DOI: http://dx.doi.org/10.14236/ewic/FDIA2015.10

Topic-centric Classification of Twitter User's Political Orientation

Anjie Fang[1], Iadh Ounis[2], Philip Habel[2], Craig Macdonald[2] and Nut Limsopatham[2]
University of Glasgow, UK
[1]*a.fang.1@research.gla.ac.uk*, [2]*{firstname.secondname}@glasgow.ac.uk*

We aim to classify people's voting intentions by the content of their Tweets about the Scottish Independence Referendum (hereafter, IndyRef). By observing the IndyRef dataset, we find that people not only discussed the vote, but raised topics related to an independent Scotland including oil reserves, currency, nuclear weapons, and national debt. We show that the views communicated on these topics can inform us of the individuals' voting intentions ("Yes" vs. "No"). In particular, we argue that an accurate classifier can be designed by leveraging the differences in the features' usage across different topics related to voting intentions. We demonstrate improvements upon a Naive Bayesian classifier using the topics enrichment method. Our new classifier identifies the closest topic for each unseen tweet, based on those topics identified in the training data. Our experiments show that our proposed Topics-Based Naive Bayesian classifier improves accuracy by 7.8% over the classical Naive Bayesian baseline.

Keywords: Classification, Topic model, Bayesian theorem, Feature selection, Twitter

1. INTRODUCTION

Twitter emerged as an especially popular platform during the IndyRef held in 2014. We propose a technique to analyse the voting intentions of users, based on data mining and machine learning approaches. The general approach we propose could also be used to understand users' voting intentions in other major elections. To analyse voting intentions, we capture two months of Twitter data related to the IndyRef. To form a ground truth, we label users based upon hashtags appearing in their tweets, and we verify the reliability of this approach using the users' followee networks. After removing the hashtags from these tweets, we then focus on the remaining terms, treating each term as a feature. However, the referendum created an evolving discourse, with different topical themes (such as *oil*, *currency*, and *debt*), which make the accurate classification of users' voting intentions more challenging. For instance, the word "change" is indicative of a "No" voter in the *currency* topic, and of a "Yes" voter in the *nuclear weapons* topic. That is, there was a significant discussion over whether Scotland would need to "change" its currency if it obtained independence, while the "Yes" camp purported that the nuclear arsenal base could "change" in an independent Scotland. The dichotomy of the term "change" in indicating voting intentions across different topics highlights the main benefit of our approach. Indeed, this paper contributes the use of topical clusters to identify the topic of discussion in a tweet and subsequently it leverages this topic to classify the user's voting intention. Our approach, called *Topics-Based Naive Bayesian* (TBNB) demonstrates marked improvements over a classical Naive Bayes (NB) classification baseline.

2. BACKGROUND AND RELATED WORK

Cohen and Ruths (2004) demonstrated that classification of political orientation was still a difficult problem and that the earlier result in Al Zamal et al. (2004) was exaggerated since it used easily classifiable political data. We focus on the content of tweets to classify the users' voting intentions. We use as a starting point a classical Naive Bayesian (NB) classifier. Since the number of features can be very large, we use several feature selection approaches in Mladenic and Grobelnik (1999). Each selection approach ranks and selects F informative features based on the training data. Of course, not every selected feature will appear in the unseen test tweets.

3. TOPICS-BASED NAIVE BAYESIAN

The IndyRef discussions on Twitter revolved around a number of topics, for which people's opinions usually reflected their vote intentions. Let us continue the example of the word "change" usage in Section 1. The difference in usage of "change" across different topics is high. Furthermore, the conditional probability of "change" in the "Yes" category is higher than in the "No" category in the "currency" topic. Typically, the feature selection approaches just select features with higher differences between categories. If a feature differs between topics (e.g. "change"), it will be treated as different features in the TBNB model. Thus TBNB can capture term dependencies between topic and user voting intentions. Our TBNB classifier leverages both the features' dissimilarities across topics and in the categories. In the training step, the topics are first detected by Latent Dirichlet Allocation (LDA). For each topic, a corresponding probability table is produced, where each feature has two associated conditional probabilities related to the two possible voting intentions ("Yes"/"No"). Consequently, during the training step, we produce as many feature tables as the number of used topics. In the testing step, we treat a user as a virtual document and this document contains the users' tweets. For each tweet in the user's virtual document, the topic that is closest to the tweet's content is selected. Terms in an unseen tweet are then examined using the probability table generated during the training step for the topic with which this tweet is associated. In this way, terms in different tweets are treated differently based on their associated topics, and the TNBN classifier applies, for each unseen tweet, those features that were learned from the corresponding topic. Note that the feature selection approaches can naturally be applied to the TBNB classifier. For example, if F is set to 1000, the top 1000 features learned from each topic are selected.

