# On Understanding the Divergence of Online Social Group Discussion

**Hemant Purohit**[⋆1], **Yiye Ruan**[⋆2], **David Fuhry**[2], **Srinivasan Parthasarathy**[2], **Amit Sheth**[1]
[⋆]Joint first authors
[1]Kno.e.sis Center, Wright State University, USA
[2]Department of Computer Science and Engineering, The Ohio State University, USA
[1]{*hemant,amit*}@*knoesis.org*
[2]{*ruan,fuhry,srini*}@*cse.ohio-state.edu*

## Abstract

We study online social group dynamics based on how group members diverge in their online discussions. Previous studies mostly focused on the link structure to characterize social group dynamics, whereas the group behavior of content generation in discussions is not well understood. Particularly, we use *Jensen-Shannon (JS) divergence* to measure the divergence of topics in user-generated contents, and how it progresses over time. We study Twitter messages (tweets) in multiple real-world events (natural disasters and social activism) with different times and demographics. We also model structural and user features with guidance from two socio-psychological theories, social cohesion and social identity, to learn their implications on group discussion divergence. Those features show significant correlation with group discussion divergence. By leveraging them we are able to construct a classifier to predict the future increase or decrease in group discussion divergence, which achieves an area under the curve (AUC) of 0.84 and an F-1 score (harmonic mean of precision and recall) of 0.8. Our approach allows to systematically study collective diverging group behavior independent of group formation design. It can help to prioritize whom to engage with in communities for specific topics of needs during disaster response coordination, and for specific concerns and advocacy in the brand management.

## Introduction

Online social networks allow Internet users all over the globe to share information, exchange thoughts, and work collaboratively. All of those activities involve more than a single user, consequently, making the dynamics of *online social groups* worthy of study. The prevalence of online social networks in the last decade has enabled computational social scientists to answer various questions of group dynamics, such as group formation, evolution and engagement (Backstrom et al. 2006; Ren, Kraut, and Kiesler 2007; Shi et al. 2009; Farzan et al. 2011; Kairam, Wang, and Leskovec 2012; Grabowicz et al. 2013). Most studies, however, investigate implications of the network structure alone in characterizing group dynamics, and they lack the insights of dynamics based on user-generated content in online social groups. In this paper, we take a new perspective on characterizing group dynamics based on divergence of group discussion topics.

Social scientists have defined the groups based on various common user characteristics. We define a *group* here as the set of users interacting in discussions about a real-world event. We refer to *group discussion divergence* as collectively diverging behavior in user-generated discussion topics, and it is quantitatively measured as the Jensen-Shannon divergence among latent topic distributions of a group's messages (tweets). Understanding of such collective behavior in discussions around events can lead to actions of prioritizing for engagement, such as whom to engage with in communities for specific needs during disaster response coordination, and for specific concerns and advocacy in brand management.

In this paper, we focus on Twitter users' discussions related with two types of real-world events: natural disaster and social activism. Particularly, we ask the following questions:

- How does the divergence of user discussion in a group change over time, within and across different phases of events?
- Do two existing theories of social group behavior, namely, social cohesion and social identity, have implications on the evolution of group discussions?
- Can we predict the change of group discussion divergence in the future by utilizing features guided by social cohesion and social identity theories?

Answers to the above questions can aid in understanding which factors contribute more in facilitating cohesion (lower divergence) in social group discussions. They also enable us to predict the change of group discussion divergence, which in turn allows fast identification of groups whose voices are showing less divergent shifts. Such techniques may be highly valuable in scenarios like natural disaster response, where a small number of less diverging, focused groups (with resource requests or information supplies) need to be identified efficiently, so that their input will not be buried under an overwhelming amount of noise in the social content stream. Moreover, understanding of these factors will help us decipher behavior of self organizing online social groups.

**Main Results and Contribution:** In this study, using Twitter as our experimental platform, we propose a systematic approach to analyze discussions in online social groups, and understand the pattern of how discussion divergence changes over time. We discover that the divergence of topics in user-generated content starts with a low value prior to the event, peaks during the event, and fades away after the event.

We also formally define structural and user features guided by social cohesion and social identity, the two socio-psychological theories on group dynamics to be discussed in the next section. We represent a group's structural features related to social cohesion using network characteristics of its friendship-follower network. User features related to social identity are modeled via self-presentation in user profiles, capturing group members' physical world identities, as well as their online identities. These features incorporate guidance from the two theoretical approaches, while capturing users' social behavior from both physical and online worlds.

Furthermore, we study the relation between group discussion divergence and proposed features via correlation analysis. We observe that a group's network density, average length of pair-wise shortest path, and entropy values of user identities are well-correlated with its discussion divergence.

Finally, we build machine learning models to predict the future increase or decrease of social groups' discussion divergence values by using features discussed above. Our classifiers are able to achieve an AUC of 0.84 and an F-1 score of 0.8, reflecting a 33% improvement from the baseline method. As discussed earlier, this work can help in various application domains, including identification of emergent concerns during disasters, and the self-organizing group behavior of discussions.

## Related Work

First, we briefly introduce two theories proposed by socio-psychologists to explain the dynamics of traditional face-to-face social groups and their behaviors. We envision that their roles in shaping user engagement in groups (Dyaram and Kamalanabhan 2005) will contribute to our understanding of group discussion divergence. Then we describe related work on online social group bonding and dynamics.

**Socio-Psychological Theories**. The social identity theory includes two closely related parts: *social identity* (Tajfel et al. 1971) and *self-categorization* (Turner et al. 1987). In (Tajfel et al. 1971), Tajfel defines the concept of social identity as "the individual's knowledge that he belongs to certain social groups together with some emotional and value significance to him of this group membership". Therefore, group membership is the result of **"shared self-identification"** rather than "cohesive interpersonal relationship", and such shared identity leads to cohesiveness and uniformity, among other features (Turner 1982). One commonly-cited piece of evidence for the social identity theory is team sports (Branscombe and Wann 1991), where teammates are representing the same organization (a school, a club, or a country) and they are well aware of desire to sustain the reputation of their associated identity. In contrast, the social cohesion theory views social groups from a different perspective. Its hypothesis is that the necessary and sufficient condition for individuals to work as a group is the **cohesive social relationships between individuals**. We adopt the definition by Lott and Lott (Lott and Lott 1965) that interprets cohesiveness as *mutual attraction* between individuals, which is slightly different from that used in (Festinger, Schachter, and Back 1950). In accordance with this definition, the positive correlation between group cohesion and performance has been reported in various types of groups (Mullen and Copper 1994; Beal et al. 2003). A social cohesion example will attribute the inter-personal friendship between teammates of a sports club as the reasoning factor for group performance and its evolution.

