

Analyzing users' sentiment towards popular consumer industries and brands on Twitter

Guoning Hu, Preeti Bhargava, Saul Fuhrmann, Sarah Ellinger and Nemanja Spasojevic
Lithium Technologies | Klout, San Francisco, CA

Email: {guoning.hu, preeti.bhargava, saul.fuhrmann, sarah.ellinger, nemanja.spasojevic}@lithium.com

Abstract—Social media serves as a unified platform for users to express their thoughts on subjects ranging from their daily lives to their opinion on consumer brands and products. These users wield an enormous influence in shaping the opinions of other consumers and influence brand perception, brand loyalty and brand advocacy. In this paper, we analyze the opinion of 19M Twitter users towards 62 popular industries, encompassing 12,898 enterprise and consumer brands, as well as associated subject matter topics, via sentiment analysis of 330M tweets over a period spanning a month. We find that users tend to be most positive towards manufacturing and most negative towards service industries. In addition, they tend to be more positive or negative when interacting with brands than generally on Twitter. We also find that sentiment towards brands within an industry varies greatly and we demonstrate this using two industries as use cases. In addition, we discover that there is no strong correlation between topic sentiments of different industries, demonstrating that topic sentiments are highly dependent on the context of the industry that they are mentioned in. We demonstrate the value of such an analysis in order to assess the impact of brands on social media. We hope that this initial study will prove valuable for both researchers and companies in understanding users' perception of industries, brands and associated topics and encourage more research in this field.

I. INTRODUCTION

Social media has become one of the major means of communication and content production. It serves as a unified platform for users to express their thoughts on subjects ranging from their daily lives to their opinion on companies and products. This, in turn, has made it a valuable resource for mining user opinion for tasks ranging from predicting the performance of movies to results of stock markets and elections.

Although most people are hesitant to answer surveys about products or services, they express their thoughts freely on social media and wield an enormous influence in shaping the opinions of other consumers. These consumer voices can influence brand perception, brand loyalty and brand advocacy. As a result, it is imperative that large enterprises pay more attention to mining user opinion related to their brands and products from social media. With social media monitoring, they will be able to tap into consumer insights to improve their product quality, provide better service, drive sales and even identify new business opportunities. In addition, they can reduce customer support costs by responding to their customers through these social media channels, as 50% of users prefer

reaching service providers on social media rather than a call center [1].

Sentiment analysis has become synonymous with opinion mining and involves the computational study of text in order to identify sentiment polarity (positive, negative or neutral), intensity and subjectivity [2]. It is an excellent tool for enterprises to analyze users' expressed opinions on social media without explicitly asking any questions as this approach often reflects their true opinions. Though it has drawbacks regarding the population sampled (as it may not represent the general public), it can be used to approximate public opinion.

In this paper, we analyze the opinion of Twitter users towards different industries, encompassing several global consumer brands, and associated subject matter topics via sentiment analysis of their tweets. Because twitter messages are short, it is quite challenging to identify the target that an opinion is expressed towards. As a preliminary study, here we take a special approach based on the observation that **when a user interacts with an industry or brand, such as a reply or a retweet, the sentiment in the related text likely shows this user's opinion toward that industry or brand.** Therefore, by aggregating such interactions from a large number of users, the sentiment distribution can reveal the general attitude about the industry or brand among twitter users.

Our dataset consists of 19M unique users, 62 popular industries (such as *Airlines*, *Automotive*, *Telecommunications* etc.), encompassing 12,898 enterprise and consumer brands¹ (such as *United*, *Volkswagen* etc.), 8030 subject matter topics and about 330M user-brand interactions over a period spanning a month. Specifically, we try to answer the following questions:

- What is the opinion of users towards different industries? What are the industries that are viewed negatively or positively?
- Within industries, what is the users' sentiment towards different brands?
- What are the frequent topics that users mention when they interact with brands over social media? What is their sentiment towards those topics?
- How does users' general behavior on social media differ from their behavior when they interact with brands?

Our main findings are:

¹*Brand* is here defined as an independently recognizable product line or business unit.

