

Complaint Identification in Social Media with Transformer Networks

Mali Jin Nikolaos Aletras

Department of Computer Science, University of Sheffield
{mjin6, n.aletras}@sheffield.ac.uk

Abstract

Complaining is a speech act extensively used by humans to communicate a negative inconsistency between reality and expectations. Previous work on automatically identifying complaints in social media has focused on using feature-based and task-specific neural network models. Adapting state-of-the-art pre-trained neural language models and their combinations with other linguistic information from topics or sentiment for complaint prediction has yet to be explored. In this paper, we evaluate a battery of neural models underpinned by transformer networks which we subsequently combine with linguistic information. Experiments on a publicly available data set of complaints demonstrate that our models outperform previous state-of-the-art methods by a large margin achieving a macro F1 up to 87.

1 Introduction

Complaining is a basic speech act, usually triggered by a discrepancy between reality and expectations towards an entity or event (Olshtain and Weinbach, 1985; Cohen and Olshtain, 1993; Kowalski, 1996). Social media has become a popular platform for expressing complaints online (Preotiuc-Pietro et al., 2019) where customers can directly address companies regarding issues with services and products. Complaint detection aims to identify a breach of expectations in a given text snippet. However, the use of implicit and ironic expressions and accompaniment of other speech acts such as suggestions, criticism, warnings and threats (Pawar et al., 2015) make it a challenging task. Identifying and classifying complaints automatically is important for: (a) improving customer service chatbots (Coussement and Van den Poel, 2008; Lailiyah et al., 2017; Yang et al., 2019a); (b) linguists to analyze complaint characteristics on large scale (Vásquez, 2011; Kakolaki and Shahrokhi, 2016); and (c) psychologists to understand the behavior of humans that express complaints (Sparks and Browning, 2010).

Previous work has focused on binary classification between complaints and non-complaints in various domains (Preotiuc-Pietro et al., 2019; Jin et al., 2013; Coussement and Van den Poel, 2008). Furthermore, some studies have performed more fine-grained complaint classification. For instance, complaints directed to public authorities have been categorized based on their topics (Forster and Entrup, 2017; Merson and Mary, 2017) or the responsible departments (Laksana and Purwarianti, 2014; Gunawan et al., 2018; Tjandra et al., 2015). Other categorizations are based on possible hazards and risks (Bhat and Cullotta, 2017) as well as escalation likelihood (Yang et al., 2019a). Most of these previous studies have used supervised machine learning models with features extracted from text (e.g. bag-of-words, topics, features extracted from psycho-linguistic dictionaries) or task-specific neural models trained from scratch. Adapting state-of-the-art pre-trained neural language models based on transformer networks (Vaswani et al., 2017) such as BERT (Devlin et al., 2018) and XLNet (Yang et al., 2019b) has yet to be explored.

In this paper, we focus on the binary classification of Twitter posts into complaints or not (2019). We adapt and evaluate a battery of pre-trained transformers which we subsequently combine with external linguistic information from topics and emotions.

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

| Domain | Complaints | Non-complaints |
|-------------|------------|----------------|
| Food | 95 | 35 |
| Apparel | 141 | 117 |
| Retail | 124 | 75 |
| Cars | 67 | 25 |
| Services | 207 | 130 |
| Software | 189 | 103 |
| Transport | 139 | 109 |
| Electronics | 174 | 112 |
| Other | 96 | 33 |
| Total | 1232 | 739 |

Table 1: Data set statistics across 9 domains.

Contributions (1) New state-of-the-art results on complaint identification in Twitter, improving macro FI by 8.0% over previous work by Preotiuc-Pietro et al. (2019); (2) A qualitative analysis of the limitations of transformers in predicting accurately whether a given text is a complaint or not.

