

Rationale-Guided Few-Shot Classification to Detect Abusive Language

Punyajoy Saha*, Divyanshu Sheth, Kushal Kedia, Binny Mathew and Animesh Mukherjee

Indian Institute of Technology, Kharagpur

Abstract. Abusive language is a concerning problem in online social media. Past research on detecting abusive language covers different platforms, languages, demographics, etc. However, models trained using these datasets do not perform well in cross-domain evaluation settings. To overcome this, a common strategy is to use a few samples from the target domain to train models to get better performance in that domain (cross-domain few-shot training). However, this might cause the models to overfit the artefacts of those samples. A compelling solution could be to guide the models toward rationales, i.e., spans of text that justify the text’s label. This method has been found to improve model performance in the in-domain setting across various NLP tasks. In this paper, we propose RGFS (Rationale-Guided Few-Shot Classification) for abusive language detection. We first build a multitask learning setup to jointly learn rationales, targets, and labels, and find a significant improvement of 6% macro F1 on the rationale detection task over training solely rationale classifiers. We introduce two rationale-integrated BERT-based architectures (the RGFS models) and evaluate our systems over five different abusive language datasets, finding that in the few-shot classification setting, RGFS-based models outperform baseline models by about 7% in macro F1 scores and perform competitively to models finetuned on other source domains. Furthermore, RGFS-based models outperform LIME/SHAP-based approaches in terms of plausibility and are close in performance in terms of faithfulness.

Disclaimer: This paper contains material that many will find offensive or hateful. However, this cannot be avoided owing to the nature of the work.

1 Introduction

Abusive language has become a perpetual problem in today’s online social media. An ever-increasing number of individuals are falling prey to online harassment, abuse and cyber-bullying as established in a recent study by Pew Research [40]. In the online setting, such abusive behaviours can lead to traumatization of the victims [37], affecting them psychologically. Furthermore, widespread usage of such content may lead to increased bias against the target community, making violence normative [21]. Many gruesome incidents like the mass shooting at Pittsburgh synagogue¹, the Charlottesville car attack², etc. are all caused by perpetrators consuming/producing such abusive content.

In response, various social media companies have implemented

policies to moderate the content on their platforms³. These are primarily handled by moderators who manually delete posts that violate community guidelines specific to the platform [18]. The main issue with such a system is the sheer volume of content to be reviewed, which in many cases does not leave enough time for the moderator to arrive at a decision⁴. Furthermore, many moderators complain about psychological effects caused due to moderation of such abusive content. To help the moderators, various social media platforms are trying to proactively filter such abusive content using recent NLP architectures like transformers⁵. Such filtering techniques also need human support, in case of complex examples, but the manual effort of moderators should reduce a lot. This will benefit the moderation system as a whole.

However, reliably detecting abusive language is a challenging problem. Efforts using lexicon-based systems [13] and machine learning [10] have been made to detect abusive speech in online social media. As highlighted by Mishra [26], one of the issues is *domain shift*, i.e., when the training and test data come from different distributions. This domain shift is common in abusive language research due to the variability in annotation, demography and topics [38]. This calls for models/pipelines which can effectively domain transfer across different datasets with zero or few annotated datapoints. Though there has been a considerable volume of research in cross-domain transfer in other NLP problems like machine translation [7], the ever-changing nature of abusive language makes it a very crucial problem [23].

Another important aspect missing in the abuse detection pipeline is explainability [27]. A lot of the current research in NLP concentrates on not just detections but providing explanations behind the detections as well [9]. A subset of such research has been involved in creating datasets that contain the annotators’ reasoning in some form, i.e., the spans of text or rationales, textual reasons, etc. [12]. We hypothesize that the nucleus of abusiveness lies in certain text spans, i.e., rationales in typical hate posts. Hence, it is essential to give additional attention to these spans compared to the whole post. This attention can also be provided by including novel architectural changes in the computational model. As highlighted in [24], such rationales, when used as a feedback to the model, can also help in improving abuse classification and reduce the unintended bias of the model toward various target communities.

In this work, we investigate if the use of such rationales can help

* Corresponding Author. Email: punyajoy@iitkgp.ac.in

¹ https://en.wikipedia.org/wiki/Pittsburgh_synagogue_shooting

² https://en.wikipedia.org/wiki/Charlottesville_car_attack

³ <https://help.twitter.com/en/rules-and-policies/abusive-behavior> (Accessed on 05/02/2021)

⁴ <https://tinyurl.com/2m4jdw54>

⁵ <https://ai.facebook.com/blog/how-ai-is-getting-better-at-detecting-hate-speech/>

better cross domain few-shot classification⁶ in various datasets. We propose RGFS – Rationale-Guided Few-Shot Classification to Detect Abusive Language, which uses an attention framework to introduce rationales, i.e., text spans which justify the classification labels, into the prediction pipeline. To train models to predict rationales in text, we utilise a multitask framework which jointly learns labels, rationales, and targets for a given text. We use this rationale predictor - called BERT-RLT (BERT-Rationale-Label-Target) - to predict rationales for datapoints in an unseen dataset and integrate it with RGFS, using a few labeled samples for training. To evaluate the pipeline, we use five different datasets having different numbers of labels, collected from different timelines and using slightly different annotation guidelines. We observe that:

- The multitask framework that jointly learns rationales, labels and targets outperforms the model learning only rationales by 6% in terms of rationale classification macro F1 score.
- In the cross domain few-shot setting, our proposed model - RGFS outperforms BERT models by about 7% in terms of macro F1 scores and performs comparably to models already fine-tuned on a similar dataset.
- Predicted rationales used in RGFS-based models are more plausible than LIME/SHAP-based rationales. The predicted rationales outperform LIME/SHAP-based rationales by around 40% and 30% in terms of AUPRC and token-F1 scores respectively.
- Further, rationale-based explanation when utilised with the RGFS-based architecture achieves comparable performance to LIME/SHAP-based explanation in terms of faithfulness.

Please refer to Appendix⁷ for additional details about the architectures, dataset and discussion on rationales. The code is added here⁸.

