

Modeling the Severity of Complaints in Social Media

Mali Jin Nikolaos Aletras

Department of Computer Science

University of Sheffield, UK

{mjin6, n.aletras}@sheffield.ac.uk

Abstract

The speech act of complaining is used by humans to communicate a negative mismatch between reality and expectations as a reaction to an unfavorable situation. Linguistic theory of pragmatics categorizes complaints into various severity levels based on the face-threat that the complainer is willing to undertake. This is particularly useful for understanding the intent of complainers and how humans develop suitable apology strategies. In this paper, we study the severity level of complaints for the first time in computational linguistics. To facilitate this, we enrich a publicly available data set of complaints with four severity categories and train different transformer-based networks combined with linguistic information achieving 55.7 macro F1. We also jointly model binary complaint classification and complaint severity in a multi-task setting achieving new state-of-the-art results on binary complaint detection reaching up to 88.2 macro F1. Finally, we present a qualitative analysis of the behavior of our models in predicting complaint severity levels.^{1,2}

1 Introduction

Complaining is a speech act that usually conveys negative emotions triggered by a discrepancy between reality and expectations towards an entity or event (Olshtain and Weinbach, 1985). Complaints play an important role in human communication for expressing dissatisfaction. Based on complainers' personalities and specific situations, expression of complaints vary from person to person (Vásquez, 2011).

In pragmatics, complaints have been classified into various levels of severity according to their

emotional intensity, the amount of face-threat that the complainer is willing to undertake and their purpose (Olshtain and Weinbach, 1985; Trosborg, 2011; Kakolaki and Shahrokhi, 2016). Complaining purposes might include the expression of dissatisfaction, to find solutions (e.g. ask for reparations) or both. Furthermore, a complaint can be categorized as implicit (i.e. without mentioning who is responsible) or explicit (i.e. accusing someone for doing something).

Recent work on modeling complaints in natural language processing (NLP) has focused on distinguishing complaints from non-complaints in social media (Preotiuc-Pietro et al., 2019; Jin and Aletras, 2020), however there is no previous study into more fine-grained complaint categories. Table 1 shows examples of social media posts expressing complaints grouped into four severity classes according to Trosborg (2011): (a) no explicit reproach; (b) disapproval; (c) accusation; and (d) blame.

Identifying and analyzing the severity of complaints is important for: (a) improving customer service by recognizing the level of dissatisfaction and understanding complainers' needs (Van Noort and Willemsen, 2012); (b) linguists to study the speech act of complaints in different levels of granularity on large scale (Tatsuki, 2000); and (c) developing downstream NLP applications such as automatic complaint response generation (Xu et al., 2017) or voting stance prediction (Tsakalidis et al., 2018).

In this paper, we present a systematic study on analyzing complaint categories with computational methods for the first time in computational linguistics. Our main contributions are as follows:

- Grounded in linguistic theory of pragmatics (Trosborg, 2011), we enrich a publicly available data set (Preotiuc-Pietro et al., 2019) with four complaint severity levels;
- We create a new classification task for identifying different severity levels of complaints;

¹Data is available here: https://archive.org/details/complaint_severity_data

²Code is available here: <https://github.com/mali726/Complaint-Severity>

Label	Example
No Explicit Reproach	<i>Are you following me? I seem unable to send you a dm.</i>
Disapproval	<i>So far , the mac graphics drivers have been another disappointing update (for both my quadro 4000 & gtx - 285),</i>
Accusation	<i>Can u stop adding the UK keyboard layout to my Italian keyboard at every update? ktnxby</i>
Blame	<i>Thanks to <USER> 's incompetence i now can't work till October 4th, when the ati card arrives.</i>

Table 1: Examples of complaint severity levels (Trosborg, 2011).

- We evaluate transformer-based classification models (Vaswani et al., 2017) combined with linguistic information on (a) complaint severity level classification; and (b) binary complaint detection in a multi-task setting achieving new state-of-the-art results.

2 Related Work

2.1 Linguistic Categories of Complaints

Previous work in linguistic theory of pragmatics has classified complaints into different levels based on their severity and directness. Olshtain and Weinbach (1985) classified complaints into five distinct categories: (a) below the level of reproach; (b) expression of annoyance or disapproval; (c) explicit complaint; (d) accusation and (e) warning, immediate threat. More recently, Trosborg (2011) proposed four major severity levels: (a) no explicit reproach; (b) disapproval; (c) accusation and (d) blame. Finally, Kakolaki and Shahrokhi (2016) classified complaints into levels of directness: (a) very direct; (b) somewhat direct and (c) indirect. Direct complaints (i.e. very direct and somewhat direct) include obvious breaches of expectations. On the other hand, indirect complaints do not explicitly mention or can imply the breach of expectations. Moreover, the difference between very direct and somewhat direct is that the former highlights the responsibility of the complaint receiver while the latter does not.