2 Complaint Prediction Task and Data

Given a text snippet (i.e. tweet), we aim to classify it as a complaint or not. For that purpose, we use the data set¹ by Preotiuc-Pietro et al. (2019) which contains tweets written in English that were manually annotated as complaints or not. It includes 1,232 complaints (62.4%) and 739 non-complaints (37.6%) over 9 domains (e.g. food, technology, etc.). Data statistics are shown in Table 1. We opted using this data set because (1) it is publicly available; and (2) it allows a direct comparison with existing methods.

We also use the data for distant supervision² collected by Preotiuc-Pietro et al. (2019). This extra ‘noisy’ data source contains 18,218 complaint tweets collected by querying Twitter API with certain complaint related hashtags (e.g. #badbusiness, #badcustomerservice, etc.) and the same amount of non-complaint tweets that were sampled randomly.

3 Transformer-based Models

Transformer architectures trained on language modeling have been recently adapted to downstream tasks demonstrating state-of-the-art performance (Weller and Seppi, 2019; Gupta and Durrett, 2019; Maronikolakis et al., 2020). In this paper, we adapt and subsequently combine transformers with external linguistic information for complaint prediction.

BERT, ALBERT and RoBERTa Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2018) learns language representations by jointly conditioning on both left and right contexts using transformers. It is trained on masked language modeling where some of the tokens are randomly masked with the aim to predict them using only the context.

We further experiment with ALBERT (Lan et al., 2019) and RoBERTa (Liu et al., 2019). ALBERT uses two parameter-reduction methods to address memory limitations and long training time of BERT: (a) factorized embedding parameterization; (b) cross-layer parameter sharing. RoBERTa is an extension of BERT trained on more data with larger batch size using dynamic masking (i.e. changeable masked tokens of each sequence during training epochs). We adapt BERT, ALBERT and RoBERTa by adding a linear layer with a sigmoid activation and then fine-tune it on the complaint classification data.

XLNet XLNet (Yang et al., 2019b) uses a similar architecture to BERT to learn bidirectional contextual information. Instead of masked tokens used in BERT, XLNet maximizes the expected log-likelihood of all possible factorization orders. We adapt and fine-tune the XLNet model for complaint prediction similar to BERT.

¹<https://github.com/danielpreotiuc/complaints-social-media>

²We use the noisy but larger distantly supervised data to first adapt all models on the complaint classification task. Then we fine-tune them using the smaller original complaint data set.

M-BERT To combine our model with external linguistic information, we adapt the Multimodal BERT (M-BERT) (Rahman et al., 2019) model structure that has been introduced for multimodal modeling (text, image, speech). Instead of cross-modal interactions, we inject extra linguistic information as alternative views of the data into the pre-trained BERT model. We use (a) **Emotion**, a 9 dimensional vector obtained by quantifying six basic emotions of Ekman (1992) for each tweet using a predictive model by Volkova and Bachrach (2016); (b) **Topics**, a 200 dimensional vector representing word frequencies in word clusters designed to identify semantic themes in tweets by Preotiuc-Pietro et al. (2015; 2015). To inject external linguistic information to M-BERT,³ we first project the linguistic information into vectors with similar size to the BERT CLS embeddings. Then we concatenate word representations obtained from BERT and the linguistic information (Emotion, Topics or Emotion+Topics) to generate combined embeddings. During concatenation, an Attention Gating Mechanism called Multimodal Shifting Gate (Wang et al., 2019) is applied to control the importance of each representation. Finally, the combined embeddings are fed to BERT for fine-tuning. The rest of the architecture is the same as BERT.

4 Experiments

Baselines We compare the transformer-based models with two previous approaches for complaint identification by Preotiuc-Pietro et al. (2019) and a transfer learning method: (1) Logistic Regression with bag-of-words trained using the original and distantly supervised complaint data (**LR-BOW + Dist. Supervision**); (2) A Long-Short Term Memory (**LSTM**) network (Hochreiter and Schmidhuber, 1997) that takes as input a tweet, maps its words to embeddings and subsequently passes them through the LSTM to obtain a contextualized representation which is finally fed to the output layer; (3) Adapting the pre-trained Universal Language Model Fine-tuning (**ULMFiT**) model (Howard and Ruder, 2018) for complaint prediction. ULMFiT uses a AWD-LSTM (Merity et al., 2017) encoder for language modeling.