2 Related work

2.1 Few-shot learning

Research in cross domain few-shot classification mostly focuses on instance selection [31]. In the semi-supervised setting, a few labelled as well as unlabelled datasets are used for the domain transfer [6]. Some of the methods used include maximising the label entropy for the unlabelled data [32] or continuing pre-training using some auxiliary task [19] before training with the labelled data.

In abusive language research literature, researchers have studied cross-domain performance across different datasets and found that the percentage of positive examples [35] and the in-domain performance of the classifier [16] are correlated with cross-domain performance. [42] used a multi-task learning framework where auxiliary tasks were learnt on datasets from different distributions to improve performance on the target task, using a simple classifier. In recent work, Fortuna [16] show that transformer-based models like BERT [11] are already more generalisable compared to the previous models.

Although cross-domain performance has been extensively studied for abusive language research literature, there is a lack of evaluation in few-shot classification setups. Aluru [1] performed an extensive study on several multilingual datasets in a few-shot setting where the authors focused on variation of abuse detection performance with respect to languages. Also, Stappen [34] introduced a new architecture

‘AXEL’ to improve cross lingual zero-shot and few-shot classification. Their study, performed across just two datasets, did not include explainability analysis.

2.2 Explainability / Interpretability

In recent years, NLP research has focused more on black box techniques at the expense of less interpretable models. As a remedy, different local post-hoc techniques like LIME [29] and TREPAN [8] have been introduced to explain the prediction of such black box models. With the advent of transformer models like BERT, a section of research is also focusing on understanding BERT’s inner working [36] by visualising its internal layers and utilising BERT’s attention to generate the explanation using reinforcement learning [28]. Once an explanation is generated, it is crucial to measure its reliability. One of the methods is to compare it with ground truth rationales [43]. DeYoung [12] compiled previous explainable works and provided several metrics to draw comparisons with ground truth rationales. In abusive language literature, a recent study [24] showed that rationales improve hate speech classification and also help in reducing unintended biases towards target communities.

In a recent work in the computer vision community, Schmidt [33] proposed a loss using a saliency map to help in-domain generalisation for image recognition. Furthermore, Zhang [45] proposed an Explain-Then-Predict framework for document level tasks where they first extracted the rationales and then used the extracted rationales only for prediction. This is different for short sentence classification, where removing the non-rationales can result in the sentence losing its meaning. There are also some research works that show improvement in low/few-shot tasks with human rationales; however, they mostly evaluate in in-domain settings. Melamud [25] study if pre-training with such rationales can help in few-shot performance. This study was further extended to include rationale-based self-training to improve few-shot performance [5]. Both these research works focus mostly on the performance of the models in light of rationales but do not discuss the explainability of these models.

In this work, we study if rationales can help in improving cross domain few-shot classification for abusive language detection if a target dataset has a few labelled samples but does not have annotated rationales. Furthermore, while previous studies have only utilised the rationale prediction as an auxiliary task, we use rationale predictions in an attention-based framework to improve upon base models in terms of both performance and explainability.

3 Datasets

In this section, we describe the datasets used in our work. To train the rationale extractor, we used the **HX: HateXplain** [24] dataset, which contains the classification label, rationales, and targets annotated per post. To provide the model with rationales as feedback, the rationales by each annotator are converted into Boolean vectors. Values in these Boolean vectors are 1 when the corresponding token (word) in the text is a part of a rationale. To create the *ground truth rationales*, each token in the text is denoted as a rationale if at least two annotators have highlighted it as a rationale. The final ground truth rationales are Boolean vectors, considering the above constraint.

3.1 What are rationales?

Rationales are sets of words or phrases in text that justify the text’s classification into a label or category. We consider spans of consecutive words marked as rationales by annotators on the HateXplain

⁶ For our case, cross domain few-shot classification is a variant of few-shot classification where we have a few samples from the target domain and can use some information from the source.

⁷ <https://rb.gy/kbwqz>

⁸ https://github.com/punayajoy/RGFS_ECAI

(HX) dataset for our analysis. The top 5 frequent rationales seen include common slur/derogatory terms like ‘ni**er, bi**h, k**e, ass. We note that about 93% of the 6021 unique rationale phrases have a count of 1. Examples of some rationales for classifying a text as hate speech include “evil mu*rat cult”, “can not speak properly lack basic knowledge of biology”, “jew is just a ni**er turned inside out”, etc. The average length of rationale phrases is about 6 words— many rationales are seen to have multiple words in them.

3.2 Datasets for evaluation

To evaluate our RGFS models, we use five popular abusive language datasets - **DA: Davidson** [10], **FA: Founta** [17], **OD: Olid** [44], **BA: Basile** [3] and **WH: Waseem & Hovy** [41] and use the same labels as present in these datasets. These datasets differ in their choice of class labels, methods of annotation, geographies & diversities of source users in their posts, communities & groups targeted and periods of data collection. Details about the datasets are presented in Table 1 (see Appendix for more details about the datasets).

3.3 A note about the datasets

Biases in annotations of these datasets have been noted by earlier authors. The definition of abuse varies across different datasets [30] and oftentimes these definitions are incompatible [15]. Awal [2] noted inconsistencies among three popular abusive language datasets: Davidson, Founta and Waseem [41]. However, with such a subjective and difficult annotation process, such inconsistencies are unavoidable.

4 Methodology

4.1 Base model

As a base model, we use BERT [11] pre-trained on English data.⁹ We perform uniform pre-processing by normalising usernames and links in datasets. Finally, we pass the inputs through the text processing pipeline — *ekphrasis* [4]. As a baseline for classification, we attach a classifier layer on top of BERT (henceforth, BERT-L model).

