

‘Who would have thought of that!’: A Hierarchical Topic Model for Extraction of Sarcasm-prevalent Topics and Sarcasm Detection

Aditya Joshi^{1,2,3}

Prayas Jain⁴

Pushpak Bhattacharyya¹

Mark James Carman²

¹Indian Institute of Technology Bombay, India, ²Monash University, Australia

³IITB-Monash Research Academy, India, ⁴IIT-BHU (Varanasi), India,

{adityaj, pb}@cse.iitb.ac.in, prayas.jain.csel4@iitbhu.ac.in

mark.carman@monash.edu

Abstract

Topic Models have been reported to be beneficial for aspect-based sentiment analysis. This paper reports a simple topic model for sarcasm detection, a first, to the best of our knowledge. Designed on the basis of the intuition that sarcastic tweets are likely to have a mixture of words of both sentiments as against tweets with literal sentiment (either positive or negative), our hierarchical topic model discovers sarcasm-prevalent topics and topic-level sentiment. Using a dataset of tweets labeled using hashtags, the model estimates topic-level, and sentiment-level distributions. Our evaluation shows that topics such as ‘work’, ‘gun laws’, ‘weather’ are sarcasm-prevalent topics. Our model is also able to discover the mixture of sentiment-bearing words that exist in a text of a given sentiment-related label. Finally, we apply our model to predict sarcasm in tweets. We outperform two prior work based on statistical classifiers with specific features, by around 25%.

1 Introduction

Sarcasm detection is the computational task of predicting sarcasm in text. Past approaches in sarcasm detection rely on designing classifiers with specific features (to capture sentiment changes or incorporate context about the author, environment, etc.) (Joshi et al., 2015; Wallace et al., 2014; Rajadesingan et al., 2015; Bamman and Smith, 2015), or model conversations using the sequence labeling-based approach by Joshi et al. (2016c). Approaches, in addition to this statistical classifier-based paradigm are: deep learning-based approaches as in the case of Silvio Amir et al. (2016) or rule-based approaches such as Riloff et al. (2013; Khattri et al. (2015)).

This work *employs a machine learning technique that, to the best of our knowledge, has not been used for computational sarcasm. Specifically, we introduce a topic model for extraction of sarcasm-prevalent topics and as a result, for sarcasm detection.* Our model based on a supervised version of the Latent Dirichlet Allocation (LDA) model (Blei et al., 2003) is able to discover clusters of words that correspond to sarcastic topics. The goal of this work is to discover sarcasm-prevalent topics based on sentiment distribution within text, and use these topics to improve sarcasm detection. The key idea of the model is that (a) some topics are more likely to be sarcastic than others, and (b) sarcastic tweets are likely to have a different distribution of positive-negative words as compared to literal positive or negative tweets. Hence, distribution of sentiment in a tweet is the central component of our model.

Our sarcasm topic model is learned on tweets that are labeled with three sentiment labels: literal positive, literal negative and sarcastic. In order to extract sarcasm-prevalent topics, the model uses three latent variables: a topic variable to indicate words that are prevalent in sarcastic discussions, a sentiment variable for sentiment-bearing words related to a topic, and a switch variable that switches between the two kinds of words (topic and sentiment-bearing words). Using a dataset of 166,955 tweets, our model is able to discover words corresponding to topics that are found in our corpus of positive, negative and sarcastic tweets.

We evaluate our model in two steps: a **qualitative evaluation** that ascertains sarcasm-prevalent topics based on the ones extracted, and a **quantitative evaluation** that evaluates sub-components of the model. We also demonstrate how it can be used for sarcasm detection. To do so, we compare our model with two prior work, and observe a significant improvement of around 25% in the F-score.

The rest of the paper is organized as follows. Section 2 discusses the related work. Section 3 presents our motivation for using topic models for automatic sarcasm detection. Section 4 describes the design rationale and structure of our model. Section 5 describes the dataset and the experiment setup. Section 6 discusses the results in three steps: qualitative results, quantitative results and application of our topic model to sarcasm detection. Section 7 concludes the paper and points to future work.