Hyper-parameters We use BERT, ALBERT and RoBERTa Base uncased models; fine-tuning them with learning rate $l = 1e-5$, $l \in \{1e-4, 1e-5, 2e-5, 1e-6\}$. We use the Base cased pre-trained XLNet tuning the learning rate over the same range as for BERT models. For ULMFiT, we use AWD-LSTM trained on Wikitext-103. We simplify the default fine-tuning by only unfreezing the last 1 layer, the last 2 layers and all layers with learning rates $l_1 = \frac{1e-4}{2.6^4}$, $l_2 = \frac{1e-4}{2.6^3}$ and $l_3 = 1e-3$ respectively. For M-BERT, the size of feature embeddings (Emotion, Topics and Emotion+Topics) is $h = 200$, $h \in \{200, 400, 768\}$ with dropout $d = 0.1$, $d \in \{.1, .5\}$ using the same parameters as BERT. The maximum sequence length is set to 49 covering 95% of tweets in the training set.

Evaluation Following Preotiuc-Pietro et al. (2019), we use a nested 10-fold cross-validation approach to conduct our experiments for complaint prediction. In the outer 10 loops, 9 folds are used for training and one for testing; while in the inner loops, a 3-fold cross-validation method is applied where 2 folds are used for training and one for validation. During training, an early stopping method is applied based on the validation loss. We measure predictive performance using the mean Accuracy, Precision, Recall and macro F1 over 10 folds (we also report the standard deviations).

Predictive Results Table 2 shows results of the Transformer-based models as well as the baselines on the complaint prediction task. All transformer-based models (BERT, ALBERT, RoBERTa and XLNet) perform better than the previous feature-based (LR-BOW + Dist. Supervision) and the non-transformer transfer learning baseline (ULMFiT), indicating a better capability on capturing idiosyncrasies of complaints syntax and semantics. BERT outperforms other models overall across all metrics reaching a macro F1 up to 87, which is 8% higher than the previous state-of-the-art (Preotiuc-Pietro et al., 2019). The results of RoBERTa are close to BERT with 86.6 macro F1 while ALBERT and XLNet achieve lower performance.

Distant supervision is beneficial only to ULMFiT and M-BERT while BERT and other transformer models perform worse, which are consistent with results of Bataa and Wu (2019) for sentiment analysis.

³We experiment by injecting emotion (**M-BERT - Emotion**) or topical information (**M-BERT - Topics**) and their combination (**M-BERT - Emotion+Topics**)