4.2 Rationale extraction

To detect rationales, we add a token classifier layer [11] to classify each token based on whether it is a rationale or not. We use binary cross entropy between the predicted and the ground truth rationales to calculate the loss for token classification, denoted by $L_{rationale}$. This is different from the method illustrated in the following paper [24], where they attempt to directly change the attention weights inside the model. Along with the rationale classifier, we add two more parallel classification layers— one to classify labels (loss L_{label}), and another to classify targets (loss L_{target}). The classification of labels is a multi-class problem, whereas target classification is a multi-label problem. Furthermore, we also consider three classes (hate speech, offensive speech, and normal) and two classes (abusive or not abusive) as two variants for the label classification task. The final loss of the model (L_{total}) is shown in equation 1, where β controls the impact of the rationales and γ controls the impact of the targets. We denote this as the BERT-Rationale-Label-Target (BERT-RLT) classifier.

$$L_{total} = L_{label} + \beta \cdot L_{rationale} + \gamma \cdot L_{target} \quad (1)$$

⁹ We use the bert-base-uncased model having 12-layers, 768-hidden, 12-heads, 110M parameters.

4.3 Classification using rationales

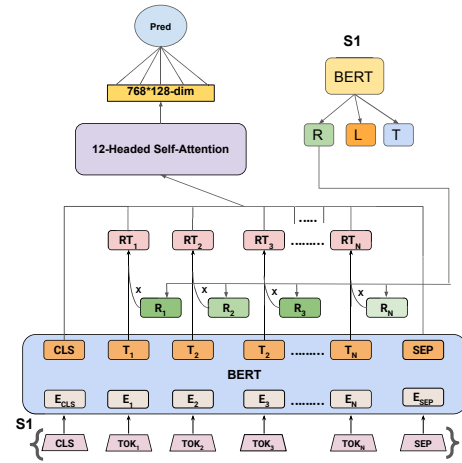


Figure 1. The RGFS-SA architecture. BERT-RLT is the model that predicts the rationales in a sentence $S1$, which are then added onto the BERT model using a self attention layer.

RGFS-SA: Input sentences are first run through BERT-RLT (BERT-Rationale-Label-Target) to obtain probability scores for each token in the sentence being a positive rationale, i.e., a rationale belonging to the hate/abusive speech class. The inputs are then passed through BERT, and the last hidden state (LHS) output from BERT is updated by multiplying each token’s output by its corresponding rationale score from the rationale-extractor, thereby assigning relative importance to each token with respect to other tokens for classification. A multi-headed self-attention layer is then constructed on top of the updated LHS; the updated LHS is taken as the query, key, and value vectors for attention. The resulting output is passed through a fully connected layer to generate final predictions. The weights of BERT-RLT (the rationale-extractor) are frozen and are not updated during the training of the classifier. We denote this classifier as RGFS-SelfAttention-Classifier (RGFS-SA). Figure 1 shows a schematic of this architecture.

RGFS-CA: Instead of using self-attention on the updated LHS in the RGFS-SA model, we introduce a cross-attention layer between the CLS-pooled output obtained from BERT and the updated LHS here. We aim to have the model learn to get a representation of the complete sentence in its CLS-pooled output that works well with the architecture we add over BERT for the task. This CLS-pooled output is taken as the query vector, and the LHS*(token-wise rationale scores) is used as the key and value vectors. The weights of the rationale-extractor are kept frozen as in the previous case. This classifier is denoted as RGFS-CrossAttention-Classifier (RGFS-CA). Figure 2 shows the schematic.

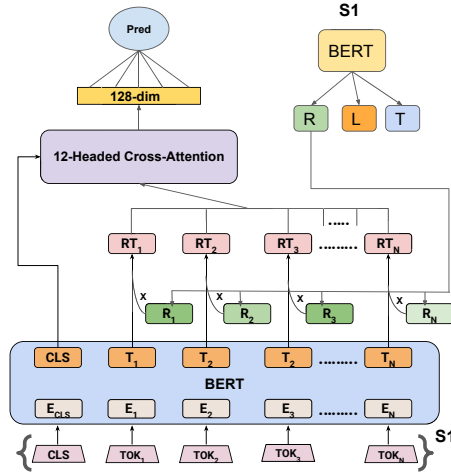
Note that the RGFS model was not trained on the past dataset or **HX**. RGFS only uses a rationale predictor (which was trained on **HX**) to predict the rationales on the cross-domain dataset. The classifier module of the RGFS model was trained with the target dataset mentioned in section 3.2.

4.4 Metrics

To evaluate the models, we rely on classification performance and explainability metrics. We use macro F1-score to measure classification

Table 1. The total dataset size and the number of datapoints available per label for each dataset being used.

Dataset	Abbv.	Labels(#numbers of datapoints)	Size
HateXplain	HX	<i>Hate speech</i> (5,935), <i>Offensive</i> (5,480), <i>Normal</i> (7,814)	19,229
Founta et al.	FA	<i>Hateful</i> (4,948), <i>Abusive</i> (27,037), <i>Normal</i> (53,790)	85,775
Davidson et al.	DA	<i>Hate speech</i> (1,430), <i>Offensive</i> (19,190), <i>Normal</i> (4,163)	24,783
OLID	OD	<i>Offensive</i> (4,640), <i>Not Offensive</i> (9,460)	14,099
Basile et al.	BA	<i>Hateful</i> (5,390), <i>Non-hateful</i> (7,415)	12,805
Waseem & Hovy	WH	<i>Racism</i> (13), <i>Sexism</i> (2855) <i>Normal</i> (8050)	10,018

**Figure 2.** The RGFS-CA architecture, where the BERT-RLT model predicts the rationales for a sentence $S1$, which are then added onto the BERT model using a cross attention layer.

performance, which is a standard metric for imbalanced datasets.

For explainability, we use the framework given by DeYoung [12], we evaluate plausibility and faithfulness of the models. *Plausibility* refers to how convincing the model’s interpretation is to humans, while *faithfulness* aims to measure the reasoning of the model to arrive at a prediction [20].

