



OPEN

SUBJECT AREAS:
PSYCHOLOGY AND
BEHAVIOUR
NETWORK TOPOLOGY
STATISTICS

Received
28 March 2014

Accepted
3 October 2014

Published
27 October 2014

Correspondence and
requests for materials
should be addressed to
X.L. (xin.lu@
flowminder.org)

* These authors
contributed equally to
this work.

Network Structure and Community Evolution on Twitter: Human Behavior Change in Response to the 2011 Japanese Earthquake and Tsunami

Xin Lu^{1,2,3,4*} & Christa Brelsford^{5,6*}

¹College of Information System and Management, National University of Defense Technology, 410073 Changsha, China,
²Flowminder Foundation, 17177 Stockholm, Sweden, ³Department of Public Health Sciences, Karolinska Institutet, 17177 Stockholm, Sweden, ⁴Department of Sociology, Stockholm University, 10691 Stockholm, Sweden, ⁵School of Sustainability, Arizona State University, Tempe, Arizona, USA, ⁶Santa Fe Institute, Santa Fe, New Mexico, USA.

To investigate the dynamics of social networks and the formation and evolution of online communities in response to extreme events, we collected three datasets from Twitter shortly before and after the 2011 earthquake and tsunami in Japan. We find that while almost all users increased their online activity after the earthquake, Japanese speakers, who are assumed to be more directly affected by the event, expanded the network of people they interact with to a much higher degree than English speakers or the global average. By investigating the evolution of communities, we find that the behavior of joining or quitting a community is far from random: users tend to stay in their current status and are less likely to join new communities from solitary or shift to other communities from their current community. While non-Japanese speakers did not change their conversation topics significantly after the earthquake, nearly all Japanese users changed their conversations to earthquake-related content. This study builds a systematic framework for investigating human behaviors under extreme events with online social network data and our findings on the dynamics of networks and communities may provide useful insight for understanding how patterns of social interaction are influenced by extreme events.

Understanding how human social behavior changes in response to extreme events like earthquakes, tsunamis or terrorist attacks is key to emergency response and recovery^{1,2}. In this paper, we consider the subset of human behavior that is demonstrated in the patterns, context and content of interactions between individuals, especially those mediated by internet enabled devices. Observations of this behavior during extreme events may provide insights into general behavior patterns during times of stress because links between unobservable decisions and the observable actions taken in response to those decisions are more likely to be temporally close due to the time sensitive nature of many extreme events. However, the study of human behavior during extreme events has historically been hindered due to the limitation of available information during such scenarios¹.

Studies of behavior during extreme events have primarily been made retrospectively, e.g., through field surveys or census sampling after the event³. For example, Elliott and Pais studied Hurricane Katrina's influences on a wide array of survivors' responses, including evacuation timing and emotional support to housing and employment situations and plans to return to prestorm communities⁴. Souza et al examined changes in the mental health status of vulnerable communities after the 2004 Indian Ocean earthquake and tsunami⁵.

While retrospective studies are critical for understanding human behavior, limitations are also obvious: first, retrospective studies are often made long after the event and cannot provide timely support for decision making when the need is most urgent; and second, data is collected through surveys or interviews, introducing bias due to reporting errors, recollection bias, and challenges identifying and accessing disaster-affected populations. To overcome these issues, researchers have recently used more objective and timely data, generated from sensor networks such as cell phone towers, to track individual mobility and population flow for large populations in real time, providing a unique solution for disaster response and relief management^{2,6,7}.



Retrospective surveys and time-location based data improves our ability to investigate human psychological and physical behavior under disasters. However, they do not provide insight into patterns of human interactions and social behavior. For example, when a disaster happens, what do people do as part of a social group? What is the role of each individual and how do they interact with people around them? How are patterns of interactions during extreme events different from those during more typical conditions? The rapid development of online social networking services has enabled researchers to answer these questions: with internet-enabled devices, any user can post and share news, events, thoughts, and the like, providing not only the content of individual psychological and physical activities, but also the context of human interactions such as the target and topic of interaction. When the regional physical telecommunication system still functions after an adverse event, the use of these online tools increases enormously, demonstrating their salience for real world interactions and information sharing behavior. Although the population that regularly uses internet-enabled devices to communicate is not representative of the population as a whole, the records of those interactions do represent a window into social behavior that has not previously been observable. Thus, studies relying on real-time records of interactions complement but do not replace existing retrospective analyses. These real-time perspectives provide more breadth of data on patterns and content of interaction behaviors, with the limitation that the only data available are those interactions that are mediated by internet-enabled devices and so are not representative of all social interactions.

Studies of online social media under adverse events were pioneered by utilizing data from Twitter, a leading micro-blogging service which allows users to rapidly communicate information in up to 140 characters (tweet) on a one-to-one, specified group or global basis⁸. The temporal, spatial and social dynamics of Twitter activity have intrigued many researchers in developing applications to facilitate early event detection and increase situational awareness^{9,10}. The behavior observable through Twitter is not representative of all social interactions but we believe that it does have the potential to show how extreme events influence patterns of interactions, despite the challenges in obtaining representative samples.

The US Geological Survey (USGS) has found that when an earthquake is felt by a population that uses Twitter, tweets reporting the incidence of an earthquake are published online sooner than the 2 to 20 minutes it takes the USGS to publicly distribute instrumentally derived estimates of location and magnitude¹¹. In an evaluation implemented by Paul et al, the authors found that tweets referencing earthquakes may be useful information for detecting earthquakes in poorly instrumented regions¹². Instant change of user activities and appearance of event-related keywords have been seen in a variety of adverse events such as earthquakes, cross border attacks, and wild fires¹³. For example, in the MW 4.3 earthquake in Morgan Hill, CA, 2009, the tweet frequency in the epicenter region quickly rose to about 150 per minute, in comparison to the background level of less than one per hour before the earthquake¹¹. Information regarding the location and specific details of events was reported on Twitter within seconds following the first explosion of the April 15th, 2013 bombing at the Boston Marathon¹⁴.

Understanding changes in patterns of user activities during extreme events requires dynamic analysis of how individuals communicate with others before and during an extreme event and the communities they interact with¹⁵. Interactions can be mapped into a network structure where nodes are individual users and edges are records of interaction between two users. Community structure can be observed through patterns of coherent and sustained interactions between groups of individuals. Despite many studies focusing on tweeting activity and text mining, studies analyzing behavioral interactions under extreme events have been rare. Preliminary studies have been made, for example, by Gupta et al., who proposed a

clustering algorithm in detecting user communities and applied the method in analyzing Twitter users in three major crisis events of 2011 (Hurricane Irene, Riots in England, and Earthquake in Virginia). They found that top users represent the topics and opinions of all the users in the community with 81% accuracy on an average¹⁶. In the 2010 Chile earthquake, Mendoza et al. analyzed information propagation on Twitter network and found that rumors tend to be questioned more than news by the Twitter community¹⁷. In a recent study, Chatfield and Brajawidagda conducted a social network analysis of Twitter information flow among the central disaster warning agency's (BMKG) Twitter followers during the 2012 Indonesia Earthquake and found that even with less than 0.1% of all followers re-tweeting the initial reports, tsunami early warnings released by BMKG could potentially reach 4 million users in 15 minutes¹⁸.

There are preliminary analyses of human interactions through Twitter under extreme events from a network perspective in the existing scientific literature. However, current research lacks a systematic framework to characterize the mechanisms and pattern changes observed in social interactions during adverse events. To investigate the dynamics of social networks and the formation and evolution of communities in response to extreme events, we collected three unique datasets from the weeks before and after the March 11th, 2011 Japanese earthquake: 1% of all tweets posted worldwide, sampled approximately uniformly from Twitter using the "Spritzer" API, which we refer to as the Global dataset; tweets written in Japanese that were sampled from the "Topsy" API, which we refer to as the TP-JP dataset; and TP-EN: tweets written in English and sampled from the "Topsy" API (See Materials and Methods for details). Hereafter, we refer to people who communicate online in Japanese as Japanese speakers, but they are not a representative sample of all people who speak Japanese, or all people who were affected by the earthquake. Nonetheless, we believe that changes in the behavior of this Japanese speaking population in response to the earthquake and tsunami can provide some insight about changes in social behavior in Japan in response to the earthquake and tsunami. We consider the TP-JP dataset to represent a treatment group, because a larger proportion of this group's population was affected by the earthquake than the populations of the other two datasets. Thus, changes in the structure of the TP-JP network before and after the earthquake that did not also occur in the TP-EN or Global dataset are likely related to the events that influenced Japanese speakers on twitter.

We are using a broad dataset intended to preferentially select the population of people who are affected by the earthquake, and examine at how their social interaction patterns change for the broad population affected by the disaster. In this sense, we are looking at population level changes in patterns of social interactions, rather than using records of online interactions to trace details of events, information discovery, coordination efforts, and emergency response^{19–21}. The detail-oriented approach of this crisis-informatics research is crucial for developing effective real time emergency response strategies and planning efforts. Understanding changes in social interaction patterns in response to an adverse event may provide insight into how patterns of interaction behaviors change in response to stress, which could inform planning disaster response efforts that deal with behavior at a population level; for example evacuation procedures or information dissemination.