2.2 Complaint Analysis

Most of the existing studies on complaint classification in NLP have explored different approaches to the complaint identification task (identifying complaints from non-complaints) in various domains, starting with feature-based machine learning models (Coussement and Van den Poel, 2008; Preotiuc-Pietro et al., 2019) and deep learning methods (Jin and Aletras, 2020). Coussement and Van den Poel (2008) used boosting ensemble models with linguistic style features to identify complaints in com-

pany emails. Preotiuc-Pietro et al. (2019) applied logistic regression with a broad range of features to detect complaints in Twitter. More recent, Jin and Aletras (2020) explored a battery of transformer-based architectures combined with sentiment and topic information for complaint identification in social media. Also, previous work has classified complaints into detailed topical categories (Forster and Entrup, 2017; Merson and Mary, 2017) or responsible departments (Laksana and Purwarianti, 2014; Gunawan et al., 2018; Tjandra et al., 2015). Furthermore, other complaint related categorizations are based on product hazards and risks (Bhat and Culotta, 2017), service failure (Jin et al., 2013) and escalation likelihood (Yang et al., 2019).

2.3 Emotion Detection

Most related to complaint severity is emotion detection and its intensity which have been extensively studied in NLP (Danisman and Alpkocak, 2008; Volkova and Bachrach, 2016; Zhang et al., 2018). More recently, Alejo et al. (2020) explored cross-lingual transfer approaches to predict emotion intensity in Twitter. Similarly, Akhtar et al. (2020) evaluated a series of feature-based machine learning models for both emotion and sentiment intensity prediction in social and news media.

3 Task & Data

We define complaint severity prediction as a multi-class classification task. Given a text snippet T , defined as a sequence of tokens $T = \{t_1, \dots, t_n\}$, the aim is to classify T as one of the four predefined severity labels.

We use an existing complaints data set developed by Preotiuc-Pietro et al. (2019), which consists of 1,235 complaints (35.8%) and 2,214 non-complaints (64.2%) in English. We opted using this data set because it is publicly available with annotated complaints collected from Twitter in 9 general domains (i.e. Food, Apparel, Retail, Cars, Service, Software, Transport, Electronics and Other).

3.1 Complaint Severity Categories

For complaint severity annotation, we adopt the four categories defined by Trosborg (2011) because it is considered as the ‘standard’ in pragmatics literature (see examples in Table 1):

- **No explicit reproach:** there is no explicit mention of the cause and the complaint is not offensive;
- **Disapproval:** express explicit negative emotions such as dissatisfaction, annoyance, dislike and disapproval;
- **Accusation:** asserts that someone did something reprehensible;
- **Blame:** assumes the complaineer is responsible for the undesirable result.

Note that the severity levels categorize complaints by type instead of intensity. Classes are disjoint according to Trosborg (2011). More specifically, ‘No explicit reproach’ is a suggestive strategy, where the complaineer is usually not mentioned in the statement. ‘Disapproval’ expresses negative sentiment or unsatisfying state only. The statement may imply the complainer holds the complaineer responsible but avoid mentioning it, which is the key component of identifying ‘Disapproval’ and ‘Accusation’/‘Blame’. The main difference between ‘Accusation’ and ‘Blame’ is in the latter one the complainer presupposes the complaineer is guilty of the offense.

3.2 Complaint Severity Annotation

Following the definitions above, each tweet was labeled by three annotators independently. In case of ties, the final decision was made by the authors through consensus. We recruited 35 native English speaking annotators from the volunteers list of our institution.^{3,4} The inter-annotator agreement between the three original annotations for each tweet is $k = 0.64$ Fleiss’ Kappa⁵ (Fleiss, 1971) which belongs to substantial agreement (Artstein and Poesio, 2008).

Table 2 shows the distribution of tweets across classes: 435 tweets belong to ‘No Explicit Reproach’ (35.2%), 378 belong to ‘Disapproval’

³We have received approval from the Ethics Committee of our institution.

⁴Annotators are provided with an introduction of the task including definitions and examples of each category.

⁵We randomize the order of three annotations for each tweet three times and compute the average Fleiss’ Kappa.

Labels	Amount	Percentage
No Explicit Reproach	435	35.2
Disapproval	378	30.6
Accusation	225	18.2
Blame	197	16.0
Total	1235	100

Table 2: Number of tweets in different complaint severity levels and class distribution.

(30.6%), 225 belong to ‘Accusation’ (18.2%); and 197 belong to ‘Blame’ (16.0%). The class distributions over 5 domains (Car, Retail, Service, Software, Transport) are similar to the overall distribution while 4 domains (Food, Apparel, Electronics, Other) differ from Table 2. In domains with different distribution, differences appear especially in ‘No Explicit Reproach’ and ‘Accusation’, which might result from domain specific complaint requests.

