

Collective Emotions Online

Anna Chmiel, Julian Sienkiewicz, Georgios Paltoglou, Kevan Buckley, Marcin Skowron, Mike Thelwall, Arvid Kappas, and Janusz A. Hołyst

Abstract This chapter analyzes patterns in messages posted to several Internet discussion forums from the perspective of the sentiment expressed in them and the collective character of observed emotions. A large set of records describing comments expressed in diverse cyber communities—blogs, forums, IRC channels, and the Digg community—was collected, and sentiment classifiers were used to estimate the emotional valence (positive, negative, or neutral) of each message. A comparison with simple models showed that the data included clusters of comments with the same emotional valence that were much longer than similar clusters created by a random process. This shows that there are emotional interactions between participants so that future posts tend to have the same valence as previous posts. Threads starting from a larger number of negative comments also last longer so negative emotions can be treated as a kind of discussion fuel; when the fuel (negativity) is used up in the discussion, it may finish. Moreover, the amount of user activity in a particular thread correlates positively with the presence of negative emotions expressed by the individual user in the thread. In summary, the analyses describe individual and collective patterns of emotional activities of Web forum users and suggest that negativity is needed to fuel important discussions.

A. Chmiel • J. Sienkiewicz • J.A. Hołyst (✉)

Center of Excellence for Complex Systems Research, Faculty of Physics, Warsaw University of Technology, Koszykowa 75, 00-662 Warsaw, Poland
e-mail: jholyst@if.pw.edu.pl

G. Paltoglou • K. Buckley • M. Thelwall

Statistical Cybermetrics Group, School of Technology, University of Wolverhampton, Wulfruna Street, Wolverhampton WV1 1LY, UK

M. Skowron

Interaction Technologies Group, Austrian Research Institute for Artificial Intelligence, Freyung 6/3/1a, 1010 Vienna, Austria

A. Kappas

School of Humanities and Social Sciences, Jacobs University Bremen, Campus Ring 1, 28759 Bremen, Germany

1 Introduction

The notion of *collective action* is usually associated with the responses of group members taken in order to maintain or improve the group's conditions (Wright et al. 1990). Research in this area frequently touches on the issue of *prediction* of the internal conditions that have to be fulfilled in order for the action to be undertaken [e.g., moral convictions (van Zomeren et al. 2012), injustice, efficacy, identity (van Zomeren et al. 2008)]. However, it seems that there is also a key component—emotions—that can play a pivotal role as an inhibitor of collective action, both directly (Stürner and Simon 2004) and indirectly (Taylor 1995). Moreover, the affective component can lead to the emergence or suppression of collective behavior (Sabucedo et al. 2011).

The Internet can be treated as a system of human behavior in which social dynamics are evident and measurable (Onnela and Reed-Tsochas 2010; Barabási 2005; Huberman et al. 1998; Sobkowicz and Sobkowicz 2010; Mitrović et al. 2010; Szell et al. 2010; Castellano et al. 2009; Kujawski et al. 2007; Chmiel et al. 2009). Communication in this medium displays different activity patterns compared to traditional communication (Radicchi 2009). People increasingly spend time using sites like MySpace, Facebook, Twitter, and blogs, and hence e-communities (Walther and Parks 2002) have become widespread and important. It is also obvious that collective actions take place in the Internet (Postmes and Brunsting 2002)—this medium can serve as both creator and transmitter of this notion [as in the case of Spanish “15M movement” (Borge-Holthoefer et al. 2012)]. The analysis of emotional interactions in e-communities is crucial for obtaining a comprehensive insight into social relations. In this study, we focus on collective emotional behavior, modeling its emergence, and the conditions that are necessary for it to happen.

Although emotions are typically expressed using a variety of nonlinguistic mechanisms—such as laughing, smiling, vocal intonation, and facial expression—textual communication can be just as rich and can be augmented by expressive textual methods—such as emoticons and slang (Gamon et al. 2005). Taking advantage of this, sentiment analysis, a research field in computational linguistics and computer science, has evolved rapidly in the last 10 years in response to a growing recognition of the importance of emotions in business and the increasing availability of masses of text in the social Web. The development of a number of algorithms to detect positive and negative sentiment has also made large-scale online text sentiment research possible, such as predicting elections by analyzing sentiment in Twitter (Tumasjan et al. 2010) and diagnosing trends for happiness in society via blogs (Dodds and Danforth 2010) and Facebook status updates (Pang and Lee 2008).

In this chapter, we discuss the impact of emotional expressions of Internet users on the vitality of online debates. We focus on (1) measuring the transfer of emotions between participants, (2) the influence of emotions on a thread's life span, and (3) the relationship between user behaviors and the emotionality of a discussion.