Plausibility: To measure plausibility, we consider both discrete and soft selection metrics. As for discrete metrics, we report the *IOU F1* and *token F1* scores, and for soft selection we use *AUCPRC* scores [12]. The token F1 is derived from token-wise precision and recall scores between predicted and ground truth rationales. DeYoung [12] defines IOU at the token level – for two spans, it is the size of the overlap of the tokens they cover divided by the size of their union. In IOU F1 metric, a prediction is considered a match if the overlap with any of the ground truth rationales is more than 0.5. These partial matches are used to calculate the IOU F1 score. Thus, while the token-level F1 score (token F1) measures the token level matching, IOU F1-score also awards credits to partial matches [14].

Faithfulness: To measure faithfulness, we report two complementary metrics: *comprehensiveness* and *sufficiency* [12].

- **Comprehensiveness:** To measure comprehensiveness, we create a contrasting example \tilde{x}_i , for each post x_i , where \tilde{x}_i is calculated by removing the predicted rationales r_i^{10} from x_i . Let $m(x_i)_j$ be the original prediction probability provided by a model m for the predicted class j . $m(x_i \setminus r_i)_j$ is then defined as the predicted probability of \tilde{x}_i ($= x_i \setminus r_i$) by the model m for the class j . We would expect the model prediction to be lower on remov-

ing the rationales. We can measure this as: *comprehensiveness* $= m(x_i)_j - m(x_i \setminus r_i)_j$. A high value of comprehensiveness implies that the rationales in the text are influential in the model’s prediction.

- **Sufficiency:** Measures the degree to which extracted rationales are adequate for a model to make a prediction. This can be measured as: *sufficiency* $= m(x_i)_j - m(r_i)_j$. A low sufficiency score would imply that the model’s performance on text containing just the rationales is close to the performance of the model on the complete text.

4.5 Rationale annotation

In order to evaluate the predicted rationales on the target domain/dataset, we sample 50 datapoints from the abusive class (hateful/offensive) from the test splits of the five target datasets (FA, DA, OD, BA, and WH) and annotate the rationales in them.

Six annotators participated in the process, including 2 bachelors students and 4 PhD students. All the annotators were aged between 20-30 years, and they all had experience in the domain of abusive language research. We used Docanno¹¹, an open source annotation platform to perform the annotation task. Each annotator was given a secure account and an interface where the posts were shown (see Appendix). Each post was labelled by two annotators. For each post, the annotators were required to classify the post as abusive or normal. If they found the post abusive, they had to mark rationales in the form of phrases in the text following the guidelines given by past research [24]. We consider a post abusive when both annotators have marked it as abusive.

The average Jaccard overlap between the annotators across different datasets is 0.64. We also simulate random rationales and find the average Jaccard overlap to be around 0.30 (see in Appendix).

When tokenizing the post (having rationales), we divide the post into rationale and non-rationale phrases based on the rationale annotation. Then these phrases are individually tokenized. Now for all the tokens in rationale phrases, the rationale label is 1 while this is 0 for the tokens in non rationale phrases. In this way we generate a vector of 1’s and 0’s for each phrase, i.e., rationale vector. Finally, we concatenate the tokenized phrases and their corresponding rationale vectors.

4.6 Experimental setup

The rationale extraction performance on the HateXplain dataset as shown in Table 2 is compared using the same *train:development:test* splits of 8:1:1 as used by Mathew [24]. For the evaluation datasets, we maintain a stratified split of *train:development:test* in the ratio 7:1:2, similar to previous research on cross domain evaluation [35]¹².

¹¹ <https://github.com/doccano/doccano>

¹² The amount of data per split is noted in Appendix.

¹⁰ We select the top 5 tokens as the rationales as per Mathew [24]

We use samples from the training dataset to train the few-shot detection models, the validation dataset is used to find the best model while training. Finally, we report the result on the test data.

We set the token length to 128 for reducing model size¹³. All our results are reported on the test set of the corresponding dataset (either source or target). We also highlight the best performance using **bold** font in all the tables. In Tables 2 and 3, we also show the second best using underline. For cross domain few-shot evaluation, we use 50, 100, 150, and 200 training datapoints from each class to train the models in the new domains¹⁴. We create five such different random sets of 50, 100, 150, and 200 datapoints for each target dataset to make our evaluations robust and we report average performance across these sets.

For all the models trained, we vary the learning rate as the main hyperparameter, taking up the following values: $1e^{-5}$, $3e^{-5}$ and $5e^{-5}$. With regards to equation 1, we vary β and γ through 1, 2, 5, 10, 100, achieving maximum performance on the validation set with $\beta = 2$ and $\gamma = 10$ for the BERT-RLT model. Once the best model is found for in-domain performance, we fix the model for all cross-domain evaluations. The learning rate for the cross domain few-shot classification models is fixed at $1e^{-5}$. Note that we did not explicitly provide a comparison with the earlier papers, since most of them use inferior models compared to BERT.

5 Results

5.1 Rationale extraction

We show the different variations to train the rationale classifier in Table 2. At first, we consider the rationale classifier alone. This results in a macro F1 score of 0.71. If we consider random rationales, we obtain a macro F1 score of 0.49 which highlights the difficulty of predicting rationales. Next, we consider two variants of the label classification problem - (i) in the two class variant, we consider toxic and non-toxic as the final labels. (ii) In the three class variant, we consider - hate speech, offensive and normal - similar to the HateXplain paper.

We do not see an improvement in rationale classification performance when we add in the label classification task (model denoted by BERT-RL) along with rationale classification. We then add the target classification task as well, the new model being denoted by BERT-RLT. In the three class variant, the BERT-RLT model performs slightly better than the original rationale classification model. We get the best performance for BERT-RLT in the two class variant. We observe that BERT-RLT outperforms the base rationale classifier by 4 F1 points. Henceforth, we will use this BERT-RLT model as our rationale predictor. Across all the models, the rationale accuracy (R-Acc) is > 0.95 since it is a highly imbalanced task.

The performance on the label classification task remains almost the same for both the variants (two class and three class). For the three class variant, the macro F1 score (L-F1 in Table 2) is in the range 0.67-0.69, whereas for the two class variant it is 0.77. The targets classification task is extremely difficult with 22 different classes (see in Appendix). For both two and three class variants of the BERT-RLT model, macro F1 score reaches around 0.21. We mostly resort to BERT models as we did not find much difference with other transformer models like RoBERTa (1% difference in rationale classification F1 score).