3.3 Text Processing

Text is processed by lower-casing, and replacing all mentions of usernames and URLs with placeholder tokens. A Twitter-aware tokenizer, DLATK (Schwartz et al., 2017), is used for text tokenization to handle emoticons and hashtags in social media text.

4 Predictive Models

Since severity complaint prediction is a new task, we first evaluate the majority class as well as three strong baselines: (1) logistic regression with bag-of-words; (2) a bidirectional recurrent neural network trained from scratch; and (3) finetuning a pre-trained transformer-based model. Furthermore, we combine linguistic information (i.e. emotion and topic information) to a transformer-based model similar to the method proposed by Jin and Aletras (2020) in the context of binary complaint classification.

4.1 Baselines

Majority Class We use Majority Class as the first baseline, where we calculate scores by labeling all the tweets with the majority class.

LR-BOW We use a linear baseline, Logistic Regression with standard bag-of-words (LR-BOW) and L2 regularization.

BiGRU-Att We also use a neural baseline trained from scratch; a bidirectional Gated Recurrent Unit

(GRU) network (Cho et al., 2014) with a self-attention mechanism (BiGRU-Att; (Tian et al., 2018)). Given a Twitter post T , a token t_i is mapped to a GloVe embedding (Pennington et al., 2014). We then apply dropout to the output of GloVe embedding layer and pass it to a bidirectional GRU with self-attention layer. Finally, the contextualized token representations are passed to an output layer using a softmax activation function for multi-class classification.

RoBERTa Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2018) is a pre-trained language model based on the Transformer architecture (Vaswani et al., 2017). It makes use of multiple multi-head attention layers to learn context information from both the left and the right side of tokens. It is trained on masked language modeling by randomly masking some of the tokens from the input aiming to predict them based on the context only. RoBERTa (Liu et al., 2019) is an extension of BERT trained on more data with different hyperparameters and has achieved better performance in social media analysis tasks (Maronikolakis et al., 2020). We fine-tune RoBERTa⁶ on complaint severity classification by adding an output dense layer with a softmax activation function.

4.2 M-RoBERTa with Linguistic Information

Multimodal-BERT (M-BERT) (Rahman et al., 2019) injects multimodal information such as image and speech into the text representations of BERT. It combines word embeddings and embeddings from other modalities (e.g. image, audio) which are then fed to a BERT encoder. M-BERT has been recently adapted by Jin and Aletras (2020) for binary complaint prediction by inducing linguistic information instead of speech and image, however it did not perform better than BERT in their setting.

We adapt M-BERT by replacing (1) the underlying BERT model with RoBERTa; and (2) the multimodal information with linguistic information. We first use a fully connected layer to project the linguistic representations into vectors with comparable size to RoBERTa’s embeddings. Then we concatenate word representations from RoBERTa and linguistic information representations using

⁶We only report the results of RoBERTa because it achieves better performance compared to BERT over all evaluation methods in our experiments.

a Multimodal Shifting Gate (Wang et al., 2019), where an attention gating mechanism is applied to control the influence of each representation. Finally, we apply layer normalization and dropout after the Multimodal Shifting Gate and pass the output to RoBERTa. We add an output layer to M-RoBERTa for classification similar to the RoBERTa model. We use M-RoBERTa with three types of linguistic features (i.e. emotion, topic and their combination):

M-RoBERTa_{Emo} We first use emotional information obtained by using a pretrained emotional classifier by Volkova and Bachrach (2016). This is 9-dimensional vector representing scores of sentiment (positive, negative and neutral) and six basic emotions of Ekman (1992) (anger, disgust, fear, joy, sadness and surprise).

M-RoBERTa_{Top} We also use topical information from a 200-dimensional vector representing the distribution of the fraction of tokens in each tweet belonging to a topic cluster (PreoŃuc-Pietro et al., 2015; Aletras and Chamberlain, 2018).

M-RoBERTa_{Emo+Top} We finally experiment with injecting both emotional and topical information to M-RoBERTa.

5 Experimental Setup

5.1 Hyperparameters

The **BiGRU-Att** model uses 200-dimensional GloVe embeddings (Pennington et al., 2014) pre-trained on Twitter data. Its hidden size is $h = 128$, $h \in \{64, 128, 256, 512\}$ with dropout $d = .2$, $d \in \{.2, .5\}$. We use Adam optimizer (Kingma and Ba, 2014) with learning rate $l = 1e-3$, $l \in \{1e-3, 5e-3, 1e-2\}$. For **RoBERTa**, we use the Base uncased model and fine-tuning it with learning rate $l = 5e-6$, $l \in \{1e-4, 1e-5, 5e-6, 1e-6\}$. The maximum sequence length is set to 50 covering 95% of tweets in the training set. For **M-RoBERTa** models, we project the linguistic features (emotions and topics) to vectors of size $l = 200$, $l \in \{200, 300, 400, 768\}$. We also use dropout $d = .5$, $d \in \{.2, .5\}$. For all models we use a categorical-cross entropy loss following a similar approach to Sun et al. (2019) which have achieved best results on fine-grained sentiment analysis (i.e. similar to the ordinal scale of complaints severity).