We especially focus on manifestations of collective emotional behavior in Sect. 3 while comparing emotional distributions with random equivalents. Section 4

touches on another interesting aspect of emotionally driven online discussions—the conditions that have to be fulfilled in order for a discussion to be fruitful (measured in the number of utterances). Finally, the BBC Forum analysis in Sect. 5 investigates the sentiments expressed by individual users within discussions at both global and local levels.

2 Data

We collected over six million comments from four prominent different interactive spaces: blogs, BBC discussion forums, popular social news Web site Digg, and #ubuntu IRC discussion channel. The texts were processed using sentiment analysis classifiers to predict their emotional valence (see Sect. 2.2), which according to a converging evidence is at the center of emotion experience and is considered one of the most important determinates of behavior from simple life-forms to humans (Lang and Davis 2006).

2.1 Datasets

The British Broadcasting Corporation (BBC)¹ (Chmiel et al. 2011a) Web site has a number of publicly open *Message Boards* covering a wide variety of topics that allow registered users to start their own discussions and post comments on existing discussions. Comments are postmoderated and anything that breaks the House Rules is subject to deletion. Our data included discussions posted on the Religion and Ethics and World/UK News message boards starting from the launch of the Web site (July 2005 and June 2005 respectively) until the beginning of the crawl (June 2009)—these were found to have interesting emotional content. The dataset comprises 100,000 discussions, 2.5 million comments and 18,000 users.

The Blogs dataset is a subset of the Blogs06 (Macdonald and Ounis 2006; Weroński et al. 2012) collection, which is an uncompressed 148 GB crawl of approximately 100,000 different blogs (more than three million Web pages) and spans 11 weeks, from 6 December 2005 to 21 February 2006. The subset was created by manually removing the HTML of the Blogs06 collection and keeping only the actual content (i.e., posts and comments). Only posts attracting more than 100 comments were extracted, as these seemed to initialize nontrivial discussions.

The Digg dataset comprises a full crawl of digg.com, one of the most popular social news Web sites. The analysis spans February to April 2009 and consists of all the stories, comments, and users that contributed to the site during this period.

¹ <http://www.bbc.co.uk>

The resulting dataset contains approximately 1.9 million stories, 1.6 million comments, and 800,000 users (Paltoglou et al. 2010; Pohorecki et al. 2013).

The Internet Relay Chat (IRC) is a medium that allows maintaining real-time multiuser discussions. The presented dataset contains information from the logs of #ubuntu discussion channels dating between 1 January 2007 and 31 December 2009. Data were preprocessed and transformed into a structure of over 90,000 one-to-one dialogues with almost 1.9 million comments (Sienkiewicz et al. 2013).

2.2 Algorithms

Sentiment analysis algorithms typically operate in three stages: (a) separate objective from subjective texts, (b) predict the polarity of the subjective texts, and (c) detect the sentiment target (Paltoglou et al. 2010). A variety of methods are used, including machine learning based upon the words used in each text, summarized in vector form (Riloff et al. 2006), and lexical approaches that start with a dictionary of known sentiment-bearing terms and apply linguistically derived heuristics to predict polarity from their occurrence and contexts (Wilson et al. 2009). One ongoing problem, however, is domain transfer—algorithms need to be tailored for each type of text that they are applied to because words tend to have different meanings in different contexts. In consequence, large sets of human-annotated data are needed to train and evaluate systems for each new application domain.

The algorithm that we used to detect and characterize the emotional content of posts is based on standard, supervised machine-learning principles (Sebastiani 2002). During a *training* phase, a corpus of documents is provided to the algorithm, i.e., a set of documents each one belonging to a specific category. The category of each document was determined during a preceding *corpus development* process by human experts who read its content and manually classified it. The algorithm extracts the characteristics of each class by analyzing the provided documents, i.e., “learns by example,” and stores this knowledge. Subsequently, during the *application/testing* phase, the algorithm applies the acquired knowledge to new, unseen documents and determines the best category under which they can be classified.

In this study, we implemented a hierarchical extension of a standard Language Model classifier (Sebastiani 2002) (abbreviated as h-LM). LM classifiers are a typical example of probabilistic classifiers, which estimate the probability that a document belongs to all of the available classes and select the one with the highest probability as the final prediction. In our hierarchical extension, a document is initially classified by the algorithm as objective or subjective, and in the latter case a second-stage classification determines its polarity, either positive or negative.

We used a manually annotated subset of Blogs06 dataset as *training corpus*. Human assessors examined approximately 34,000 documents for whether they contain factual information or positive/negative opinions about specific entities, such as people, companies, and films, and assigned a category to each document.