Table 2. Performance of different models on the HateXplain test set. The models are denoted by BERT-X where the letters in X denote the tasks that particular variant uses. R denotes rationale, L denotes label, T denotes target classification. F1 denotes the macro F1-score and Acc denotes the accuracy. L-F1: label macro F1, L-Acc: label accuracy, R-F1: rationale macro F1 and R-Acc: rationale accuracy.

Model	L-F1	L-Acc	R-F1	R-Acc
BERT-R-RANDOM	–	–	0.49	0.97
BERT-R	–	–	0.66	0.97
Two classes				
BERT-L	0.77	0.78	–	–
BERT-RL	0.77	0.78	0.66	0.96
BERT-RLT	0.77	0.78	0.70	0.98
Three classes				
BERT-L	0.69	0.70	–	–
BERT-RL	0.68	0.69	0.64	0.96
BERT-RLT	0.67	0.68	<u>0.68</u>	<u>0.97</u>

5.2 Similarity across different domains

A well-known strategy to improve cross domain few-shot classification is to use a model trained on a similar domain. Following this strategy, we aim to select the best source among the datasets (excluding the HateXplain dataset). We first calculate the normalised term distribution for all the posts in a particular dataset. Next we use the pairwise cosine similarity between the term distribution of the two datasets [31]. For a particular target dataset, we first fine-tune the model using all the training points of the best source dataset (see the Table in Appendix for best source-target pairs). We name this model as BERT-L-DOM, where the DOM refers to the best source domain for that particular target dataset. We then add a new last layer on this model to train on the training samples of the target dataset. We also compare the similarity value of the target dataset with the HateXplain (HX) dataset and find there is an average drop of 25% across different target datasets. Intuitively, this can be due to the different sources (Twitter and Gab) and a different timeline of data collection.

Table 3. Comparison of cross domain few-shot performance for BERT-L, BERT-L-DOM, RGFS-CA and RGFS-SA. Here, we train a particular model using few labeled datapoints and evaluate on the test data from the same domain. Each cell shows the average macro F1-score on the test dataset after running the model for five times with different sets of datapoints for fine-tuning. BERT-L-DOM is the best cross domain model based on the performance corresponding to each dataset.

Model	Data	No. of training datapoints (per label)				
		50	100	150	200	All
BERT-L	DA	0.590	0.609	0.671	0.674	0.768
BERT-L-DOM		0.656	<u>0.682</u>	0.689	0.684	0.766
RGFS-CA		0.636	0.681	0.700	0.699	0.766
RGFS-SA		<u>0.653</u>	0.693	<u>0.695</u>	0.711	0.752
BERT-L	OD	0.588	0.651	0.689	0.703	0.785
BERT-L-DOM		0.561	0.637	0.693	0.703	0.776
RGFS-CA		<u>0.627</u>	<u>0.668</u>	<u>0.702</u>	<u>0.712</u>	0.778
RGFS-SA		0.648	0.686	0.710	0.725	0.779
BERT-L	BA	0.634	0.647	0.671	0.673	0.795
BERT-L-DOM		<u>0.650</u>	0.662	<u>0.684</u>	<u>0.700</u>	0.781
RGFS-CA		0.648	0.673	0.683	0.690	0.787
RGFS-SA		0.658	0.673	0.686	0.701	0.791
BERT-L	FA	0.597	0.653	0.677	0.687	0.779
BERT-L-DOM		<u>0.655</u>	0.679	0.694	<u>0.702</u>	0.776
RGFS-CA		0.649	0.691	0.707	0.709	0.773
RGFS-SA		0.665	0.692	0.700	0.709	0.784
BERT-L	WH	0.692	0.733	0.757	0.775	0.857
BERT-L-DOM		0.755	0.771	0.780	0.787	0.844
RGFS-CA		<u>0.724</u>	0.735	0.752	0.767	0.858
RGFS-SA		0.719	<u>0.741</u>	<u>0.757</u>	0.772	0.846

¹³ We consider first 128 tokens when > 128 tokens (0.05% cases).

¹⁴ None of the datasets have less than 200 datapoints for any label.

Table 4. Examples where RGFS models perform better than BERT-L-DOM. The highlighted words are rationales identified within RGFS which enable it to classify correctly while BERT-L-DOM fails to do so.

GROUND TRUTH	RGFS-SA	BERT-L-DOM	TEXT
hateful	hateful	normal	Wow they tried to give me a FEMALE DOCTOR TODAY. I didn't even know those EXISTED...
hateful	hateful	normal	user i am convinced the rapefugee invasion is an integral stage leading to this. considering their victims tend to be minors.

5.3 Cross domain few-shot classification

We perform cross domain few-shot classification by considering 5 random sets of 50, 100, 150, 200 samples per label from each training dataset and use it to train different model variants. As a baseline, we consider the BERT-L model which is finetuned directly on the target dataset. We compare this model with our proposed models - RGFS-SA and RGFS-CA which use an attention framework to include rationales in the architecture.

When training with 50 samples, both the RGFS models perform better than the BERT-L model (see Table 3). In terms of macro F1 score, the RGFS-SA model outperforms the BERT-L by around 4 points. On the other hand, the RGFS-CA model outperforms the BERT-L by around 3.5 points. The highest difference between RGFS-SA/CA and BERT-L models appears for the FA dataset and the least for WH dataset. This suggests that the classification task for WH dataset might be easier compared to the FA dataset. The rationales predicted by the BERT-RLT help in the latter case to improve the classification in a cross domain few-shot setting. With increase in the number of datapoints, the difference between the RGFS models and the BERT-L models reduces but stays significant across different datasets except for the WH dataset.