2 Related Work

Topic models are popular for sentiment aspect extraction. Jo and Oh (2011) present an aspect-sentiment unification model that learns different aspects of a product, and the words that are used to express sentiment towards the aspects. In terms of using two latent variables: one for aspect and one for sentiment, they are related to our model. Mukherjee and Liu (2012a) use a semi-supervised model in order to extract aspect-level sentiment. The role of the supervised sentiment label in our model is similar to their work. Finally, McAuley and Leskovec (2013a) attempt to generate rating dimensions of products using topic models. However, the topic models that have been reported in past work have been for sentiment analysis in general. They do not have any special consideration to either sarcasm as a label or sarcastic tweets as a special case of tweets. The hierarchy-based structure (specifically, the chain of distributions for sentiment label) in our model is based on Joshi et al. (2016a) who extract politically relevant topics from a dataset of political tweets. The chain in their case is sentiment distribution of an individual and a group.

Sarcasm detection approaches have also been reported in the past (Joshi et al., 2016b; Liebrecht et al., 2013; Wang et al., 2015; Joshi et al., 2015). Wang et al. (2015) present a contextual model for sarcasm detection that collectively models a set of tweets, using a sequence labeling algorithm - however, the goal is to detect sarcasm in the last tweet in the sequence. The idea of distribution of sentiment that we use in our model is based on the idea of context incongruity. In order to evaluate the benefit of our model to sarcasm detection, we compare two sarcasm detection approaches based on our model with two prior work, namely by Buschmeier et al. (2014) and Liebrecht et al. (2013). Buschmeier et al. (2014) train their classifiers using features such as unigrams, laughter expressions, hyperbolic expressions, etc. Liebrecht et al. (2013) experiment with unigrams, bigrams and trigrams as features. To the best of our knowledge, past approaches for sarcasm detection do not use topic modeling, which we do.

3 Motivation

Topic models enable discovery of thematic structures in a large-sized corpus. The motivation behind using topic models for sarcasm detection arises from two reasons: (a) presence of sarcasm-prevalent topics, and (b) differences in sentiment distribution in sarcastic and non-sarcastic text. In context of sentiment analysis, topic models have been used for aspect-based sentiment analysis in order to discover topic and sentiment words (Jo and Oh, 2011). The general idea is that for a restaurant review, the word ‘spicy’ is more likely to describe food as against ambiance. On similar lines, the discovery that a set of words belong to a sarcasm-prevalent topic - a topic regarding which sarcastic remarks are common - can be useful as additional information to a sarcasm detection system. The key idea of our approach is that some topics are more likely to evoke sarcasm than some others. For example, a tweet about working late night at office/ doing school homework till late night is much more probable to be sarcastic than a tweet on Mother’s Day. A sarcasm detection system can benefit from incorporating this information about sarcasm-prevalent topics. The second reason is the difference in sentiment distributions. A positive tweet is likely to contain only positive words, a negative tweet is likely to contain only negative words. On the other hand, a sarcastic tweet may contain a mix of the two kind of words (for example, ‘*I love being ignored*’ where ‘*love*’ is a positive word and ‘*ignored*’ is a negative word), except in the case of hyperbolic sarcasm (for example ‘*This is the best movie ever!*’ where ‘*best*’ is a positive word and there is no negative word). Hence, in addition to sarcasm-prevalent topics, sentiment distributions for tweets also form a critical component of our topic model.

