of being a "white-box" modelling technique, enabling the modeller to understand how it arrives at its predicted values. Henceforth, we will only report on the linear regression models.

Emotion Models: We create regression models for four target values, which are the variation from one week to the next in levels for the emotions anger, fear, sadness, and joy as measured in tweets from regular players in the *#globalwarming* Twitter community. To predict these values, we use 16 independent attributes that describe the variation, week by week, in the average levels of each of the four emotions for big player tweets across the four communication categories. We name these attributes and the targets of the models using a three-part syntax to indicate the community group, the communication category, and the emotion:

$$\begin{pmatrix} big \\ reg \end{pmatrix}\_\begin{pmatrix} oter \\ rter \\ rted \\ tmed \end{pmatrix}\_\begin{pmatrix} anger \\ fear \\ sadness \\ joy \end{pmatrix}$$

Independent attributes begin with *big* as these are values representing big players. For example, big_rter_fear gives the variation in the level of fear in retweets from users in the big retweeters category, and big_tmed_joy is the variation in joy measured in tweets that include frequently-mentioned authors. Target attributes start with the community code *reg* as these are values we are predicting for the regular players. [12] As an example, the following represents a typical regression model to predict the variation in the level of anger measured in regular players' tweets for a given week with respect to the previous week:

$$\begin{aligned} reg\_oter\_anger = \quad & 0.218 \times big\_oter\_anger \\ & - 0.105 \times big\_oter\_fear \\ & - 0.056 \times big\_tmed\_fear \\ & - 0.001 \end{aligned} \tag{1}$$

Data Preparation: To prepare tracking run data for a regression model, our system runs all the tweets tagged with *#globalwarming*, collected over the

12. For targets we consider only the original tweets (*oter*) for the four base emotions.

tracking period, through the Affective Tweets Weka filter to determine their associated emotion vectors. We keep running averages for all users with respect to each of the four emotions for all four communication categories. Once we have the emotion vector for a tweet, we apply it to the user either as an original tweet (*oter*) or a retweet (*rter*). For a retweet, we also apply the emotion vector to the tweet's original author (*rted*). Finally, we search the tweet text for all mentioned users and apply the emotion vector to their accounts as well (*tmed*). Applying the vector to a user means that the system incorporates the emotion levels for the tweet into an emotion vector representing the running averages of the emotions on the day the tweet was published for the communication type being processed. The system increments the user's counter for the communication activity, used to track the user's position in the ranking for that category.

After the tweets have been processed, the system finds the big players. For N ranging from 5 to 25, we create player activity rankings for the four communication categories and select the Top N players in each case. Note that as players are ranked, there is an occasional tie in activity levels. For example, when determining the Top 10 original tweeters, the players ranked #10 and #11 may have both sent the same number of tweets. In this case we accept both players, and the Top 10 big players for one category will incorporate 11 users. As our models target the variation from one week to the next for the levels of emotion in the regular players' original tweets, the group of regular players is defined as all accounts which have authored and sent one or more tweets during the tracking run but are not included in the four big player groups.

After identifying the big players in the four communication categories and the remaining accounts that make up the regular players, our system then creates five groups (one for each player type) and calculates an average of the emotion vectors across all users in a group for each day of the tracking run. It then sorts the day groupings and bundles them into super-groups representing a week (seven consecutive day groups). The system reduces all the emotion vectors in each week-long grouping into a single vector corresponding to the average emotion levels measured in the tweets for all players in a community (*big* or *reg*) and communication category (*oter*, *rter*, *rted*, or *tmed*) over that week.