We also include the best cross domain model per dataset based on the similarity metric (see in Appendix). The best model trained on a source dataset (BERT-L-DOM) is trained again on this new dataset with a new last linear classification layer, on the number of the samples of the target dataset being used. The RGFS-SA model outperforms BERT-L-DOM in low data settings (50, 100 datapoints) for the OD, BA and FA datasets. In fact, with 100 training points, the RGFS-SA model beats the BERT-L-DOM model for all datasets except WH (some examples where BERT-L-DOM fails are noted in Table 4). In addition, we also apply the few shot detection method used in [1] where we consider all the other datasets source dataset except the target dataset. We find that on average the BERT-L-DOM approach is 2% ($\sigma = 0.11$) better than this method. Hence we use the BERT-L-DOM for further experiments. Finally, the RGFS-SA models outperform these models by 4% ($\sigma = 0.06$) on average.

Are the predicted rationales useful? Here, we attempt to understand the impact of the predicted rationales in the overall RGFS architecture. Our approach is to dilute the importance of the rationales and observe how this affects the overall performance of the model. The rationale predictor predicts a score per token (can be $+ve/-ve$), where a higher score means more propensity to become a rationale. We select a threshold for each dataset based on the higher quartile value in the score distribution. We assume tokens with scores higher than the threshold (i.e., the top 25 percentile scores) represent the rationales and update their scores to a very low value (-4)¹⁵ to reduce

their importance.

The rest of the tokens are assigned a random score based on a uniform distribution $-\mathcal{U}_{(-5,5)}$ ¹⁶. This modified vector is passed through softmax function to create the random rationale vector and the experiments with 50 datapoints in Table 3 is repeated. We observe that for this change there is an average percentage drop of around 6% in F1 score for RGFS-CA and RGFS-SA respectively, with the highest drop being observed for the Davidson dataset. We also add a qualitative analysis of how the rationales help in the Appendix.

5.4 Explainability

In order to evaluate the predicted rationales, we take the models which are domain adapted using 5 different random sets of 50 datapoints (per dataset) from the training dataset¹⁷, and the rationale annotated data (see in Appendix). To compare with models which do not predict rationales within them, we pass the model output through two explainability methods namely, LIME [29] and SHAP [22] to get importance scores for each word. After getting the raw vector of importance scores, we normalise the score between 0 and 1 using min-max normalisation. For the RGFS models, we perform the same operation but with the rationales predicted by the BERT-RLT. Next, we evaluate these rationales.

Plausibility: For plausibility, we first consider the soft token metric - AUPRC. The rationales used in RGFS models outperform the BERT-L-DOM + LIME configuration by 2 points on average and BERT-L-DOM + SHAP configuration by 4 points on average (see Table 6). The difference in some of the datasets (BA, WH, OD) is more significant than the others (DA and FA dataset). In terms of token F1, RGFS outperforms the BERT-L-DOM + LIME configuration by around 1.5 F1 points and BERT-L-DOM + SHAP around 2 F1 points (see Table 6). The results are closer for the IOU scores since it also considers partial matches but here again the predicted rationales outperform LIME/SHAP rationales across all datasets except for WH. Overall, we observe that the rationales predicted using the BERT-RLT model provide more plausible explanations than the rationales generated using LIME/SHAP across all the target datasets.

Faithfulness: We measure the faithfulness of the rationales using *sufficiency* and *comprehensiveness* to understand whether the rationales act as correct reasoning behind the model predictions. In terms of sufficiency, we observe that the RGFS-SA model outperforms the RGFS-CA across DA, OD, BA datasets (see Table 7). The RGFS based models overall outperform the BERT-L-DOM + LIME configuration across all datasets except WH but underperforms when compared to the SHAP based configuration. In terms of comprehensiveness, BERT-L-DOM is slightly superior to the RGFS-SA model. Additional artefacts in the model might be a reason for this slight underperformance. SHAP based rationales are much worse compared to other variations in terms of comprehensiveness.

6 Conclusion

In the field of abusive language, one of the major issues in building better detection is the variability in the context across demographics and the subjective nature of abuse. Further, since platforms are coming up with stronger moderation algorithms, the language of abuse is also changing to bypass such moderation. The paper [23] discussed in detail the temporal nature of hate speech. Hence, many researchers

¹⁵ The value of 4 is based on the observed distribution.

¹⁶ The value of 5 is based on the observed distribution.

¹⁷ Corresponding to column 1 in Table 3.

Table 5. Ablation study: Percentage drops in F1-scores upon using random rationales (RR) with RGFS models instead of rationales obtained from the BERT-RLT rationale predictor. Random rationale weights are assigned to non-rationales identified by BERT-RLT, while detected rationales are assigned low rationale weights.

Dataset	% drop with RR Vs. RGFS-CA	% drop with RR Vs. RGFS-SA
DA	-14.97%	-13.49%
OD	-13.10%	-5.68%
BA	-2.56%	-4.18%
FA	-3.87%	-6.12%
WH	-4.37%	-3.08%

Table 6. Average AUPRC, token-F1 and IOU-F1 scores for the rationales predicted by the models which were trained using different sets of 50 datapoints. For the models not utilising rationales in their architecture, LIME and SHAP are used to predict the rationales. BERT-L-DOM is the best cross-domain model for each dataset.

Model	Data	Plausibility		
		AUPRC	token-F1	IOU-F1
BERT-L-DOM + LIME	DA	0.77	0.52	0.21
BERT-L-DOM + SHAP		0.45	0.37	0.14
RGFS-CA/SA		0.84	0.58	0.26
BERT-L-DOM + LIME	OD	0.49	0.36	0.10
BERT-L-DOM + SHAP		0.37	0.32	0.07
RGFS-CA/SA		0.68	0.54	0.11
BERT-L-DOM + LIME	BA	0.63	0.43	0.19
BERT-L-DOM + SHAP		0.46	0.34	0.14
RGFS-CA/SA		0.76	0.55	0.23
BERT-L-DOM + LIME	FA	0.61	0.46	0.11
BERT-L-DOM + SHAP		0.48	0.36	0.06
RGFS-CA/SA		0.67	0.54	0.13
BERT-L-DOM + LIME	WH	0.57	0.42	0.01
BERT-L-DOM + SHAP		0.50	0.35	0.01
RGFS-CA/SA		0.84	0.63	0.01

Table 7. Average comprehensiveness scores (Comp) and the sufficiency scores (Suff) for the rationales predicted by the models which were trained using different sets of 50 datapoints. For the models not utilising rationales in their architecture, LIME and SHAP are used to predict the rationales. BERT-L-DOM is the best cross domain model corresponding to each dataset. For sufficiency scores, lower values are better.

