Identifying Rhetorical Questions in Social Media

Suhas Ranganath[†], Xia Hu[⋆], Jiliang Tang[‡], Suhang Wang [†], Huan Liu[†]
[†] Arizona State University, [⋆] Texas A&M University, [‡] Yahoo! Labs
{srangan8, swang187, huan.liu}@asu.edu, hu@cse.tamu.edu, jlt@yahoo-inc.com

Abstract

Social media provides a platform for seeking information from a large user base. Information seeking in social media, however, occurs simultaneously with users expressing their viewpoints by making statements. Rhetorical questions have the form of a question but serve the function of a statement and might mislead platforms assisting information seeking in social media. It becomes difficult to identify rhetorical questions as they are not syntactically different from other questions. In this paper, we develop a framework to identify rhetorical questions by modeling the motivations of the users to post them. We focus on one motivation of the users drawing from linguistic theories, to implicitly convey a message. We develop a framework from this motivation to identify rhetorical questions in social media and evaluate the framework using questions posted on Twitter. This is the first framework to model the motivations for posting rhetorical questions to identify them on social media platforms.

Introduction

Social media makes it easier for users to reach out to a large number of friends, leading people to use it to seek information from them (Morris, Teevan, and Panovich 2010). This phenomenon is prevalent in social media platforms like Twitter and Facebook and has received considerable attention in recent literature (Ranganath et al. 2015a; Paul, Hong, and Chi 2011). Information seeking in social media, however, goes hand in hand with users expressing their thoughts by making statements. Millions of social media users express their thoughts by posting rhetorical questions (Paul, Hong, and Chi 2011). Rhetorical questions are defined as "posts that have the form of a question but serve the function of a statement" (Anzilotti 1982).

Identifying rhetorical questions in social media has several applications. First, systems that assist information seeking in social media can ignore them when addressing user queries (Ranganath et al. 2015a). Second, sociological studies have shown that rhetorical questions are widely employed persuasion tools by political and commercial campaigners (Petty, Cacioppo, and Heesacker 1981). Identifying rhetorical questions, therefore, can also assist in detecting persuasive tactics by social media campaigners and inform users potentially affected by them.

This task faces several challenges. First, rhetorical questions are in the form of a question, although they are not looking for a response. They are not syntactically different from other questions. Hence features like "bag of words" might not be suitable for identifying them. Second, rhetorical questions have the function of a statement, which can be exploited to identify them. However, the function served by a standalone post of a user is not always apparent.

Linguistic theories for the motivations of users to post rhetorical questions can be used to understand their function. Rhetorical questions are stated as an indirect speech act (Schmidt-Radefeldt 1977). This means that they do not convey a message explicitly, but implies the message from its context. Consider the question, "Would somebody willingly die for a claim he knew was a lie? ". It is hard to determine if the question is rhetorical only from the text. However, his most recent status message before the question says "RT @PastorKentB: our pride keeps us from seeing who Jesus is... John 8...he Pharisees are too concerned with themselves to see the son of God!". It is clear from this that the question is rhetorical and has an implied message. This indicates that rhetorical questions are likely to share context with the most recent status message of their user.

Drawing concepts from these linguistic theories, we propose a framework to identify rhetorical questions in social media by modeling the motivations of the user to post them. We evaluate the framework on a dataset of questions from the social media platform Twitter and demonstrate that the framework is effective in identifying rhetorical questions in social media. Specifically, we address the following questions: How to model the motivations of the user for posting rhetorical questions to identify them? Are approaches based on motivations of the user useful in identifying rhetorical questions in social media?

The primary contributions of the paper are the following:

- Formally defining the problem of identifying rhetorical questions posted in social media;

- Demonstrating the utility of linguistic theories for motivations to employ rhetorical questions in social media data;

- Proposing a framework to identify rhetorical questions posted in social media by modeling user motivations; and

- Evaluating the framework using a real world dataset of questions posted in the social media platform Twitter.

Parameter	Statistics
# Questions	40,146
# Rhetorical Questions	16,058
# Randomly Sampled Questions	24,088
# Prev Status Messages	40,146
# Users	33,331

Table 1: Dataset containing questions posted in Twitter.

Problem Statement

In this section, we present the notations and formally state the problem. Matrices are denoted by boldface uppercase letters (e.g. $\mathbf{X}$) and vectors by lowercase letters (e.g. $\mathbf{x}$). $\mathbf{X}_{ij}$ signifies the element in the i^{th} row and j^{th} column of matrix $\mathbf{X}$. The Frebonius norm of matrix $\mathbf{X}$ is denoted as $||\mathbf{X}||_F = \sqrt{\sum_{i,j} \mathbf{X}_{ij}^2}$.