At this point we create a preliminary set of data instances, which each contain levels of emotion intensity for one week of tweets. As a final step, we subtract the attribute values for each instance from those of the previous one to obtain the variation in emotion intensity from week to week. [13] Thus, variations in the levels of the four emotions in the four big player groups become the 16 independent attributes (*big_ oter_ anger*, etc.). These values remain the same across all affect models for a given N. The target attribute for each model represents the weekly variation in the average affect intensity expressed by all regular players for one

13. For a given attribute, A_Δ, we compute $A_\Delta = A_i - A_{i-1}$, where A_i is the emotion level for the current week, and A_{i-1} is the level for the previous week. One might consider using instead the relative change, dividing our A_Δ by A_{i-1}, to get a percentage. This was not possible here because if a player group does not express some given emotion in a week's worth of tweets, the relative change will be undefined.

emotion. Each element of an emotion vector in the grouping for regular players (e.g., *reg_oter_fear*) represents the dependent attribute for one affect model. Its value for a given data instance is computed from a pair of consecutive weeks from the regular player group.

4 Results and Analysis

This section presents results from the models predicting the variation in level from one week to the next for anger, fear, sadness, and joy in the regular players' original tweets for the *#globalwarming* Twitter community. For each emotion we ran a series of models for the Top N big players with N ranging from 5 to 25. Our dataset contains 20 full months of tweets tagged with *#globalwarming*, beginning on January 1, 2018 and ending on August 31, 2019. For each set of experiments we ran 9 tracking runs with each run using 12 months of data, starting on midnight of the first day of month M_i, for i ranging from 1 to 9, and ending at 23:59:59.999 on the last day of month M_{i+12}. As each tracking run shifts the starting month by one, we are essentially sweeping a 12-month window across the 20 months of Twitter data.

In addition to being a natural choice with respect to the calendar, the one-year window provides sufficient data to train our models, given that each data instance represents a whole week of Twitter activity. For each 12-month period we used the first nine months (75%) for training and tested the models on the last three months (25%). We report averages of the Pearson Correlation Coefficient (PCC) across the 9 tracking runs in Table 1 for each of the four emotions. The PCC is a value between -1 and 1, with 1 indicating a total positive correlation between the model's predictions and the measured values in the test data.

To determine whether a group of big players of a given size is significant with respect to predicting emotion in the general community, we compare the results for the big players with results for reference groups of the same size. In order to build a meaningful model, however, we must ensure that users in the reference groups have a minimal level of participation. Therefore, we generate reference groups in the following way. We pick random users to form four groups (*oter*, *rter*, *rted*, and *tmed*).[14] These groups are the same size as the big player groups, and they must have at least 40 tweets in the communication category for their group. To avoid biased results wherein a few big players represent the dominant contribution, we require that users in the reference groups not be in any of the four big player groups. After creating the groups, we process the tweets and create regression models as explained above, replacing the big players with these reference groups. For each tracking run, and for each group size of N reference users (analogous to the Top N players), we repeat this procedure 20 times.

The "ref" columns in Tables 1 and 2 report the averages over these 180 models (20 sets of reference groups $\times$ 9 tracking runs). In these tables, a field containing asterisks (*****) indicates that for all 180 models, there was at least one week for

14. We use the Erlang *rand* library's implementation (*exrop*) of the Xoroshiro116+ pseudorandom number generator [28] with 58 bits of precision and a period of 2^{116-1}.

which a reference group had no tweets for one or more communication categories. This generally occurs for small values of N and for larger training periods (12 months as opposed to 9). By definition, the reference users are less active than the big players, and if there are too few of them in a given group, together their activity may not be consistent enough to cover every week in the tracking run. [15]

Table 1. Correlation (PCC) for models predicting emotion in the last 3 months.