**User-Group Bonding**. One study relevant to our work is by Grabowicz et al. (Grabowicz et al. 2013), where authors discussed methods to translate the common identity and common bond theories for group attachment into general metrics applicable to large social graphs. They also devised a method to predict whether a group is social (formation dependent on interpersonal bonds) or topical (formation based on role awareness). Prior to that, Ren et al. (Ren, Kraut, and Kiesler 2007) presented a study on the similar direction, focusing on the implications of the two theories of group attachment and link these theories with design decisions for online communities. Our differing objective here is to rather analyze a group's discussion having the characteristics of identity and cohesion features instead of predicting group type or evaluating community design decisions. In a similar spirit, Farzan et al. (Farzan et al. 2011) studied group commitment on Facebook within a controlled environment and observed that designs that encourage relationships among members or emphasize the community as an entity both increase the commitment and retention of players. Budak and Agrawal (Budak and Agrawal 2013) utilized data analytics and user survey to study factors that drive group chats on Twitter, and found that social inclusion contributes most to user retention. Our objective here is slightly different, in that it focuses on the effects of group commitment in discussion divergence in the communities emerging around real-world events.

**Group Dynamics**. Most prior work on group dynamics has focused on structural dynamics. Notably, Backstrom et al. (Backstrom et al. 2006) proposed a structure-centric model for network membership, growth and evolution by analyzing DBLP and LiveJournal social networks. Their findings show how individuals join communities and how communities grow depending on the underlying network structure, which supports cohesion-based structural features in our study. Taking a different path of a user-centric approach, Shi et al. (Shi et al. 2009) studied the user behavior of joining communities on online forums. Among other features, the authors studied the similarity between users and the similarity's relation with community overlap. They found that user similarity defined by the frequency of communication or number of common friends was inadequate to predict grouping behavior, but adding node/user-level features could improve the fit of the model. Kairam et al. (Kairam, Wang, and Leskovec 2012) analyzed long term dynamics of communities and modeled future community growth rate. They found that growth rate is correlated with current size and age of a group and the size of the largest clique is the best feature for community sustainability. Relevant efforts on un-

| Event Name | Type | Duration | #Tweets | #Users | Type |
|---|---|---|---|---|---|
| Hurricane Irene (Irene) | Disaster (D) | 08/24-09/19, 2011 | 183K | 77K | Transient |
| Hurricane Sandy (Sandy) | Disaster (D) | 10/27-11/07, 2012 | 4.9M | 1.8M | Transient |
| India Anti-Corruption (IAC) | Civil Protest (P) | 11/05-12/02, 2011 | 100K | 21K | Lasting |
| Occupy Wall Street (OWS) | Civil Protest (P) | 11/05-12/02, 2011 | 2.1M | 331K | Lasting |

Table 1: Twitter data statistics centered on diverse set of evolving events

| Event | Timeline |
|---|---|
| IAC | *During-phase* Beginning (11/24): Minister Sharad Pawar got slapped due to alleged corruption<br>*During-phase* End (11/29): No further substantial tweet w.r.t. the incident of slapping<br>Source: http://en.wikipedia.org/wiki/2011_Indian_anti-corruption_movement |
| OWS | *During-phase* Beginning (11/15): Raid of Zuccotti Park<br>*During-phase* End (11/23): President speech interrupted by protesters<br>Source: https://99.occupymediawiki.org/wiki/Timeline_of_Occupy_movement#November_2011 |
| Irene | *During-phase* Beginning (08/27): Landfall in North Carolina<br>*During-phase* End (08/30): Hurricane dissipated<br>Source: http://www.theguardian.com/world/blog/2011/aug/27/hurricane-irene-new-york-live |
| Sandy | *During-phase* Beginning (10/29): Landfall in New Jersey<br>*During-phase* End (10/31): Hurricane dissipated<br>Source: http://en.wikipedia.org/wiki/Effects_of_Hurricane_Sandy_in_New_York |

Table 2: Timeline and dates signifying the beginning and end of *during-event* phase of each event

derstanding individual-level characteristics include a study by Rao et al. (Rao et al. 2011), where authors presented an approach for automatic creation of ethnic profiling of users, focusing on names as the key factor. Pennacchiotti and Popescu (Pennacchiotti and Popescu 2011) also proposed a machine learning approach for user classification on Twitter by analyzing user's friends, user posts and profile information. These studies of group and individual characteristics provide a base for the modeling of user and structural features in our study.

# Problem Formulation

In this section we describe preliminaries including event-based discussion collection, social group identification, measure of group discussion divergence, and a formal specification of the prediction task. Feature design, experiment results and analyses are presented in subsequent sections.

## Data Collection

We focus on Twitter user-generated contents and discussions based on particular real-world events, and thus, proper filtering of the generic content stream is required.

We implemented a Twitter Streaming API-based crawler that collected an on-going tweet stream relevant to the event based on a seed keyword set, similar to (Ruan et al. 2012). For a keyword `k`, we crawl all tweets that mention `k`, `K`, `#k` or `#K`. The seed list of keywords and hashtags is kept up-to-date by first automatically extracting the top-N most frequent hashtags and keywords from the crawled tweets, and then manually selecting and adding highly unambiguous hashtags and keywords (e.g. `hurricane sandy`, `#sandy`, `#ows`). This process provides a control for contextual relevance of tweet content to the event. Tweets containing seed hashtags/keywords and their corresponding authors then become our dataset. We also store metadata associated with the dataset, such as each author's location, followers/friends, and profile description.

In this study, we choose four events (two for social activism and two for natural disasters), and collect relevant data using the mechanism described above. Table 1 summarizes basic information about each dataset. We note that events possess varying characteristics on the dimensions of activity, social significance, participant types, etc. In Table 1, we specifically show temporal feature values as 'Lasting' and 'Transient' that denotes how enduring an event is. For example, the Occupy Wall Street movement was highlighted in social media discussion for a long time frame, while Twitter users' attention to Hurricane Sandy quickly decreased significantly after it dissipated.