Because the distribution of documents per category is uneven in the specific corpus, the probability thresholds for both classification tasks were optimized on a small subset of humanly annotated BBC comments. The optimized classifier has an average F1-value of 66.68 % on the subjectivity detection task and 60.93 % on the polarity detection task on a humanly annotated BBC subset.

3 Cluster Distribution

To detect affective interactions between discussion participants, we calculated statistics for groups of comments with similar emotion levels. Each thread was transformed into a chain structure even if the the original structure was tree-like structure was present (see Fig. 1). The four datasets on which sentiment analysis was performed are characterized by different structures. In the case of Blogs and IRC dialogues, all posts were originally arranged in chains of successive comments, as shown in Fig. 1b. In other words, new posts are automatically added after the last post. On the other hand, the BBC and Digg data are arranged in a forum-like structure (see Fig. 1a), meaning that each user may make a comment to any previous post, thus starting a separate discussion. However, for the purpose of this study, each discussion (thread) in each dataset was arranged chronologically, as in Fig. 1b. In this way, it was possible to compare these different communities. Although BBC and Digg have a forum-like structure, the default view as presented to the user was chronological. Thus, a chronological simplification for analysis can be justified.

We define an emotional cluster of size n as a chain of n consecutive messages with the same sentiment orientation, i.e., negative, positive, or neutral, where before the cluster and after it there is message with the valence different from the cluster valence (see upper row of Fig. 2) (Chmiel et al. 2011b). For the purpose of comparison, we show also the shuffled data obtained from the same discussion (see bottom row of Fig. 2), which display clearly shorter clusters than those in the original data. A possible explanation of the observed fact could be the emotional interactions between the participants of the discussion in the original data.

In order to test this hypothesis, we checked the complementary cumulative distribution function (CCDF) $P^{(e)}(\geq n)$ that describes the frequency of clusters of size greater than or equal to n for the cases of negative, positive, and neutral emotions, $e = \{-1, 0, 1\}$. The results are presented as symbols in Fig. 3. For comparison, the corresponding CCDF from the independent and identically distributed (i.i.d.) random process

$$P_{\text{iid}}^{(e)}(\geq n) = p(e)^{n-1} \quad (1)$$

(dotted lines) is also plotted, where $p(e)$ is the probability of negative, positive, or neutral emotion (see Table 1). The i.i.d. random process (Feller 1968) is an effect




Fig. 1 The difference between the actual tree structure (a) present in the BBC and Digg datasets as compared to the chronological layout of the posts (b), e.g., originally present in Blogs and IRC. The *numbers* indicate the order of messages (1 being the first, 10 being the last) while *arrows* indicate that a post was given in reply to another one (e.g., post 9 is the response to post 7)




Fig. 2 An example of a discussion in the “Eastern religion” BBC Forum. The original thread that consists of 22 posts is shown in the *upper row*. Each *box* represents one post. *Red, blue, or black boxes* indicate that the comment was classified as, respectively, positive, negative, or neutral (objective). The *bottom row* presents shuffled data, i.e., the comments were arranged in a random order

of the simplest stochastic process where there is no statistical dependence between events at consecutive time-steps, and at every time-step the event probability distribution is the same. The fit diverges for large n , which means that in the data collected the probability of long clusters of the same emotional valence is large




Fig. 3 Complementary cumulative distribution function (CCDF) $P^{(e)}(\geq n)$ of the cluster size for all data used in the study. Symbols are data (blue triangles, red circles, and white squares, respectively, for negative, positive, and neutral clusters), dotted lines are i.i.d. processes given by Eq. (1), dashed lines are Markov processes given by Eq. (2), while solid lines come from Eq. (3) and represent distributions based on the preferential attraction rule. The spurious increase of $P^{(e)}(\geq n)$ for $n \geq 40$ for Blogs data is due to violation of the scaling $p(ele) = p(ele)n^\alpha$

Table 1 Datasets’ properties

	N	T	$\langle e \rangle$	Probability of emotion			Exponent α		
				$p(+)$	$p(-)$	$p(0)$	α_+	α_-	α_0
BBC	2,474,781	97,946	-0.44	0.19	0.65	0.16	0.38	0.05	0.45
Digg	1,646,153	129,998	-0.16	0.31	0.48	0.21	0.20	0.11	0.37
IRC	1,899,119	93,379	0.16	0.32	0.15	0.53	0.18	0.31	0.07
Blogs	242,057	1,219	0.14	0.35	0.22	0.43	0.23	0.19	0.16

Properties of the four datasets: number of comments N ; number of discussions T ; average valence in the dataset $\langle e \rangle$; probability of finding positive ($p(+)$), negative ($p(-)$), and neutral emotion ($p(0)$); values of α exponents for positive (α_+), negative (α_-), and neutral (α_0) clusters

compared to the probability expected for mutually independent messages. It follows that there is a tendency for emotions of the same valence to cluster together, suggesting that there may be attractive affective forces between discussion participants—posts tend to trigger follow-up posts of the same valence.