Let the set of rhetorical questions be denoted by $\mathcal{R}$ and the set of randomly sampled questions by $\mathcal{S}$. The combined set of questions is denoted by $\mathcal{F} = [\mathcal{R}, \mathcal{S}]$ and the total number of questions by Q. For each question $q \in \mathcal{F}$, we collect the most recent status message the user posted previous to the question, and denote the set as $\mathcal{M}$.

We construct a dictionary of words, $\mathcal{W}$, used in the questions and the most recent message of the users posting it, with W words. We then construct the question word matrix $\mathbf{F} \in \mathbb{R}^{Q \times W}$ from the question set $\mathcal{F}$, whose each row contains the word-frequency vector of the corresponding question. We then construct the message word matrix $\mathbf{M} \in \mathbb{R}^{Q \times W}$ from the set $\mathcal{M}$. The i^{th} row of the question word matrix $\mathbf{F}_i$ has a corresponding row $\mathbf{M}_i$ containing the word frequencies of the most recent status message. We concatenate $\mathbf{F}$ and $\mathbf{R}$ vertically to form the matrix $\mathbf{P} \in \mathbb{R}^{2Q \times W}$.

The problem can then be formally stated as follows: *"Given the question set $\mathcal{Q}$, consisting of a set of known rhetorical questions $\mathcal{R}$ and randomly sampled questions $\mathcal{S}$, and the post word matrix $\mathbf{P}$, determine if a new question q is a rhetorical question".*

Data

The dataset consists of a set of questions collected from the social media platform Twitter. To collect rhetorical questions, we use questions which the user has labeled as rhetorical with appropriate hashtags following (Ma et al. 2014), where users have shown to employ hashtags to label their intention behind the tweet. We collect questions containing the hashtags "#rhetoricalquestion" or "#dontanswerthat" along with "?" appended to each hashtag from the Twitter Streaming API, denoting them as positive examples. Tweets containing "?" have been shown to be questions with high precision in (Cong et al. 2008). To construct negative examples, we randomly sample tweets containing "?", with a ratio of positive-negative examples following (Paul, Hong, and Chi 2011). For each question, we collect the most recent status message of the user and construct the matrix $\mathbf{P}$ from the questions and most recent status messages. Some statistics of the dataset are given in Table 1.

The questions that constitute the negative examples might have some rhetorical questions, so we use human assess-ments to validate them. The number of questions are large, and evaluating the entire set might be expensive. To address this we apply the mark and recapture technique used in population estimation methodologies (Brower, Zar, and Von Ende 1998). We draw two random samples and assign the probability of finding errors in the random samples r_1 and r_2 be p_{r_1} and p_{r_2}. The total number of errors in the samples are $e_{r_1} = p_{r_1} N_e$ and $e_{r_2} = p_{r_2} N_e$, where N_e. The number of errors in the intersection of the two samples are then $e_{r_1 r_2} = p_{r_1} p_{r_2} N_e$. The number of errors in the dataset, accounting for overestimation are

$$N_e = \frac{(e_{r_1} + 1) \times (e_{r_2} + 1)}{e_{r_1 r_2} + 1} - 1 \qquad (1)$$

with a measure of uncertainty computed as

$$SE = \sqrt{\frac{(e_{r_1} + 1)(e_{r_2} + 1)(e_{r_2} - e_{r_1 r_2})(e_{r_1} - e_{r_1 r_2})}{(e_{r_1 r_2} + 1)^2 (e_{r_1 r_2} + 2)}}$$

We draw two random samples of 1% of the size of the list and use human evaluators from Amazon Mechanical Turk to verify them. The definition of rhetorical question is given to them as "posts that have the form of a question, but serve the function of a statement". The evaluators mark 1 if they think the question is a rhetorical question and 0 otherwise. The 95% confidence intervals are computed as $Int_e = N_e \pm 1.96 SE$. We compute the confidence values as $92.51\% \pm 8.3\%$, thus validating the negative examples.

Identifying Rhetorical Questions

We first propose postulates to quantify the motivations of the user to post rhetorical questions drawing concepts from linguistic theories (Schmidt-Radefeldt 1977). A possible motivation for the user is to imply a message from the context of the question. We derive the context of the question from the most recent status message of the asker. We state the following postulate to operationalize this theory "Rhetorical questions share more context with the most recent post of their user than randomly sampled questions share with the most recent post of their user".