To enable temporal analysis and reasoning, tweets are grouped into three phases (*pre-*, *during-*, and *post-event*). Our categorization of phases for each event is aligned with its real-world timeline, and Table 2 shows the occurrences leading to phase division.

## Identifying Social Groups

Social groups can be defined in many ways. Our focus here lies on those groups of people who interact (and potentially emerge) in the times of evolving real-world events.

Therefore, given all users in a community formed around discussions of an event, it is necessary to identify appropriate social groups on which quantitative analyses will be performed to understand the dynamics of group discussion divergence. Resultant social groups should reflect online interaction among users that is beyond simply using the same word in their tweets. Moreover, the grouping criterion needs to be independent of any feature of social structure and user characteristics due to some of our features being based on social cohesion and identity phenomena (defined in the following sections), so that the results are not biased.

To that end, we propose an approach of clustering users based on their interactions, which can be either *retweet*, *reply* or *mention*. An interaction graph is created to represent those relationships during each phase of the event, where vertices stand for users and edges indicate at least one interaction between two users through the phase. We apply Markov clustering (Satuluri and Parthasarathy 2009), a commonly-used community detection algorithm to identify social groups. Only groups that have at least 10 members and are active (that is, at least one member posts a relevant tweet by mentioning event-related keyword(s)) for at least two days are retained. Again, while there exist other choices of identifying latent online user groups without ground truth labels, we believe our simple approach can effectively capture online interactions and yield meaningful groupings of users. Table 3 summarizes the information of each dataset's social groups.

| Event | # Groups | # Users | Average Group Size |
|---|---|---|---|
| Irene | 137 | 22,068 | 161 |
| Sandy | 4,947 | 284,062 | 57 |
| IAC | 76 | 7,907 | 104 |
| OWS | 6,202 | 296,279 | 48 |

Table 3: Information of social groups

## Defining Group Discussion Divergence

We use Jensen-Shannon divergence (JS-divergence) to quantify the divergence of group discussions. Compared with other information-theoretic measures such as Kullback-Leibler divergence, JS-divergence is always defined, bounded, and can be generalized to more than two distributions (Lin 1991). JS-divergence has long been employed in computational linguistics (Lin et al. 2006; Louis and Nenkova 2013), though its usage in social network analytics has been limited.

In order to calculate the JS-divergence, we first construct a dynamic topic model (Blei and Lafferty 2006) and infer the topics of discussion. Input into the topic model is a collection of vocabulary vectors, each of which represents one event-related tweet and is indexed by discrete time-stamps. The vocabulary includes words and phrases pertaining to the event, as well as hashtags with the leading '#' symbol stripped. The dynamic topic model has the advantage of modeling a systematic topic shift (due to event's progress) automatically, which allows us to investigate the true difference of an individual member's topic distribution to the corresponding group's topic distribution at any given time.

For topic inference, we use the `dtm` package[1] with default parameters. We evaluated results from 2 to 5 latent topics, and found that topics become similar and redundant after 3. For expository simplicity we use 3 as the default number of topics and report the top vocabulary in the different event phases for two events (Hurricane Sandy and Occupy Wall Street) in Table 4.

[1] https://code.google.com/p/princeton-statistical-learning/downloads/detail?name=dtm_release.tgz

| Hurricane Sandy | | | |
|---|---|---|---|
| | **Pre-event** | **During-event** | **Post-event** |
| Topic 1 | tropical storm | red cross | red cross |
| | east coast | jersey shore | staten island |
| | canada | caused | mexico |
| | path | staten island | caused |
| Topic 2 | new york | new york | new york |
| | state | new jersey | new jersey |
| | google | hurricane katrina | states |
| | android | media | hurricane katrina |
| Topic 3 | frankenstorm | frankenstorm | frankenstorm |
| | halloween | fema | knicks |
| | east coast | halloween | fema |
| | atlantic | mitt romney | nyc |
| **Occupy Wall Street** | | | |
| | **Pre-event** | **During-event** | **Post-event** |
| Topic 1 | occupy | occupy | occupy |
| | protest | n17 | oo |
| | movement | nypd | occupyla |
| | occupytogether | brooklyn bridge | movement |
| Topic 2 | movement | nypd | nypd |
| | us | movement | movement |
| | bahrain | protest | anonymous |
| | occupy movement | time | protest |
| Topic 3 | occupy | occupy | p2 |
| | oo | p2 | tcot |
| | p2 | tcot | republican |
| | tcot | oo | teaparty |

Table 4: Top vocabulary representing the latent topics of discussions at each event phase

The inference process of the topic model returns a latent topic distribution for each tweet $t$, denoted as $\beta_t$. A group $g$'s mean topic distribution at phase $s$ over all its users' tweets ($T_g^s$) can then be calculated as:

$$\beta_g^s(i) = \frac{\sum_{t \in T_g^s} \beta_t(i)}{|T_g^s|}, \forall i = 1, \cdots, \text{number of topics} \quad (1)$$

and $g$'s JS-divergence at phase $s$ is defined as

$$JS(g^s) = H(\beta_g^s) - \frac{\sum_{t \in T_g^s} H(\beta_t)}{|T_g^s|} \quad (2)$$

where $H(\bullet)$ is the Shannon entropy function (with log base 2) (Lin 1991). Intuitively, JS-divergence here gauges the divergence among topic distributions of a group's tweets. **The greater the $JS$ value, the larger the difference and the stronger indication of a group lacking conformity in discussion.**

## Prediction Problem Statement

Our goal is to solve a learning problem to predict the increase or decrease in the divergence of a group's discussion topics, measured by its discussion divergence, over an event's three phases: *pre-*, *during-*, and *post-event* (however, our analysis approach is applicable in general beyond the three phases of interests here). Specifically:

*Given a real-world event $E$, a collection of $N$ Twitter users discussing about $E$, and an assignment of them into $K$ non-overlapping user groups $g_i (1 \leq i \leq K)$ based on interactions, predict the change of each group's discussion divergence $JS(g_i)$ between two consecutive event phases (that is,* *from pre-event to during-event or from during-event to post-event).*

## Feature Design

In this section, we describe the feature design driven by socio-psychological theories.