By comparing network and community structure before and after the earthquake, we were able to examine the effects of the earthquake on networks of communication in social media. We find that while almost all users demonstrate increased activity after the earthquake, Japanese speakers expanded their network of interactions to a much greater degree. Japanese speakers are more likely to be directly affected by the earthquake than a person who uses another language. By applying the community detection algorithm of information mapping (Infomap)^{22,23}, we find that the behavior of joining or leaving a

**Table 1 | Data description and basic network characteristics**

Data	Period	Raw data			Extracted network		
		#users	#tweets	#links	#nodes	\bar{D}	s
TP-JP	before	36,893	42,678	23,444	3,647	2.8	2,791
	after	36,893	114,247	112,004	3,647	5.7	3,191
TP-EN	before	36,234	49,095	36,394	4,037	2.6	3,261
	after	36,234	64,507	49,142	4,037	2.9	3,417
Global	before	7,403,050	6,184,122	4,191,050	87,939	2.0	40,568
	after	7,403,050	8,062,823	5,545,562	87,939	2.1	38,552

D: average degree; s: size of giant connected component or giant strongly connected network; c: clustering coefficient.

community is far from random. In the week following the earthquake, Japanese speakers and users globally were more likely to have remained in the community they were in before the earthquake than would be expected if community structure were randomly generated. Additionally, users were much less likely to join new communities from a solitary pre earthquake state or shift to other communities from their current community. Additionally, we find that while non-Japanese speakers did not significantly change the content of their conversations relative to the pre-earthquake topics, nearly all Japanese speakers changed their main conversation topic to earthquake-related content. This study builds a systematic framework for investigating human behaviors and social interactions during extreme events with online social network data. The findings on network dynamics and community evolution may provide useful insight for our understanding of collective human behavior in future studies.

Methods

Network data. Global. Using Twitter's 1% sampling API ("spritzer") collected from the same sources as Mandel et al²⁴ and Culotta²⁵ we sample approximately uniformly from all messages from February 28th to March 7th, before the Tōhoku earthquake on March 11th 2011, and from March 14th to March 21st, after the earthquake. On Wednesday, March 2nd there was a server error and no data was collected. In order to maintain day of week symmetry in the before and after dataset, we also exclude Wednesday, March 16th. This dataset results in 14.2 million messages and is hereafter referred to as Global. Unfortunately, the data was recorded in ASCII, which is sufficient for content analysis when users write in the Latin alphabet, but does not permit content analysis or language estimation when users write in language that does not use the Latin alphabet, especially Japanese. For this reason, the following dataset is also used for a more detailed analysis of content changes in Japanese language tweets.

Topsy. We used the API interface to the website www.topsy.com to search for the largest possible subset of tweets in English and Japanese for the seven days before and after the earthquake. We searched on the 6 most commonly used hiragana characters: うんいしょん and じこ to retrieve Japanese tweets (TP-JP) and the six most common words in English language to retrieve English tweets (TP-EN): "the", "be", "to", "of", "and", and "a".

For the Japanese language search, this method returns between 100 and 400 tweets per hour before the earthquake, and 300 to 1500 tweets per hour after the earthquake. For the English language search, between 300 and 600 tweets per hour were returned for the entire study period.

Inferring the geographic origin of individual tweets is challenging²⁶. We rely on the geographic concentration of Japanese speakers in Japan to develop a sample of tweets written by people who are likely to be affected by the earthquake. Ninety-nine percent of all people for whom Japanese is a first language live in Japan^{27,28}, and so searching just by the Japanese language will preferentially select people who were affected by the earthquake, especially in comparison to people speaking English on Twitter, as in our TP-EN control group. Those who use Japanese to communicate on Twitter are likely to be socially connected to those in Japan, and thus influenced by the earthquake either directly or indirectly through social ties to people in Japan. We consider the TP-JP dataset to represent a treatment group, because a larger proportion of this groups population was affected by the earthquake than the populations of the other two datasets. Thus, changes in the structure of the TP-JP network before and after the earthquake that did not also occur in the TP-EN or Global dataset may be related to the events during the study period that influenced Japanese speakers more significantly than people who speak other languages, especially the earthquake and tsunami.

Network formation. Each node in the network represents a single Twitter user who sent or received tweets in the sample dataset. Each user is uniquely identified by a Twitter handle, which must contain only Latin letters and underscores. The first

character of all Twitter usernames is the @ symbol. A directed link from A to B is formed when user A writes user B's username in the posted content by either mentioning B directly or by forwarding a tweet written by B, called a retweet. The number of times A mentions B in all tweets from the study period is used as the link weight.

In the end, 36,893, 36,234 and 7,403,050 unique users were obtained from the TP-JP, TP-EN and Global datasets, respectively. This is shown in Table 1. Data is sampled from posted tweets rather than from a full list of users and so it is important to follow the behavior of each user individually to evaluate the effect of the earthquake. We thereby focus on users who were included in the sample both before and after the earthquake. After this filtering, we obtained 3,647, 4,037, and 87,939 nodes, respectively. Due to the limited sampling rate in TP data, the resulting networks are relatively sparse. For the purpose of network and community analysis, we dismiss link direction and weight in the TP-JP and JP-EN networks, collapsing the data to an undirected unweighted network. We maintain the weighted and directed network for the Global dataset. In the Global dataset, outdegree refers to the number of tweets sent by a given user that refer to another user by name.

Community detection. We apply the Infomap method^{22,23} to detect the underlying community structures in the TP-JP, TP-EN networks (undirected unweighted) and the Global networks (directed weighted). The Infomap method is built based on optimizing the minimum description length of the random walk on the network. It takes advantage of the duality between finding community structure and minimizing description length of a random walker's movements on a network. With a random walker as a proxy for real flow, the minimization over all possible network partitions reveals important aspects of network structure with respect to the dynamics on the network.

Infomap has the advantage of being flexible for finding community structures on both undirected and directed, weighted and unweighted networks, it has also shown that it is one of the most efficient, reliable and accurate community detecting method in comparison to a range of other models^{29–31}. For these reasons, it has been becoming the state of the art and has been increasingly used in the detection of network communities in empirical studies^{30,32,33}.

Categorize network dynamics. **Community dynamics.** We propose a framework for modeling the dynamics of communities using the following five processes. In the time between t and $t + 1$, a community can dissolve, form, survive, split or merge. Let s_i^t be the size of community i at time t , and s_{ij}^{t+1} be the number of nodes, or users, who shift from community i to community j between time t and $t + 1$. Then:

Community i survives if a sufficiently large proportion of its nodes remain in the same community at time $t + 1$. Formally, any community j ($1 \leq j \leq m^{t+1}$) has survived from time t to $t + 1$ if $\min(s_{ij}^{t+1} / s_i^t, s_{ij}^{t+1} / s_j^{t+1}) \geq k_1$, where m^{t+1} is the total number of communities detected at time $t + 1$, and k_1 is the constant threshold of the proportion of nodes that must remain active in community j for that community to be defined as having survived from time t to $t + 1$. In this case, we define all destination communities as having also survived.

Community j dissolves between time t and $t + 1$ if only a small proportion of its nodes remain in the same community at time $t + 1$. Formally, for any community j at time $t + 1$, $\max(s_{ij}^{t+1} / s_i^t, s_{ij}^{t+1} / s_j^{t+1}) < k_2$, where k_2 is the constant for setting the threshold of the minimum proportion of nodes that must remain in community j for that community to remain active and avoid dissolving. Thus, k_2 must be less than k_1 .

Finally, if a community at time t does not survive or dissolve in time $t + 1$, it has split, and if a community in time $t + 1$ does not survive or form, it has merged. In our analysis, a community can dissolve, survive or split in time t , and can form, survive, or merge in time $t + 1$. To avoid including many communities with very few nodes, we limit our analysis to communities with at least s nodes.

Node dynamics. Similarly to community dynamics, the dynamics of individual nodes can be modeled according to their status in a community. There are five possible behaviors for each node between time t and time $t + 1$. A node may stay-solitary if it does not belong to any community in time t or time $t + 1$; a solitary node at time t may join a community at time $t + 1$; a node that belongs to a community at time t may stay-social, by remaining in the same community at time $t + 1$, jump to another

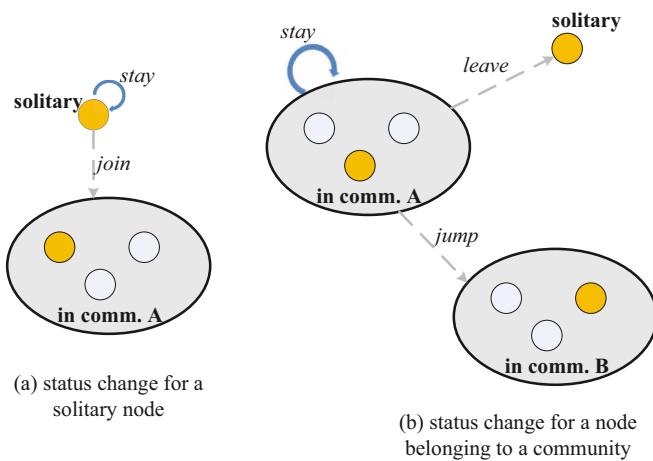


Figure 1 | Illustration on the dynamics of nodes.

community, or leave the community and become solitary. These behaviors are shown graphically in Fig. 1.