Model	Acc	P	R	F1
Majority Class	35.2	8.8	25.0	13.0
LR-BOW	46.7 ± .03	44.3 ± .06	43.6 ± .03	43.5 ± .03
BiGRU-Att	46.1 ± .03	43.6 ± .03	42.7 ± .02	43.5 ± .03
RoBERTa	58.7 ± .03	55.8 ± .05	55.4 ± .03	54.7 ± .04
M-RoBERTa _{Emo}	59.8 ± .02	56.6 ± .03	55.7 ± .03	55.7[†] ± .03
M-RoBERTa _{Top}	59.0 ± .03	55.9 ± .04	55.6 ± .03	55.2 ± .04
M-RoBERTa _{Emo+Top}	59.4 ± .03	56.5 ± .03	56.2 ± .03	55.5 ± .02

Table 3: Accuracy (Acc), Precision (P), Recall (R) and macro F1-Score (F1) for complaint severity level prediction (\pm std. dev.). Best results are in bold. \dagger indicates statistically significant improvement over RoBERTa (t-test, $p < 0.05$).

5.2 Training and Evaluation

We run all models using a nested 10-fold cross validation approach, which consists of 2 nested loops as in Preotiuc-Pietro et al. (2019). In the outer loop, 9-folds are used for training and one for testing; while in the inner loop, a 3-fold cross validation is applied on the data from the nine folds (in the outer loop), where 2-folds are used for training and one for validation. During training, we choose the model with the smallest validation loss over 30 epochs. We measure predictive performance using the mean Accuracy, Precision, Recall and macro F1 over 10 folds (we also report the standard deviations).

6 Results

Table 3 shows the performance of all models including baselines and M-RoBERTa combined with linguistic information on complaint severity level prediction.

Overall, M-RoBERTa models with linguistic features achieve best results. M-RoBERTa_{Emo} outperforms all other models and reaches macro F1 up to 55.7. This confirms our hypothesis that injecting extra emotion information helps improve the performance of complaint severity level prediction. This is also in line with Trosborg (2011) who states that the expression of complaints is relevant to different emotional states. The results of M-RoBERTa_{Top} and M-RoBERTa_{Emo+Top} are comparable with 55.2 and 55.5 macro F1 respectively. RoBERTa performs competitively but worse than the M-RoBERTa models. We also notice that BiGRU-Att does not perform well in our task (43.5 macro F1), which may result from the fact that it has not been pretrained.

Figure 1 presents the confusion matrix of our




Figure 1: Confusion matrix of the best performing model (M-RoBERTa_{Emo}).

best model (i.e. M-RoBERTa_{Emo}). The confusion matrix is normalized over the actual values (rows). The ‘No Explicit Reproach’ category has the highest percentage (77.2%) of correctly classified data points by the model, followed by label ‘Disapproval’ with 59.0%. These are also the two most frequent classes in the data set. On the other hand, results on ‘Accusation’ are the lowest (32.9%) which is confused with adjacent categories (‘Disapproval’ and ‘Blame’). Furthermore, the differences between mis-classifications and correct classification are relatively large for ‘Blame’. We speculate that this is because of the unique linguistic characteristic of the ‘Blame’ category which gives emphasis on someone’s responsibility. Finally, a category is more likely, in general, to be mis-classified to its adjacent severity categories. For example, when predicting ‘Disapproval’, the number of model mis-classifications as ‘No Ex-




Figure 2: Confusion matrix of human agreement.

licit Reproach’ and ‘Accusation’ is larger than ‘Blame’. This hints that tweets belonging to neighboring levels share more semantic, syntactic and stylistic similarities.

We also compare the performance of our best model (i.e. M-RoBERTa_{Emo}) with human agreement for each class (Figure 2). In general, the results of the model (shown in Figure 1) correlate to human agreement. In other words, the model and humans agree in the categories they confuse. For instance, it is easy for both of them to confuse ‘Accusation’ with ‘Disapproval’ (32.9% vs. 31.1% for the model and 43.6% vs. 31.0% for humans). However, we observe that annotators are better at distinguishing high severity complaints from ‘No Explicit Reproach’, where 21.2% ‘Disapproval’ and 12.4% ‘Accusation’ are wrongly classified as ‘No Explicit Reproach’ by the model while the corresponding values are 18.5% and 8.9% by humans respectively. We argue that this is because annotators are able to identify subtle language (More details will be discussed in Section 7). Also, we notice that the model achieves better performance when predicting ‘Blame’, indicating a better capability on capturing the main characteristics of this class compared to humans.