### Structural Features Guided by Social Cohesion

To study the structural features driven by cohesion of social groups in a quantitative manner, we extract information from Twitter users' follower network. For each social group, we construct its corresponding node-induced sub-graph from the follower network. Because the follower relation is directional, there are three groups of features:

- *Reciprocal*:
  An undirected edge will be created between two users only when both of them are following each other. This choice directly reflects the assumption of mutual interpersonal attraction in the social cohesion theory. Features here include density, transitivity[2], average clustering coefficient[3], and maximum average length of pairwise shortest paths over all connected components (short-named "average shortest path length").
- *Undirected*:
  An undirected edge will be created between two users if either of them is following the other. The underlying assumption is that one-way interpersonal attraction is sufficient to keep the social group sustained. The same group of features as in the reciprocal sub-graph are computed.
- *Directed*:
  We also compute density and transitivity on the directed sub-graph for each social group, without converting it to an undirected graph.

The range for all cohesion features is $[0, 1]$, except for the average shortest path length. Note that in existing sociology literature (Moody and White 2003; White and Harary 2001) the term "structural cohesion" is a specific measure, defined as the minimum number of nodes one needs to remove from a graph to disconnect it. We do not include this feature as we find that almost all (more than 97% of total) social groups contain at least one fringe node (whose degree is one) or singleton, meaning that the value of this feature for most social groups will be at most one.

### User Features Guided by Social Identity

To quantify the social identity-based features, we extract user's profile information as well as activity, as we note that social behavior tends to associate the user with established identities (regional, organizational, etc.) via self-representation and with incentive-based identity via user actions in the cyber-world. For example, 'New Yorker' in a user's profile is an indicative signal of his location-based identity, and a profile containing 'professional NBA player' or 'Emergency Management' is highly suggestive of the user's occupational expertise. A user's action of adding such indicative terms into the profile suggests his self-awareness of the identity. Moreover, recently emerged social analytics services show online identities of users such as 'celebrity' on Klout, 'Mayor of a place' on Foursquare, etc., and users often tend to identify with them (Cramer, Rost, and Holmquist 2011). Thus, we are living with various social identities in both our *physical* as well as *cyber* world. We use location and description metadata in user profiles in addition to user actions (status updating, interacting, etc.) to extract the following types of social identities. Each identity type is modeled as a discrete attribute and for each social group under study, we compute the class distribution entropy for each identity and serve them as user features for the analysis. The range of identity features is from 0 to $\ln(C)$, where $C$ is the number of unique classes in an identity type.

- *Regional Identity feature*:
  Using location information in user profiles, we map users to regional classes that is sometimes used to represent self-identification in our daily lives — state-based (e.g., 'Ohio' for Ohioans) and nation-based (e.g., 'Brazil' for Brazilians). For creating feature value, we choose a user's state identity if it belongs to the host nation of the event (e.g., user from Buffalo will have 'NY' as the identity value in the OWS event), otherwise, we choose the user's nation identity (e.g., user from London will have 'UK' as the identity value in the OWS event). We use the Geonames dataset on Linked Open Data (LOD) cloud and Google Maps API to convert user profile locations into latitude-longitude, and then state and nation identity. We note that this simple model of two regional levels (state and nation) for self-identity can be expanded further.
- *Expertise Identity feature*:
  Users generally write their interests, expertise and affiliations in the description on Twitter user profiles. It is an example of self-representation of social identity (e.g., artist, researcher, etc.). Therefore, we first derive expertise classes by 2 steps: a) collect occupation categories and titles from trusted knowledge sources — Wikipedia and the US department of Labor Statistics reports, and b) classify the resulting occupation lexicon into ten broad classes, inspired by the domain classification on news websites and also from the super classes in the knowledge bases:
  {*ACADEMICS, BUSINESS, POLITICS, TECHNOLOGY, BLOGGING, JOURNALISM, ART, SPORTS, MEDICAL, OTHERS* }
  For expertise identity assignment to a user, we first create N-grams from the description metadata in the profile by tokenizing on punctuations, and filter out those not containing any of the occupation lexicon terms. From the remaining N-gram set, each N-gram is associated with one of the ten classes, and its weight is determined by its position in the description text. Because users tend to place terms that are more socially identifying and important to them at the beginning, due to self-awareness of identity representation. Finally, the user is assigned to the highest-weighted identity class.

[2] transitivity = $\frac{3 \times \text{number of triangles}}{\text{number of connected triples of vertices}}$

[3] clustering coefficient of node $i$ = $\frac{2 \times \text{number of triangles in } i\text{'s neighborhood}}{degree(i) \times (degree(i)-1)}$

| | Irene | Sandy | IAC | OWS |
|---|---|---|---|---|
| **Directed Structural Features** | | | | |
| Density | $0.04 \pm 0.07$ | $0.06 \pm 0.08$ | $0.02 \pm 0.03$ | $0.05 \pm 0.04$ |
| Transitivity | $0.23 \pm 0.20$ | $0.21 \pm 0.23$ | $0.10 \pm 0.18$ | $0.19 \pm 0.23$ |
| **Reciprocal Structural Features** | | | | |
| Density | $0.03 \pm 0.07$ | $0.04 \pm 0.07$ | $0.01 \pm 0.02$ | $0.03 \pm 0.04$ |
| Transitivity | $0.16 \pm 0.19$ | $0.18 \pm 0.24$ | $0.07 \pm 0.20$ | $0.14 \pm 0.24$ |
| Average Clustering Coefficient | $0.06 \pm 0.10$ | $0.08 \pm 0.12$ | $0.02 \pm 0.05$ | $0.05 \pm 0.09$ |
| Average Shortest Path Length | $2.25 \pm 1.19$ | $1.83 \pm 1.10$ | $1.06 \pm 0.99$ | $1.56 \pm 0.76$ |
| **Undirected Structural Features** | | | | |
| Density | $0.05 \pm 0.09$ | $0.07 \pm 0.09$ | $0.04 \pm 0.04$ | $0.06 \pm 0.05$ |
| Transitivity | $0.16 \pm 0.16$ | $0.19 \pm 0.22$ | $0.08 \pm 0.15$ | $0.16 \pm 0.21$ |
| Average Clustering Coefficient | $0.14 \pm 0.13$ | $0.13 \pm 0.15$ | $0.05 \pm 0.09$ | $0.10 \pm 0.12$ |
| Average Shortest Path Length | $2.72 \pm 0.90$ | $2.36 \pm 1.06$ | $2.01 \pm 0.82$ | $2.07 \pm 0.64$ |
| **User Features** | | | | |
| Regional Entropy | $2.71 \pm 0.78(5.28)$ | $2.24 \pm 0.73(5.74)$ | $2.06 \pm 0.45(4.94)$ | $2.12 \pm 0.62(5.65)$ |
| Expertise Entropy | $1.79 \pm 0.26(2.30)$ | $1.08 \pm 0.46(2.30)$ | $1.56 \pm 0.31(2.30)$ | $1.50 \pm 0.27(2.30)$ |
| Online Entropy | $0.97 \pm 0.21(2.08)$ | $1.03 \pm 0.21(2.08)$ | $1.24 \pm 0.24(2.08)$ | $1.18 \pm 0.23(2.08)$ |