We investigate the distributions of the five forms of community evolution and five forms of node behavior by varying the minimum size s for a community to be analyzed, with $k_1 = 50\%$ and $k_2 = 10\%$.

Content analysis. In order to estimate the most important and relevant conversational topics for the largest communities in each time period, tweets were stripped of all punctuation except Twitter relevant indicators: #, @ and _, and the

frequency of all remaining words was counted. For non-Japanese tweets (Global and TP-EN), three categories of words are filtered out: common English words, common Twitter phrases, and common words in non-English languages. First, the 50 most common English language words are excluded. Second, the 25 most common words in the TP-EN dataset are excluded, most of which are fragments of web addresses and other meaningless words. Finally, the 50 most common words in the Global dataset are excluded. In the Global dataset, because only ASCII characters were retained, most of the content for users writing in a language that does not use the Latin alphabet was lost, particularly Korean and Japanese content. Because Twitter usernames can only contain the 26 letters of the Latin alphabet and the underscore character, the network structure is still observable for all language groups. One consequence of this data collection challenge is that the most frequently used words remaining in the dataset of communities using a non-Latin alphabet are usernames.

Written Japanese does not have spaces between words. As a result, tokenization to break sentences into words is challenging. We used MeCab³⁴, which is a morphological analysis system for Japanese language, to parse the Japanese tweets³⁵. After tokenization, the 50 most common words in the TP-JP dataset were removed. This still left a large number of word fragments and meaningless words as some of the most frequently used words in each cluster. With assistance from a native Japanese speaker, the top twenty most frequently used meaningful words were retained for each cluster in time one, before the earthquake, and time two, after.

Results

Basic network characteristics. Characterizing the Japanese network before and after the earthquake shows clear increases in Twitter use, directed communication, and the global clustering coefficient that is not observed to the same extent in the TP-EN or Global dataset (See Table 1 and Fig. 2. In all metrics, changes in the TP-EN dataset closely mirror the changes in the Global dataset, which supports the validity of Topsy's sampling method. The number of tweets

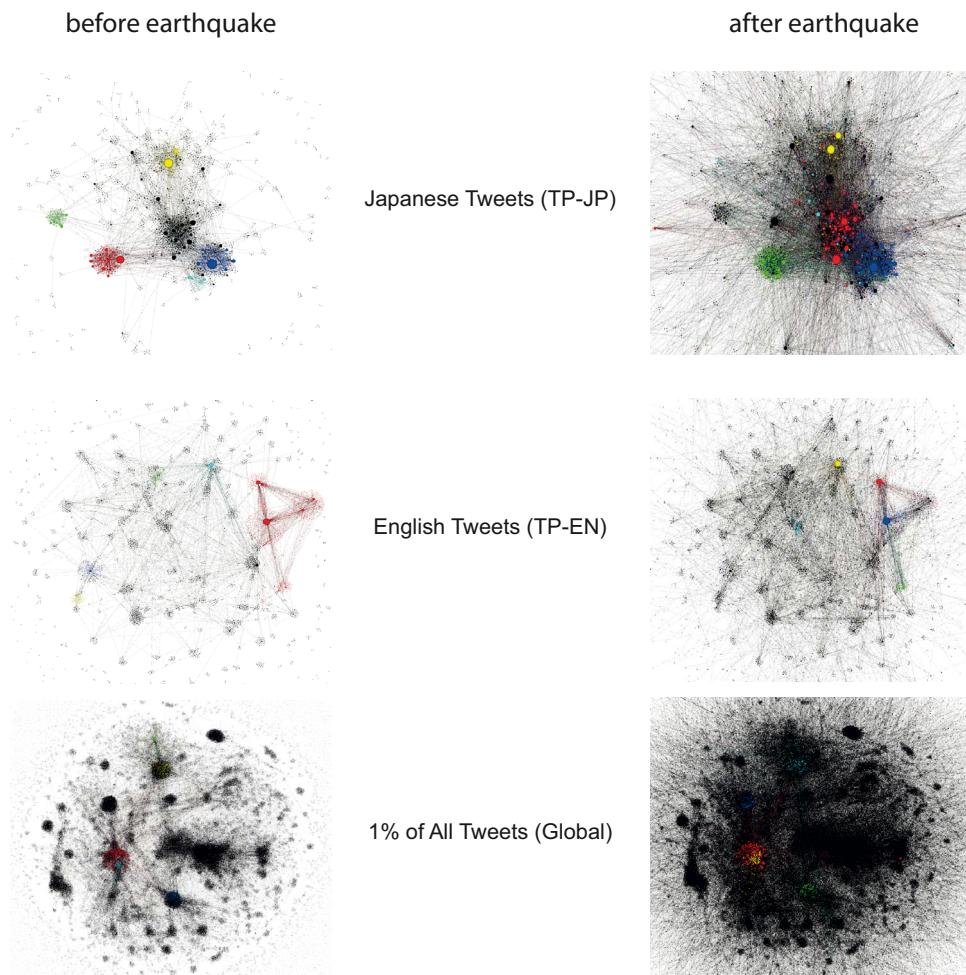


Figure 2 | Network visualization of Japanese and English tweets before and after the earthquake. Each node has the same x-y coordinates before and after. Node size is proportional to its outdegree. Nodes in the five largest communities detected by Infomap are colored, in order from the largest to the fifth largest: red, blue, green, yellow, and light blue.

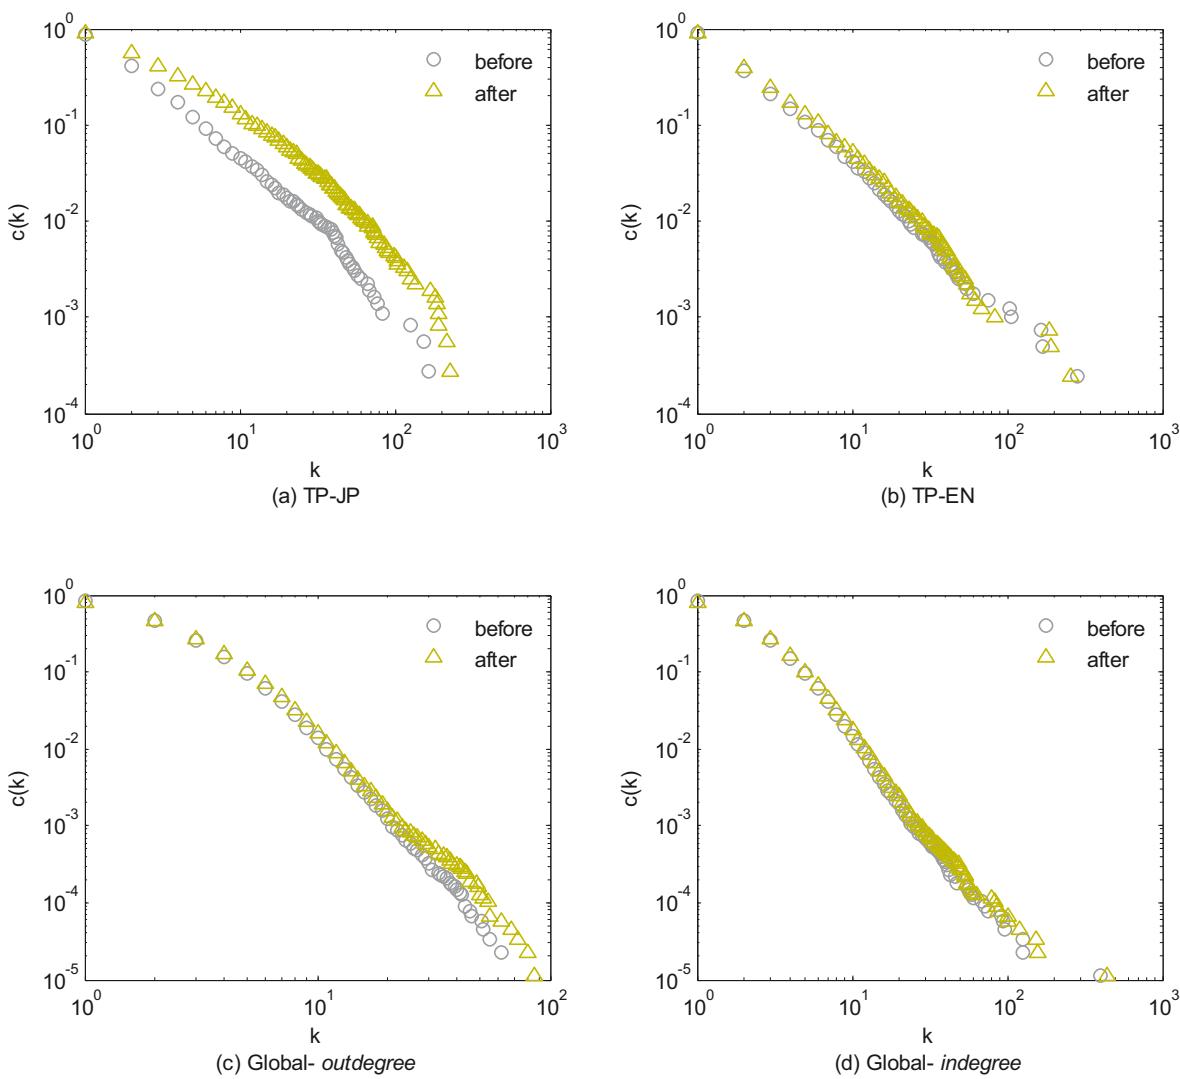


Figure 3 | Cumulative degree distribution for the extracted networks.