Then, could emotions in the discussions be described by a Markov process? A Markov chain (Norris 1997) is a basic stochastic process with one memory step

when the probability of the next time state depends only on the previous one by corresponding conditional probabilities. In this case,

$$P_M^{(e)}(\geq n) = p(e|e)^{n-1} \quad (2)$$

(dotted lines in Fig. 3), where $p(ele)$ is the conditional probability that two consecutive messages have the same emotion. It is defined as $p(ele) = p(ee)/p(e)$, where $p(ee)$ is the joint probability of the pair ee that is measured as a number of occurrences of the two consecutive messages with the same valence e divided by the number of all appearing pairs. The fit is better than for the i.i.d., but there is still divergence for large n .

Now, let us consider conditional probability $p(elne)$ that after n comments with the same emotion valence the next comment will have the same valence. The data reveal the relation $p(ele) < p(elee) < \dots < p(elne) \approx p(ele)n^\alpha$ with the characteristic exponent α representing the strength of the sublinear preferential process and α in $[0;1]$. This relationship means that finding a positive message after seven positive comments is more likely than after six. It holds true for $n < 10$, but then saturation follows, finally decreasing to zero for long clusters (see Fig. 4). Preferential processes are common in complex systems with positive loop dynamics, and they are responsible for the emergence of fat-tailed distributions, including power-law scaling (Barabási and Albert 1999; Krapivsky and Redner 2001). The influence of preferential attraction is visible in the CCDF in Fig. 3. Its decay is much slower than in the case of random or even one-step Markov processes. In order to find an analytical approximation to the cluster distribution, extending the scaling relation $p(elne)$ leads to the following approximation of the CCDF:

$$P_\alpha^{(e)}(\geq n) \approx p(e|e)^{n-1} [(n-1)!]^\alpha \quad (3)$$

This approximation is presented in Fig. 3 (solid lines). The fit with the data is far better than in the case of the i.i.d. and one-step Markov distributions, especially for large n . The differences between the analytical assumption and the real data come from the artificial extension of the scaling relation $p(elne)$, which results in underestimation of the norm of probability and, consequently, overestimation of the number of clusters. The range of applicability of this analytical result is limited since for large n the function $P_\alpha^{(e)}(\geq n)$ possesses a minimum depending on the parameters $p(ele)$ and α , and it diverges as n goes to infinity.

Figure 3 confirms that the occurrence of emotional posts cannot be described by the i.i.d. process, and there are specific correlations between emotions in consecutive posts. These correlations result from emotional interactions between discussion participants via their messages. The interactions possess an attractive character because clusters of posts with the same emotional valence are longer than clusters from random distributions. The emotion expressed by a participant depends on the emotions in previous posts—he/she tends to express emotions that have been recently used in the discussion. This observation is consistent with general ideas regarding functions of emotions (e.g., Frijda 1986). Moreover, the relations