N	Anger BIG	Anger ref.	Fear BIG.	Fear ref.	Sadness BIG	Sadness ref.	Joy BIG	Joy ref.
5	0.1376	*****	0.0129	*****	-0.0992	*****	0.1130	*****
6	0.1660	*****	-0.0818	*****	-0.1369	*****	0.1473	*****
7	0.1465	0.3301	0.0509	0.2432	-0.1539	-0.2223	0.1680	-0.0852
8	0.3415	0.0751	0.0429	-0.1945	-0.0758	-0.3385	0.0170	0.1602
9	0.2727	0.0090	-0.0337	-0.1281	-0.1429	0.2542	0.1183	-0.0504
10	0.2294	-0.0753	0.0204	0.1587	-0.1407	-0.0119	0.1291	0.1278
11	0.2534	0.0951	0.0031	0.0112	-0.2119	0.0155	0.1066	0.0938
12	0.3189	0.0595	-0.0009	-0.0340	-0.2380	-0.0701	0.1058	0.0816
13	0.3035	0.0144	0.0197	0.0425	-0.1542	0.0436	0.0935	0.0766
14	0.3024	-0.0786	0.1568	0.0128	-0.1876	0.0682	0.0854	-0.0219
15	0.2793	-0.0599	0.2358	0.0334	-0.1989	-0.0854	0.0115	-0.0606
16	0.2770	-0.0229	**0.2972**	-0.0332	-0.1221	0.0127	0.0632	-0.0007
17	**0.3869**	0.0139	0.2566	0.0099	-0.1812	0.0468	0.0945	0.0639
18	0.3787	0.0580	0.1282	0.0530	-0.2420	-0.0236	-0.0110	0.0829
19	0.3402	-0.0351	-0.0134	0.0053	-0.1857	-0.0609	0.2517	-0.0044
20	0.2764	0.0046	0.0389	0.0204	-0.1725	-0.0063	0.0016	0.0721
21	0.2917	0.0233	0.0914	0.0677	-0.0572	0.0015	0.0531	-0.0348
22	0.2042	0.0524	0.0934	-0.0289	-0.1444	0.0032	0.0942	-0.0314
23	0.1969	-0.0066	0.1732	0.0245	-0.1179	0.0222	0.1064	-0.0126
24	0.1816	-0.0012	0.1331	0.0177	-0.0533	0.0185	0.1061	0.0595
25	0.1201	0.0552	0.0594	0.0241	-0.0897	-0.0442	0.0522	0.0115

Table 2. Correlation (PCC) for models evaluated using 10-fold cross-validation.

N	Anger BIG	Anger ref.	Fear BIG.	Fear ref.	Sadness BIG	Sadness ref.	Joy BIG	Joy ref.
5	0.3521	*****	0.0721	*****	0.3203	*****	0.1580	*****
6	0.3620	*****	0.2865	*****	0.3309	*****	0.1559	*****
7	0.4067	*****	0.3893	*****	0.3138	*****	0.2326	*****
8	0.4127	-0.0685	0.4050	0.1910	0.2369	0.1227	0.2317	0.5470
9	0.4576	0.3461	0.3075	0.2587	0.2970	-0.0342	0.2418	0.3372
10	0.4545	0.1041	0.2906	0.2317	0.3216	0.1892	0.2215	0.2865
11	0.4322	0.2651	0.3815	0.2895	0.3102	0.2942	0.2428	0.3487
12	0.4327	0.1883	0.3369	0.2678	0.3000	0.1924	0.1974	0.3406
13	0.3800	0.1766	0.3189	0.1528	0.2297	0.1536	0.2233	0.3262
14	0.4299	0.2770	0.3341	0.3137	0.2917	0.4133	0.2367	0.2675
15	0.4364	0.2358	0.3238	0.3180	0.2610	0.3765	0.2432	0.2508
16	0.4098	0.2712	0.3625	0.2686	0.2416	0.3788	0.1705	0.3115
17	0.3740	0.2361	0.2448	0.2638	0.2651	0.2903	0.1393	0.2719
18	0.3823	0.3284	0.2670	0.2845	0.2697	0.2582	0.1425	0.2410
19	0.3550	0.2701	0.3248	0.2804	0.2871	0.2609	0.1194	0.2724
20	0.3955	0.2926	0.2482	0.3146	0.2509	0.2686	0.0858	0.2840
21	0.4236	0.3049	0.3583	0.2927	0.3052	0.3601	0.1160	0.3011
22	0.3716	0.2828	0.3044	0.2811	0.2900	0.3102	0.1401	0.2419
23	0.4306	0.3191	0.3106	0.2623	0.3086	0.3263	0.1364	0.3048
24	0.4248	0.3189	0.3235	0.3178	0.3080	0.2934	0.1673	0.3004
25	0.4201	0.3382	0.2678	0.2700	0.2740	0.3054	0.1909	0.2849

Comparing columns 2 and 3 from Table 1, we see the PCC for anger (variation) predictions from big players closer to 1 and higher than the PCC for the reference groups. Interestingly, as we follow N, looking at larger groups of big players, the PCC increases to 0.3869 before decreasing again for even larger groups. The best results are obtained with a group of around 17 big players.