collected in the TP-JP dataset increased by 68% after the earthquake, while the number of tweets increased by 31% and 30%, respectively, for the TP-EN and Global datasets. While Twitter use increased for all networks, it increased to a greater degree for Japanese speakers than it did for English speakers and those in the Global dataset.

In addition to the simple rate of Twitter use, the rate of directed communication (the number of links) increased by 477% in the TP-JP dataset and only 32% and 35% for the TP-EN and Global datasets respectively. This shows that for non-Japanese speakers, the rate of directed tweets, those mentioning another person by name, compared to all tweets stayed consistent in both the TP-EN dataset and the Global dataset before and after the earthquake. For Japanese speakers, the rate of tweets that are directed at a particular user, out of all tweets, increased from 54% of all tweets to 98% of all tweets, indicating that the number of people they were in direct contact with expanded to a much greater extent than did their raw increase in posts.

For all three datasets, the increase in the number of tweets and number of links contributes to an increase in the size of the giant connected component. However, the clustering coefficient remained constant for the English tweets, and increased from 0.06 to 0.09 for Japanese data. Thus, after the earthquake, the Japanese social network was much more densely connected than it had been before the earthquake even though there was no measurable change in the density of the other two networks. This continues to suggest that

the observed changes in the Japanese social network were caused by the earthquake and related events, and implies changes in patterns of interacting behavior that are independent of just increased communication volume.

Dynamics of online activity. We start with the degree distribution, as shown in Fig. 3, to investigate the dynamics of users' online activity. We can see that on all networks the degree distributions are skewed, with a small percentage of nodes possessing a large number of links in the network. In both the Global and English datasets, the change in degree distribution is very small, while the Japanese network shows a significant increase in the degree at all levels of the distribution: the increase occurred across all types of nodes and included both those that were initially very connected, and those that were less connected. The average degree increased from 2.8 to 5.7 (see Table 1).

The changes in interaction patterns at the individual node level before and after the earthquake are presented in Fig. 4. We can see that there is no meaningful change for the Global indegree (number of links pointing to a node), outdegree (number of links starting from a node), and for the TP-EN network. However, the Japanese network shows significantly increase in degree after the earthquake. This is consistent with the aggregate behavior demonstrated in Fig. 3. No matter how active a Japanese user was in communicating directly with others on twitter before the earthquake, s/he is likely to have

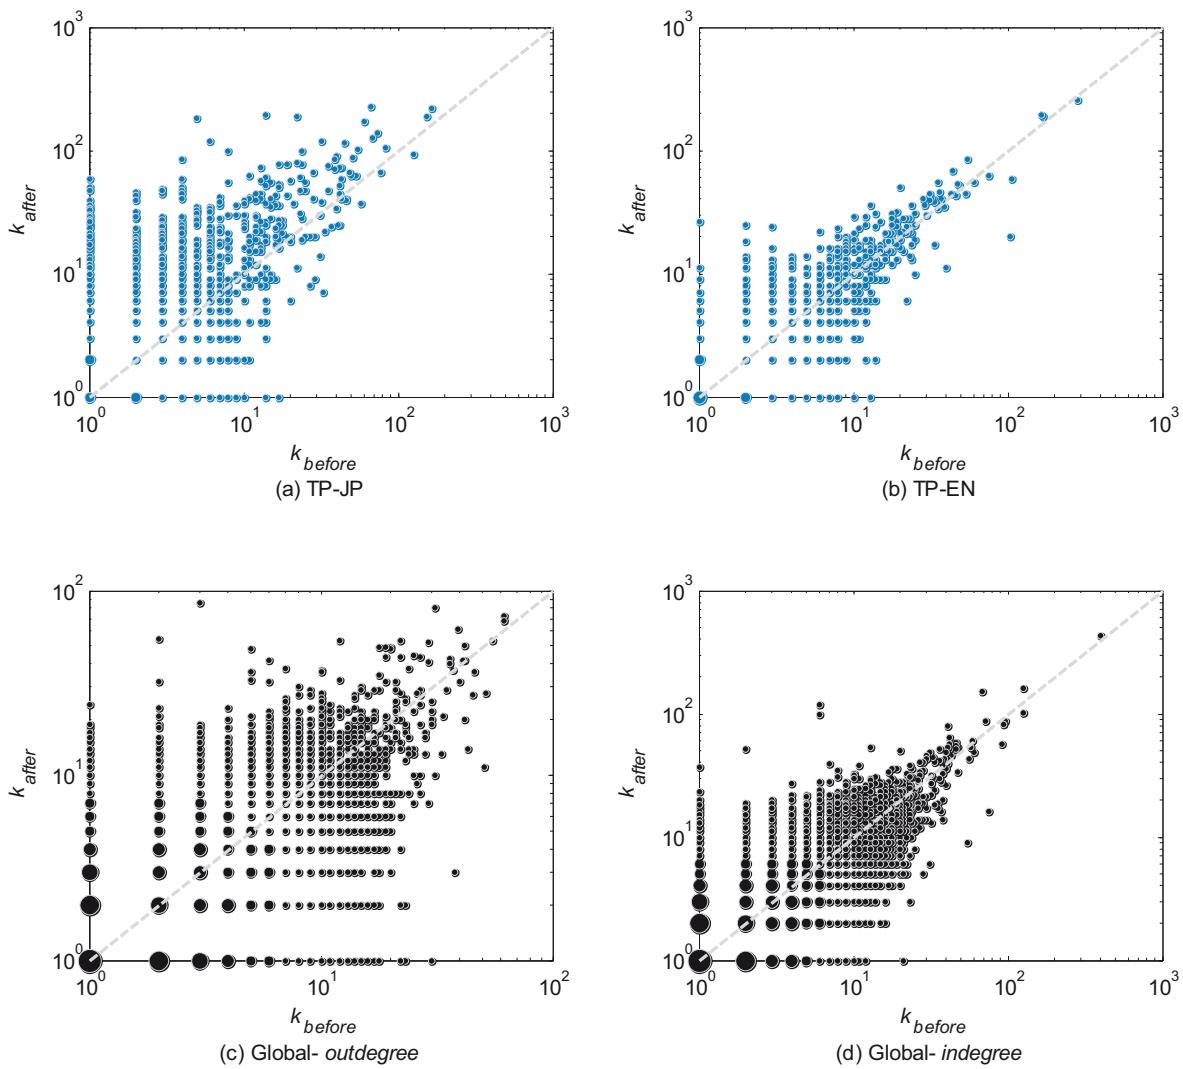


Figure 4 | Individual changes in degree before and after the earthquake. The size of each point indicates the log of the number of unique nodes with that pair of degrees in the before and after cases.

increased the number of people s/he communicated with after the earthquake.

Community evolution. Dynamics of communities. The above analysis on network density and node degree shows that individuals in the Japanese network became more active in posting tweets and interacting with others after the earthquake. To further investigate these changes in social behavior, in this section we evaluate the dynamics of communities: the social groups which are formed among intensely interactive users and are detected by the Infomap algorithm^{22,23}. To identify significant patterns, we generate a null hypothesis scenario in which the size distribution of communities is fixed both before and after the earthquake, but each user joins some detected community at random after the earthquake. Based on changes in each users community between the truly observed community in the before case and the null hypothesis scenario in the after case, the number of communities that survive, dissolve or form, and split or merge is then plotted in Fig. 5.

We can see that in the true data, for all networks, a number of communities survived after the earthquake, a phenomenon which cannot be observed under the random participation assumption. In communities that survived, users stayed in contact with a similar group of people, despite the strong effect of earthquake. The

Japanese network had fewer communities that survived than the English network did: In the Japanese network, only three communities survived after the earthquake under a community size threshold, s , of 10, under a community size threshold of 60, only the largest community survived. The relatively small number of sustained communities in the Japanese dataset indicates that the earthquake imposed a strong force in disrupting online user activities. It is possible that after the earthquake, users were likely to change their conversations towards topics that were socially universal at that time, such as news transfer, disaster response and relief work. Additionally, the expansion of personal networks due to highly increased volume of mentioning and retweeting other users could also lead to the failure of members to remain in their own isolated communities.

