

A Computational Approach to Automatic Prediction of *Drunk-Texting*

Aditya Joshi^{1,2,3} Abhijit Mishra¹ Balamurali AR⁴
Pushpak Bhattacharyya¹ Mark James Carman²

¹IIT Bombay, India, ²Monash University, Australia

³IITB-Monash Research Academy, India ⁴Aix-Marseille University, France

{adityaj, abhijitmishra, pb}@cse.iitb.ac.in

balamurali.ar@lif.univ-mrs.fr, mark.carman@monash.edu

Abstract

Alcohol abuse may lead to unsociable behavior such as crime, drunk driving, or privacy leaks. We introduce automatic drunk-texting prediction as the task of identifying whether a text was written when under the influence of alcohol. We experiment with tweets labeled using hashtags as distant supervision. Our classifiers use a set of N-gram and stylistic features to detect drunk tweets. Our observations present the first quantitative evidence that text contains signals that can be exploited to detect drunk-texting.

1 Introduction

The ubiquity of communication devices has made social media highly accessible. The content on these media reflects a user’s day-to-day activities. This includes content created under the influence of alcohol. In popular culture, this has been referred to as ‘*drunk-texting*’¹. In this paper, we introduce automatic ‘*drunk-texting prediction*’ as a computational task. Given a tweet, the goal is to automatically identify if it was written by a drunk user. We refer to tweets written under the influence of alcohol as ‘*drunk tweets*’, and the opposite as ‘*sober tweets*’.

A key challenge is to obtain an annotated dataset. We use hashtag-based supervision so that the authors of the tweets mention if they were drunk at the time of posting a tweet. We create three datasets by using different strategies that are related to the use of hashtags. We then present SVM-based classifiers that use N-gram and stylistic features such as capitalisation, spelling errors, etc. Through our experiments, we make subtle points related to: (a) the *performance of our features*, (b) *how our approach compares against*

human ability to detect drunk-texting, (c) *most discriminative stylistic features*, and (d) *an error analysis* that points to future work. To the best of our knowledge, this is a first study that shows the feasibility of text-based analysis for drunk-texting prediction.

2 Motivation

Past studies show the relation between alcohol abuse and unsociable behaviour such as aggression (Bushman and Cooper, 1990), crime (Carpenter, 2007), suicide attempts (Merrill et al., 1992), drunk driving (Loomis and West, 1958), and risky sexual behaviour (Bryan et al., 2005). Merrill et al. (1992) state that “*those responsible for assessing cases of attempted suicide should be adept at detecting alcohol misuse*”. Thus, a drunk-texting prediction system can be used to identify individuals susceptible to these behaviours, or for investigative purposes after an incident.

Drunk-texting may also cause regret. Mail Goggles² prompts a user to solve math questions before sending an email on weekend evenings. Some Android applications³ avoid drunk-texting by blocking outgoing texts at the click of a button. However, to the best of our knowledge, these tools require a user command to begin blocking. An ongoing text-based analysis will be more helpful, especially since it offers a more natural setting by monitoring stream of social media text and not explicitly seeking user input. Thus, automatic drunk-texting prediction will improve systems aimed to avoid regrettable drunk-texting. To the best of our knowledge, ours is the first study that does a quantitative analysis, in terms of prediction of the drunk state by using textual clues.

Several studies have studied linguistic traits associated with emotion expression and mental

¹Source: <http://www.urbandictionary.com>

²<http://gmailblog.blogspot.in/2008/10/new-in-labs-stop-sending-mail-you-later.html>

³<https://play.google.com/store/apps/details?id=com.oopsapp>

health issues, suicidal nature, criminal status, etc. (Pennebaker, 1993; Pennebaker, 1997). NLP techniques have been used in the past to address social safety and mental health issues (Resnik et al., 2013).