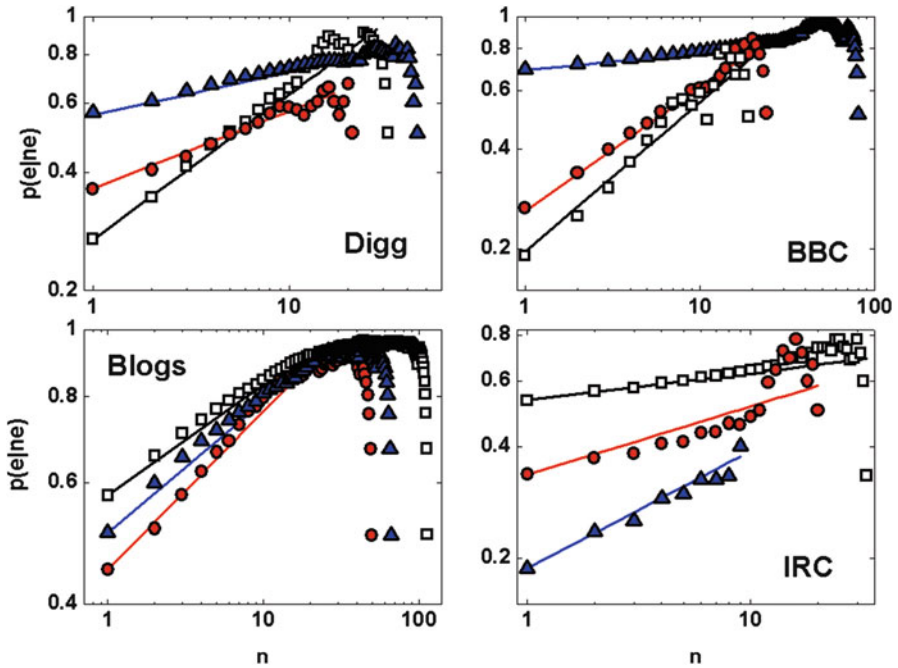


Fig. 4 The conditional probability $p(e|ne)$ of the next comment occurring having the same emotion for Digg, BBC, Blogs, and IRC data. Symbols are data (blue triangles, red circles, and white squares, respectively, for negative, positive, and neutral clusters) and lines reflect the fit to the preferential attraction relation $p(e|ne) = p(ene)n^\alpha$

observed in both Figs. 3 and 4 indicate that the behavior of the participants can be regarded as a collective one—the more the emotional posts submitted the more emotional will be the next one.

Note that the collective effect happens regardless of emotion valence—in other words, the affective attractive forces that produce a snowball effect in a set of consecutive messages are not sensitive to the emotion type. The prerequisite for a significant value of such a force is a large *number of messages* in the preceding cluster of homogeneous emotions. This observation supports recent findings by Sabucedo et al. (2011) who showed that both positive and negative emotions (respectively, enthusiasm and anger in their study) can be responsible for triggering collective actions (political demonstrations).

4 Life-Span of the BBC Forum and Digg Communities

The influence of emotions on the duration of BBC forum and Digg discussions was investigated as follows. Threads of the same size were grouped together, and a moving average of the emotion type of the last ten comments was calculated for




Fig. 5 Time dependence of emotions in BBC Forum (a, b, c) and Digg (d) threads. (a) Average emotion valence in the thread (moving average of the previous 10 messages in the thread). Four groups of threads of lengths 20, 40, 60, and 80 are represented by different symbols (respectively *circles*, *squares*, *triangles*, and *diamonds*). Shorter threads start from emotional levels closer to zero. (b) Emotional level (valence) at the beginning of a thread minus the emotional level at the end as a function of thread length (*grey symbols*). *Black triangles* display binned data. Longer threads use more emotional “fuel” over time. (c) Average length of the thread as a function of the absolute value of the average emotion valence of the first 10 comments. Emotional thread starts, whether positive or negative, usually lead to longer discussions. (d) Average length of thread in Digg data as a function of the average valence of the first 10 comments. Here, only negative start of thread leads to longer discussions

each point. As seen in Fig. 5a, shorter threads tend to start from a lower (i.e., less negative) emotional level than longer ones. On the other hand, threads end with a similar mean emotional valence value regardless of their lengths—the last point of each data series in 3a (circles, squares, triangles, and diamonds) is at almost the same level, i.e., about -0.42 . This phenomenon is echoed in Fig. 5b where the average emotional valence of the first ten comments minus the average emotional valence of the last ten comments is plotted, showing that longer threads have bigger eventual decreases in negative valence. Figure 5c also suggests that the initial emotional content (whether positive or negative) may be used as an indicator of the expected length of a thread—low absolute average emotion valences lead to shorter discussions. A possible heuristic explanation is that the first few posts in a

thread may give it the potential (emotional fuel) to propel further discussion. Once the emotions driving the discussion dry out, the thread is no longer of interest to its participants and it may die. For the threads possessing higher initial levels of emotion, it takes more comments to resolve the emotional issue, resulting in longer threads. A similar although slightly different phenomenon is spotted in the Digg data. Here, as seen in Fig. 5d, only negative start of the thread prolongs the discussion.

5 Users Impact on the Discussion in the BBC Forum

Here we consider user activity a_i defined to be the total number of posts written by user i in all discussion threads during the observation period. For simplicity, this quantity will also be referred to as a . The maximum observed activity in the dataset is $a_{\max} = 18274$, i.e., one user authored more than 18,000 messages, while the average activity is $\langle a \rangle = 137$ and the median $m_a = 3$. The number of occurrences of a is presented in Fig. 6a (red triangles) and it is well fitted by the power-law relationship $h_a \sim a^{-\beta}$, $\beta = 1.4$ (black line in Fig. 6a). The relatively small value of the exponent β suggests a high number of very active users of this forum. Since all discussions in the forum are split into separate threads j , we define d_{ij} (or d for short) to be the local activity of user i in thread j measured by the number of posts that this user submitted in the discussion. Whereas both its maximum and average values (respectively $d_{\max} = 1582$ and $\langle d \rangle = 2.84$) are lower than in the case of a , the number of occurrences of d shown as black circles in Fig. 6a still follows a similar relationship $h_d \sim d^{-\gamma}$ with exponent $\gamma = 2.9$ (red line in Fig. 6a), which is double that of a .