For all three networks, the number of communities that survived is larger than the null model, and the number of communities that dissolved or formed is substantially smaller than the null model. This shows that even though interactions between users were significantly increased under these extreme events, most users remained in their original communities, and did not shift to other communities, which demonstrates the coherence of online social groups.

Although the direction of difference between data and the null model is the same for the Japanese, English, and Global datasets, the magnitude of difference is larger for the English and Global datasets than for the Japanese dataset. This may show that while

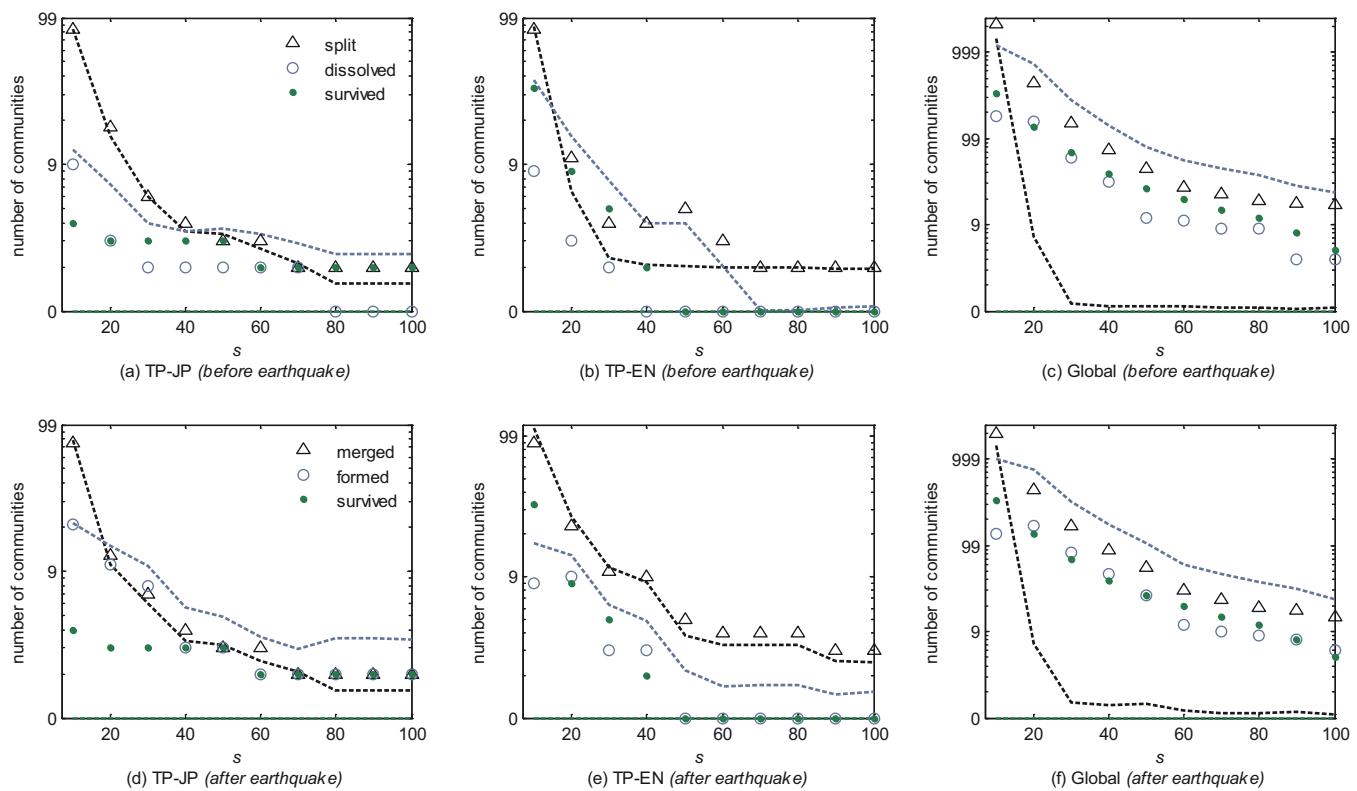


Figure 5 | Community dynamics. (a)–(c) show the number of survived, dissolved and split communities before the earthquake and (d)–(f) show the number of survived, formed and merged communities after the earthquake. Dashed lines represent corresponding communities from the null model.

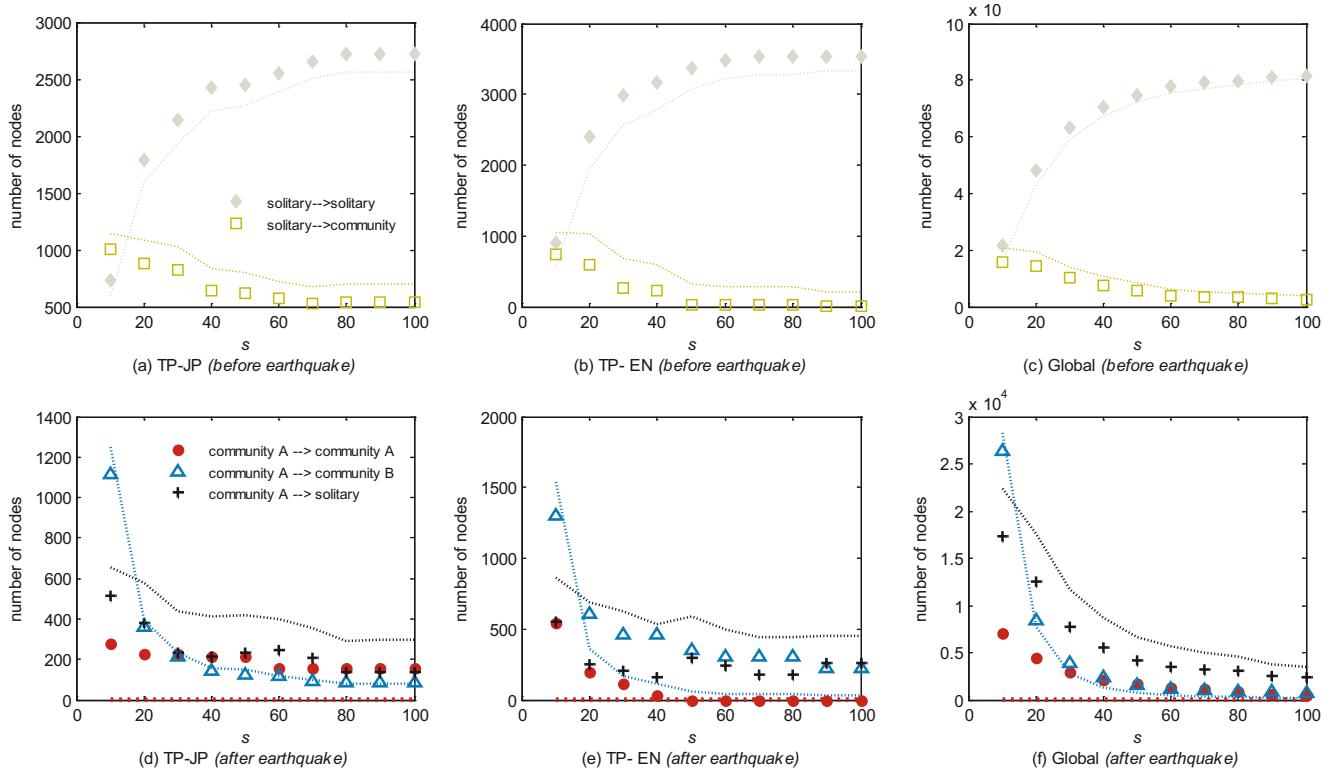


Figure 6 | Dynamics of individuals within communities. (a)–(c) show status changes of solitary users who either stayed solitary or joined a community after the earthquake, and (d)–(f) show status changes of users who were in a community before the earthquake. Dashed lines represent corresponding communities from the null model.