| Model | Acc | P | R | F1 |
|---|-------------------|-------------------|-------------------|-------------------|
| Preotiuc-Pietro et al. (2019) | | | | |
| LR-BOW + Dist. Supervision | 81.2 | - | - | 79.0 |
| LSTM | 80.2 | - | - | 77.0 |
| Transfer Learning Baseline | | | | |
| ULMFiT | 82.4 ± .04 | 81.1 ± .04 | 81.8 ± .04 | 81.2 ± .05 |
| ULMFiT + Dist. Supervision | 83.3 ± .05 | 82.5 ± .05 | 81.8 ± .04 | 81.9 ± .05 |
| Transformers | | | | |
| BERT | 88.0 ± .03 | 87.1 ± .03 | 87.3 ± .03 | 87.0 ± .03 |
| ALBERT | 85.9 ± .03 | 84.8 ± .03 | 84.6 ± .03 | 84.6 ± .03 |
| RoBERTa | 87.6 ± .03 | 86.6 ± .03 | 86.9 ± .03 | 86.6 ± .03 |
| XLNet | 83.9 ± .04 | 83.2 ± .04 | 82.3 ± .03 | 82.4 ± .05 |
| M-BERT - Emotion | 87.3 ± .03 | 86.5 ± .04 | 86.0 ± .03 | 86.1 ± .04 |
| M-BERT - Topics | 87.5 ± .03 | 86.7 ± .04 | 86.5 ± .03 | 86.4 ± .03 |
| M-BERT - Emotion+Topics | 87.1 ± .03 | 86.4 ± .03 | 85.6 ± .03 | 85.9 ± .03 |
| BERT + Dist. Supervision | 87.8 ± .03 | 87.0 ± .04 | 86.7 ± .03 | 86.7 ± .04 |
| ALBERT + Dist. Supervision | 83.9 ± .04 | 82.6 ± .04 | 82.7 ± .04 | 82.6 ± .04 |
| RoBERTa + Dist. Supervision | 85.2 ± .04 | 84.4 ± .05 | 84.0 ± .04 | 84.0 ± .04 |
| XLNet + Dist. Supervision | 82.1 ± .05 | 81.7 ± .05 | 79.9 ± .05 | 80.1 ± .05 |
| M-BERT - Emotion + Dist. Supervision | 87.7 ± .04 | 86.9 ± .04 | 87.2 ± .03 | 86.8 ± .04 |
| M-BERT - Topics + Dist. Supervision | 87.6 ± .05 | 87.0 ± .05 | 86.9 ± .04 | 86.7 ± .05 |
| M-BERT - Emotion+Topics + Dist. Supervision | 87.8 ± .04 | 87.1 ± .05 | 87.0 ± .04 | 86.9 ± .04 |

Table 2: Accuracy (Acc), Precision (P), Recall (R) and F1-Score (F1) for complaint prediction (\pm std. dev.) Best results are in bold.

Also, results of M-BERT models are comparable to BERT, among which M-BERT - Topics is slightly better with 86.4 macro F1. We notice that injecting external linguistic information in BERT’s structure for fine-tuning does not help in our case without substantially hurting performance. We speculate that modifying BERT embeddings by injecting extra linguistic information is not complementary to BERT’s text representations.

Error Analysis We also investigate the limitations in predicting capacity of our best performing model (BERT). We randomly analyze 100 cases in predictive results, where 50 cases were misclassified as non-complaints and another 50 cases were misclassified as complaints. In cases where complaints were misclassified as non-complaints, 26% errors are due to implicit expressions while 14% errors are because complaints contain irony. In the former situation, complaints express weak emotional intensity without explicit reproach, where complainers imply their dissatisfaction instead of directly complaining or mentioning the cause (Trosborg, 2011). The following tweet is a typical example:

It started yesterday , but I try again it could work normal. But since last night its just like this <url>

Such expressions rarely include words related to complaints (e.g. ‘disappointed’, ‘bad service’) and are therefore difficult to be correctly classified. In the latter situation, complaints are expressed in an ironic way using terms such as ‘congratulations’, ‘thank you’ and ‘brilliant’. For instance, the following text was wrongly classified as a non-complaint:

Thank you so much for making a box that shreds apart even when carried by both handles.

In cases where non-complaints were misclassified as complaints, errors can be roughly divided into four categories: (1) 26% errors are because certain terms appear frequently in complaints during training such as ‘thank you’, ‘dm’, ‘lost’, ‘work’. The following non-complaint was wrongly classified as a complaint:

*BTW <user> – <user> did me right, and replaced my two failed batteries under warranty. I’m **happy** :) **thanks** <user>!*

It contains similar words with the following complaint in the same fold (similarities highlighted in bold):