Observed Variables and Distributions

w Word in a tweet
 l Label of a tweet; takes values: positive, negative, sarcastic)

Distributions

η_w Distribution over switch values given a word w

Latent Variables and Distributions

z Topic of a tweet
 s Sentiment of a word in a tweet; takes values: positive, negative
 is Switch variable indicating whether a word is a topic word or a sentiment word; takes values: 0, 1

Distributions

θ_l Distribution over topics given a label l , with prior α
 ϕ_z Distribution over words given a topic z and switch =0 (topic word), with prior γ
 χ_s Distribution over words given sentiment s and switch=1 (sentiment word), with prior δ_1
 χ_{sz} Distribution over words given a sentiment s and topic z and switch=1 (sentiment word), with prior δ_2
 ψ_l Distribution over sentiment given a label l and switch =1 (sentiment word), with prior β_1
 ψ_{zl} Distribution over sentiment given a label l and topic z and switch =1 (sentiment word), with prior β_2

Table 1: Glossary of Variables/Distributions used

4 Sarcasm Topic Model

4.1 Design Rationale

Our topic model requires sentiment labels of tweets, as used in Ramage et al. (2009). This sentiment can be positive or negative. However, in order to incorporate sarcasm, we re-organize the two sentiment values into three: literal positive, literal negative and sarcastic. The observed variable l in our model indicates this sentiment label. *For sake of simplicity, we refer to the three values of l as positive, negative and sarcastic, in rest of the paper.*

Every word w in a tweet is either a topic word or a sentiment word. A topic word arises due to a topic, whereas a sentiment word arises due to combination of topic and sentiment. This notion is common to several sentiment-based topic models from past work (Jo and Oh, 2011). To determine which of the two (topic or sentiment word) a given word is, our model uses three latent variables: a tweet-level topic label z , a word-level sentiment label s , and a switch variable is . Each tweet is assumed to have a single topic indicated by z . The single-topic assumption is reasonable considering the length of a tweet. At the word level, we introduce two variables is and s . For each word in the dictionary, is denotes the probability of the word being a topic word or a sentiment word. Thus, the model estimates three sets of distributions: (A) Probability of a word belonging to topic (ϕ_z) or sentiment-topic combination (χ_{sz}), (B) Sentiment distributions over label and topic (ψ_{zl}), and (C) Topic distributions over label (θ_l). The switch variable is is sampled from η_w , the probability of the word being a topic word or a sentiment word. We thus allow a word to be either a topic word or a sentiment word.¹

4.2 Plate Diagram

Our sarcasm topic model to extract sarcasm-prevalent topics is based on supervised LDA (Blei et al., 2003). Figure 1 shows the plate diagram while Table 1 details the variables and distributions in the model. Every tweet consists of a set of observed words w and one tweet-level, observed sentiment label l . The label takes **three** values: positive, negative or sarcastic. The third label value ‘sarcastic’ indicates a scenario where a tweet appears positive on the surface but is implicitly negative (hence, sarcastic). z is a tweet-level latent variable, denoting the topic of the tweet. The number of topics, Z is experimentally determined. is is a word-level latent variable representing if a word is a topic word or a sentiment word,

¹Note that η_w is not estimated during the sampling but learned from a large-scale corpus, as will be described later.

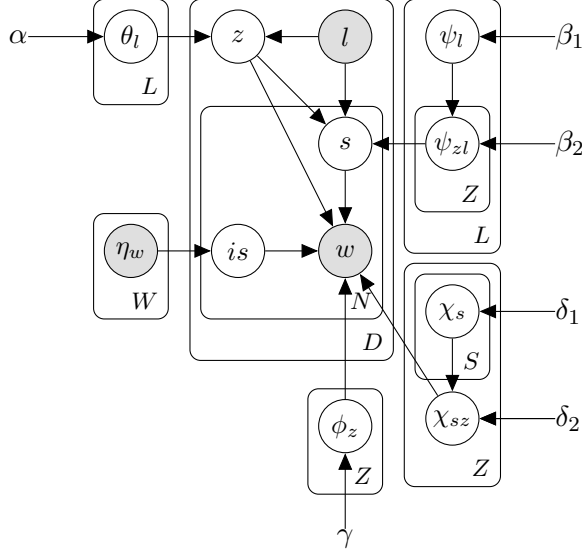


Figure 1: Plate Diagram of Sarcasm Topic Model

similar to Mukherjee and Liu (2012c). If the word is a sentiment word, the word-level latent variable s represents the sentiment of that word. It can take S unique values. Intuitively, S is set as 2.