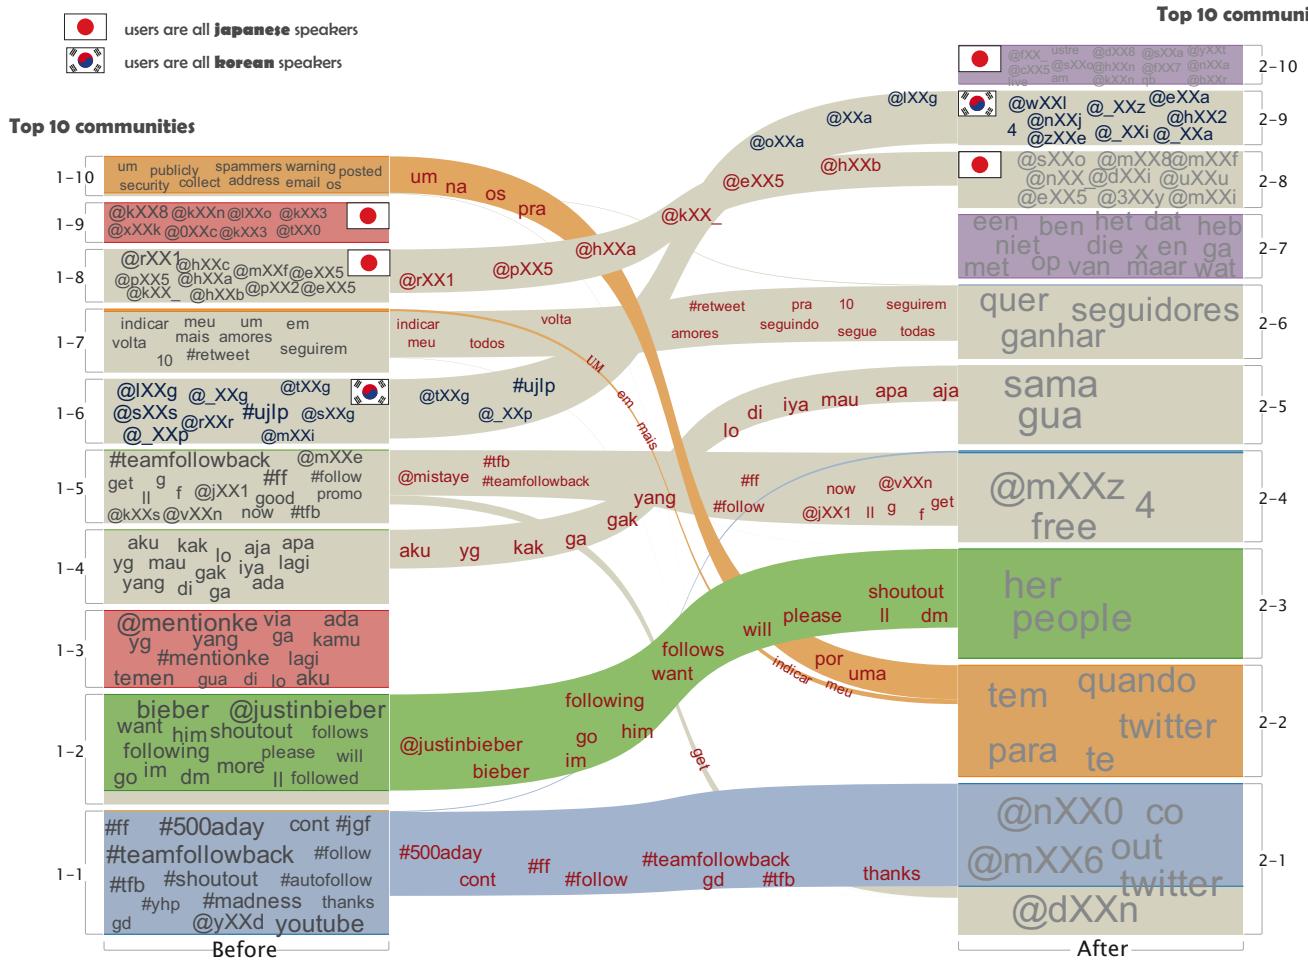


Figure 7 | Alluvial diagram for mapping changes in the Global network. The top ten communities detected by Infomap are ordered and shown as blocks. For each community, the most frequent words are shown with a word cloud in which each word's font size is related to the frequency of appearance of that word. The height of a streamline between a community identified in the before dataset and a community identified in the after dataset is proportional to the number of users who were active in both communities. Words shown in the streamline are common to those two communities. Contents of surviving communities are colored dark blue.

online social groups do cohere during adverse events, users are more likely to change social groups during an extreme event than they are when the social conditions are stable.

Dynamics of individuals. Community evolution is a consequence of the behavioral choices of many individuals. In Fig. 6, we show the number of users whose status in a community changed. Users who were solitary prior to the earthquake were more likely than expected from the null model to remain in a solitary state after the earthquake, and users who belonged to a community before the earthquake were more likely to remain in that community after the earthquake.

This shows common patterns of community and individual dynamics under either disastrous events or stable social conditions. At the community level, groups of users tend to form stable social groups. The higher number of split and merged communities also implies that users in communities behave as sub-groups and may shift between broad communities in smaller clusters. At the individual level, our analysis reveals that users prefer to keep their social status unchanged: they either stay solitary, or remain in the same community.

Content analysis. In Fig. 7, 8 and 9 we show the user movements among the top 10 communities detected before and after the earthquake in each dataset with an alluvial diagram. Communities are ordered by size, and the content of each community is shown

with a word cloud in which the size of font approximates the frequency of appearance.

Although the sizes of communities changed between the before and after cases, Fig. 7 shows that membership is quite stable for the largest clusters in the Global dataset. Communities that retain a similar population of users in the two time-steps also have consistent topics of conversation. Further, communities with similar topics of conversation exchanged measurable numbers of users. For example, of the four communities that used the hashtag #teamfollowback with high frequency, 1-1, 1-5, 2-1 and 2-4, streamlines are visible from each of the communities in time one to each of the communities in time two. The streamline between 1-1 and 2-1 represents more than half of the population of both communities, as does the streamline between community 1-5 and 2-4. The streamline from 1-1 to 2-4 is small but observable, as is the streamline between 1-5 and 2-1.

Additionally, the likely conversational language can be identified for each community. In cases of geographically isolated languages, this could be used to facilitate a broad geographic identification procedure. In Fig. 7, communities 1-3, 1-4 and 2-5 appear to be speaking Indonesian or Malaysian. Communities 1-6 and 1-9 are speaking Korean, while communities 1-8, 1-9, 2-8 and 2-10 are speaking Japanese. Finally, communities 1-7, 1-10, 2-2, and 2-6 are Portuguese speaking. Fig. 2 colors the five largest communities in each time-step, and those colors match up with communities which



Top 10 communities

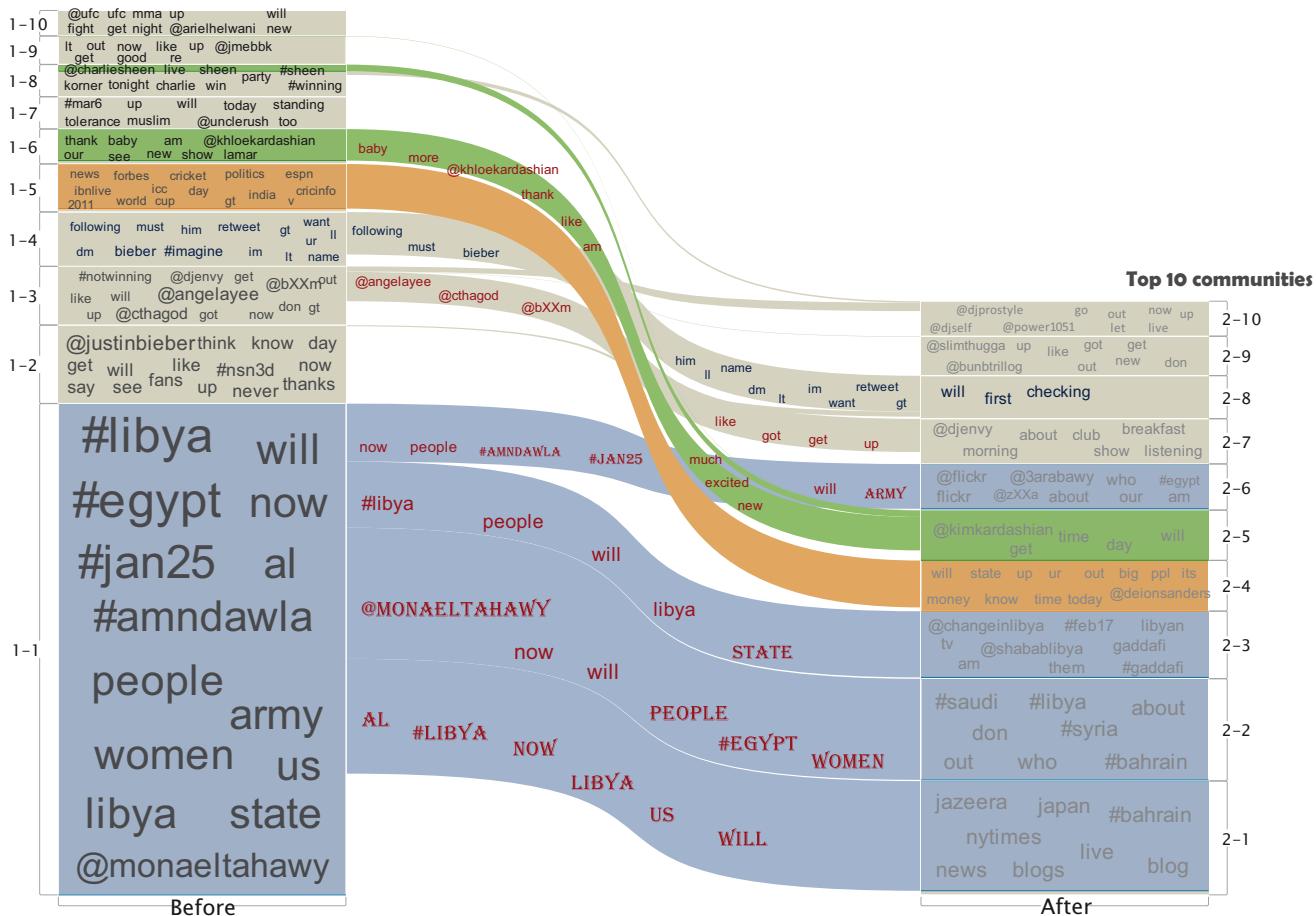


Figure 8 | Alluvial diagram for mapping change in the English Topsy network. The top ten communities detected by Infomap are ordered and shown as blocks. Contents of surviving communities are colored dark blue.

have significantly overlapping populations in the before and after time periods, seen in Fig. 7.