Implying a Message

To verify the postulate, we use the topic closeness of the question and the most recent status message to measure shared context between them. We first obtain the topics in the messages of $\mathbf{P}$ using LDA introduced in (Blei, Ng, and Jordan 2003) to construct the document-topic matrix $\mathbf{Y} \in \mathbb{R}^{2Q \times T}$, where T is the total number of topics. For a given rhetorical question $q \in \mathcal{R}$ and the most recent status message from the user posting it, the topic distributions can be obtained from the corresponding rows of $\mathbf{Y}$ and is denoted as $\mathbf{q}$ and $\mathbf{x}$ respectively. We measure the closeness in topic distributions between the question and the most recent message by computing the Euclidean similarity of $\mathbf{q}$ and $\mathbf{x}$ and assign it to the vector $\mathbf{s}_r$. We then randomly select a question from the set of negative examples $\mathcal{S}$ and measure its topic closeness with the most recent status messages of the user posting it and assign it to vector $\mathbf{s}_n$. We repeat this procedure for all the rhetorical questions in $\mathcal{R}$. The null hypothesis is given by $\mathcal{H}_0 : \mathbf{s}_r = \mathbf{s}_n$ and the alternate hypothesis is given by $\mathcal{H}_1 : \mathbf{s}_r > \mathbf{s}_n$. A paired t-test shows

that rhetorical questions share context with the most recent status messages of the user posting it significantly more than randomly sampled questions do with their most recent status message with $p < 0.0001$, thus verifying the postulate.

Modeling Shared Context

We represent the latent dimension representation of each row in $\mathbf{P}$ by the matrix $\mathbf{U} \in \mathbb{R}^{2Q \times K}$, where Q is the number of questions, and K is the number of latent dimensions. The first Q rows of $\mathbf{U}$ contain the latent dimension representation of the questions and the last Q rows contain the latent dimension representation of the most recent status message. To capture the shared context between the question and the most recent message of the user posting it, we make their latent dimensions close to each other. We formulate this as minimizing the loss function

$$\mathcal{F}_1 = \frac{1}{2} \sum_{i=1}^{2Q} \sum_{j=1}^{2Q} ||\mathbf{M}_{ij}\big(\mathbf{U}(i,*) - \mathbf{U}(j,*)\big)||_2^2, \quad (2)$$

where each element $\mathbf{M}_{ij}$ is written as follows

$$\mathbf{M}_{ij} = \begin{cases} 1 & \text{if } |j - i| = Q \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

This loss function proposes a penalty if the latent dimensions of the question are far from the latent dimensions of its most recent status message. It can be rewritten as

$$\mathcal{F}_1 = \sum_{k=0}^{2Q} \mathbf{U}_k \mathcal{L} \mathbf{U}_k = tr(\mathbf{U}^T \mathcal{L} \mathbf{U}) = ||\mathbf{U}^T \mathcal{L}^{1/2}||_F^2,$$

where $\mathcal{L}$ is the laplacian of matrix $\mathbf{M}$.

Incorporating Shared Context

We first give a semantic meaning to the latent matrices $\mathbf{U}$ and $\mathbf{V}$. Let us represent the post-word matrix as a product of latent dimension matrices by minimizing

$$\min_{\mathbf{U} \geq 0, \mathbf{V} \geq 0} ||\mathbf{P} - \mathbf{U}\mathbf{V}^T||_F^2. \quad (4)$$

We use the latent dimensions of the question from $\mathbf{U}$ to identify rhetorical questions. A random fraction of questions is labeled for training. Least squares perform classification on the latent dimension matrix $\mathbf{U}$ by minimizing

$$\min_{\mathbf{W}} ||\mathbf{I}(\mathbf{U}\mathbf{W} - \mathbf{Y})||_F^2. \quad (5)$$

Here $\mathbf{I} \in \mathbb{R}^{Q \times Q}$ is a diagonal matrix where each diagonal element $\mathbf{I}_{ii} = 1$, if the i^{th} question is labeled or 0 otherwise. $\mathbf{W} \in \mathbb{R}^{K \times N}$ is the weight matrix containing the weights given to each latent feature, where N is the number of classes. We assign $N = 2$ here as we have two classes corresponding to whether the question is rhetorical or not. $\mathbf{Y} \in \mathbb{R}^{2Q \times N}$ is the output of the classifier, and we label together the question and the most recent status message of the user. Each row of $\mathbf{Y}$, $\mathbf{Y}_i$ is labeled as $\{1, 0\}$ if the question is rhetorical, $\{0, 1\}$ if it is not rhetorical and $\{0, 0\}$ otherwise.