Table 5: Mean and standard deviation of structural and user features. Identity entropy upper bounds are listed in brackets.

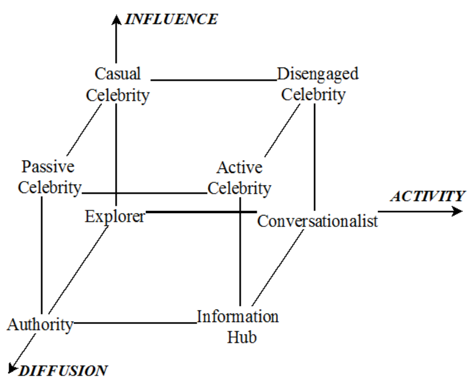


Figure 1: Online Identity based on three action measures (Influence, Diffusion, Activity)

- *Online Identity feature*:
  Based on user actions on the platform (Twitter here), we use three metrics following the work of expertise presentation in (Purohit et al. 2012) and influence and passivity in (Romero et al. 2011) that contribute to building a user's incentive-based identity (e.g., 'Celebrity' on Twitter) of cyber-world — an online identity in contrast to real-world identities by capturing user activity, influence and diffusion strength. We model the activity metric by number of posts of the user, influence metric by number of mentions of the user, and diffusion strength by number of retweets of the user's posts during event time-frame. We compute scores on each of the three metric dimensions for all users and then consider the basic 50th percentile threshold to create two levels on each of the dimensions, yielding 8 user classes as shown in Figure 1. The computation on number of mentions, number of retweets, and number of posts here is different from the step of identifying social groups in the interaction-only network, because here node-centric features (a *local* viewpoint) are taken for identity measure, and not the connection-centric feature set, (a *global* viewpoint), which is the basis of clustering.

  In contrast with regional and expertise identities, which are meaningful in the physical world, online identities exclusively define behavior in the cyber realm. From our knowledge, few attempts have been made to study the impact of both online and offline identities on group dynamics in online social networks.

  In Table 5 we summarize the basic statistical information of each of the features related to social cohesion and identity. The upper bounds of entropy values for user features are included in brackets. From the assumptions of social cohesion and social identity theories, we hypothesize the following:

- A more structurally cohesive social group has less diverse discussion. Therefore, groups with higher density, transitivity, clustering coefficient, or lower shortest path length are expected to have lower discussion divergence.
- Groups whose members are similar in identities (those having lower entropy for identity features) are speculated to have low discussion divergence, as motivated by the social identity theory.

## Analyses of Group Features and Discussion Divergence

In this section, we present the characteristics of structural and user features described in the previous section on our dataset and their correlation with group discussion divergence. It rationalizes the choice of features for the prediction task discussed in the next section.

### User & Structural Feature Statistics

We identify several interesting trends in the results reported in Table 5. First, in general the entropy values[4] are higher for

[4]Note, it is important to normalize these values against the maximum entropy possible for each case.

| | Irene | Sandy | IAC | OWS |
|---|---|---|---|---|
| **Directed Structural Features** | | | | |
| Density | $[-0.37, -0.06]$ | $[-0.22, -0.16]$ | $[-0.38, 0.07]$ | $[-0.03, 0.05]$ |
| **Reciprocal Structural Features** | | | | |
| Density | $[-0.36, -0.06]$ | $[-0.20, -0.15]$ | $[-0.29, 0.06]$ | $[-0.01, 0.07]$ |
| Shortest Path | $[0.27, 0.52]$ | $[0.10, 0.15]$ | $[-0.14, 0.21]$ | $[0.10, 0.16]$ |
| **Undirected Structural Features** | | | | |
| Density | $[-0.36, -0.05]$ | $[-0.22, -0.17]$ | $[-0.43, 0.10]$ | $[-0.05, 0.04]$ |
| Shortest Path | $[0.31, 0.56]$ | $[0.16, 0.21]$ | $[0.02, 0.37]$ | $[0.09, 0.13]$ |
| **User Features** | | | | |
| Regional Entropy | $[0.23, 0.50]$ | $[0.25, 0.30]$ | $[0.07, 0.52]$ | $[0.09, 0.14]$ |
| Expertise Entropy | $[0.11, 0.51]$ | $[0.45, 0.50]$ | $[0.37, 0.66]$ | $[0.01, 0.06]$ |
| Online Entropy | $[0.45, 0.69]$ | $[0.20, 0.25]$ | $[0.11, 0.57]$ | $[0.26, 0.31]$ |

Table 6: 95% confidence intervals of correlation coefficients between structure/user-based features and group discussion divergence

the Occupy Wall Street (OWS) and India Anti-Corruption (IAC) events, the two on-the-ground social activism events, possibly because the offline interactions heavily involved in those events are not captured by online social identity features. Such distinction is most pronounced when comparing online identity entropy values of those two events with respect to the other two events. The social groups in these two events tend to revolve around opinion leaders who often help direct and orchestrate the movement (such individuals likely will have high online identity values). Therefore social groups formed in those events generally have more diverse online identity composition, reflecting the presence of opinion leaders as well as followers in groups.

Another finding from Table 5 is that groups have great divergence in terms of their memberships from different regions. This may simply be a reflection of the times and the fact that online social networks are bringing people closer together and almost all events have had significant media attention.