From columns 4 and 5 of the same table we observe that while the fear models for big players do not always outperform the reference models, they do show a better linear correlation for N between 14 and 18, where the PCC for the big players reaches values up to 0.2972. The big players perform better than the reference groups near the middle of the table, but not when we create models using smaller groups or larger ones. For anger and fear, one may conclude that the big players are indeed significant as a group. Furthermore, the PCC is maximal for N in the range 16–18, showing that interested researchers can focus their attention on a reasonable number of the top players.

For sadness and joy, the results are much less clear as Table 1 shows. Models for sadness do not show a significant level of correlation for the big players, nor for the reference groups. As for joy, one notable PCC of 0.2517 for $N = 19$ does not seem significant.

15. Setting a higher minimum tweet limit would help to correct this problem; however, raising the minimum past 40 means there may not be a large enough pool of candidate users to fill the reference groups for larger values of N.

In order to determine if this lack of correlation for sadness and joy might be a consequence of the limitations of linear regression, we experimented with other learners as mentioned in Section 3. These include Gaussian processes, decision lists, random forests, and support-vector machines with first and second degree polynomial kernels. None of these algorithms outperformed the linear model.

We also used the affect models to predict variation in levels of emotion expressed in the Twitter community *during* the twelve-month period itself, rather than predicting over the last three months. To test this scenario, we evaluated the models using 10-fold cross-validation for each of the nine 12-month periods. From Table 2 one can see that the PCC does not vary significantly across values of N, and therefore it is difficult to identify a value of N of particular interest. Even more significant is the fact that the PCC obtained by the big players and those for the reference groups are similar. This indicates that the big players are comparable to other groups of the same size with respect to their predictive value within the twelve-month periods. This finding is in stark contrast to their predictive value for the last 3 months of these periods.

As an additional test, we repeated the experiment, but rather than using cross-validation, we created an independent test dataset by randomly selecting three-month's worth of data instances throughout the 12-month period, removing those instances from the training dataset. This method is of interest since it more closely parallels the methodology we followed when using the final three months of the period for model evaluation. The results with this independent test dataset were similar to those we obtained using cross-validation.

5 Discussion

When examining ways in which our results may help to further research on climate change, three hypotheses give likely interpretations of the correlation the models show for anger and fear:

1. Emotion expressed by a big player is representative of the larger community.
2. A big player's tweets are influencing the emotional state of the community. (High Twitter activity may indicate a user is seeking to gain an online presence or communicate a specific message to a perceived audience [3,18].)
3. A big player and the community are each influencing each other's emotional state. (Users interact online mainly with like-minded individuals [9,30].)

In each case, the big player may potentially be of high interest. The fact that the model identifies a relatively small number of big players greatly reduces the work effort involved in looking up user profiles and following specific chains of tweets. Furthermore, since models take into account different types of big player activity as well as a set of base emotions, they may be useful for organizations aiming to evaluate various types of high-level participation in order to improve communication methods which use emotion-based message framing.

Furthermore, researchers are exploring the relation between emotion at a social level and people's response to the dangers of climate change. Anger and fear are of particular interest, and our results show that big players can be a group of significant interest in the context of *#globalwarming* on Twitter. For instance, studies are looking into how fear affects people's reactions to information about climate change [13] and the role fear can play when framing messages intended to promote climate change advocacy [22].