We now constrain the factorization framework with the models of the linguistic theories, to obtain

$$\min_{\mathbf{U}, \mathbf{V} \geq 0, \mathbf{W}} ||\mathbf{P} - \mathbf{U}\mathbf{V}^T||_F^2 + ||\mathbf{U}^T \mathcal{L}^{1/2}||_F^2$$
$$+ ||\mathbf{I}(\mathbf{U}\mathbf{W} - \mathbf{Y})||_F^2 + \eta(||\mathbf{U}||_F^2 + ||\mathbf{V}||_F^2),$$

Baseline	AUC	Accu	F1
Random	50.32	50.44	50.41
Phrase (Podgorny et al. 2015)	58.44	51.43	51.94
Topics	64.28	58.67	59.30
BOW	65.87	61.92	61.93
InfNeed (Zhao and Mei 2013)	66.01	62.41	62.42
Qweet (Li et al. 2011)	66.10	62.70	62.57
PrMsg (Bhattasali et al. 2015)	66.18	62.75	62.51
RhetId	**68.83**	**63.43**	**64.04**

Table 2: Performance evaluation of the algorithm.

We find optimal solutions for the three variables $\mathbf{U}$, $\mathbf{V}$, and $\mathbf{W}$ by minimizing the objective function. After obtaining the latent variables, we compute the estimated value of $\mathbf{Y}$ as $\hat{\mathbf{Y}} = \mathbf{U}\mathbf{W}$. We then select the rows of $\mathbf{Y}$ pertaining to the unlabeled questions to construct $\hat{\mathbf{Y}}_{test} \in \mathbb{R}^{Q \times 2}$. For each row, we compare the values in the two columns and assign it as rhetorical if the value in the first column is greater than the second column and not rhetorical otherwise.

Performance Evaluation

We now evaluate the performance of the algorithm using Accuracy, AUC, and F1 measure, and compare it with state-of-the-art baselines. We set the number of latent dimensions as 50. We randomly select 50% of the candidate questions for training. We illustrate the results in Table 2.

From Table 2, we see that (Podgorny et al. 2015) performs slightly better than random assignment, showing that the characteristics of poorly phrased questions are distinct from rhetorical questions. The baseline **Topics** shows an improvement in the performance but performs worse than **BOW**. This indicates that there is not a large difference in the topics of rhetorical questions and randomly sampled questions.

Algorithms for identifying information seeking questions (Zhao and Mei 2013) and (Li et al. 2011) give a small improvement over **BOW** indicating that linguistic characteristics for categorizing them are not similar to the characteristics of rhetorical questions. Hence, we need to use concepts unique to rhetorical questions to identify them. Algorithms employing context of the question (Bhattasali et al. 2015), gives an improvement over **BOW**, showing the importance of contextual information for the identification of rhetorical questions. The significant improvement in **RhetId** demonstrates that modeling motivations of rhetorical questions are useful for identifying them.

In summary, the results show that modeling the motivations of the user to post rhetorical questions drawing concepts from linguistic theories is effective in identifying them.

Related Work

Information seeking behavior of social media users has received considerable attention in research communities. The authors of (Morris, Teevan, and Panovich 2010; Paul, Hong, and Chi 2011) analyze the types of questions people ask in the social media platform Twitter. A large-scale analysis of question types in Twitter was conducted by (Paul, Hong, and Chi 2011), and rhetorical questions were identified as one of the primary types of questions. Responder identification for different categories of questions have been compared in

(Ranganath et al. 2015a) and the performance for rhetorical questions have been shown to be the lowest. These papers give insights into the different categories of questions in social media, but they manually label the categories.

Automatic identification of questions expressing an information need been addressed in (Zhao and Mei 2013) and (Li et al. 2011). (Li et al. 2011) state that rhetorical questions are most prone to get misclassified. They use linguistic, lexical and parts of speech features to identify information seeking questions. Identifying poorly phrased questions using the grammatical structure of questions have been addressed in (Podgorny et al. 2015). The authors in (Bhattasali et al. 2015) identify rhetorical questions by directly combining contextual information. We model the motivations of the user to post rhetorical questions by utilizing specific relations between the question and its context.