Model	Data	Faithfulness	
		Suff.(↓)	Comp.
BERT-L-DOM + LIME	DA	-0.03	0.67
BERT-L-DOM + SHAP		-0.29	-0.14
RGFS-CA		-0.06	0.11
RGFS-SA		-0.08	0.25
BERT-L-DOM + LIME	OD	-0.02	0.09
BERT-L-DOM + SHAP		-0.10	-0.09
RGFS-CA		-0.04	0.04
RGFS-SA		-0.08	0.29
BERT-L-DOM + LIME	BA	-0.05	0.34
BERT-L-DOM + SHAP		-0.38	-0.38
RGFS-CA		-0.04	0.06
RGFS-SA		-0.07	0.19
BERT-L-DOM + LIME	FA	-0.02	0.40
BERT-L-DOM + SHAP		-0.09	0.09
RGFS-CA		-0.13	0.09
RGFS-SA		-0.11	0.26
BERT-L-DOM + LIME	WH	-0.07	0.20
BERT-L-DOM + SHAP		-0.09	-0.09
RGFS-CA		0.01	0.03
RGFS-SA		0.06	0.06

are aiming to create dynamic datasets [39], since static datasets cannot serve as consistent benchmarks for abusive language research. In addition to such dynamic datasets, we need models that can learn using a few samples from target domains and are more explainable in order to identify the drawbacks of the models in a new domain.

In this paper, we show that models that utilise rationales can perform better in cross domain few-shot classification than the models

without them. These models also provide competitive performance as compared to the best model pre-trained on a source dataset (BERT-L-DOM), which has already been trained using more abusive examples. Further, the rationale predictors of models like RGFS provide more plausible explanations compared to traditional LIME/SHAP-based explanations, whilst slightly under-performing in terms of faithfulness. In future, we would focus on creating a typology of the errors in classification and rationale extraction task. We also plan to extend this framework to other languages.

References

- [1] Sai Saketh Aluru, Binny Mathew, Punyajoy Saha, and Animesh Mukherjee, 'A deep dive into multilingual hate speech classification', in *ECML PKDD 2020*, pp. 423–439. Springer International Publishing, (2021).
- [2] Md Rabiul Awal, Rui Cao, Roy Ka-Wei Lee, and Sandra Mitrović, 'On analyzing annotation consistency in online abusive behavior datasets', *arXiv preprint arXiv:2006.13507*, (2020).
- [3] Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti, 'SemEval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter', in *SemEval 2019*, pp. 54–63. Association for Computational Linguistics, (June 2019).
- [4] Christos Baziotis, Nikos Pelekis, and Christos Doukeridis, 'Datastories at semeval-2017 task 4: Deep lstm with attention for message-level and topic-based sentiment analysis', in *SemEval-2017*, pp. 747–754, Vancouver, Canada, (August 2017). Association for Computational Linguistics.
- [5] Meghana Moorthy Bhat, Alessandro Sordani, and Subhabrata Mukherjee, 'Self-training with few-shot rationalization: Teacher explanations aid student in few-shot nlu', *arXiv preprint arXiv:2109.08259*, (2021).
- [6] Ziqiu Chi, Zhe Wang, Mengping Yang, Wei Guo, and Xinlei Xu, 'Better embedding and more shots for few-shot learning', in *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, ed., Lud De Raedt, pp. 2874–2880. International Joint Conferences on Artificial Intelligence Organization, (7 2022). Main Track.
- [7] Chenhui Chu and Rui Wang, 'A survey of domain adaptation for machine translation', *Journal of Information Processing*, **28**, 413–426, (2020).
- [8] Roberto Confalonieri, Tillman Weyde, Tarek R Besold, and Fermín Moscoso del Prado Martín, 'Trepan reloaded: a knowledge-driven approach to explaining artificial neural networks', *arXiv preprint arXiv:1906.08362*, (2019).
- [9] Marina Danilevsky, Kun Qian, Ranit Aharonov, Yannis Katsis, Ban Kawas, and Prithviraj Sen, 'A survey of the state of explainable ai for natural language processing', *arXiv preprint arXiv:2010.00711*, (2020).
- [10] Thomas Davidson, Dana Warmesley, Michael Macy, and Ingmar Weber, 'Automated hate speech detection and the problem of offensive language', *ICWSM '17*, pp. 512–515, (2017).
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, 'BERT: Pre-training of deep bidirectional transformers for language understanding', in *NAACL*, pp. 4171–4186, Minneapolis, Minnesota, (June 2019). ACL.
- [12] Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C Wallace, 'Eraser: A benchmark to evaluate rationalized nlp models', in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4443–4458, (2020).