User brand interaction	Sentiment
Things I have learned today: we own way too many dishes, @comcast is literally the WORST, and moving is the main cause of drinking problems	Negative
Beautiful morning flying out of the Bay with @southwestair . #baytola #sfbay #family #fun #flying https://t.co/bnogVW2X8H	Positive
@united A measure of the harm inflicted by disrupting people’s travel plans: no one here felt \$800 came close to making them whole.	Negative
SUPER excited and grateful to be at the #smashbox event @menloparkmall with @sephora! Thanks for ... https://t.co/CMexZF3pLr	Positive

TABLE I: Examples of users’ interaction with a brand (highlighted) and the interaction’s sentiment

- There is a positive correlation between users’ general sentiment and their sentiments toward brands. Users who generally express one sentiment class (positive, negative or neutral) on social media tend to express that same sentiment toward brands.
- Among industries, the most negative industries tend to be those providing services (such as *Airlines*, *Mail and Shipping*, and *Telecommunications*) whereas the most positive industries tend to be those manufacturing and selling consumer goods such as *Household appliances*. We also find that sentiment towards brands within an industry varies greatly and we demonstrate this using two industries as use cases.
- Finally, we find that there is no strong correlation between topic sentiments of different industries, demonstrating that topic sentiments are highly dependent on the context of the industry that the topics are mentioned in.

To the best of our knowledge, no other work has attempted to analyze sentiment of users towards industries, consumer brands and topics at such a large scale.

II. RELATED WORK

Several recent papers have focused on analyzing social media sentiment for a variety of purposes such as public opinion mining, stock market prediction and prediction of election results. In this paper, we do not intend to cover all work on sentiment analysis of social media exhaustively. Instead, we focus on discussing a few relevant works.

Balasubramanyan et al. [3] found high correlation between public opinion (in terms of consumer confidence and political opinion), measured from polls, with sentiment measured from tweets. Hu et al. [4] analyzed public events by automatically characterizing segments and event topics in terms of the aggregate sentiments they elicited on Twitter. Bollen et al. [5] analyzed tweets to find correlation between overall public mood and social, economic and other major events. Si et al. [6] used the co-occurrences of stock tickers in tweets to propagate sentiment of stocks in order to predict stock values. Mittal and Goel [7] mined public mood from Twitter and used it along with previous days’ stock market values to predict future stock market movements. Ceron et al. [8] used sentiment analysis on tweets to understand the policy preferences of Italian and French citizens. Tumasjan et al. [9] analyzed the role of political sentiment and mentions of political parties in Tweets as indicators of election results.

In contrast to this rich landscape, very few studies have focused on investigating users’ sentiment towards

major worldwide brands. Adeborna and Siau [10] propose a sentiment topic recognition model based on correlated topics models with variational expectation-maximization algorithm using airline data from Twitter. They validate their approach on a dataset of 1046 tweets concerning three major airlines by computing their airline quality rating. Mostafa [11] used a random sample of 3516 tweets to evaluate consumers’ sentiment towards 16 global well-known brands.

Our work differs in the scale of the Twitter dataset size as well as the number of industries and brands we analyze. In addition, we look at users’ sentiment towards topics. Finally, we derive some meaningful insights from this analysis and demonstrate the value of such an analysis in order to assess the impact of brands on social media.

III. DATA SET

A. Industries and brands dataset

For our purposes, a brand is defined as an independently recognizable product line or business unit, often with its own logo, trademark, slogan, etc. Examples of brands in our dataset include *Cadbury*, *AAA*, and *23andme*. A brand may be synonymous with a given company or business, but we make no attempt to represent company ownership or hierarchies (e.g. *Fanta* is considered a separate brand from *Coca-Cola*, although both brands are owned by the corporate entity *The Coca-Cola Company*). Brands in our data set can contain one or more Twitter profiles and belong to one or more industries.

An industry is defined as a high-level category of brands that perform similar functions or provide similar products. In this analysis, we use a subset of 62 industries drawn from LinkedIn Industry Codes², the Fortune 500³, the Forbes Global 2000⁴, and Hoovers⁵. See Figure 2 for a list of included industries.

The brand dataset was initially bootstrapped by collecting companies’ names and industries from the LinkedIn profiles of Klout⁶ users. We then attempted to match the company names to Twitter handles via the Bing Search API. However, this approach had some limitations:

- LinkedIn included many companies that had shut down, changed names, etc.