Among the distributions, η_w is an observed distribution that is estimated beforehand. It denotes the probability of the word w being a topic word or a sentiment word. Distribution θ_l represents the distribution over z given the label of the tweet as l . ψ_l and ψ_{zl} are an hierarchical pair of distributions. ψ_{zl} represents the distribution over sentiment of the word given the topic and label of the tweet and that the word is a sentiment word. χ_s and χ_{sz} are an hierarchical pair of distributions, where χ_{sz} represents distribution over words, given the word is a sentiment word with sentiment s and topic z . ϕ_z is a distribution over words given the word is an topic word with topic z . The generative story of our model is:

1. For each label l , select $\vec{\theta}_l \sim \text{Dir}(\alpha)$
2. For each label l , select $\vec{\psi}_l \sim \text{Dir}(\beta_1)$
For each topic z , select $\vec{\psi}_{l,z} \sim \text{Dir}(\beta_2 \vec{\psi}_l)$
3. For each topic z and sentiment s , select $\vec{\chi}_s \sim \text{Dir}(\delta_1)$, and $\vec{\chi}_{s,z} \sim \text{Dir}(\delta_2 \vec{\chi}_s)$
4. For each topic z select $\vec{\phi}_z \sim \text{Dir}(\gamma)$
5. For each tweet k select
 - (a) topic $z_k \sim \vec{\theta}_{l_k}$
 - (b) switch value for all words, $i_{s_{k,j}} \sim \vec{\eta}_j$
 - (c) sentiment for all sentiment words, $s_{k,j} \sim \vec{\psi}_{z_k, l_k}$
 - (d) all topic words, $w_{k,j} \sim \vec{\phi}_{z_k}$
 - (e) all sentiment words, $w_{k,j} \sim \vec{\chi}_{s_{k,j}, z_k}$

We estimate these distribution using Gibbs sampling. The joint probability over all variables is decomposed into these distributions, based on dependencies in the model. Estimation details have not been included due to lack of space.

5 Experiment Setup

We create a dataset of English tweets for our topic model. We do not use datasets reported in past work (related to classifiers) because topic models typically require larger datasets than classifiers. The tweets are downloaded from twitter using the twitter API² using hashtag-based supervision. Hashtag-based supervision is common in sarcasm-labeled datasets (Joshi et al., 2015). Tweets containing hashtags #happy, #excited are labeled as positive tweets. Tweets with #sad, #angry are labeled as negative tweets. Tweets with #sarcasm and #sarcastic are labeled as sarcastic tweets. The tweets are converted to lowercase, and the hashtags used for supervision are removed. Function words³, punctuation, hashtags, author names

²<https://dev.twitter.com/rest/public>

³www.sequencepublishing.com

and hyperlinks are removed from the tweets. Duplicate tweets (same tweet text repeated for multiple tweets) and re-tweets (tweet text with the ‘RT’ added in the beginning) are discarded. Finally, words which occur less than three times in the vocabulary are also removed. As a result, the tweets that have less than 3 words are removed. This results in a dataset of 166,955 tweets. Out of these, 70,934 are positive, 20,253 are negative and the remaining 75,769 are sarcastic. A total of 35398 tweets are used for testing, out of which 26,210 are of positive sentiment, 5535 are of negative sentiment and 3653 are sarcastic. We repeat that these labels are determined based on hashtags, as stated above.

The total number of distinct labels (L) is 3, and the total number of distinct sentiment (S) is 2. The total number of distinct topics (Z) is experimentally determined as 50. We use block-based Gibbs sampling to estimate the distributions. The block-based sampler samples all latent variables together based on their joint distributions. We set asymmetric priors based on sentiment word-list from McAuley and Leskovec (2013b).