*Was **happy** to find out <user> had an app to watch all their shows, until 6 episodes in it stops working. **Thanks!** <user>*

| Test Train | F | A | R | C | Se | So | T | E | O |
|-------------|------------------|------------------|------------------|------------------|------------------|------------------|------------------|------------------|------------------|
| Food | - | 69.7 49.8 | 76.5 53.2 | 85.7 61.8 | 73.3 56.8 | 76.2 59.2 | 71.2 52.7 | 74.4 61.1 | 83.0 48.7 |
| Apparel | 75.0 69.7 | - | 74.2 81.5 | 78.8 74.3 | 72.7 76.3 | 73.2 81.8 | 66.7 74.1 | 75.7 75.9 | 81.6 84.2 |
| Retail | 74.9 75.9 | 72.8 80.0 | - | 80.0 73.0 | 75.4 75.9 | 75.9 80.7 | 70.0 74.9 | 72.7 75.3 | 86.6 79.5 |
| Cars | 76.1 57.1 | 70.2 62.1 | 75.5 65.1 | - | 70.2 51.6 | 75.1 70.6 | 71.3 62.3 | 74.4 61.4 | 82.1 71.8 |
| Services | 79.8 64.7 | 71.1 82.4 | 77.2 76.5 | 77.5 75.4 | - | 73.8 78.6 | 73.4 76.3 | 75.5 79.4 | 78.9 83.0 |
| Software | 74.6 69.3 | 70.4 80.0 | 73.5 77.5 | 79.1 78.0 | 77.9 76.4 | - | 73.4 75.2 | 76.7 76.2 | 81.7 82.0 |
| Transport | 72.5 62.2 | 70.5 73.5 | 77.1 80.0 | 80.0 79.2 | 74.4 76.3 | 75.8 75.3 | - | 72.4 70.4 | 82.0 82.6 |
| Electronics | 69.9 72.9 | 72.2 78.5 | 69.1 78.9 | 73.0 75.4 | 77.0 78.0 | 71.0 72.4 | 69.8 69.7 | - | 82.1 80.8 |
| Other | 65.9 64.8 | 75.2 74.8 | 79.2 72.2 | 81.7 69.2 | 76.3 69.9 | 76.5 77.9 | 70.6 70.6 | 72.8 70.8 | - |
| All | 48.9 77.5 | 67.5 87.7 | 72.6 85.8 | 73.9 80.9 | 72.0 81.1 | 65.8 85.1 | 64.9 81.4 | 67.6 82.0 | 81.9 88.2 |

Table 3: F1-score of models in Preotiuc-Pietro et al. (2019) (left) and BERT (right) trained from one domain and tested on other domains. Domains include Food (F), Apparel (A), Retail (R), Cars (C), Services (Se), Software (So), Transport (T), Electronics (E) and Other (O). The All line shows results on training on all categories except the category in testing. Best results are in bold.

(2) 22% errors due to interrogative tone, which is common in complaints. An example is “Folks , what is cost of text message to a us number?” (3) 22% errors are from negation words such as “No luck with pc or phone.” (4) 12% errors are because texts contain negative sentiment such as “This would be a terrible idea <url>” are likely to be classified as complaints incorrectly since words such as ‘terrible’ are widely used to express dissatisfaction. However, there are not enough cues to indicate violation of expectations. According to the statistics, the proportion of complaints misclassified as non-complaints (15.22%) is higher than that of non-complaints misclassified as complaints (10.25%) indicating implicit and figurative expressions as well as unknown factors in complaints are more challenging to identify.