²<http://developer.linkedin.com/docs/reference/industry-codes>

³<http://fortune.com/fortune500/list/>

⁴<http://www.forbes.com/global2000/>

⁵<http://www.hoovers.com/industry-analysis/industry-directory.html>

⁶<https://klout.com>

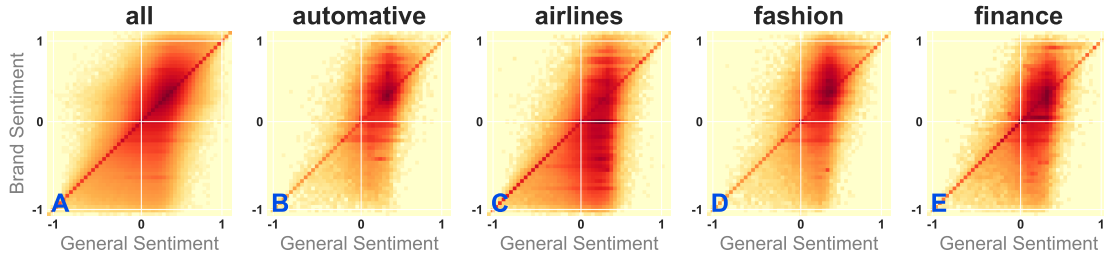


Fig. 1: Heatmap of users’ brand and general sentiments

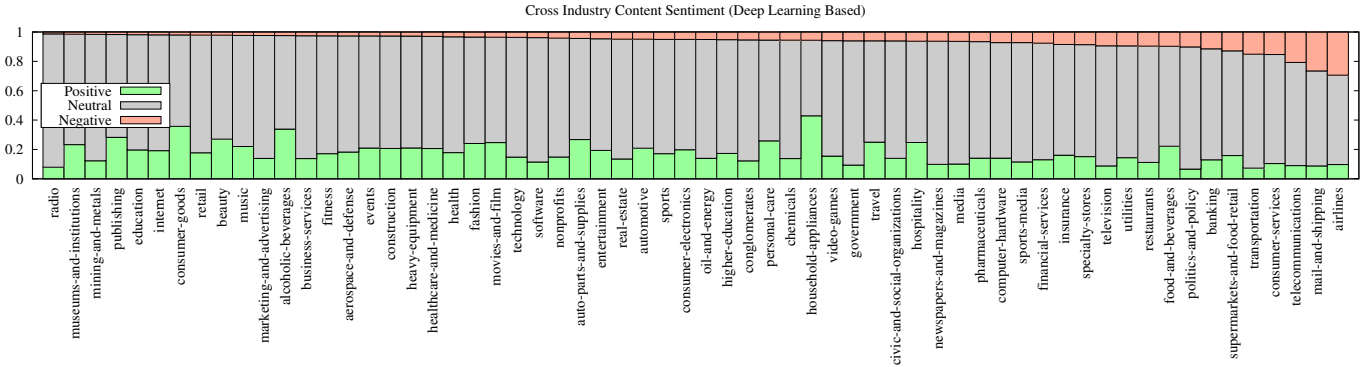


Fig. 2: Sentiment distribution by industry, ordered by negative sentiment (best viewed in color)




Fig. 3: Sentiment distribution by industry, ordered by positive sentiment (best viewed in color)

- Even with the industry name appended, the Bing Search API often returned non-brand Twitter accounts with similar names to the target brands.
- Bing Search API provided only one Twitter handle per brand, whereas many brands have several specialized handles (for customer service vs. corporate communications, in different languages, etc.).

Thus, in order to guarantee the quality of the data, we chose to manually curate the brand dataset. Candidate handles were collected both from the initial bootstrap data and by mining the Twitter Firehose stream⁷, looking at signals such as reply volume and average time to reply (which can indicate a handle staffed by professional social media agents), follower and mention volume, and the use of enterprise-level publishing and reply tools. Each candidate handle was reviewed to determine if it belonged to a brand. If yes, it was assigned a new or existing brand node and

⁷<http://support.gnip.com/apis/firehose/overview.html>

given one or more industry labels. At this time, the brand data set contains 12,898 brands and 16,137 handles.

B. Twitter dataset

For each brand, we collected *interactions* (including replies to tweets, mentions, retweets, and cc) of Twitter users with the brand for the month of April 2017, via the Twitter Firehose stream. Each tweet was processed through the Lithium NLP pipeline [12] to extract information such as the tweet language, named entities etc. For each Twitter user in our dataset, we collected all their English tweets for April 2017.

The complete dataset contains more than 19M unique users that interacted with the 12,898 brands. These users generated more than 330M user-brand interactions. Among these interactions, there are more than 56M unique tweets in English. Overall, these users posted about 282M unique tweets in English. Table I shows some examples of user brand interactions in our dataset along with their sentiment.