A key parameter of the model is η_w since it drives the split of a word as a topic or a sentiment word. SentiWordNet (Baccianella et al., 2010) is used to learn the distribution η_w prior to estimating the model. We average across multiple senses of a word. Based on the SentiWordNet scores to all senses of a word, we determine this probability.

6 Results

Work	Party	Jokes	Weather
day	life	Quote	Snow
morning	friends	Jokes	Today
night	night	Humor	Rain
today	drunk	Comedy	Weather
work	parties	Satire	Day
Women	School	Love	Politics
Women	tomorrow	love	Ukraine
Wife	school	feeling	Russia
Compliment(s)	work	break-up	again
Fashion	morning	day/night	deeply
Love	night	sleep	raiders

Table 2: Topics estimated when the topic model is learned on only sarcastic tweets

6.1 Qualitative Evaluation

The goal of this section is to present topics discovered by our sarcasm topic model. We do so in two steps. We first describe the topics generated when only sarcastic tweets from our corpus are used to estimate the distributions, followed by the ones when the full corpus is used. In case of the former, since only sarcastic tweets are used, the topics generated here indicate words corresponding to sarcasm-prevalent topics. In case of the latter, the sentiment-topic distributions in the model capture the prevalence of sarcasm.

The model estimates the ϕ and χ distributions corresponding to topic words and sentiment words. Top five words for a subset of topics (as estimated by ϕ) are shown in Table 2. The headings in boldface are manually assigned⁴. Sarcasm-prevalent topics, as discovered by our topic model, are work, party, weather, women, etc. The corresponding sentiment topics for each of these sarcasm-prevalent topics (as estimated by χ) are given in Table 3. The headings in boldface are manually assigned. For topics corresponding to ‘party’ and ‘women’, we observe that the two columns contain words from opposing sentiment polarities. An example sarcastic tweet about work is ‘Yaay! Another night spent at office! I love working late night’.

The previous set of topics are all from sarcastic text. We now show the topics extracted by our model from the full corpus. These topics will indicate peculiarity of topics for each of the three labels, allowing

⁴This is a common practice in topics model papers, in order to interpret topics. (Mukherjee and Liu, 2012b; Joshi et al., 2016a; Kim et al., 2013)

Work		Party		Jokes		Weather	
Love	Great	Lol	Hate	Funny	Lol	Love	nice
Good	Sick	Attractive	Allergic	Liar	Fucks	Glad	wow
Awesome	Seriously	Love	Insulting	Hilarious	Like	Fun	really
Women		School		Love		Politics	
Compliment(s)	Talents	excited	fun	best	love	Losing	issues
Thrilled	Sorry	love	omg	awesome	ignored	lies	weep
Recognized	Bad	really	awesome	greatest	sick	like	really

Table 3: Sentiment-related topics estimated when the topic model is learned on only sarcastic tweets

Music	work/school	Orlando Incident	Holiday	Quotes	Food
pop	work	orlando	summer	quote(s)	food
country	sleep	shooting	weekend	morning	lunch
rock	night	prayers	holiday	inspiration	vegan
bluegrass	morning	families	friends	motivation	breakfast
beatles	school	victims	sun,beach	mind	cake
Stock(s)/Commodities	Father	Gun	Pets	Health	
silver	father(s)	gun(s)	dog	fitness	
gold	dad	orlando	cat	gym	
index	daddy	trump	baby	run	
price	family	shooting	puppy	morning	
consumer	work	muslim	pets	health	

Table 4: Topics estimated when the topic model is learned on full corpus

us to infer what topics are sarcasm-prevalent. Table 4 shows the top 5 topic words for the topics discovered (as estimated in ϕ) from the full corpus (*i.e.*, containing tweets of all three tweet-level sentiment labels: positive, negative and sarcastic). Table 5 shows the top 3 sentiment words for each sentiment (as estimated by χ) of each of the topics discovered. Like in the previous case, the heading in boldface is manually assigned. One of the topic discovered was ‘Music’. The top 5 topic words for the topic ‘Music’ are Pop, Country, Rock, Bluegrass and Beatles. The corresponding sentiment words for Music are ‘love’, ‘happy’, ‘good’ on the positive side and ‘sad’, ‘passion’ and ‘pain’ on the negative side.