While user behavior shows a strong tendency to be scale-invariant, this is not so clear for the thread statistics shown in Fig. 6b. Here, we consider thread length L and the number of unique users U posting at least one comment in the thread. Histograms of both quantities h_L (black circles) and h_U (red triangles) display power-law tails for $U, L > 20$. This is most prominent in the case of h_U , which is also characterized by a rather large exponent $\eta = 4.9$ (black line in Fig. 6a).

To understand the impact exerted by the most frequent users on the length of a thread, consider the dependence between the normalized number of unique users in a single thread defined by $u = U/L$ and thread length (Fig. 6c). For short threads (L between 1 and 10), u is about 0.6–1 while for threads larger than 400 comments, it drops below 0.1. A good fit is $u(L) = A(L + b)^{-0.58}$ (the blue line in Fig. 6c); thus, the number of unique users grows more slowly than linearly with thread lengths. This suggests that mutual discussions between specific users rather than a large number of independent comments submitted by many users sustain thread life.

The following quantities describe the emotions of individual debaters and discussions threads. The average (global) emotion of a user e_a is the sum of all emotions e in posts written by the user i divided by his/her activity a_i . The average emotion of a thread e_L is the sum of all emotions in the thread j divided by its length




Fig. 6 (a) Histogram of user activity a (*triangles*); histogram of user activity in thread d (*circles*). Lines are fits to the data and they follow relations $h_a \sim a^{-\beta}$ and $h_d \sim d^{-\gamma}$ with $\beta = 1.4$ and $\gamma = 2.9$. (b) Histogram of threads with length L (*circles*); histogram of the number of unique users U making a comment in the thread (*triangles*). The *black line* is a fit to the tail of the distribution and it follows the relation $h_U \sim U^{-\eta}$ with $\eta = 4.9$. (c) The normalized number of unique users u making a comment in a thread of length L (*dots*—original data, *triangles*—data binned logarithmically). The *blue line* corresponds to relation $u = A(L+b)^{-\delta}$ with fitted parameters $\delta = 0.58$, $A = 3.72$ and $b = 8.6$





Fig. 7 (a) Probability distribution of users' global average emotions $\langle e \rangle_a$. (b) Users' global average emotion $\langle e \rangle_a$ versus users' global activity a . (c) Users' average global activity $\langle a \rangle$ ($\langle e \rangle_a$) versus their global average emotion $\langle e \rangle_a$: *red bars*—empirical data, *black bars*—shuffled data. (d) Relationship between users' average emotion in a thread $\langle e \rangle_d$ and users' activity in the thread d . *Grey circles* are original data, *black triangles* are binned data, and the *red curve* corresponds to equation $\langle e \rangle_d = A_1 + B_1 \ln(d+b)$ with $A_1 = -0.31$, $B_1 = -0.054$ and $b = 8.6$

L_j . The third value e_d is the average emotional expression of the user i in the thread j . The main features of the distribution $p(e_a)$, presented in Fig. 7a, are peaks for $e_a = -1, 0, 1$ which are a straightforward effect of the large number of users with $a = 1$ and threads with $L = 1$ (see Fig. 6a, b). The local maximum around $e_a = -0.5$ is a specific attribute of the BBC Forum because it possesses a strong bias toward negative emotions, with an average value of $\langle e \rangle = -0.44$. We observe similar distribution shapes for $p(e_l)$ and $p(e_d)$.

So far we have treated user activities and emotions as mutually independent variables, but we now consider the relationship between them. Figure 7b plots users' global average emotions $\langle e \rangle_a$ versus global activity a . Neglecting fluctuations for large values of a caused by small numbers of very active users, there is a

Fig. 8 Relationship between the average emotion in discussions with fixed thread length L (diamonds), circles are binned data, the line corresponds to the relationship $\langle e \rangle_L = A' + B' \ln(L)$ with fitted parameters $A' = -0.34$, $B' = -0.03$. The inset shows the same plot for the shuffled data



constant mean emotion that is around the forum's average value $\langle e \rangle$. Hence, on average, the user activity level a does not influence his/her average emotions $\langle e \rangle$. In Fig. 7c, the reversed relationship is plotted, i.e., the average global activity versus users' average emotions (red bars). For comparison, we present shuffled data (black bars) where the emotional values of posts were randomly interchanged between users. Whereas the second distribution follows a Gaussian-like function, the original set is characterized by a broad maximum stretching across almost all of the negative part of the plot and some minor fluctuations in the positive part. This means that although there are the same mean emotions for groups of users of various activities (see Fig. 7b), there are different average activities for users of various mean emotions.