Lastly, we point out that the average directed transitivity (global clustering coefficient) is at least 82% higher than that of the whole follower network (not shown in the table), and results based on the reciprocal and undirected definitions are similar, indicating that there is likely a community structure embedded in the social groups we have identified.

## Correlation Between Features & Group Discussion Divergence

To investigate the relation between structural/user features and group discussion divergence, we first compute their statistical correlation. Particularly, we use bootstrap method (sampling with replacement) to construct the 95% confidence interval of correlation coefficients. In Table 6, we report a subgroup of features whose correlation with group discussion divergence is considered significant.

**User features statistics:** We note in Table 6 that user features (especially regional identity entropy and online identity entropy) have a moderate to high positive correlation with group discussion divergence, for the first three events. This finding agrees with our hypothesis that group discussion divergence rises when group members' identities become less distinct and thus identity entropy values rise. Correlation values for Occupy Wall Street are less significant, possibly due to some intrinsic characteristics of its conversation (Conover et al. 2013).

For social groups with a stronger regional concentration, in-group discussions tend to be more location-specific and consistent, leading to a smaller degree of member-wise discussion divergence, compared with groups whose members' locations are more dispersed. Similarly, the presence of users with similar expertise or interest domain in a social group tends to keep the scope of discussions more focused.

For the online identity feature, we note that it is reflective of user actions. Therefore, we speculate that for the sake of maintaining their incentive-based action identity via lesser change in their actions, users are likely to maintain a pattern of focused topic discussions in the groups.

**Structural features statistics:** For structural features, we find that patterns of correlation with group discussion divergence can be categorized into two types:

- Density features have a moderate correlation with group discussion divergence for Hurricane Irene and Hurricane Sandy, indicating that a better-connected social group tends to have a more cohesive discussion.
  We ask an event-type specific question, why is the correlation weaker for Occupy Wall Street and the India anti-corruption movements? As mentioned earlier, both of them are long-lasting events accompanied by an arguably more engaged offline component, whose information is not captured in cohesion features. Therefore, the density of online social groups is low (see Table 5), making it less indicative of sustainability for those two events.
- Average shortest path length (especially the undirected version) shows consistency in its positive correlation with group discussion divergence, which also agrees with our hypothesis. Compared with others structural features that reflect the tightness of a social group, average shortest path length shows clearer dispersion in values, making the result from its correlation analysis more meaningful.
- When comparing correlation strengths with content-divergence by reciprocal features and undirected features, we find that they are often comparable. In fact, a one-sided binomial test rejects the alternative hypothesis that "reciprocal features have stronger correlation with group discussion divergence than undirected features" with a p-value of 0.89. This finding is particularly interesting as the key premise of reciprocal structural features is *mutual* interpersonal attractions (social cohesion theory), an assumption that undirected structural features do not make. This leads to the question of whether mutual attraction is still a necessary condition for *online* communities to form and last, and we believe it requires more research attention in the future.

## Contrasting High & Low Divergent Groups

We performed a case study of 10 highest and lowest divergent groups in each event, where we analyzed their content

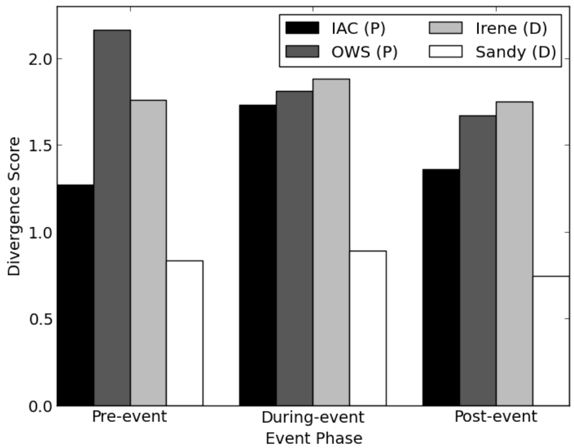


Figure 2: Average discussion divergence of groups in each of the phases for various events.

to check if there is contrast between the content practices. Specifically, we compared the frequency of using hashtags, retweets (RT), mentions, URL links, and emoticons in the content of candidate group members. An interesting contrast was that the least divergent group members use practice of RT heavily, while the most divergent groups use hashtags heavily, indicating diverging nature of user classified topics. Therefore, we suspect content practices also play a role in predicting trend of divergence.

### Effects of Event Characteristics

From Table 6 we note that transient events (Hurricane Irene and Hurricane Sandy) have stronger correlations with user features than with structural features. We conjecture it is due to the fact that groups in such volatile events form in an ad-hoc setting, where groups are less likely to have existing cohesively connected users, undermining the effects of structural features. Therefore, discussions can be highly dependent on the characteristics of participants of the group, their personal behavior and identities.

Furthermore, Figure 2 shows the general pattern of lower topical divergence in the *pre-event* phase, while increasing in the *during-event* phase and then again decreasing to lower value in the *post-event* phase. OWS is an outlier here likely due to high number of incidents even prior to the *pre* phase of the event in our dataset.

## Predicting Trend of Group Discussion Divergence

In this section, we present the methodology and results for our main task: to predict the trend of social groups' discussion divergence. We plan to leverage observations from previous sections, including 1) statistical correlations between features and group discussion divergence, and 2) disparities of a subgroup of feature values between groups of high versus low group discussion divergence.

More precisely, our goal is to solve a learning problem where the label is whether the discussion divergence of a group of users will increase or decrease over time. Since each event is divided into three phases, there are two transitions: pre-event to during-event, during-event to post-event. Features selection are guided by the statistical analyses and case studies in previous sections.

### Feature Sets and Learning Instances

We consider three main categories of features to use in the prediction problem. First, structural features focus on the cohesion and connectivity of each group's follower network. Second, user features emphasize the conformity of group users' offline and online identities. We have defined a family of those features in previous sections, and their significance varies. Lastly, content features capture the content practices of user-generated content. Based on the analyses in previous sections, we select different subsets of features from all of them, in order to reduce redundancy and improve prediction performance. The subsets are as follows:

- $Divergence$: Discussion divergence of the group at the current phase.
- $Structure_{sub}$: Directed density, reciprocal density, undirected density, reciprocal average shortest path length, undirected average shortest path length.
- $Structure_{all}$: All structural features described in the Feature Design section.
- $User_{all}$: Location entropy, occupation entropy, and online entropy.
- $Content_{sub}$: Average numbers of retweets and hashtags.
- $Content_{all}$: $Content_{sub}$ and average numbers of mentions, URLs and emoticons.