The **remaining sections present results when the model is learned on the full corpus.**

6.2 Quantitative Evaluation

In this section, we answer three questions: (A) What is the likely sentiment label, if a user is talking about a particular topic? (Section 6.2.1), (B) We hypothesize that sarcastic text tends to have mixed-polarity




Figure 2: Distribution of word-level sentiment labels for tweet-level labels

Stock(s)/ Commodities		Father		Gun		Pets		Health		Music	
gains	risks	happy	lol	like	sad	happy	small	love	tired	love	sad
happiness	dipped	love	little	good	hate	love	sad	fun	sick	happy	passion
unchanged	down	best	bless	wow	angry	cute	miss	laugh	unfit	good	pain
Work/School		Orlando Incident		Holiday		Quotes		Food			
great	sick	love	tragedy	love	beauty	positive	simple	happy	foodie		
fun	hate	like	hate	smile	hot	happy	kind	yummy	seriously		
yay	ugh	want	heartbroken	fun	sexy	happiness	sad	healthy	perfect		

Table 5: Sentiment-related topics estimated when the topic model is learned on full corpus

Topics P(l/z)	Positive	Negative	Sarcastic
Holiday	0.9538	0.0140	0.0317
Father	0.9224	0.0188	0.0584
Quote	0.8782	0.0363	0.0852
Food	0.8100	0.0331	0.1566
Music	0.7895	0.0743	0.1363
Fitness	0.7622	0.0431	0.1948
Orlando Incident	0.0130	0.9500	0.0379
Gun	0.1688	0.3074	0.5230
Work	0.1089	0.0354	0.8554
Humor	0.0753	0.1397	0.7841

Table 6: Probability of sentiment label for various discovered topics

words. Does it hold in case of our model? (Section 6.2.2), and (C) How can sarcasm topic model be used for sarcasm detection? (Section 6.2.3).

6.2.1 Probability of sentiment label, given topic

We compute the probability $p(l/z)$ based on the model. Table 6 shows these values for a subset of topics. Topics with a majority positive sentiment are Father’s Day (0.9224), holidays (0.9538), etc. The topic with the highest probability of a negative sentiment is the Orlando shooting incident (0.95). Gun laws (0.5230), work and humor are where sarcasm is prevalent.

6.2.2 Distribution of sentiment words for tweet-level sentiment labels

Figure 2 shows the proportion of word-level sentiment labels, for the three tweet-level sentiment labels, **as estimated by our model**. The X-axis indicates percentage of positive sentiment words in a tweet, while Y-axis indicates percentage of tweets which indicate a specific value of percentage. More than 60% negative tweets (bar in red) have 0% positive content words. The ‘positive’ here indicates the value of s for a word in a tweet. In other words, the said red bar indicates that 60% tweets have 0% words sampled with s as positive.

It follows intuition that negative tweets have low percentage of positive words (red bar on the left part of the graph) while positive tweets have high percentage of positive words (blue bar on the right part of the graph). The interesting variations are observed in case of sarcastic tweets. It must be highlighted that *the sentiment labels considered for these proportions are as estimated by our topic model*. Many sarcastic tweets contain very high percentage of positive sentiment words. Similarly, the proportion of tweets with around 50% positive sentiment words is around 20%, as expected. Thus, the model is able to capture the sentiment mixture as expected in the three tweet-level sentiment labels: (literal) positive, (literal) negative and sarcastic.