3 Definition and Challenges

Drunk-texting prediction is the task of classifying a text as drunk or sober. For example, a tweet ‘*Feeling buzzed. Can’t remember how the evening went*’ must be predicted as ‘*drunk*’, whereas, ‘*Returned from work late today, the traffic was bad*’ must be predicted as ‘*sober*’. The challenges are:

1. **More than topic categorisation:** Drunk-texting prediction is similar to topic categorisation (that is, classification of documents into a set of categories such as ‘*news*’, ‘*sports*’, etc.). However, Borrill et al. (1987) show that alcohol abusers have more pronounced emotions, specifically, anger. In this respect, drunk-texting prediction lies at the confluence of topic categorisation and emotion classification.
2. **Identification of labeled examples:** It is difficult to obtain a set of sober tweets. The ideal label can be possibly given only by the author. For example, whether a tweet such as ‘*I am feeling lonely tonight*’ is a drunk tweet is ambiguous. This is similar to sarcasm expressed as an exaggeration (for example, ‘*This is the best film ever!*’), where the context beyond the text needs to be considered.
3. **Precision/Recall trade-off:** The goal that a drunk-texting prediction system must chase depends on the application. An application that identifies potential crimes must work with high precision, since the target population to be monitored will be large. On the other hand, when being used to avoid regrettable drunk-texting, a prediction system must produce high recall in order to ensure that a drunk message does not pass through.

4 Dataset Creation

We use hashtag-based supervision to create our datasets, similar to tasks like emotion classification (Purver and Battersby, 2012). The tweets are downloaded using Twitter API (<https://dev>.

twitter.com/). We remove non-Unicode characters, and eliminate tweets that contain hyperlinks⁴ and also tweets that are shorter than 6 words in length. Finally, hashtags used to indicate drunk or sober tweets are removed so that they provide labels, but do not act as features. The dataset is available on request. As a result, we create three datasets, each using a different strategy for sober tweets, as follows:



Figure 1: Word cloud for drunk tweets

1. **Dataset 1** (2435 drunk, 762 sober): We collect tweets that are marked as drunk and sober, using hashtags. Tweets containing hashtags #drunk, #drank and #imdrunk are considered to be drunk tweets, while those with #notdrunk, #imnotdrunk and #sober are considered to be sober tweets.
2. **Dataset 2** (2435 drunk, 5644 sober): The drunk tweets are downloaded using drunk hashtags, as above. The list of users who created these tweets is extracted. For the negative class, we download tweets by these users, which do not contain the hashtags that correspond to drunk tweets.
3. **Dataset H** (193 drunk, 317 sober): A separate dataset is created where drunk tweets are downloaded using drunk hashtags, as above. The set of sober tweets is collected using both the approaches above. The resultant is the held-out test set *Dataset-H that contains no tweets in common with Datasets 1 and 2.*

The drunk tweets for Datasets 1 and 2 are the same. Figure 1 shows a word-cloud for these drunk tweets (with stop words and forms of the word ‘*drunk*’ removed), created using

⁴This is a rigid criterion, but we observe that tweets with hyperlinks are likely to be promotional in nature.

Feature	Description
N-gram Features	
Unigram & Bigram (Presence)	Boolean features indicating unigrams and bigrams
Unigram & Bigram (Count)	Real-valued features indicating unigrams and bigrams
Stylistic Features	
LDA unigrams (Presence/Count)	Boolean & real-valued features indicating unigrams from LDA
POS Ratio	Ratios of nouns, adjectives, adverbs in the tweet
#Named Entity Mentions	Number of named entity mentions
#Discourse Connectors	Number of discourse connectors
Spelling errors	Boolean feature indicating presence of spelling mistakes
Repeated characters	Boolean feature indicating whether a character is repeated three times consecutively
Capitalisation	Number of capital letters in the tweet
Length	Number of words
Emoticon (Presence/Count)	Boolean & real-valued features indicating unigrams
Sentiment Ratio	Positive and negative word ratios

Table 1: Our Feature Set for Drunk-texting Prediction

WordItOut⁵. The size of a word indicates its frequency. In addition to topical words such as ‘bar’, ‘bottle’ and ‘wine’, the word-cloud shows sentiment words such as ‘love’ or ‘damn’, along with profane words.

Heuristics other than these hashtags could have been used for dataset creation. For example, timestamps were a good option to account for time at which a tweet was posted. However, this could not be used because user’s local times was not available, since very few users had geolocation enabled.

5 Feature Design

The complete set of features is shown in Table 1. There are two sets of features: (a) N-gram features, and (b) Stylistic features. We use unigrams and bigrams as N-gram features- considering both presence and count.