Our models for sadness and joy, however, do not show any significant correlation. This does not seem to be an artifact of linear regression as the other algorithms do not produce better results. Hence, techniques to model sadness and joy may differ sharply from those for anger and fear. Further research on sadness in the context of high-activity online users is certainly warranted because this emotion is an important aspect in studies of human reactions to climate change. One example is the study by Farbotko and McGregor [7] exploring the influence which sadness can have on shaping international policy on climate change.

Our initial interpretations indicate that models analyzing sadness and joy may need to handle additional complications. We would not generally expect joy in particular to be a clear, unblurred emotion in messages about global warming. For instance, Sulis et al. [27] demonstrate that high levels of sadness may be found in tweets expressing irony, while joy occurs frequently in tweets expressing sarcasm. We would expect these emotions to be particularly difficult to model for communications on climate change, and it is intriguing to speculate on the extent to which irony and sarcastic remarks are influencing our models.

Limitations: With statistical models we must remember that finding a correlation does not mean we understand the causes behind the phenomenon we are studying. We present a method for predicting emotion levels in tweets about global warming, but we cannot say that the elements that we are considering as big player activity is causing the expression of emotion. We must also exercise a measure of restraint as we interpret our results. Tweets are a noisy, extremely informal, and non-standard use of language that traditional NLP techniques often find problematic [8]. Users may repeatedly send the same tweet (or retweet) numerous times; they may alter the original author's text when retweeting [3]; and they may use the "@" sign for purposes other than addressing another user. Additionally, emotions are only a part of the complex system that is human cognition. When using affective models to study how best to talk to people to inform them about climate change and work with them to mitigate its effects, we must continually be conscious of the underlying complexities and, as much as possible, avoid oversimplifying human understanding and behaviour [4].

6 Conclusion

This work shows that top high-activity users in the *#globalwarming* community on Twitter do not demonstrate a general predictive capacity compared to other groups of the same size. However, high-activity users do show a distinct predictive capability when predicting for the three months following the training period for anger and fear, two particularly relevant emotions with respect to climate change. Furthermore, this correlation occurs for the Top N players in groups small enough to allow researchers to follow up on them if needed. In contrast, this is not the case for sadness and joy, indicating that modelling these emotions is not a completely straightforward process in the context of online communications about global warming.

References

1. Aleahmad, A., Karisani, P., Rahgoza, M., Oroumchian, F.: Olfinder: Finding opinion leaders in online social networks. Journal of Information Science **42**(5), 659–674 (2016)
2. Bigonha, C., Cardoso, T., Moro, M., Gonçalves, M., Almeida, V.: Sentiment-based influence detection on Twitter. Journal of the Brazilian Computer Society **18**(3), 169–183 (2012)
3. Boyd, D., Golder, S., Lotan, G.: Tweet, tweet, retweet: Conversational aspects of retweeting on Twitter. In: Proceedings of the 43rd Hawaii International Conference on System Sciences. pp. 1–10 (2010)
4. Chapman, D., Lickel, B., Markowitz, E.: Reassessing emotion in climate change communication. Nature Climate Change **5**(12), 850–852 (2017)
5. Cody, E.M., Reagan, A.J., Mitchell, L., Dodds, P.S., Danforth, C.M.: Climate change sentiment on Twitter: An unsolicited public opinion poll. PLoS: ONE **10**(8), e0136092 (2015)
6. Eliacik, A.B., Erdogan, N.: Influential user weighted sentiment analysis on topic based microblogging community. Expert Systems With Applications **92**, 403–418 (2018)
7. Farbotko, C., Mcgregor, H.V.: Copenhagen, climate science and the emotional geographies of climate change. Australian Geographer **41**(2), 159–166 (2010)
8. Farzindar, A., Inkpen, D.: Natural Language Processing for Social Media. Morgan & Claypool, San Rafael, CA (2015)
9. Fersini, E.: Sentiment analysis in social networks: A machine learning perspective. In: Pozzi, F.A., Fersini, E., Messina, E., Liu, B. (eds.) Sentiment Analysis in Social Networks, pp. 91–111. Morgan Kaufmann (2017)
10. Frank, E., Hall, M.A., Witten, I.H.: The WEKA Workbench. Online Appendix for "Data Mining: Practical Machine Learning Tools and Techniques". Morgan Kaufmann, 4th edn. (2016)
11. Gimpel, K., Schneider, N., O'Connor, B., Das, D., Mills, D., Eisenstein, J., Heilman, M., Yogatama, D., Flanigan, J., Smith, N.A.: Part-of-speech tagging for Twitter: Annotation, features, and experiments. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2. pp. 42–47 (2011)
12. Halse, S.E., Tapia, A., Squicciarini, A., Caragea, C.: An emotional step toward automated trust detection in crisis social media. Information, Communication & Society **21**(2), 288–305 (2018)
13. Haltinner, K., Sarathchandra, D.: Climate change skepticism as a psychological coping strategy. Sociology Compass **12**(6), 84–85 (2018)
14. Honeycutt, C., Herring, S.C.: Beyond microblogging: Conversation and collaboration via Twitter. In: Proceedings of the 42rd Hawaii International Conference on System Sciences (2009)