Cross Airline Industry Content Sentiment (Deep Learning Based)




Fig. 4: Sentiment distribution for Airlines, ordered by negative sentiment (best viewed in color)

Cross Automotive Industry Content Sentiment (Deep Learning Based)




Fig. 5: Sentiment distribution for Automotive, ordered by positive sentiment (best viewed in color)

C. Topics dataset

Each tweet in our dataset is associated with subject matter topics that are most relevant to the tweet text. For instance, the tweet “Anti-net neutrality bots are swarming the FCC’s comments <http://engt.co/2q4Kd7l>” gets associated with topics such as *Net Neutrality*, *Federal Communications Commission* and *Telecommunications* etc. These topics belong to the Klout Topic Ontology⁸ (KTO) [13] - a manually curated ontology built to capture social media users’ interests and expertise scores, in different topics, across multiple social networks. As of April 2017, it consists of roughly 8,030 topic nodes and 13,441 edges encoding hierarchical relationships among them. On Klout, we use these topics to model users’ interests [14] and expertise [15] by building their topical profiles. These user topics are also available via the GNIP PowerTrack API⁹.

The topics are inferred by mapping the recognized entities in a tweet to KTO topics via the Lithium NLP pipeline [12] which uses a weighted ensemble of several semi-supervised learning models that employ entity co-

occurrences, word embeddings, Freebase hierarchical relationships and Wikipedia.

IV. METHOD

To analyze the sentiment of a tweet and label it as one of the 3 sentiment classes (‘Positive’, ‘Negative’ or ‘Neutral’), we used a Deep Learning based approach as this field has recently enabled strong advances in NLP [16]. In particular, we use Recurrent Neural Networks (RNNs) which are a class of neural networks that are able to model relationships in time. Unlike traditional neural networks, RNNs use units with internal states that can persist information about previous events. This makes them suitable for problems that require sequential information, such as text processing tasks. We also experimented with a lexicon based approach but found deep learning to be more accurate given empirical evidence ([17], [18]).

In our implementation, we use a conventional Long Short Term Memory (LSTM) network [17]. We first pre-process the text by tokenizing and normalizing it. We then feed this text to the LSTM, and use the labels to perform supervised learning on the messages. We use a vocabulary of 40k words, 128 embeddings units, 32 units, and a

⁸https://github.com/klout/opendata/tree/master/klout_topic_ontology

⁹<http://support.gnip.com/enrichments/klout.html>




Fig. 8: Topic Sentiment across different industries (best viewed in color)

negative sentiment. Figure 3 shows the same data except that industries are sorted by positive sentiment.

Clearly, sentiment distributions vary significantly across different industries. Interestingly, the most negative industries tend to be those providing services to customers, such as *Airlines*, *Mail and Shipping*, and *Telecommunications*, whereas the most positive industries tend to be those manufacturing and selling consumer goods, such as *Household appliances*. Additionally, the most polarized industries are surprisingly not *Politics* and *Sports* (as one might expect).

We also analyzed the sentiment within specific industries¹⁰, taking *Airlines* (which has the highest negative sentiment) and *Automotive* (which has a high positive sentiment) as use cases. Figure 4 shows the sentiment distribution of individual brands in the airline industry, sorted by negative sentiment. Clearly, sentiment distribution varies significantly across different brands. As shown, users are quite negative towards *Spirit*, *Delta*, and *United*, while quite positive towards *Air-Mauritius*, *Sun-Country*, and *Thomas Cook*.

Figure 5 show the sentiment distribution of individual brands in the *Automotive* industry. Again, we observe significant variation across different brands. The most positively viewed brands include luxury brands such as *Aston Martin*, *Lamborghini* and *BMW*, while the most negatively viewed brands include rental car companies (such as *Enterprise* and *National*), *Hyundai* and *Volkswagen*.

In addition, we noticed that brands with a higher volume of interactions tend to have more neutral sentiment. As an illustration, Figure 6 shows the average user-brand interactions across industries. Among the top 10 industries with the most interactions, 4 of them are among the top 10 industries with the most neutral sentiment, and 8 of them are among either the bottom 10 industries with the most positive or negative sentiment. We suspect this is because

these industries get a lot of interactions and coverage from media, whose messages generally tend to be neutral.

C. Topic sentiment

We also analyzed the sentiment towards topics associated with users’ brand interaction tweets. Here we assume that when users interact with brands belonging to an industry, they will mention certain topics in the tweet. By assigning the tweet’s sentiment score to these topics, we can roughly estimate the average sentiment of those topic within those industries. Figure 7 shows the word cloud representing the average sentiment of the 100 most frequent topics in users’ interactions across all industries. As shown, highly negative topics include *Fox News* while highly positive topics include *Video Games* and *Music*.