Cross Domain Experiments Finally, we use BERT to train models on one domain and test on another as well as training on all domains except the one that the model is tested on. Table 3 shows the performance of models in Preotiuc-Pietro et al. (2019) (left) and BERT (right) across 9 domains. We first observe that BERT results in nearly half of the cases when training on a single domain are lower than LR-BOW (especially ‘Food’, ‘Car’ and ‘Other’) while BERT trained on all domains performs better across all testing domains, achieving a macro F1 up to 88.2 when tested on ‘Other’. This indicates that, fine-tuning BERT on a small training data set (‘Food’, ‘Car’ and ‘Other’ are three of the domains with the smaller amount of data) is not enough to make it perform well. In contrast, it achieves better performance consistently on larger data sets (All). We also notice that BERT performs robustly for domain pairs where the domains are either used for training or testing. For example, training on ‘Apparel’ achieves high performance when testing on ‘Software’ (81.8 F1) and vice versa (80.0 F1). Furthermore, domain relevance affect predictive performance. For example, BERT trained on ‘Transport’ achieves 79.2 F1 when tested on ‘Car’, which is the highest performance compared to other training domains since these two domain share common vocabulary (see ‘Car’ column for BERT).

5 Conclusion

We evaluated a battery of transformer networks on the Twitter complaint identification task and obtained 87 macro F1, which outperforms the previous state-of-the-art results of Preotiuc-Pietro et al. (2019). We further presented a thorough analysis of the limitations of our models in predicting complaints. In future work, we intend to explore more in how we can combine other sources of linguistic information with transformers as well as information from other modalities such as images.

Acknowledgements

Nikolaos Aletras is supported by ESRC grant ES/T012714/1.

References

- Enkhbold Bataa and Joshua Wu. 2019. An Investigation of Transfer Learning-Based Sentiment Analysis in Japanese. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4652–4657. Association for Computational Linguistics.
- Shreesh Kumara Bhat and Aron Culotta. 2017. Identifying leading indicators of product recalls from online reviews using positive unlabeled learning and domain adaptation. In *Eleventh International AAAI Conference on Web and Social Media*.
- Andrew D Cohen and Elite Olshtain. 1993. The Production of Speech Acts by EFL Learners. *Tesol Quarterly*, 27(1):33–56.
- Kristof Coussement and Dirk Van den Poel. 2008. Improving customer complaint management by automatic email classification using linguistic style features as predictors. *Decision Support Systems*, 44(4):870–882.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, page 4171–4186.
- Paul Ekman. 1992. An Argument for Basic Emotions. *Cognition & Emotion*, 6(3-4):169–200.
- J Forster and B Entrup. 2017. A Cognitive Computing Approach for Classification of Complaints in the Insurance Industry. In *IOP Conference Series: Materials Science and Engineering*, volume 261, page 012016. IOP Publishing.
- D Gunawan, RP Siregar, RF Rahmat, and A Amalia. 2018. Building automatic customer complaints filtering application based on Twitter in Bahasa Indonesia. In *Journal of Physics: Conference Series*, volume 978, page 012119. IOP Publishing.
- Aditya Gupta and Greg Durrett. 2019. Effective Use of Transformer Networks for Entity Tracking. *arXiv preprint arXiv:1909.02635*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780.
- Jeremy Howard and Sebastian Ruder. 2018. Universal Language Model Fine-tuning for Text Classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339. Association for Computational Linguistics.
- Jiahua Jin, Xiangbin Yan, You Yu, and Yijun Li. 2013. Service Failure Complaints Identification in Social Media: A Text Classification Approach.
- Leila Nasiri Kakolaki and Mohsen Shahrokhi. 2016. Gender Differences in Complaint Strategies among Iranian Upper Intermediate EFL Students. *Studies in English Language Teaching 4 (1)*, pages 1–15.
- Robin M Kowalski. 1996. Complaints and complaining: Functions, antecedents, and consequences. *Psychological Bulletin*, 119(2):179.
- M Lailiyah, S Sumpeno, and IK E Purnama. 2017. Sentiment Analysis of Public Complaints Using Lexical Resources Between Indonesian Sentiment Lexicon and Sentiwordnet. In *2017 International Seminar on Intelligent Technology and Its Applications (ISITIA)*, pages 307–312. IEEE.
- Janice Laksana and Ayu Purwarianti. 2014. Indonesian Twitter Text Authority Classification For Government in Bandung. In *2014 International Conference of Advanced Informatics: Concept, Theory and Application (ICAICTA)*, pages 129–134. IEEE.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. *arXiv preprint arXiv:1909.11942*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.