15. Leiserowitz, A.: Climate change risk perception and policy preferences: The role of affect, imagery, and values. Climatic Change **77**(1), 45–72 (2006)
16. Liu, B.: Sentiment Analysis: Mining Opinions, Sentiments, and Emotions. Cambridge University Press, Cambridge (2015)
17. Lu, H., Schuldt, J.P.: Exploring the role of incidental emotions in support for climate change policy. Climatic Change **131**(4), 719–726 (2015)
18. Marwick, A.E., Boyd, D.: I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. New Media & Society **13**(1), 114–133 (2011)
19. Mitchell, L., Frank, M.R., Harris, K.D., Dodds, P.S., Danforth, C.M.: The geography of happiness: Connecting Twitter sentiment and expression, demographics, and objective characteristics of place. PLoS: ONE **8**(5), e64417 (2013)
20. Mobley, C., Vagias, W.M., DeWard, S.L.: Exploring additional determinants of environmentally responsible behavior: The influence of environmental literature and environmental attitudes. Environment and Behavior **42**(4), 420–447 (2010)
21. Mohammad, S.M., Bravo-Marquez, F.: Emotion intensities in tweets. In: Proceedings of the Joint Conference on Lexical and Computational Linguistics. pp. 65–77 (2017)
22. Nabi, R.L., Gustafson, A., Jensen, R.: Framing climate change: Exploring the role of emotion in generating advocacy behavior. Science Communication **40**(4), 442–468 (2018)
23. Plutchik, R.: The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. American Scientist **89**(4), 344–350 (2001)
24. Porter, M.F.: An algorithm for suffix stripping. Program **40**(3), 211–218 (1980/2006)
25. Preoţiuc-Pietro, D., Volkova, S., Lampos, V., Bachrach, Y., Aletras, N.: Studying user income through language, behaviour and affect in social media. PLoS ONE **10**(9), e0138717 (2015)
26. Stocker, T., Qin, D., Plattner, G.K., Tignor, M., Allen, S., Boschung, J., Nauels, A., Xia, Y., Bex, V., Midgley, P. (eds.): Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA (2013)
27. Sulis, E., Irazú Hernández Farías, D., Rosso, P., Patti, V., Ruffo, G.: Figurative messages and affect in Twitter: Differences between #irony, #sarcasm and #not. Knowledge-Based Systems **108**, 132–143 (2016)
28. Vigna, S.: Further scramblings of marsaglia's xorshift generators. Journal of Computational and Applied Mathematics **315**, 175–181 (2016)
29. Weber, E.: Experience-based and description-based perceptions of long-term risk: Why global warming does not scare us (yet). Climate Change **77**(1), 103–120 (2006)
30. 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)
31. Witten, I.H., Frank, E.: Data Mining: Practical Machine Learning Tools and Techniques. Morgan Kaufmann, 2nd edn. (2005)