6.2.3 Application to Sarcasm Detection

We now use our sarcasm topic model to detect sarcasm, and compare it with two prior work. The task here is to classify a tweets as either sarcastic or not. We use the topic model for sarcasm detection using two methods:

1. **Log-likelihood based:** The topic model is first learned using the training corpus where the distributions in the model are estimated. Then, the topic model performs sampling for a pre-determined number of samples, in three runs - once for each label. For each run, the log-likelihood of the tweet given the estimated distributions (in the training phase) and the sampled values of the latent variables (for this tweet) is computed. The label of the tweet is returned as the one with the highest log-likelihood.
2. **Sampling-based:** Like in the previous case, the topic model first estimates distributions using the training corpus. Then, the topic model is learned again where the label l is assumed to be latent, in addition to the tweet-level latent variable z , and word-level latent variables s , and is . The value of l as learned by the sampler is returned as the predicted label.

We compare our results with two previously existing techniques, Buschmeier et al. (2014) and Liebrecht et al. (2013). We ensure that our implementations result in performance comparable to the reported papers. The two rely on designing sarcasm-level features, and training classifiers for these features. For these classifiers, the positive and negative labels are combined as non-sarcastic. As stated above, the test set is separate from the training set. The results of these two past methods compared with the two based on topic models are shown in Table 7. The values are averaged over the two classes. Both prior work show poor F-score (around 18-19%) while our sampling based approach achieves the best F-score of 46.80%. The low values, in general, may be because our corpus is large in size, and is diverse in terms of the topics. Also, features in Liebrecht et al. (2013) are unigrams, bigrams and trigrams which may result in sparse features.

Approach	P (%)	R (%)	F (%)
(Buschmeier et al., 2014)	10.41	100.00	18.85
(Liebrecht et al., 2013)	11.03	99.88	19.86
Topic Model: Log Likelihood	46.40	46.56	46.48
Topic Model: Sampling	45.94	47.70	46.80

Table 7: Comparison of Various Approaches for Sarcasm Detection

7 Conclusion & Future Work

We presented a novel topic model that discovers sarcasm-prevalent topics. Our topic model uses a dataset of tweets (labeled as positive, negative and sarcastic), and estimates distributions corresponding to prevalence of a topic, prevalence of a sentiment-bearing words. We observed that topics such as work, weather, politics, etc. were discovered as sarcasm-prevalent topics. We evaluated the model in three steps: (a) Based on the distributions learned by our model, we show the most likely label, for all topics. This is to understand sarcasm-prevalence of topics when the model is learned on the full corpus. (b) We then show distribution of word-level sentiment for each tweet-level sentiment label as estimated by our model. *Our intuition that sentiment distribution in a tweet is different for the three labels: positive, negative and sarcastic, holds true.* (c) Finally, we show how topics from this topic model can be harnessed for sarcasm detection. We implement two approaches: one based on most likely label as per log likelihood, and another based on last sampled value during iteration. In both the cases, we are able to significantly outperform two prior work based on feature design by F-Score of around 25%.

The current model is limited because of its key intuition about sentiment mixture in sarcastic text. Instances such as hyperbolic sarcasm go against the intuition. The current approach relies only on bag of words which may be extended to n-grams since a lot of sarcasm is expressed through phrases with implied sentiment. This work, being an initial work in topic models for sarcasm, sets up the promise of topic models for sarcasm detection, as also demonstrated in corresponding work in aspect-based sentiment analysis.