7 Error Analysis

We perform an error analysis to shed light on the limitations of our best performing model (M-RoBERTa_{Emo}) on complaint severity level classification.

Firstly, we observe that most errors happen when the differences of tweets belonging to ‘Accusation’ are blurred with ‘Disapproval’ and ‘Blame’. The

following two tweets are typical examples for ‘Accusation’ being mis-classified as ‘Disapproval’ and ‘Blame’ respectively:

*<USER>, thank you ! Clear guidelines here, but **not at all** what your advisor on the phone stated!*

*The new <USER> stinks ...10mins to **take my order** and another 15 to get it. And **stop asking** my name like we’re friends <URL>*

This is because some tweets belonging to ‘Accusation’ also contain negation (e.g. ‘not at all’) or negative terms (e.g. ‘disappointed’), which appear frequently in ‘Disapproval’. Also, consistent with the definition by Trosborg (2011) (directly or indirectly accuses someone for causing the problem), tweets belonging to ‘Accusation’ may involve doing something and contain terms like ‘<USER>’ or ‘you’, which is similar to complaints labeled as ‘Blame’ such as:

*Thanks <USER> for **selling** expired beer #fail <USER> <URL>*

Secondly, the model struggles with complaints expressed in more subtle ways. In the following two examples, tweets belonging to ‘Disapproval’ and ‘Accusation’ are mis-classified as ‘No Explicit Reproach’ respectively:

Think someone at <USER> had been drinking the stuff before they put the label on

Just opened a fresh bud light that was filled with water. Please explain <USER>.

Such complaints do not contain terms that are typical of any specific complaint severity category (e.g. negation and negative terms in ‘Disapproval’, person pronouns and terms describing undesirable results in ‘Blame’) thus predicting them correctly needs more contextual understanding.

Finally, compared to other categories, the model is more likely to confuse tweets belonging to ‘No Explicit Reproach’ and ‘Disapproval’. This happens because some tweets express weak dissatisfaction, which is difficult to identify. The following tweet is mis-classified as ‘No Explicit Reproach’:

Dearest <USER>: there really needs to be an easier method to report names that are inappropriate <URL>

The model might need to learn more contextual information about such tweets instead of capturing certain relevant terms. Also, these two labels contain more similar terms such as ‘dm’, ‘please help’, ‘can’t work’ and interrogative tone. Examples of a ‘No Explicit Reproach’ and ‘Disapproval’ are the following (where similarities are in bold):

Hey guys, I love this product featured on <USER> today **but** don't see a price? **Help** a girl out? <URL>

So it's going to cost \$7000 to fix the exhaust on my <USER> 2009 jetta, and **only** \$300 is covered under warranty. **Help** <USER>?

8 Multi-task Learning for Binary Complaint Prediction

We further experiment with multi-task learning (MTL) (Caruana, 1997) for using severity categories to improve binary complaint prediction (i.e. complaint or non-complaint). MTL enables two or more tasks to be learned jointly by sharing information and parameters of a model.

We explore whether or not the severity level of a complaint helps in complaint identification. We use the same data set as Preotiuc-Pietro et al. (2019), where each tweet is annotated as a complaint or not and our severity level annotations.⁷

8.1 Predictive models

We first adapt three multi-task learning models based on bidirectional recurrent neural networks recently proposed by Rajamanickam et al. (2020) for jointly modeling abusive language detection and emotion detection. We also adapt our M-RoBERTa_{Emo} model in a multi-task setting using two variants. We use the severity complaint prediction as an auxiliary task and the binary complaint prediction as the main task to train different MTL models. All models are trained on the two tasks and updated at the same time with a joint loss:

$$L = (1 - \alpha)L_{com} + \alpha L_{sev}$$

where L_{com} and L_{sev} are the losses of complaint identification and severity level classification tasks respectively. α is a parameter to control the importance of each loss.

MTL-Hard Sharing We adapt the MTL-Hard Sharing model of Rajamanickam et al. (2020), where a single encoder is shared and updated by both tasks. We first pass GloVe embedding representations to a shared stacked BiGRU encoder. Then the output of the shared encoder is fed to two different BiGRU-Att models specific to each task (complaint detection and severity level identification) separately. Finally, we add an output layer with a sigmoid and a softmax activation function for binary and multi-class prediction respectively.

⁷For a tweet that is a non-complaint, we assign an extra class for severity (i.e. 'No Complaint Severity').