Figure 8 shows the average sentiment of the 20 most frequent topics (in users’ interactions across all industries) within each industry, sorted by the average sentiment. It depicts a weak correlation between industry and topic sentiment. But there are many exceptions to this, such as topic *CNN* in industry *Heavy Equipment* and topic *Russia* in industry *Entertainment*. Moreover, average sentiments towards these topics vary significantly across different industries and are highly dependent on the context of the industry that the topics are mentioned in. For instance, from Figure 7 we can infer that *Fox News* is a highly negative topic and we might expect it to be negative across all industries. However, we can see in Figure 8 that this topic actually has positive sentiment in industries such as *Events*, *Music* and *Sports*.

Figure 9 shows the correlation between industries represented by their topic sentiment vectors. Each industry is represented as a 100 dimensional vector of the most frequent topics (those shown in Figure 7) and the value of each dimension represents the average sentiment towards that topic within that industry. We centered these sentiment values around zero in order to mitigate any bias in the data. We then computed cosine similarity between

¹⁰Due to lack of space, it is not possible to show sentiment towards brands in all the 62 industries




Fig. 9: Relationship between Industries represented by topic sentiment vectors (best viewed in color)

the vectors. Overall, there is no strong correlation between topic sentiments of different industries, demonstrating that topic sentiments are highly industry-dependent.

VI. LIMITATIONS

While we have attempted to cover as much ground as possible with respect to industries, brands and topics as well as users' interactions and sentiment towards them, we note that this study has some limitations:

- We limit our dataset to Twitter alone since it is one of the most widely used platforms for self expression.
- Our work mainly focuses on the most popular industry and brands.
- Due to resource constraints, we analyzed Twitter data for only a month.
- The popular topics within industries are often based on current context and can change with time.
- This study analyzed sentiment only at the message level. Accurate sentiment analysis at the aspect level will lead to more insightful results.

VII. APPLICATIONS

As mentioned earlier, these insights about users' sentiment towards consumer brands can help brands provide

better online customer service, assess the impact of social media marketing campaigns or public relations incidents, identify perceived flaws in their products, and even identify new sales or business opportunities. They also help brands gauge the consumers' perception of them.

Another direct application of analyzing consumers' sentiment is to rank a brand relative to its competitors, allowing that brand to contextualize its performance on social media. This is easily demonstrated by the Social Impact Ranking tool¹¹ that Lithium has developed and which employs the sentiment of users' tweets, in addition to several other metrics, to assess the impact of different competing brands in various industries on social media. As an illustration, Figures 10 and 11 show screenshots of the impact ranking of airlines on different dates separated by 3 weeks. Figure 10 coincides with an unfortunate incident involving a passenger on United Airlines¹². As evident, United has the highest mentions as well as strong negative sentiment towards it on Twitter. However, the number of mentions have decreased 3 weeks later (Figure 11).

¹¹<https://klout.com/impact-ranking/>

¹²https://en.wikipedia.org/wiki/United_Express_Flight_3411_incident



Fig. 10: Airlines ranking retrieved on April 24, 2017

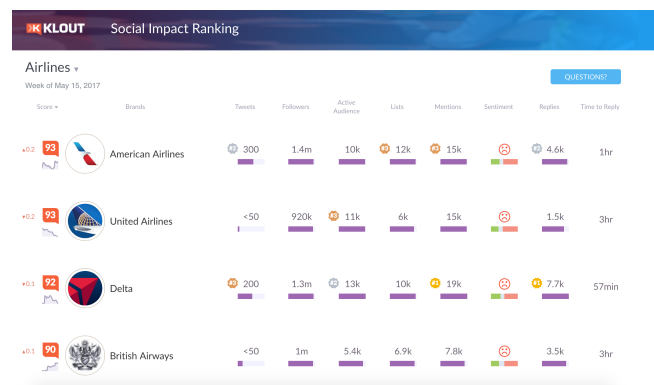


Fig. 11: Airlines ranking retrieved on May 15, 2017

Brands are very interested in understanding customers' attitudes towards them. There have been many demands on quantitative measures of overall attitude of users towards a brand as well as their underlying business components, such as product, website, support and customer service. Sentiment analysis of brands, topics and their relationship plays a key role in building such measures.