For each event, we identify pairs of social groups that are overlapping (Jaccard similarity[5] is above 0.5) before and after transition between two phases. There are 69 instances of group pairs meeting this criterion, and for 35 pairs their group discussion divergence values increase. We assign a label of 'increase' or 'decrease' to each group pair, depending on the change of its group discussion divergence value.

### Experiment Setup

For each pair of social groups of consideration, we use its features *before* the transition for the prediction task. Both SVM[6] (*SVM*) and logistic regression (*logistic*) are used.

We also create another baseline method (referred to as *baseline*), which relies its classification on the current phase. In the preliminary analysis of content divergence above, it is observed that groups' content divergence in general increases from *pre-event* to *during-event*, and decreases from *during-event* to *post-event*. Therefore, *baseline* always predicts a group's discussion divergence to 'increase' if it is currently in the *pre-event* phase, and 'decrease' if it belongs to the *during-event* phase.

### Learning performance

To evaluate the performance of group discussion divergence prediction, we perform a five-fold cross validation on *SVM*

[5]The Jaccard similarity between two sets $A$ and $B$ is $\frac{|A \cap B|}{|A \cup B|}$.

[6]RBF kernel with $\gamma$ value set to 0.5.

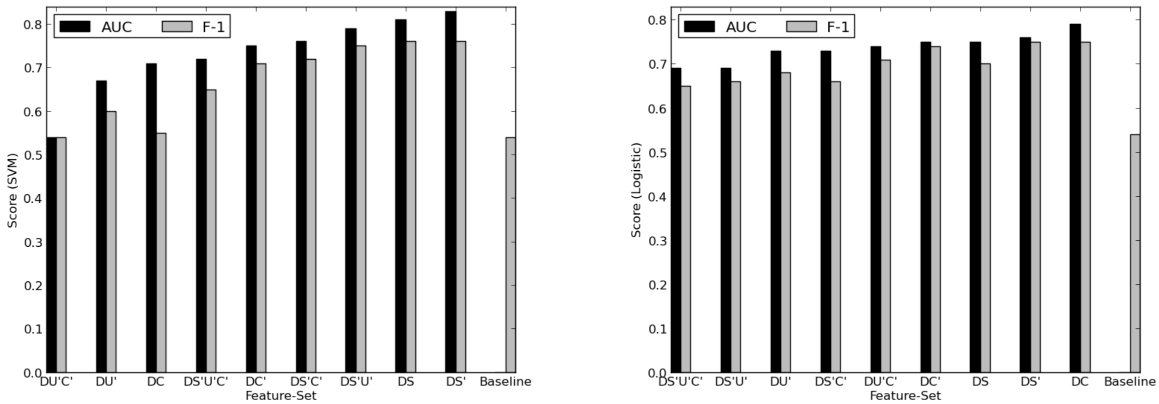


(a) SVM

(b) Logistic Regression

Figure 3: AUC and F-1 of prediction for SVM and logistic regression, organized by feature set and sorted by AUC. D=$Divergence$, U'=$User_{all}$, S=$Structure_{sub}$, S'=$Structure_{all}$, C=$Content_{sub}$, C'=$Content_{all}$.

and *logistic*. For *baseline*, we directly compute its F-1 score (0.54). Figure 3 shows the performance of various feature sets and learning models, measured by area under the curve (AUC) and F-1 score.

It is demonstrated from Figure 3 that classification based on features described in previous sections are significantly more accurate than the baseline method (F-1 of SVM using structural and user features is 0.75, a 39% improvement). Furthermore, performance of classifiers varies according to the selection of features to use. While user features have shown high correlation with *static* group discussion divergence, our results suggest that structural features contribute most to accurately predicting the *dynamic* change of content divergence. Using structural features only, SVM achieves the best AUC (0.83) and F-1 score (0.76).

## Discussion

We performed qualitative study on the content of the overlapping groups by transition of phase (e.g., mid to post), and the divergence shift (e.g., decrease) using the Linguistic Inquiry Word Count (LIWC) software. We observe that groups who tend to diverge in their discussions write more of general reporting type content based on past incidents. While the groups with decreasing diverging behavior write more social and future action related content, likely due to users being organized to inform the fellow members about updates on the situation. For example, we found in the overlapping candidate groups of Hurricane Sandy event that a group with decreasing diverging behavior was highly focused on the updates of flight statuses of different airlines, first delays and cancellation, and later on the resuming parts. Such focused and active topic-specific groups will be valuable to engage with by the response coordinators.

To summarize our main contribution, we present an approach to understand factors that drive the shift of collective diverging behavior in the group discussion topics, and illustrate by a prediction model to show that these factors can help track the behavior of group discussion divergence. Its application can be in several domains, such as in brand management, or disaster response coordination. We can identify groups of audience that are active and concerned about specific issues. In the massive social media community after disasters, identifying reliable sources for engagement to coordinate about specific needs is a daunting task and the proposed approach also helps in identifying reliable sources of groups with specific information of needs. Another application of the proposed approach is for deciphering the self-organizing behavior of groups by learning the collective diverging trends.

Summarizing limitations about our study, we note that other group formation methods can be used and evaluated. We also limit ourselves to three phases in the prediction model experiment, namely *pre-*, *during-* and *post-event*, based on the real-world incidents on the event timeline. However, more phases may be considered for longer events, as they could also possess long-term impact. Extended evaluation needs to be performed across more events of diverse types in the future to validate the work's generalizability. We also did not consider other types of group behaviors due to first time analyzing event-oriented group discussion for collective behavior and thus, future studies can expand on that.

For our future work, we plan to extend our features of social identity and cohesion, including ethnic and religious social relationships, and structural properties from Twitter List subscriptions. We shall also validate models into other social networks, such as Facebook, Google+, LinkedIn, and the DBLP co-authorship network, to see if they show a similar social phenomena of group dynamics. Finally, we are also interested in detecting transition point of group discussion divergence over time, which may corresponds to the phase change from *storming* to *norming* in the group developmental sequence theory (Tuckman 1965).