Users can take part in many threads; thus, their local and global activities as well as corresponding local emotions can be very different. But how are users' emotions $\langle e \rangle_d$ expressed in a thread connected to the activity level d in it? Figure 7d shows the average emotions of a user in a thread as a function of the user's local activity. In this case, an increase in activity in a particular thread leads to more negative average emotions in the thread. Recall that there was no relationship between a user's global activity and his/her emotions, as shown in Fig. 7b. For longer discussions, there is a more homogeneous group of users (see Fig. 6c); thus, on average one user writes a larger number of posts $\langle d \rangle(L) = 1/u(L)$. As shown in Fig. 7d, the average emotions for users locally more active decreases logarithmically. These two effects cause the longer threads to possess, on average, more negative emotions. In fact in Fig. 8, there is a logarithmic decay in mean thread emotions $\langle e \rangle_L$ as a function of thread length L . To confirm the statistical validity of this phenomenon, we randomly shuffled emotions between various threads. The inset in Fig. 8 proves that in this case mean thread emotion is independent of thread length. The qualitative outcome of Fig. 8 resembles Fig. 5b with respect to the idea of emotional (negative) fuel that has to be included in the discussion in order to sustain it.

6 Conclusions

On the basis of automatic sentiment detection methods applied to huge datasets, we have shown that Internet users' messages correlate at the simplest emotional level—positive, negative, or neutral messages tend to provoke similar responses. The collective character of the emotions expressed was evident in several different types of e-community—it was observed for BBC forums and Digg (mainly negative emotions), for the Blogs (mainly positive comments), and also in IRC dialogues (neutral). The strength of emotional interactions can be indirectly measured by the parameter α expressing the influence of the most recent emotional cluster on the probability that the next post has the same emotion. The results indicate the presence of online collective behavior among users that creates longer discussion threads.

We also found patterns in individual users' emotional behaviors in online BBC Forums. We observed a scale-free distribution of users' activities in the whole forum and in singular threads as well as power law tails for the distribution of thread lengths and the number of unique users in a thread. At the level of the entire forum, negative emotions boost users' activities; participants with more negative emotions write more posts. At the level of individual threads, users that are more active in a specific thread tend to express more negative emotions in it and seem to be the key agents for sustaining discussion. As result, longer threads possess more negative emotional content. Overall, then, negativity is the key sentiment to start and sustain online discussions, at least in the forums investigated here.

Acknowledgments This work was supported by a European Union grant by the 7th Framework Programme, Theme 3: Science of complex systems for socially intelligent ICT. It is part of the CyberEmotions (*Collective Emotions in Cyberspace*) project (contract 231323). J.A.H, A.Ch. and J.S. acknowledge support from Polish Ministry of Science Grant 1029/7.PR UE/2009/7.

References

- Barabási A-L (2005) The origin of bursts and heavy tails in human dynamics. *Nature* 207:435–433
- Barabási A-L, Albert R (1999) Emergence of scaling in random networks. *Science* 286:509–519
- Borge-Holthoefer J, Rivero A, Moreno Y (2012) Locating privileged spreaders on a online social network. *Phys Rev E* 85:066123
- Castellano C, Fortunato S, Loreto V (2009) Statistical physics of social dynamics. *Rev Mod Phys* 81:591–646
- Chmiel A, Kowalska K, Hołyst JA (2009) Scaling of human behavior during portal browsing. *Phys Rev E* 80:066122
- Chmiel A, Sobkowicz P, Sienkiewicz J, Paltoglou G, Buckley K, Thelwall M, Hołyst JA (2011a) Negative emotions boost user activity at BBC forum. *Phys A* 390:2936–2944
- Chmiel A, Sienkiewicz J, Thelwall M, Paltoglou G, Buckley K, Kappas A, Hołyst JA (2011b) Collective emotions online and their influence on community life. *PLoS ONE* 6(7):e22207
- Dodds PS, Danforth CM (2010) Measuring the happiness of large-scale written expression: songs, blogs, and presidents. *J Happiness Stud* 11:441–456