MTL-Double Encoder Instead of sharing a single encoder, the MTL-Double Encoder model (Rajamanickam et al., 2020) utilizes two stacked BiGRU encoders, where one is task-specific (complaint detection only) and the another one is shared by both tasks. We pass the output of the shared encoder to a BiGRU-Att model for severity level prediction. We also concatenate the output of the task-specific and shared encoder and pass it to another BiGRU-Att model for complaint prediction. The rest of the architecture is the same as the MTL-Hard Sharing model.

MTL-Gated Double Encoder The MTL-Gated Double Encoder model (Rajamanickam et al., 2020) has the same architecture as the MTL-Double Encoder. The outputs from two stacked BiGRU-Att encoders are concatenated by assigning a weight to each representation $[(1 - \beta)$ for the output of the task-specific encoder layer and β for the output of the shared one] that controls the importance of the two representations.

MTL-M-RoBERTa_{Emo} We also adapt our best performing model in the severity prediction task (M-RoBERTa_{Emo}) to support multi-task learning by adding an extra output layer for binary complaint prediction (MTL-M-RoBERTa_{Emo}).

MTL-M-RoBERTa_{Emo}-DE We pass the M-RoBERTa_{Emo} embedding to two separate RoBERTa encoders, i.e. double encoder (DE), followed by two classifiers for binary complaint and severity level prediction (MTL-M-RoBERTa_{Emo}-DE).

8.2 Experimental Setup

Baselines We compare MTL models with the following baselines on binary complaint identification: (1) Logistic Regression with bag-of-words using distant supervision⁸ (**LR-BOW + DS**) (Preotiuc-Pietro et al., 2019); (2) A standard BiGRU-Att model (**BiGRU-Att**); (3) A RoBERTa base model without combining linguistic information (**RoBERTa**) (Jin and Aletras, 2020); (4) A BERT base model without combining linguistic information (**BERT**) which has been shown to achieve state-of-the-art results in binary complaint identification (Jin and Aletras, 2020); and (5) replacing M-RoBERTa_{Emo} with BiGRU-Att in the MTL-M-RoBERTa_{Emo} and MTL-M-

⁸Distant supervision method uses larger 'noisy' data to further boost the performance of the model

Model	Acc	P	R	F1
Single-task Learning				
LR-BOW+DS (Preotiuc-Pietro et al., 2019)	81.2	-	-	79.0
BiGRU-Att	79.2 ± .05	79.2 ± .06	74.5 ± .05	74.5 ± .05
RoBERTa (Jin and Aletras, 2020)	87.6 ± .03	86.6 ± .03	86.9 ± .03	86.6 ± .03
BERT (Jin and Aletras, 2020)	88.0 ± .03	87.1 ± .03	87.3 ± .04	87.0 ± .03
Multi-task Learning				
MTL-BiGRU-Att	77.2 ± .05	75.4 ± .04	75.7 ± .04	75.4 ± .04
MTL-BiGRU-Att-DE	75.7 ± .05	74.1 ± .05	74.6 ± .04	74.1 ± .05
Rajamanickam et al. (2020)				
MTL-Hard Sharing	75.2 ± .04	73.5 ± .05	71.5 ± .04	72.1 ± .05
MTL-Double Encoder	74.6 ± .03	72.7 ± .04	71.7 ± .03	72.0 ± .04
MTL-Gated Double Encoder	74.7 ± .03	73.4 ± .04	70.4 ± .03	71.1 ± .03
Ours				
MTL-M-RoBERTa _{Emo}	89.0 ± .04	88.2 ± .03	88.4 ± .03	88.2[†] ± .03
MTL-M-RoBERTa _{Emo} -DE	88.9 ± .04	88.3 ± .04	88.3 ± .03	88.1 ± .04

Table 4: Accuracy (Acc), Precision (P), Recall (R) and macro F1-Score (F1) for binary complaint prediction (\pm std. dev.). Best results are in bold. \dagger indicates statistically significant improvement over BERT (Jin and Aletras, 2020) in STL (t-test, $p < 0.05$).

RoBERTa_{Emo}-DE models (**MTL-BiGRU-Att** and **MTL-BiGRU-Att-DE**).

Hyperparameters We train⁹ the **MTL-BiGRU-Att** and **MTL-BiGRU-Att-DE** model with the same hyperparameters as BiGRU-Att in complaint severity prediction. For the **MTL-Hard Sharing**, **MTL-Double Encoder** and **MTL-Gated Double Encoder** model, the hidden size of the stacked BiGRU encoder(s) and BiGRU-Att models is $h = 128$, $h \in \{64, 128, 256, 512\}$. We set β in **MTL-Gated Double Encoder** and the remaining parameters in three models to be the same as Rajamanickam et al. (2020). We train **MTL-M-RoBERTa_{Emo}** and **MTL-M-RoBERTa_{Emo}-DE** with a learning rate $l = 1e-6$, $l \in \{1e-5, 5e-6, 1e-6\}$. The rest of the parameters is the same as M-RoBERTa_{Emo} in the complaint severity prediction. The parameter α which controls the importance of the two losses is set to .1, $\alpha \in \{.001, .01, .1, .3, .5\}$.