Table 1 shows the complete set of stylistic features of our prediction system. POS ratios are a set of features that record the proportion of each POS tag in the dataset (for example, the proportion of nouns/adjectives, etc.). The POS tags and named entity mentions are obtained from NLTK (Bird, 2006). Discourse connectors are identified based on a manually created list. Spelling errors are identified using a spell checker by Aby (2014). The repeated characters feature captures a situation in which a word contains a letter that is repeated three or more times, as in the case of

happy. Since drunk-texting is often associated with emotional expression, we also incorporate a set of sentiment-based features. These features include: count/presence of emoticons and sentiment ratio. Sentiment ratio is the proportion of positive and negative words in the tweet. To determine positive and negative words, we use the sentiment lexicon in Wilson et al. (2005). To identify a more refined set of words that correspond to the two classes, we also estimated 20 topics for the dataset by estimating an LDA model (Blei et al., 2003). We then consider top 10 words per topic, for both classes. This results in 400 LDA-specific unigrams that are then used as features.

	A (%)	NP (%)	PP (%)	NR (%)	PR (%)
Dataset 1					
N-gram	85.5	72.8	88.8	63.4	92.5
Stylistic	75.6	32.5	76.2	3.2	98.6
All	85.4	71.9	89.1	64.6	91.9
Dataset 2					
N-gram	77.9	82.3	65.5	87.2	56.5
Stylistic	70.3	70.8	56.7	97.9	6.01
All	78.1	82.6	65.3	86.9	57.5

Table 2: Performance of our features on Datasets 1 and 2

⁵www.worditout.com

6 Evaluation

Using the two sets of features, we train SVM classifiers (Chang and Lin, 2011)⁶. We show the five-fold cross-validation performance of our features on Datasets 1 and 2, in Section 6.1, and on Dataset H in Section 6.2. Section 6.3 presents an error analysis. *Accuracy, positive/negative precision and positive/negative recall are shown as A, PP/NP and PR/NR respectively. ‘Drunk’ forms the positive class, while ‘Sober’ forms the negative class.*

Top features		
#	Dataset 1	Dataset 2
1	POS_NOUN	Spelling_error
2	Capitalization	LDA_drinking
3	Spelling_error	POS_NOUN
4	POS_PREPOSITION	Length
5	Length	LDA_tonight
6	LDA_Llife	Sentiment_Ratio
7	POS_VERB	Char_repeat
8	LDA_today	LDA_today
9	POS_ADV	LDA_drunken
10	Sentiment_Ratio	LDA_lmao

Table 3: Top stylistic features for Datasets 1 and 2 obtained using Chi-squared test-based ranking

6.1 Performance for Datasets 1 and 2

Table 2 shows the performance for five-fold cross-validation for Datasets 1 and 2. In case of Dataset 1, we observe that N-gram features achieve an accuracy of 85.5%. We see that our stylistic features alone exhibit degraded performance, with an accuracy of 75.6%, in the case of Dataset 1. Table 3 shows top stylistic features, when trained on the two datasets. Spelling errors, POS ratios for nouns (POS_NOUN)⁷, length and sentiment ratios appear in both lists, in addition to LDA-based unigrams. However, negative recall reduces to a mere 3.2%. This degradation implies that our features capture a subset of drunk tweets and that there are properties of drunk tweets that may be more subtle. When both N-gram and stylistic features are used, there is negligible improvement. The accuracy for Dataset 2 increases from

⁶We also repeated all experiments for Naïve Bayes. They do not perform as well as SVM, and have poor recall.

⁷POS ratios for nouns, adjectives and adverbs were nearly similar in drunk and sober tweets - with the maximum difference being 0.03%

77.9% to 78.1%. Precision/Recall metrics do not change significantly either. The best accuracy of our classifier is 78.1% for all features, and 75.6% for stylistic features. This shows that text-based clues can indeed be used for drunk-texting prediction.

	A1	A2	A3
A1	-	0.42	0.36
A2	0.42	-	0.30
A3	0.36	0.30	-

Table 4: Cohen’s Kappa for three annotators (A1-A3)

	A (%)	NP (%)	PP (%)	NR (%)	PR (%)
Annotators	68.8	71.7	61.7	83.9	43.5
Training Dataset	Our classifiers				
Dataset 1	47.3	70	40	26	81
Dataset 2	64	70	53	72	50

Table 5: Performance of human evaluators and our classifiers (trained on all features), for Dataset-H as the test set

6.2 Performance for Held-out Dataset H

Using held-out dataset H, we evaluate how our system performs in comparison to humans. Three annotators, A1-A3, mark each tweet in the Dataset H as drunk or sober. Table 4 shows a moderate agreement between our annotators (for example, it is 0.42 for A1 and A2). Table 5 compares our classifier with humans. Our human annotators perform the task with an average accuracy of 68.8%, while our classifier (with all features) trained on Dataset 2 reaches 64%. The classifier trained on Dataset 2 is better than which is trained on Dataset 1.