- [13] Mai ElSherief, Vivek Kulkarni, Dana Nguyen, William Yang Wang, and Elizabeth Belding, 'Hate lingo: A target-based linguistic analysis of hate speech in social media', in *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12, (2018).
- [14] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman, 'The pascal visual object classes (voc) challenge', *International journal of computer vision*, **88**(2), 303–338, (2010).
- [15] Paula Fortuna, Juan Soler, and Leo Wanner, 'Toxic, hateful, offensive or abusive? what are we really classifying? an empirical analysis of hate speech datasets', in *LREC*, pp. 6786–6794, Marseille, France, (May 2020). European Language Resources Association.
- [16] Paula Fortuna, Juan Soler-Company, and Leo Wanner, 'How well do hate speech, toxicity, abusive and offensive language classification models generalize across datasets?', *Information Processing and Management*, **58**(3), 102524, (2021).
- [17] Antigoni-Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis, 'Large scale crowdsourcing and characterization of twitter abusive behavior', in *ICWSM 2018*. AAAI Press, (2018).
- [18] Tarleton Gillespie, *Custodians of the Internet: Platforms, content moderation, and the hidden decisions that shape social media*, Yale University Press, 2018.
- [19] Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith, 'Don't stop pretraining: Adapt language models to domains and tasks', in *Proceedings of ACL*, (2020).
- [20] Alon Jacovi and Yoav Goldberg, 'Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness?', in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4198–4205, Online, (2020). Association for Computational Linguistics.
- [21] Aliza Luft. Dehumanization and the normalization of violence: It's not what you think – items. <https://tinyurl.com/s4mxz7j2>, May 2019. (Accessed on 05/02/2021).
- [22] Scott M Lundberg and Su-In Lee, 'A unified approach to interpreting model predictions', in *Advances in Neural Information Processing Systems 30*, eds., I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, 4765–4774, Curran Associates, Inc., (2017).
- [23] Binny Mathew, Anurag Illendula, Punyajoy Saha, Soumya Sarkar, Pawan Goyal, and Animesh Mukherjee, 'Hate begets hate: A temporal study of hate speech', *Proc. ACM Hum.-Comput. Interact.*, **4**(CSCW2), (oct 2020).
- [24] Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee, 'Hatexplain: A benchmark dataset for explainable hate speech detection', in *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 14867–14875, (2021).
- [25] Oren Melamud, Mihaela Bornea, and Ken Barker, 'Combining unsupervised pre-training and annotator rationales to improve low-shot text classification', in *EMNLP-IJCNLP*, pp. 3884–3893, Hong Kong, China, (November 2019). ACL.
- [26] Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova, 'Tackling online abuse: A survey of automated abuse detection methods', *arXiv preprint arXiv:1908.06024*, (2019).
- [27] Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova, 'The user behind the abuse: A position on ethics and explainability', *arXiv preprint arXiv:2103.17191*, (2021).
- [28] Runliang Niu, Zhepei Wei, Yan Wang, and Qi Wang, 'Attexplainer: Explain transformer via attention by reinforcement learning', in *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, ed., Lud De Raedt, pp. 724–731. International Joint Conferences on Artificial Intelligence Organization, (7 2022). Main Track.
- [29] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin, '"why should I trust you?": Explaining the predictions of any classifier', in *ACM SIGKDD*, pp. 1135–1144, (2016).
- [30] Björn Roß, Michael Rist, Guillermo Carbonell, Benjamin Cabrera, Nils Kurosky, and Michael Maximilian Wojatzki, 'Measuring the reliability of hate speech annotations: The case of the european refugee crisis, Nov 2016.
- [31] Sebastian Ruder and Barbara Plank, 'Learning to select data for transfer learning with bayesian optimization', *arXiv preprint arXiv:1707.05246*, (2017).
- [32] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko, 'Semi-supervised domain adaptation via minimax entropy', *ICCV*, (2019).
- [33] Robin M Schmidt, 'Explainability-aided domain generalization for image classification', *arXiv preprint arXiv:2104.01742*, (2021).
- [34] Lukas Stappen, Fabian Brunn, and Björn Schuller, 'Cross-lingual zero- and few-shot hate speech detection utilising frozen transformer language models and axel', *arXiv preprint arXiv:2004.13850*, (2020).
- [35] Steve Durairaj Swamy, Anupam Jamatia, and Björn Gambäck, 'Studying generalisability across abusive language detection datasets', in *CoNLL*, pp. 940–950, Hong Kong, China, (November 2019). Association for Computational Linguistics.
- [36] Ian Tenney, James Wexler, Jasmijn Bastings, Tolga Bolukbasi, Andy Coenen, Sebastian Gehrmann, Ellen Jiang, Mahima Pushkarna, Carey Radebaugh, Emily Reif, and Ann Yuan, 'The language interpretability tool: Extensible, interactive visualizations and analysis for NLP models', in *EMNLP*, pp. 107–118, Online, (October 2020). Association for Computational Linguistics.
- [37] Janikke Solstad Vedeler, Terje Olsen, and John Eriksen, 'Hate speech harms: A social justice discussion of disabled norwegians? experiences', *Disability & Society*, **34**(3), 368–383, (2019).
- [38] Bertie Vidgen and Leon Derczynski, 'Directions in abusive language training data: Garbage in, garbage out', *arXiv preprint arXiv:2004.01670*, (2020).
- [39] Bertie Vidgen, Tristan Thrush, Zeerak Waseem, and Douwe Kiela, 'Learning from the worst: Dynamically generated datasets to improve online hate detection', *arXiv preprint arXiv:2012.15761*, (2020).
- [40] Emily A Vogels. The state of online harassment | pew research center. <https://www.pewresearch.org/internet/2021/01/13/the-state-of-online-harassment/>, 2021. (Accessed on 05/02/2021).
- [41] Zeerak Waseem and Dirk Hovy, 'Hateful symbols or hateful people? predictive features for hate speech detection on Twitter', in *Proceedings of the NAACL Student Research Workshop*, pp. 88–93, San Diego, California, (June 2016). Association for Computational Linguistics.
- [42] Zeerak Waseem, James Thorne, and Joachim Bingel, *Bridging the Gaps: Multi Task Learning for Domain Transfer of Hate Speech Detection*, 29–55, Springer International Publishing, Cham, 2018.
- [43] Omar Zaidan, Jason Eisner, and Christine Piatko, 'Using "annotator rationales" to improve machine learning for text categorization', in *Human Language Technologies 2007: NAACL*, pp. 260–267, Rochester, New York, (April 2007). ACL.
- [44] Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar, 'Predicting the type and target of offensive posts in social media', in *NAACL*, pp. 1415–1420, Minneapolis, Minnesota, (June 2019). ACL.
- [45] Zijian Zhang, Koustav Rudra, and Avishek Anand, 'Explain and predict, and then predict again', in *WSDM*, WSDM '21, p. 418–426, New York, NY, USA, (2021). Association for Computing Machinery.