## Conclusion

This study focuses on characterizing the online social group dynamics using content of group discussion in contrast to structural properties studied earlier, and proposes a measure of *group discussion divergence*. We include structural and user features guided by two socio-psychological theories of group bonding and attachment — social identity and social cohesion. Leveraging these features in addition to content features, our classifiers accurately predict the future change of collectively diverging behavior in the group discussions. The classifiers achieve F-1 scores of up to 0.8, which is a 33% relative improvement from the baseline method. This study provides a framework to further research about collective behavior in online social groups.

## Acknowledgements

We thank Valerie Shalin, Christopher Browning, and the reviewers for their useful comments and NSF for sponsoring SoCS grants IIS-1111118 and IIS-1111182 titled as 'Social Media Enhanced Organizational Sensemaking in Emergency Response'.

## References

Backstrom, L.; Huttenlocher, D.; Kleinberg, J.; and Lan, X. 2006. Group formation in large social networks: membership, growth, and evolution. In *SIGKDD'06*, 44–54. ACM.

Beal, D.; Cohen, R.; Burke, M.; and McLendon, C. 2003. Cohesion and performance in groups: a meta-analytic clarification of construct relations. *Journal of Applied Psychology* 88(6):989.

Blei, D., and Lafferty, J. 2006. Dynamic topic models. In *ICML'06*, 113–120. ACM.

Branscombe, N. R., and Wann, D. L. 1991. The positive social and self concept consequences of sports team identification. *Journal of Sport & Social Issues* 15(2):115–127.

Budak, C., and Agrawal, R. 2013. On participation in group chats on twitter. In *WWW'13*, 165–176.

Conover, M. D.; Davis, C.; Ferrara, E.; McKelvey, K.; Menczer, F.; and Flammini, A. 2013. The geospatial characteristics of a social movement communication network. *PloS one* 8(3):e55957.

Cramer, H.; Rost, M.; and Holmquist, L. E. 2011. Performing a check-in: emerging practices, norms and'conflicts' in location-sharing using foursquare. In *MobileHCI'11*, 57–66. ACM.

Dyaram, L., and Kamalanabhan, T. 2005. Unearthed: the other side of group cohesiveness. *Journal of Social Science* 10(3):185–90.

Farzan, R.; Dabbish, L. A.; Kraut, R. E.; and Postmes, T. 2011. Increasing commitment to online communities by designing for social presence. In *CSCW'11*, 321–330.

Festinger, L.; Schachter, S.; and Back, K. 1950. The spatial ecology of group formation. *Social pressure in informal groups* 33–60.

Grabowicz, P. A.; Aiello, L. M.; Eguíluz, V. M.; and Jaimes, A. 2013. Distinguishing topical and social groups based on common identity and bond theory. In *WSDM'13*.

Kairam, S.; Wang, D.; and Leskovec, J. 2012. The life and death of online groups: Predicting group growth and longevity. In *ICWSM'12*, 673–682. ACM.

Lin, C.-Y.; Cao, G.; Gao, J.; and Nie, J.-Y. 2006. An information-theoretic approach to automatic evaluation of summaries. In *NAACL HLT'06*, 463–470. Association for Computational Linguistics.

Lin, J. 1991. Divergence measures based on the shannon entropy. *Information Theory, IEEE Transactions on* 37(1):145–151.

Lott, A., and Lott, B. 1965. Group cohesiveness as interpersonal attraction: A review of relationships with antecedent and consequent variables. *Psychological bulletin* 64(4):259.

Louis, A., and Nenkova, A. 2013. Automatically assessing machine summary content without a gold standard. *Computational Linguistics* 39(2):267–300.

Moody, J., and White, D. 2003. Structural cohesion and embeddedness: A hierarchical concept of social groups. *American Sociological Review* 103–127.

Mullen, B., and Copper, C. 1994. The relation between group cohesiveness and performance: An integration. *Psychological Bulletin; Psychological Bulletin* 115(2):210.

Pennacchiotti, M., and Popescu, A. 2011. A machine learning approach to twitter user classification. In *ICWSM'11*.

Purohit, H.; Dow, A.; Alonso, O.; Duan, L.; and Haas, K. 2012. User taglines: Alternative presentations of expertise and interest in social media. In *Proceedings of the first ASE International Conference on Social Informatics*. IEEE.

Rao, D.; Paul, M.; Fink, C.; Yarowsky, D.; Oates, T.; and Coppersmith, G. 2011. Hierarchical bayesian models for latent attribute detection in social media. *ICWSM'11* 598–601.

Ren, Y.; Kraut, R.; and Kiesler, S. 2007. Applying common identity and bond theory to design of online communities. *Organization Studies* 28(3):377–408.

Romero, D. M.; Galuba, W.; Asur, S.; and Huberman, B. A. 2011. Influence and passivity in social media. In *ECML/PKDD'11*, 18–33.

Ruan, Y.; Purohit, H.; Fuhry, D.; Parthasarathy, S.; and Sheth, A. 2012. Prediction of topic volume on twitter. In *WebSci'12*, 397–402. ACM.

Satuluri, V., and Parthasarathy, S. 2009. Scalable graph clustering using stochastic flows: applications to community discovery. In *SIGKDD'09*, 737–746. ACM.

Shi, X.; Zhu, J.; Cai, R.; and Zhang, L. 2009. User grouping behavior in online forums. In *SIGKDD'09*, 777–786. ACM.

Tajfel, H.; Billig, M. G.; Bundy, R. P.; and Flament, C. 1971. Social categorization and intergroup behaviour. *European Journal of Social Psychology* 1(2):149–178.

Tuckman, B. W. 1965. Developmental sequence in small groups. *Psychological bulletin* 63(6):384.

Turner, J.; Hogg, M.; Oakes, P.; Reicher, S.; and Wetherell, M. 1987. *Rediscovering the social group: A self-categorization theory.* Basil Blackwell.

Turner, J. 1982. Towards a cognitive redefinition of the social group. *Social identity and intergroup relations* 15–40.

White, D., and Harary, F. 2001. The cohesiveness of blocks in social networks: Node connectivity and conditional density. *Sociological Methodology* 31(1):305–359.