6.3 Error Analysis

Some categories of errors that occur are:

1. **Incorrect hashtag supervision:** The tweet ‘Can’t believe I lost my bag last night, literally had everything in! Thanks god the bar man found it’ was marked with ‘#Drunk’. However, this tweet is not likely to be a drunk tweet, but describes a drunk episode in retrospective. Our classifier predicts it as sober.

2. **Seemingly sober tweets:** Human annotators as well as our classifier could not identify whether ‘*Will you take her on a date? But really she does like you*’ was drunk, although the author of the tweet had marked it so. This example also highlights the difficulty of drunk-texting prediction.
3. **Pragmatic difficulty:** The tweet ‘*National dress of Ireland is one’s one vomit.. my family is lovely*’ was correctly identified by our human annotators as a drunk tweet. This tweet contains an element of humour and topic change, but our classifier could not capture it.

7 Conclusion & Future Work

In this paper, we introduce automatic drunk-texting prediction as the task of predicting a tweet as drunk or sober. First, we justify the need for drunk-texting prediction as means of identifying risky social behavior arising out of alcohol abuse, and the need to build tools that avoid privacy leaks due to drunk-texting. We then highlight the challenges of drunk-texting prediction: one of the challenges is selection of negative examples (sober tweets). Using hashtag-based supervision, we create three datasets annotated with drunk or sober labels. We then present SVM-based classifiers which use two sets of features: N-gram and stylistic features. Our drunk prediction system obtains a best accuracy of 78.1%. We observe that our stylistic features add negligible value to N-gram features. We use our heldout dataset to compare how our system performs against human annotators. While human annotators achieve an accuracy of 68.8%, our system reaches reasonably close and performs with a best accuracy of 64%.

Our analysis of the task and experimental findings make a case for drunk-texting prediction as a useful and feasible NLP application.

References

Aby. 2014. Aby word processing website, January.

Steven Bird. 2006. Nltk: the natural language toolkit. In *Proceedings of the COLING/ACL on Interactive presentation sessions*, pages 69–72. Association for Computational Linguistics.

David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022.

Josephine A Borrill, Bernard K Rosen, and Angela B Summerfield. 1987. The influence of alcohol on judgement of facial expressions of emotion. *British Journal of Medical Psychology*.

Angela Bryan, Courtney A Rocheleau, Reuben N Robbins, and Kent E Hutchinson. 2005. Condom use among high-risk adolescents: testing the influence of alcohol use on the relationship of cognitive correlates of behavior. *Health Psychology*, 24(2):133.

Brad J Bushman and Harris M Cooper. 1990. Effects of alcohol on human aggression: An integrative research review. *Psychological bulletin*, 107(3):341.

Christopher Carpenter. 2007. Heavy alcohol use and crime: Evidence from underage drunk-driving laws. *Journal of Law and Economics*, 50(3):539–557.

Chih-Chung Chang and Chih-Jen Lin. 2011. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27.

Ted A Loomis and TC West. 1958. The influence of alcohol on automobile driving ability: An experimental study for the evaluation of certain medicological aspects. *Quarterly journal of studies on alcohol*, 19(1):30–46.

John Merrill, GABRIELLE MILKER, John Owens, and Allister Vale. 1992. Alcohol and attempted suicide. *British journal of addiction*, 87(1):83–89.

James W Pennebaker. 1993. Putting stress into words: Health, linguistic, and therapeutic implications. *Behaviour research and therapy*, 31(6):539–548.

James W Pennebaker. 1997. Writing about emotional experiences as a therapeutic process. *Psychological science*, 8(3):162–166.

Matthew Purver and Stuart Battersby. 2012. Experimenting with distant supervision for emotion classification. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 482–491. Association for Computational Linguistics.

Philip Resnik, Anderson Garron, and Rebecca Resnik. 2013. Using topic modeling to improve prediction of neuroticism and depression. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language*, pages 1348–1353. Association for Computational Linguistics.

Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the conference on human language technology and empirical methods in natural language processing*, pages 347–354. Association for Computational Linguistics.