- Feller W (1968) An introduction to probability: theory and its applications. Wiley, Hoboken, NJ
- Frijda NH (1986) The emotions. Cambridge University Press, Cambridge, MA
- Gamon M, Aue A, Corston-Oliver S, Ringger E (2005) Lect Notes Comput Sci 3646:121–132
- Huberman BA, Pirolli PLT, Pitkow JE, Lukose RJ (1998) Strong regularities in World Wide Web surfing. *Science* 280:95–97
- Krapivsky P, Redner S (2001) Organization of growing random networks. *Phys Rev E* 63:066123
- Kujawski B, Hołyst JA, Rodgers GJ (2007) Growing trees in internet news groups and forums. *Phys Rev E* 76:036103
- Lang PJ, Davis M (2006) Emotion, motivation, and the brain: reflex foundations in animal and human research. *Prog Brain Res* 156:3–34
- Macdonald C, Ounis I (2006) The TREC Blogs06 collection: creating and analyzing a blog test collection (Technical Report TR-2006-224). Department of Computer Science, University of Glasgow, Glasgow
- Mitrović M, Paltoglou G, Tadić B (2010) Networks and emotion-driven user communities at popular blogs. *Eur Phys J B* 77:597–609
- Norris JR (1997) Markov chains. Cambridge University Press, Cambridge, MA
- Onnela J-P, Reed-Tsochas F (2010) Spontaneous emergence of social influence in online systems. *Proc Natl Acad Sci USA* 107:18375–18380
- Paltoglou G, Thelwall M, Buckely K (2010) Online textual communication annotated with grades of emotion strength. In: Proceedings of the third international workshop on EMOTION (satellite of LREC): corpora for research on emotion and affect, Valletta, Malta, pp 25–31
- Pang B, Lee L (2008) Opinion mining and sentiment analysis. *Found Trends Inf Retr* 1(1–2):1–135
- Pohorecki P, Sienkiewicz J, Mitrović M, Paltoglou G, Hołyst JA (2013) Statistical analysis of emotions and opinions at Digg website. *Acta Phys Pol* 123:604–614
- Postmes T, Brunsting S (2002) Collective action in the age of the internet mass communication and online mobilization. *Soc Sci Comput Rev* 20:290–301
- Radicchi F (2009) Human activity in the web. *Phys Rev E* 80:026118
- Riloff E, Patwardhan S, Wiebe J (2006) Feature subsumption for opinion analysis. In: Proceedings of the conference on empirical methods in natural language processing, Morristown, NJ, USA, pp 440–448
- Sabucedo JM, Durán M, Alzate M, Barreto I (2011) Emotions, ideology and collective political action. *Univ Psychol* 10:27–34
- Sebastiani F (2002) Machine learning in automated text categorization. *ACM Comput Surv* 34(1):1–47
- Sienkiewicz J, Skowron M, Paltoglou G, Hołyst JA (2013) Entropy-growth-based model of emotionally charged dialogues. *Adv Complex Syst* 16:1350026
- Sobkowicz P, Sobkowicz A (2010) Dynamics of hate based Internet user networks. *Eur Phys J B* 73:633–643
- Stürmer S, Simon B (2004) Collective action: towards a dual pathway model. *Eur Rev Soc Psychol* 15:59–99
- Szell M, Lambiotte R, Turner S (2010) Multirelational organization of large-scale social networks in an online world. *Proc Natl Acad Sci USA* 107:13636–13641
- Taylor V (1995) Watching for vibes: bringing emotions into the study of feminist organizations. In: Ferree MM, Martin PY (eds) *Feminist organizations: harvest of the new women's movement*. Temple University Press, Philadelphia, pp 223–233
- Tumasjan A, Sprenger TO, Sandner PG, Welpel IM (2010) In: Proceedings of the fourth international AAAI conference on weblogs and social media. AAAI Press, Menlo Park, CA, pp 178–185
- van Zomeren M, Postmes T, Spears R (2008) Toward an integrative social identity model of collective action: a quantitative research synthesis of three socio-psychological perspectives. *Psychol Bull* 134:504–535
- van Zomeren M, Postmes T, Spears R (2012) On conviction's collective consequences: integrating moral conviction with the social identity model of collective action. *Br J Soc Psychol* 51:52–71

- Walther J, Parks M (2002) In: Knapp M, Daly J, Miller G (eds) *The handbook of interpersonal communication*. Sage, Thousand Oaks, CA, pp 529–563
- Weroński P, Sienkiewicz J, Paltoglou G, Buckley K, Thelwall M, Hołyst JA (2012) Emotional analysis of blogs and forums data. *Acta Phys Pol A* 121:B128–B132
- Wilson T, Wiebe J, Hoffman P (2009) Recognizing contextual polarity: an exploration of features for phrase-level sentiment analysis. *Comput Linguist* 35(3):399–433
- Wright SC, Taylor DM, Moghaddam FM (1990) Responding to membership in a disadvantaged group: from acceptance to collective protest. *J Person Soc Psychol* 58:994–1003