References

- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *LREC*, volume 10, pages 2200–2204.
- David Bamman and Noah A Smith. 2015. Contextualized sarcasm detection on twitter. In *Ninth International AAAI Conference on Web and Social Media*.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Konstantin Buschmeier, Philipp Cimiano, and Roman Klinger. 2014. An impact analysis of features in a classification approach to irony detection in product reviews. In *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 42–49.
- Yohan Jo and Alice H Oh. 2011. Aspect and sentiment unification model for online review analysis. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 815–824. ACM.
- Aditya Joshi, Vinita Sharma, and Pushpak Bhattacharyya. 2015. Harnessing context incongruity for sarcasm detection. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, volume 2, pages 757–762.
- Aditya Joshi, Pushpak Bhattacharyya, and Mark Carman. 2016a. Political issue extraction model: A novel hierarchical topic model that uses tweets by political and non-political authors. In *7th Workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis*, pages 82–90.
- Aditya Joshi, Pushpak Bhattacharyya, and Mark James Carman. 2016b. Automatic sarcasm detection: A survey. *arXiv preprint arXiv:1602.03426*.
- Aditya Joshi, Vaibhav Tripathi, Pushpak Bhattacharyya, and Mark Carman. 2016c. Harnessing sequence labeling for sarcasm detection in dialogue from tv series ‘friends’. *CoNLL 2016*, page 146.
- Anupam Khattri, Aditya Joshi, Pushpak Bhattacharyya, and Mark James Carman. 2015. Your sentiment precedes you: Using an author’s historical tweets to predict sarcasm. In *6TH WORKSHOP ON COMPUTATIONAL APPROACHES TO SUBJECTIVITY, SENTIMENT AND SOCIAL MEDIA ANALYSIS WASSA 2015*, page 25.
- Suin Kim, Jianwen Zhang, Zheng Chen, Alice H Oh, and Shixia Liu. 2013. A hierarchical aspect-sentiment model for online reviews. In *AAAI*.
- CC Liebrecht, FA Kunneman, and APJ van den Bosch. 2013. The perfect solution for detecting sarcasm in tweets# not.
- Julian McAuley and Jure Leskovec. 2013a. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 165–172. ACM.
- Julian John McAuley and Jure Leskovec. 2013b. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. In *Proceedings of the 22nd international conference on World Wide Web*, pages 897–908. International World Wide Web Conferences Steering Committee.
- Arjun Mukherjee and Bing Liu. 2012a. Aspect extraction through semi-supervised modeling. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 339–348. Association for Computational Linguistics.
- Arjun Mukherjee and Bing Liu. 2012b. Aspect extraction through semi-supervised modeling. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 339–348. Association for Computational Linguistics.
- Arjun Mukherjee and Bing Liu. 2012c. Modeling review comments. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 320–329. Association for Computational Linguistics.
- Ashwin Rajadesingan, Reza Zafarani, and Huan Liu. 2015. Sarcasm detection on twitter: A behavioral modeling approach. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, WSDM ’15*, pages 97–106, New York, NY, USA. ACM.
- Daniel Ramage, David Hall, Ramesh Nallapati, and Christopher D. Manning. 2009. Labeled lda: A supervised topic model for credit attribution in multi-labeled corpora. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1 - Volume 1, EMNLP ’09*, pages 248–256, Stroudsburg, PA, USA. Association for Computational Linguistics.

- Ellen Riloff, Ashequl Qadir, Prafulla Surve, Lalindra De Silva, Nathan Gilbert, and Ruihong Huang. 2013. Sarcasm as contrast between a positive sentiment and negative situation. In *EMNLP*, volume 13, pages 704–714.
- Byron C Silvio Amir, Wallace, Hao Lyu, and Paula Carvalho Mário J Silva. 2016. Modelling context with user embeddings for sarcasm detection in social media. *CoNLL 2016*, page 167.
- Byron C. Wallace, Do Kook Choe, Laura Kertz, and Eugene Charniak. 2014. Humans require context to infer ironic intent (so computers probably do, too). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 512–516, Baltimore, Maryland, June. Association for Computational Linguistics.
- Zelin Wang, Zhijian Wu, Ruimin Wang, and Yafeng Ren. 2015. Twitter sarcasm detection exploiting a context-based model. In *Web Information Systems Engineering–WISE 2015*, pages 77–91. Springer.