VIII. CONCLUSION AND FUTURE WORK

In this paper, we analyzed the opinion of 19M Twitter users towards 62 popular industries, encompassing 12,898 enterprise and consumer brands, as well as associated subject matter topics via sentiment analysis of 330M tweets over a period spanning a month. To the best of our knowledge, no other work has attempted to analyze sentiment of users towards industries, consumer brands and topics at such a large scale. We found that users tend to be most positive towards manufacturing and most negative towards service industries. In addition, they tend to be more positive or negative when interacting with brands than generally on Twitter. We also found that sentiment towards brands within an industry varies greatly. In addition, we discovered that there is no strong correlation between topic sentiments of different industries, demonstrating that topic sentiments are highly dependent on the context of the industry that they are mentioned in. We also demonstrated the value of such an analysis in order to assess the impact of brands on social media using several metrics.

This preliminary study acts as a first step towards understanding the users' perception of industries, brands and associated topics. We hope that our analysis and the insights that we derive from it will prove valuable for both researchers and companies in understanding user behavior and encourage more research in this field. In future, we will work on employing these datasets results for sentiment forecasting. We also plan to conduct surveys among Twitter users in order to validate the insights we have derived from our analysis.

REFERENCES

[1] A. Okeleke and S. Bali, "Integrating social media into CRM for next generation customer experience," Ovum, Tech. Rep., 05 2014.

[2] B. Pang, L. Lee *et al.*, "Opinion mining and sentiment analysis," *Foundations and Trends® in Information Retrieval*, vol. 2, no. 1–2, pp. 1–135, 2008.

[3] R. Balasubramanyan, B. R. Routledge, and N. A. Smith, "From tweets to polls: Linking text sentiment to public opinion time series." *ICWSM*, pp. 122 – 129, 2010.

[4] Y. Hu, F. Wang, and S. Kambhampati, "Listening to the crowd: Automated analysis of events via aggregated twitter sentiment." in *IJCAI*, 2013.

[5] J. Bollen, H. Mao, and A. Pepe, "Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena." *ICWSM*, vol. 11, pp. 450–453, 2011.

[6] J. Si, A. Mukherjee, B. Liu, S. J. Pan, Q. Li, and H. Li, "Exploiting social relations and sentiment for stock prediction." in *EMNLP*, vol. 14, 2014, pp. 1139–1145.

[7] A. Mittal and A. Goel, "Stock prediction using twitter sentiment analysis," *Stanford University, CS229* (<http://cs229.stanford.edu/proj2011/GoelMittal-StockMarketPredictionUsingTwitterSentimentAnalysis.pdf>).

[8] A. Ceron, L. Curini, S. M. Iacus, and G. Porro, "Every tweet counts? how sentiment analysis of social media can improve our knowledge of citizens' political preferences with an application to italy and france," *New Media & Society*, vol. 16, no. 2, pp. 340–358, 2014.

[9] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welpe, "Predicting elections with twitter: What 140 characters reveal about political sentiment." *ICWSM*, vol. 10, no. 1, pp. 178–185, 2010.

[10] E. Adeborna and K. Siau, "An approach to sentiment analysis-the case of airline quality rating." in *PACIS*, 2014, p. 363.

[11] M. M. Mostafa, "More than words: Social networks' text mining for consumer brand sentiments," *Expert Systems with Applications*, vol. 40, no. 10, pp. 4241–4251, 2013.

[12] P. Bhargava, N. Spasojevic, and G. Hu, "Lithium nlp: A system for rich information extraction from noisy user generated text on social media," in *Proc. of the 3rd Workshop on Noisy User-generated Text*, 2017, pp. 131–139.

[13] S. Ellinger, P. Bhattacharyya, P. Bhargava, and N. Spasojevic, "Klout topics for modeling interests and expertise of users across social networks," in *under review*.

[14] N. Spasojevic, J. Yan, A. Rao, and P. Bhattacharyya, "Lasta: Large scale topic assignment on multiple social networks," in *Proc. of ACM KDD*, 2014, pp. 1809–1818.

[15] N. Spasojevic, P. Bhattacharyya, and A. Rao, "Mining half a billion topical experts across multiple social networks," *Social Network Analysis and Mining*, vol. 6, no. 1, pp. 1–14, 2016.

[16] T. Young, D. Hazarika, S. Poria, and E. Cambria, "Recent trends in deep learning based natural language processing," in *arXiv preprint arXiv:1708.02709*, 2017.

[17] A. Rao and N. Spasojevic, "Actionable and political text classification using word embeddings and lstm," *arXiv preprint arXiv:1607.02501*, 2016.

[18] N. Spasojevic and A. Rao, "Identifying actionable messages on social media," in *IEEE International Conference on Big Data*, ser. IEEE BigData '15, 2015.