Fig. 8 shows an alluvial diagram overlaid with the most frequent words for the TP-EN dataset. Like the Global Infomap, there is overlap in content for communities that have overlapping populations between time one and time. This occurs in all cases except communities 1-5 and 2-4, which share nearly all of their users, but none of the most frequently used words. The most dramatic change in this figure is the breakup of community 1-1 into four separate communities. In the English row of Fig. 2, the largest cluster in time one is colored red, and we can see that in time two, it has split to make up three of the five largest clusters which are colored red, green and blue. The broad topic of conversation for cluster 1-1, and also clusters 2-1, 2-2, 2-3, and 2-6, which are primarily composed of users who were a member of cluster 1-1 is clearly the events of the Arab spring. However, the four communities in time two have unique conversational foci. Cluster 2-1 seems to take a current events perspective. Two major international news sources, Al Jazeera and the New York Times are mentioned, and the Japanese Earthquake is also referenced. Cluster 2-2 refers to many of the countries with significant political turmoil during the Arab spring, but does not reference any traditional media sources. Cluster 2-3 refers only to events in Libya, and the usernames that are some of the most frequent words are also Libya-based groups. Cluster 2-6, the last of the clusters, refers to bloggers, activists and users based in Egypt. This shows that the algorithm is able to pick up on changes in interaction patterns that occur as current events change.

The TP-JP network shows very different patterns in community stability when compared to the TP-EN and Global network (see Fig. 9). In datasets not effected by disaster, the largest clusters in time two were nearly completely represented by nodes from a single cluster in time one. Collections of nodes that joined from another cluster are very small, and there were only a few other source clusters, so few of the largest communities merged or split. In the Japanese alluvial diagram there is a great deal of mixing among the different communities. Users and small groups of users appear significantly more likely to jump from one community to another between the two time-steps, meaning that the communities are more likely to merge or split under disaster conditions than in normal social circumstances. This is consistent with our observations of community dynamics shown in Fig. 5.

This mixing between communities is also evident in the topics of conversation in each cluster through time. Before the earthquake, the topics of conversation in each community varied from the Japanese political scandal of the week, which was a potentially inappropriate campaign donation from a South Korean person to the political campaign of the Mr. Maehara, the Minister of Foreign Affairs- to the events of the Arab spring, and Twitter's typical self-reflective discussions around obtaining more followers. After the earthquake, the main topic of conversation in every single cluster centers around the earthquake. A notable difference between the before conversations and the after conversations is that the after conversations contain many more words that are typically related to specifics, which is evidence that new information is being transmitted. For example, the characters “市”(city) and “県”(prefecture), which are modifiers and

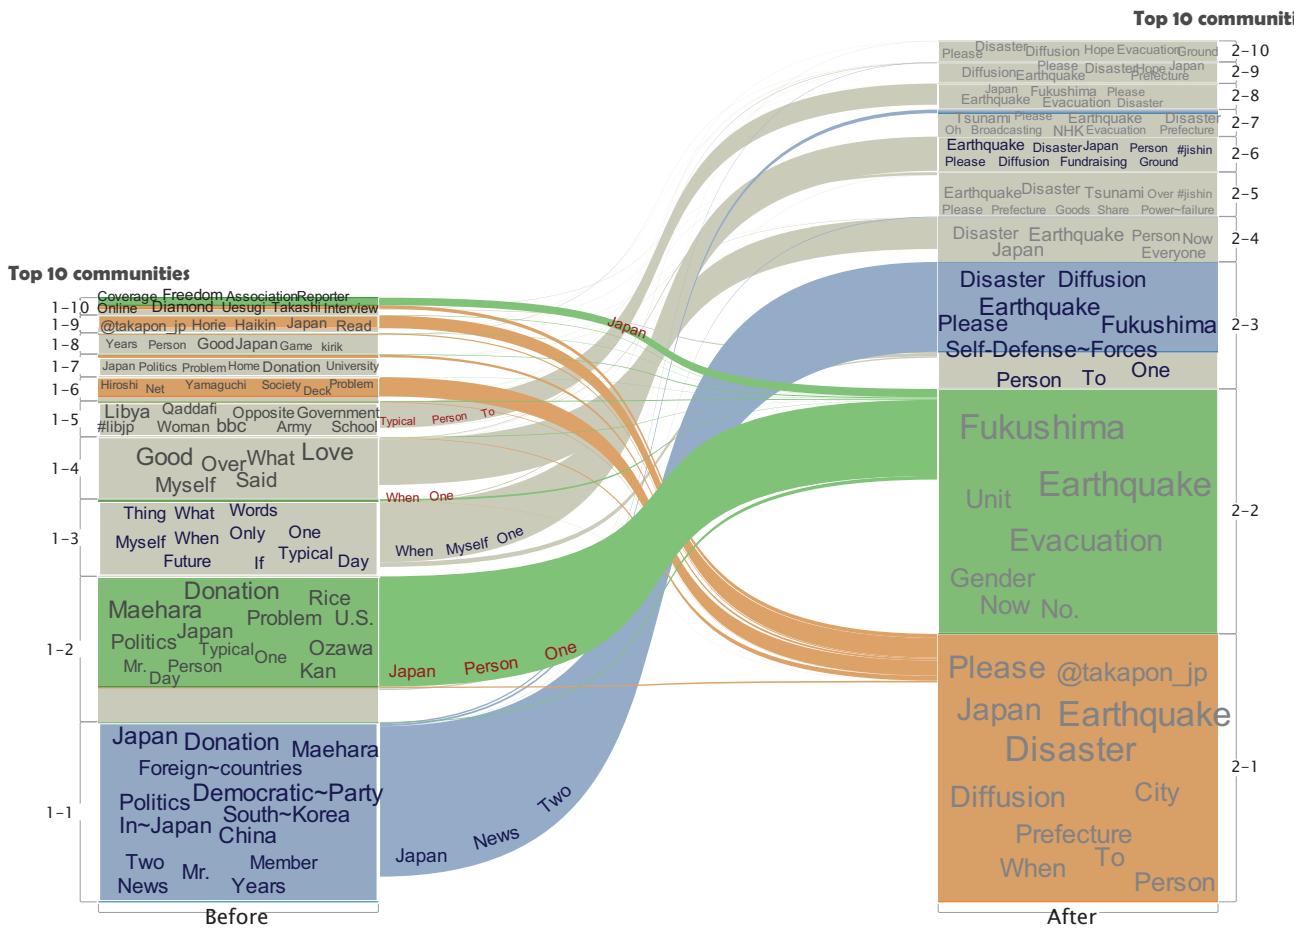


Figure 9 | Alluvial diagram for mapping change in Japanese Topsy network. Contents of surviving communities are colored dark blue.

would usually be attached to another character denoting a specific city or prefecture are much more common in the after case than before. Additionally, the higher rate of numerical words is also evidence of specific pieces of information being transmitted through the social network.

In the Global dataset, of the top 10 communities, only community 1–6 survived from time one to time two, becoming community 2–9 in time two. The members of this community are all speaking Korean. In the 10 largest communities for the TP-EN dataset, only community 1–4 survived into time two, become community 2–8. The members of this community were primarily discussing Justin Bieber. Community 1–2 was also discussing Justin Bieber, especially his “Never Say Never” film and album, both of which were released in February 2011, but there was little cross-over between members of the two different communities. In the TP-JP dataset, two out of the top ten communities survived from time period one to time period two. Community 1–1, initially discussing Japanese politics and current events was still the third largest community in time two, while community 1–3 survived to become community 2–6, both surviving communities switched to discussing the earthquake in the post earthquake period.