4. REFERENDUM DATA AND EXPERIMENTS

Our IndyRef dataset was collected from Twitter by searching for a number of referendum-specific hashtags and keywords using the Twitter API from August 1, 2014 to September 30, 2014. In our dataset, certain "Yes" hashtags (e.g. #YesBecause) were associated with a "Yes" vote, and "No" hashtags (e.g. #NoBecause) with a "No" vote. To generate our ground truth, we assume that if a user's tweets are only tagged by "No" hashtags, this user is labeled as a "No" voter. Similarly, if a user's tweets contain only "Yes" hashtags, this user is labeled as a "Yes" supporter, favoring independence. Using this method, we find 5326 "Yes" users and 2011 "No" users. Together these 7337 users account for more than 420K tweets. After labelling, all "Yes" and "No" hashtags are removed from their original tweet text. The resulting tweets constitute our classification dataset. Without the hashtags, the classification task is naturally more challenging, but importantly, the resulting generalisable classifier does not require the presence of hashtags. We verify our ground-truth's reliability using the users' followee networks. In particular, If a user mainly follows Conservative politicians ("No" campaign supporters), this person is likely to be a "No" voter. If a user follows Scottish National Party politicians ("Yes" campaign supporters), their vote intention is more likely to be "Yes". We then examined the networks of the 7337 users in our dataset, and identified who these users follow among the 536 public Twitter accounts corresponding to Members of the British or Scottish Parliaments. We find that, of the 7337 users, 87% can be verified into "Yes" or "No" voters, demonstrating that our ground-truth produced by the hashtags labeling method is reasonable and reliable.

We use our IndyRef dataset to compare the performances of the NB and TBNB classifiers. We vary the number of selected features F and the deployed feature selection approach for both NB and TBNB. We also vary the number of topics T in the TBNB classifier. We use a 10-fold cross validation process over the 7337 users and use accuracy to measure the performance. Our results show that all TBNB classifiers markedly outperform the NB baseline when F ranges from 10K to 50K. The highest accuracy of TBNB (90.4%) is achieved when applying the weighted odds ratio feature selection approach with T=10 and F =30K, while the accuracy of the baseline is 82.6%. In an additional experiment aiming to check the generalisation of our conclusions, we obtained similar results using a different IndyRef dataset (collected from different period) with the same aforementioned T and F values.

5. CONCLUSIONS AND FUTURE WORK

We classified the users' voting intentions on Twitter during the IndyRef. We noted that the users tended to focus their discussions on topics, reflecting their voting intentions. We proposed to enrich the Naive Bayes classifier by leveraging the underlying topics covered in the tweets. Our proposed approach leverages the difference of the features across the topics and voting categories to increase the classification confidence. Our results demonstrate the effectiveness of our resulting TBNB classifier on two datasets. In the future, we plan to analyse the effect of the evolving discussions on the users' voting intentions over time.

6. ACKNOWLEDGEMENTS

We thank ACM SIGIR for the awarded scholarship for participating in ESSIR 2015 and the FDIA Symposium.

REFERENCES

Cohen, Raviv and Ruths, Derek (2013) *Classifying Political Orientation on Twitter: It's Not Easy!*, ICWSM.

Al Zamal, Faiyaz and Liu, Wendy and Ruths, Derek (2012) *Homophily and Latent Attribute Inference: Inferring Latent Attributes of Twitter Users from Neighbors*, ICWSM.

Mladenic, Dunja and Grobelnik, Marko (1999) *Feature* *selection for unbalanced class distribution and Naive Bayes*, ICML.