- Antonios Maronikolakis, Danae Sánchez Villegas, Daniel Preotiuc-Pietro, and Nikolaos Aletras. 2020. Analyzing political parody in social media. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4373–4384.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2017. Regularizing and Optimizing LSTM Language Models. *arXiv preprint arXiv:1708.02182*.
- Febina Merson and Roseline Mary. 2017. A Text Mining Approach to Identify and Analyse Prominent Issues from Public Complaints. *International Journal of Advanced Research in Computer and Communication Engineering*, 6(3).
- E Olshtain and L Weinbach. 1985. Complaints: A study of speech act behavior among native and non-native speakers of Hebrew.
- Sachin Pawar, Nitin Ramrakhiani, Girish K Palshikar, and Swapnil Hingmire. 2015. Deciphering Review Comments: Identifying Suggestions, Appreciations and Complaints. In *International Conference on Applications of Natural Language to Information Systems*, pages 204–211. Springer.
- Daniel Preotiuc-Pietro, Vasileios Lampos, and Nikolaos Aletras. 2015. An analysis of the user occupational class through Twitter content. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1754–1764.
- Daniel Preotiuc-Pietro, Mihaela Gaman, and Nikolaos Aletras. 2019. Automatically Identifying Complaints in Social Media. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5019.
- Daniel Preotiuc-Pietro, Svitlana Volkova, Vasileios Lampos, Yoram Bachrach, and Nikolaos Aletras. 2015. Studying user income through language, behaviour and affect in social media. *PLOS ONE*, 10(9):1–17, 09.
- Wasifur Rahman, Md Kamrul Hasan, Amir Zadeh, Louis-Philippe Morency, and Mohammed Ehsan Hoque. 2019. M-BERT: Injecting Multimodal Information in the BERT Structure. *arXiv preprint arXiv:1908.05787*.
- Beverly A Sparks and Victoria Browning. 2010. Complaining in Cyberspace: The Motives and Forms of Hotel Guests’ Complaints Online. *Journal of Hospitality Marketing & Management*, 19(7):797–818.
- Suhatasi Tjandra, Amelia Alexandra Putri Warsito, and Judi Prajetno Sugiono. 2015. Determining Citizen Complaints to The Appropriate Government Departments using KNN Algorithm. In *2015 13th International Conference on ICT and Knowledge Engineering (ICT & Knowledge Engineering 2015)*, pages 1–4. IEEE.
- Anna Trosborg. 2011. *Interlanguage Pragmatics: Requests, Complaints, and Apologies*, volume 7. Walter de Gruyter.
- Camilla Vásquez. 2011. Complaints online: The case of TripAdvisor. *Journal of Pragmatics*, 43(6):1707–1717.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Svitlana Volkova and Yoram Bachrach. 2016. Inferring Perceived Demographics from User Emotional Tone and User-Environment Emotional Contrast. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1567–1578.
- Yansen Wang, Ying Shen, Zhun Liu, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. 2019. Words Can Shift: Dynamically Adjusting Word Representations Using Nonverbal Behaviors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7216–7223.
- Orion Weller and Kevin Seppi. 2019. Humor Detection: A Transformer Gets the Last Laugh. *arXiv preprint arXiv:1909.00252*.
- Wei Yang, Luchen Tan, Chunwei Lu, Anqi Cui, Han Li, Xi Chen, Kun Xiong, Muzi Wang, Ming Li, Jian Pei, et al. 2019a. Detecting Customer Complaint Escalation with Recurrent Neural Networks and Manually-Engineered Features. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Industry Papers)*, pages 56–63.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019b. XLNet: Generalized Autoregressive Pretraining for Language Understanding. In *Advances in Neural Information Processing Systems*, pages 5754–5764.