Conclusion and Discussion

Understanding human behavior during adverse events is critical for disaster preparation, warning, response, and recovery. In this study, we have built a systematic framework for the analysis of online communities in response to natural disasters. By investigating and comparing the structure of interacting networks, the evolution of online communities, and the content of communication before and after an adverse event, we have shown distinctive changes in patterns of

interactions in online communities that have been affected by a natural disaster compared to communities that were not affected. While almost all users in all datasets increased their online communications after the Japan earthquake, the Japanese speakers, more likely to be directly affected by the event than users in the other two datasets, expanded their personal network of communications to a higher degree. By applying the community detection algorithm of information mapping, we find that the behavior of joining or quitting a community is far from random: users tend to stay in their current social status and are less likely to join new communities from a solitary state or to shift to other communities from their original community. In close relation to the evolution of communities, the content of the information flowing within these networks has shown distinctive differences between Japanese speakers and other users: the topics conversation for non-Japanese speakers did not change significantly in the two time periods, while the diversity of topics for Japanese speakers disappeared after the earthquake, and Japanese speakers changed their main topic of conversation to earthquake-related content.

It is worth noting that even though we used different data sources to increase the validity and reliability of the study, neither the Global nor Topsy datasets are the complete datastream from Twitter, and Topsy.com does not publish their specific methods for subsetting and returning data. Because of the unknown sampling method, care should be used when extrapolating conclusions from the Topsy data to the broader Twitter community, or from Twitter use to social interactions that are not digitally mediated. Nevertheless, the consistent findings between the Topsy English tweets and the Global tweets suggest that our findings on the community of Japanese speaking Twitter users drawn from the Topsy dataset are represent-



ative of Japanese speakers on Twitter generally. Future researchers are encouraged to use a well-defined sampling strategy or integrate a larger fraction of tweets to increase the validity of their analyses.

1. Kenett, D. Y. & Portugali, J. Population movement under extreme events. *Proc Natl Acad Sci* **109**, 11472–11473 (2012).
2. Lu, X., Bengtsson, L. & Holme, P. Predictability of population displacement after the 2010 haiti earthquake. *Proc Natl Acad Sci* **109**, 11576–11581 (2012).
3. Norris, F. H. Disaster research methods: Past progress and future directions. *J Trauma Stress* **19**, 173–184 (2006).
4. Elliott, J. R. & Pais, J. Race, class, and hurricane katrina: Social differences in human responses to disaster. *Soc Sci Res* **35**, 295–321 (2006).
5. Souza, R., Bernatsky, S., Reyes, R. & Jong, K. Mental health status of vulnerable tsunami - affected communities: A survey in aceh province, indonesia. *J Trauma Stress* **20**, 263–269 (2007).
6. Bengtsson, L., Lu, X., Thorson, A., Garfield, R. & Von Schreeb, J. Improved response to disasters and outbreaks by tracking population movements with mobile phone network data: a post-earthquake geospatial study in haiti. *PLOS Med* **8**, e1001083 (2011).
7. Lu, X., Wetter, E., Bharti, N., Tatem, A. J. & Bengtsson, L. Approaching the limit of predictability in human mobility. *Sci. Rep.* **3**, 2932; DOI:10.1038/srep02923 (2013).
8. Williams, S. A., Terras, M. M. & Warwick, C. What do people study when they study twitter? classifying twitter related academic papers. *J Doc* **69**, 384–410 (2013).
9. Aramaki, E., Maskawa, S. & Morita, M. Twitter catches the flu: detecting influenza epidemics using twitter. In *Proc Conf on Empir Methods Nat Lang Process*, 1568–1576 (ACI, 2011).
10. Crooks, A., Croitoru, A., Stefanidis, A. & Radzikowski, J. # earthquake: Twitter as a distributed sensor system. *Trans GIS* **17**, 124–147 (2013).
11. Earle, P. *et al.* Omg earthquake! can twitter improve earthquake response? *Seismol Res Lett* **81**, 246–251 (2010).
12. Earle, P. S., Bowden, D. C. & Guy, M. Twitter earthquake detection: earthquake monitoring in a social world. *Ann Geophys* **54**, 708–715 (2012).
13. Bahir, E. & Peled, A. Identifying and tracking major events using geo-social networks. *Soc Sci Comput Rev* **31**, 458–470 (2013).
14. Cassa, C. A., Chunara, R., Mandl, K. & Brownstein, J. S. Twitter as a sentinel in emergency situations: lessons from the boston marathon explosions. *PLOS Curr* **5**, doi:10.1371/currents.dis.ad70cd1c8bc585e9470046cde334ee4b. (2013).
15. Sundaram, H., Lin, Y.-R., De Choudhury, M. & Kelliher, A. Understanding community dynamics in online social networks: a multidisciplinary review. *IEEE Trans Sig Process* **29**, 33–40 (2012).
16. Gupta, A., Joshi, A. & Kumaraguru, P. Identifying and characterizing user communities on twitter during crisis events. In *Proc 2012 workshop on Data-driven user behavioral modelling and mining from social media*, 23–26 (ACM, 2012).
17. Mendoza, M., Poblete, B. & Castillo, C. Twitter under crisis: can we trust what we rt? In *Proc 1st workshop on social media analytics*, 71–79 (ACM, 2010).
18. Chatfield, A. & Brajavidagda, U. Twitter tsunami early warning network: a social network analysis of twitter information flows. In *Proc 23rd Australas Conf on Inf Syst* (ACIS, 2012).
19. Sarcevic, A. *et al.* Beacons of hope in decentralized coordination: Learning from on-the-ground medical twitterers during the 2010 Haiti earthquake. In *Proceedings of the ACM 2012 conference on computer supported cooperative work*, 47–56 (ACM, 2012).
20. Palen, L., Vieweg, S., Liu, S. B. & Hughes, A. L. Crisis in a networked world features of computer-mediated communication in the April 16, 2007, Virginia Tech event. *Social Science Computer Review* **27**, 467–480 (2009).
21. Starbird, K. & Palen, L. (How) will the revolution be retweeted? Information diffusion and the 2011 Egyptian uprising. In *Proceedings of the ACM 2012 conference on computer supported cooperative work*, 7–16 (ACM, 2012).
22. Rosvall, M. & Bergstrom, C. T. Mapping change in large networks. *PLOS ONE* **5**, e8694 (2010).
23. Rosvall, M. & Bergstrom, C. T. Maps of random walks on complex networks reveal community structure. *Proc Natl Acad Sci* **105**, 1118–1123 (2008).
24. Mandel, B. *et al.* A demographic analysis of online sentiment during hurricane irene. In *Proc 2nd Workshop on Lang Soc Media*, 27–36 (ACL, 2012).
25. Culotta, A. Towards detecting influenza epidemics by analyzing twitter messages. In *Proc 1st workshop on social media analytics*, 115–122 (ACM, 2010).
26. Friedhorsky, R., Culotta, A. & Del Valle, S. Y. Inferring the origin locations of tweets with quantitative confidence. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW ’14, 1523–1536 (ACM, New York, NY, USA, 2014). http://doi.acm.org/10.1145/2531602.2531607.
27. Miyagawa, S. The japanese language. (1999) Available at: <http://aboutworldlanguages.com/japanese>. [Date of access: 24/10/14].
28. Thompson, I. Japanese. (2013). Available at: <http://aboutworldlanguages.com/japanese>. [Date of access: 20/02/14].
29. Lancichinetti, A. & Fortunato, S. Community detection algorithms: a comparative analysis. *Phys Rev E* **80**, 056117 (2009).
30. Fortunato, S. Community detection in graphs. *Phys Rep* **486**, 75–174 (2010).
31. Aldecoa, R. & Marin, I. Exploring the limits of community detection strategies in complex networks. *Sci. Rep.* **3**, 2216; DOI:10.1038/srep02216 (2013).
32. Ahn, Y.-Y., Bagrow, J. P. & Lehmann, S. Link communities reveal multiscale complexity in networks. *Nature* **466**, 761–764 (2010).
33. Opsahl, T. & Panzarasa, P. Clustering in weighted networks. *Soc Networks* **31**, 155–163 (2009).
34. Kudo, T. Mecab: Yet another part-of-speech and morphological analyzer. available: <http://mecab.sourceforge.net/> (2005).
35. Kyoto University (2013) Available at: <https://code.google.com/p/mecab/downloads/list>. [Date of access: 24/10/14].

Acknowledgments

The authors would like to thank Joshua Pestka and Yumiko Pestka for Japanese translation assistance. The authors would also like to thank Simon DeDeo and Ryan James for helpful discussions. This work was partially supported by the National Natural Science Foundation of China (no. 71301165, 91024030, 71101013, 71331008), the Bill and Melinda Gates Foundation (OPP1076282), the John D. and Catherine T. MacArthur Foundation (no. 13-105749-000-MCI) and by the Army Research Office Minerva Program (no. W911NF1210097).

Author contributions

X.L. and C.B. designed research; X.L. and C.B. performed research; X.L. and C.B. analyzed data; X.L. and C.B. wrote the paper.

Additional information

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Lu, X. & Brelsford, C. Network Structure and Community Evolution on Twitter: Human Behavior Change in Response to the 2011 Japanese Earthquake and Tsunami. *Sci. Rep.* **4**, 6773; DOI:10.1038/srep06773 (2014).



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder in order to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>