Discussion

We now discuss the possible implications of identifying rhetorical questions in social media data. One implication of our work is in improving information seeking systems in social media. People post questions in their status messages to seek subjective information that is better obtained from their social circles than a search engine (Morris, Teevan, and Panovich 2010). Social media platforms provide timely information and hence is used by people in seeking time-critical replies during natural disasters (Ranganath et al. 2015b). However, rhetorical questions form a large part of questions posted in social media platforms (Paul, Hong, and Chi 2011), and this can mislead information seeking systems. Our algorithms filter out rhetorical questions to enable better performance of information seeking systems.

Social media is used in socio-political campaigns owing to it's wide reach and easy access. Examples can be advocacy groups in election campaigns or attempts of radicalization by groups like the ISIS. The main purpose of the campaigns is to influence users for a cause, and persuasive tools are widely employed (Gass and Seiter 2015). Detecting these tools will warn users targeted by the campaigns. Rhetorical questions are important means of persuasion (Petty, Cacioppo, and Heesacker 1981), and algorithms for identifying them can play a crucial part in tracking the behavior of social media campaigners.

Conclusion and Future Work

Social media provides a new platform for people seeking information through their status messages. However, information seeking in social media goes hand in hand with users posting statements. Rhetorical questions that have the syntactic form of a question and the function of a statement can be misleading to users. In this paper, we develop a framework to identify rhetorical questions by modeling the motivations of the user for posting them. We focus on two motivations: the need for implying a message without explicitly stating it and to strengthen or mitigate a previous statement. We evaluate the framework on questions posted in Twitter and find that motivations for a user posting a rhetorical question is effective in identifying rhetorical questions.

Our work opens up interesting directions for future research. Modeling the context of a social media question beyond using the previous status message of the asker can help us in better model the motivations of the asker. The replies to rhetorical questions and the conversation dynamics driven by them can be used to study how social media respond to persuasive tactics. The traits of users posting rhetorical questions can be analyzed to better identify accounts used for persuasive social and political campaigns.

Acknowledgements

References

Anzilotti, G. I. 1982. The rhetorical question as an indirect speech device in english and italian. *Canadian Modern Language Review* 38(2):290–302.

Bhattasali, S.; Cytryn, J.; Feldman, E.; and Park, J. 2015. Automatic identification of rhetorical questions. In *ACL*.

Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3:993–1022.

Brower, J. E.; Zar, J. H.; and Von Ende, C. N. 1998. *Field and laboratory methods for general ecology*. WCB McGraw-Hill Boston, Massachusetts.

Cong, G.; Wang, L.; Lin, C.-Y.; Song, Y.-I.; and Sun, Y. 2008. Finding question-answer pairs from online forums. In *SIGIR*. ACM.

Gass, R. H., and Seiter, J. S. 2015. *Persuasion: Social influence and compliance gaining*. Routledge.

Li, B.; Si, X.; Lyu, M. R.; King, I.; and Chang, E. Y. 2011. Question identification on twitter. In *CIKM*. ACM.

Ma, Z.; Sun, A.; Yuan, Q.; and Cong, G. 2014. Tagging your tweets: A probabilistic modeling of hashtag annotation in twitter. In *CIKM*. ACM.

Morris, M. R.; Teevan, J.; and Panovich, K. 2010. What do people ask their social networks, and why?: a survey study of status message q&a behavior. In *CHI*. ACM.

Paul, S. A.; Hong, L.; and Chi, E. H. 2011. Is twitter a good place for asking questions? a characterization study. In *ICWSM*.

Petty, R. E.; Cacioppo, J. T.; and Heesacker, M. 1981. Effects of rhetorical questions on persuasion: A cognitive response analysis. *Journal of personality and social psychology* 40(3):432.

Podgorny, I. A.; Cannon, M.; Gielow, C.; and Goodyear, T. 2015. Real time detection and intervention of poorly phrased questions. In *CHI EA*. ACM.

Ranganath, S.; Tang, J.; Hu, X.; Sundaram, H.; and Liu, H. 2015a. Leveraging social foci for information seeking in social media. In *AAAI*.

Ranganath, S.; Wang, S.; Hu, X.; Tang, J.; and Liu, H. 2015b. Finding time-critical responses for information seeking in social media. In *ICDM*. ACM.

Schmidt-Radefeldt, J. 1977. On so-called rhetorical questions. *Journal of Pragmatics* 1(4):375–392.

Zhao, Z., and Mei, Q. 2013. Questions about questions: An empirical analysis of information needs on twitter. In *WWW*. ACM.