8.3 Results

Table 4 shows results of the single-task learning (STL) and multi-task learning (MTL) models on the complaint identification task. Overall, we observe that all MTL models using M-RoBERTa_{Emo} perform better than the majority of STL models, indicating severity detection improves binary complaint identification. MTL-M-RoBERTa_{Emo} outperforms all other models achieving 88.2 macro F1, followed by MTL-M-RoBERTa_{Emo}-DE with

⁹We experiment with all MTL models using the same training and evaluation method as in the complaint severity prediction task.

88.1 F1. This confirms our hypothesis that complaint identification can be benefited by the complaint severity level information when jointly learning these two tasks simultaneously. Also, MTL-BiGRU-Att performs better than BiGRU-Att in STL achieving 75.4 F1 while the results of BiGRU-Att (74.5 F1) and MTL-BiGRU-Att-DE (74.1 F1) are comparable. We notice that the models proposed by Rajamanickam et al. (2020) (i.e. MTL-Hard sharing, MTL-Double Encoder and MTL-Gated Double Encoder) achieve low performance with only the MTL-Hard Sharing model performing slightly better than the others with 72.1 macro F1. We speculate that adding one or more extra BiGRU encoders before the BiGRU-Att model is an overly complex structure for our data set.

8.4 Analysis

We investigate the influence of recognizing severity levels of complaints on binary complaint identification in our MTL setting. We analyze predictive results by inspecting predictions from the previous best performing model BERT (STL) and MTL-M-RoBERTa_{Emo} models in a random fold (out of 10 CV folds). We observe that 9.8% of predictions flip, where the number of complaints flipping to non-complaints is noticeably larger (88.2%) than that of non-complaints flipping to complaints (11.8%). Similarly, we also compare predicted results between BiGRU-Att (STL) and MTL-BiGRU-Att in the same fold. The flipping percentage (6.9%) is lower than BERT and MTL-M-RoBERTa_{Emo} while the proportions of one class flipping to an-

Tweet	BERT	MTL-M-RoBERTa _{Emo}	Actual Label	Severity Label
What’s your secret to poaching eggs? Mine never look that good.	Complaint	Non-complaint	Non-complaint	No Complaint Severity
<URL> How bad do you really want a ps4 this year? Get a pre-owned playstation 4 at a very low dis <URL>	Complaint	Non-complaint	Non-complaint	No Complaint Severity
So, I’m now having to check my <USER> forester’s oil each month. Put 4 quarts in today, got about 2 out. #smh	Non-complaint	Complaint	Complaint	Disapproval
Is this how you fix the exhaust of your <USER> in #belarus? <URL>	Non-complaint	Complaint	Complaint	Blame

Table 5: Complaint classification examples by BERT (Jin and Aletras, 2020) and our MTL-M-RoBERTa_{Emo} compared to the actual labels.

other are consistent (83.4% and 16.6% respectively). These indicate that complaint severity information encapsulates complementary information for the model to predict non-complaints accurately.

Table 5 shows flipping examples from BERT (STL) and MTL-M-RoBERTa_{Emo}. From the first two rows, we see that the MTL model is not affected by negation (e.g. ‘never’) and negative terms (e.g. ‘bad’, ‘very low’) using the extra knowledge provided by the severity level prediction task. Also, in the last two examples, complaints are expressed in a more subtle way that rarely contains typical complaint-related terms. This indicates the MTL model is able to detect this type of complaints correctly because the severity level information encourages the model to learn to distinguish between such stylistic idiosyncrasies.

We further observe that 11.2% of wrong predictions remain the same for the two models, where complaints and non-complaints account for 59.0% and 41.0% respectively which means severity features benefit more posts that are complaints to be classified accurately. On the other hand, the model still has difficulty in predicting some non-complaint posts which might happen because of the lower performance of severity detection¹⁰ when used as an auxiliary task in the MTL setting.

9 Conclusion

We presented the first study on severity level of complaints in computational linguistics. We developed a publicly available data set of tweets labeled with four categories based on theory of pragmatics. We modeled complaint severity level prediction as a new multi-class classification task and conducted experiments using different transformer-based networks combined with linguistic features reaching

¹⁰Severity prediction is less accurate in MTL than in a single task setting.

up to 55.7 macro F1. We further used a multi-task learning setting to jointly model binary complaint prediction and complaint severity classification as an auxiliary task achieving new state-of-the-art performance on complaint detection (88.2 macro F1). In the future, we plan to apply our methods on a multilingual setting across different platforms.

Acknowledgments

We would like to thank Areej Alokaili, Daniel Preoȃiuc-Pietro and Danae Sanchez Villegas for their helpful feedback. NA is supported by ESRC grant ES/T012714/1.

Ethics Statement

Our work has received approval from the Ethics Committee of our institution (Ref. No 031504) and complies with Twitter data policy for research.¹¹

References

- Md Shad Akhtar, Asif Ekbal, and Erik Cambria. 2020. [How intense are you? Predicting intensities of emotions and sentiments using stacked ensemble](#). *IEEE Computational Intelligence Magazine*, 15(1):64–75.
- Irean Navas Alejo, Toni Badia, and Jeremy Barnes. 2020. [Cross-lingual Emotion Intensity Prediction](#). *arXiv preprint arXiv:2004.04103*.
- Nikolaos Aletras and Benjamin Paul Chamberlain. 2018. [Predicting twitter user socioeconomic attributes with network and language information](#). In *Proceedings of the 29th ACM International Conference on Hypertext and Social Media*, pages 20–24.
- Ron Artstein and Massimo Poesio. 2008. [Inter-Coder Agreement for Computational Linguistics](#). *Computational Linguistics*, 34(4):555–596.

¹¹<https://developer.twitter.com/en/developer-terms/agreement-and-policy>

- 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*.
- Rich Caruana. 1997. [Multitask Learning](#). *Machine learning*, 28(1):41–75.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. *arXiv preprint arXiv:1406.1078*.
- 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.
- Taner Danisman and Adil Alpkocak. 2008. Feeler: Emotion Classification of Text Using Vector Space Model. In *AISB 2008 Convention Communication, Interaction and Social Intelligence*, volume 1, page 53.
- 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.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- 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.
- Jiahua Jin, Xiangbin Yan, You Yu, and Yijun Li. 2013. Service Failure Complaints Identification in Social Media: A Text Classification Approach.
- Mali Jin and Nikolaos Aletras. 2020. Complaint identification in social media with transformer networks. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1765–1771.
- Leila Nasiri Kakolaki and Mohsen Shahrokhi. 2016. Gender Differences in Complaint Strategies among Iranian Upper Intermediate EFL Students.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*.
- 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.
- 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 Preoțiuc-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.
- 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 Nonnative Speakers of Hebrew. The Prag-matic Perspective.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- 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 Preoțiuc-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.
- 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*.

- Santhosh Rajamanickam, Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2020. Joint Modelling of Emotion and Abusive Language Detection. *arXiv preprint arXiv:2005.14028*.
- H Andrew Schwartz, Salvatore Giorgi, Maarten Sap, Patrick Crutchley, Lyle Ungar, and Johannes Eichstaedt. 2017. Dlatk: Differential Language Analysis Toolkit. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 55–60.
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to Fine-Tune BERT for Text Classification? In *China National Conference on Chinese Computational Linguistics*, pages 194–206. Springer.
- Donna Hurst Tatsuki. 2000. [If my complaints could passions move: An interlanguage study of aggression](#). *Journal of Pragmatics*, 32(7):1003–1017.
- Zhengxi Tian, Wenge Rong, Libin Shi, Jingshuang Liu, and Zhang Xiong. 2018. Attention Aware Bidirectional Gated Recurrent Unit Based Framework for Sentiment Analysis. In *International Conference on Knowledge Science, Engineering and Management*, pages 67–78. Springer.
- Suhataati 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.
- Adam Tsakalidis, Nikolaos Aletras, Alexandra I Cristea, and Maria Liakata. 2018. Nowcasting the stance of social media users in a sudden vote: The case of the greek referendum. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 367–376.
- Guda Van Noort and Lotte M Willemsen. 2012. [Online Damage Control: The Effects of Proactive Versus Reactive Webcare Interventions in Consumer-generated and Brand-generated Platforms](#). *Journal of Interactive Marketing*, 26(3):131–140.
- 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.
- Anbang Xu, Zhe Liu, Yufan Guo, Vibha Sinha, and Rama Akkiraju. 2017. [A New Chatbot for Customer Service on Social Media](#). In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, pages 3506–3510.
- Wei Yang, Luchen Tan, Chunwei Lu, Anqi Cui, Han Li, Xi Chen, Kun Xiong, Muzi Wang, Ming Li, Jian Pei, et al. 2019. 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.
- Yuxiang Zhang, Jiamei Fu, Dongyu She, Ying Zhang, Senzhang Wang, and Jufeng Yang. 2018. [Text Emotion Distribution Learning via Multi-Task Convolutional Neural Network](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 4595–4601.