

Hate Speech Detection: A Solved Problem? The Challenging Case of Long Tail on Twitter

Ziqi Zhang*

Information School, University of Sheffield, Regent Court, 211 Portobello, Sheffield, S1 4DP, UK

E-mail: ziqi.zhang@sheffield.ac.uk

Lei Luo

College of Pharmaceutical Science, Southwest University, Chongqing, 400716, China

E-mail: drluolei@swu.edu.cn

Editors: First Editor, University or Company name, Country; Second Editor, University or Company name, Country

Solicited reviews: First Solicited Reviewer, University or Company name, Country; Second Solicited Reviewer, University or Company name, Country

Open reviews: First Open Reviewer, University or Company name, Country; Second Open Reviewer, University or Company name, Country

Abstract. In recent years, the increasing propagation of hate speech on social media and the urgent need for effective counter-measures have drawn significant investment from governments, companies, and researchers. A large number of methods have been developed for automated hate speech detection online. This aims to classify textual content into non-hate or hate speech, in which case the method may also identify the targeting characteristics (i.e., types of hate, such as race, and religion) in the hate speech. However, we notice significant difference between the performance of the two (i.e., non-hate v.s. hate). In this work, we argue for a focus on the latter problem for practical reasons. We show that it is a much more challenging task, as our analysis of the language in the typical datasets shows that hate speech lacks unique, discriminative features and therefore is found in the ‘long tail’ in a dataset that is difficult to discover. We then propose Deep Neural Network structures serving as feature extractors that are particularly effective for capturing the semantics of hate speech. Our methods are evaluated on the largest collection of hate speech datasets based on Twitter, and are shown to be able to outperform state of the art by up to 6 percentage points in macro-average F1, or 9 percentage points in the more challenging case of identifying hateful content.

Keywords: hate speech, classification, neural network, CNN, GRU, skipped CNN, deep learning, natural language processing

1. Introduction

The exponential growth of social media such as Twitter and community forums has revolutionised communication and content publishing, but is also increasingly exploited for the propagation of hate speech and the organisation of hate-based activities [1, 3]. The anonymity and mobility afforded by such media has made the breeding and spread of hate speech – eventually leading to hate crime – effortless in a virtual land-

scape beyond the realms of traditional law enforcement.

The term ‘hate speech’ was formally defined as ‘any communication that disparages a person or a group on the basis of some characteristics (to be referred to as **types of hate** or **hate classes**) such as race, colour, ethnicity, gender, sexual orientation, nationality, religion, or other characteristics’ [27]. In the UK, there has been significant increase of hate speech towards the migrant and Muslim communities following recent events including leaving the EU, the Manchester and the London attacks [16]. In the EU, surveys and reports focusing on young people in the EEA (European Economic Area) region show rising hate speech

*Correspondence author

and related crimes based on religious beliefs, ethnicity, sexual orientation or gender, as 80% of respondents have encountered hate speech online and 40% felt attacked or threatened [2]. Statistics also show that in the US, hate speech and crime is on the rise since the Trump election [28]. The urgency of this matter has been increasingly recognised, as a range of international initiatives have been launched towards the qualification of the problems and the development of counter-measures [12].

Building effective counter measures for online hate speech requires as the first step, identifying and tracking hate speech online. For years, social media companies such as Twitter, Facebook, and YouTube have been investing hundreds of millions of euros every year on this task [13, 17, 21], but are still being criticised for not doing enough. This is largely because such efforts are primarily based on manual moderation to identify and delete offensive materials. The process is labour intensive, time consuming, and not sustainable or scalable in reality [5, 13, 37].

A large number of research has been conducted in recent years to develop automatic methods for hate speech detection in the social media domain. These typically employ semantic content analysis techniques built on Natural Language Processing (NLP) and Machine Learning (ML) methods, both of which are core pillars of the Semantic Web research. The task typically involves classifying textual content into non-hate or hateful, in which case it may also identify the types of the hate speech. Although current methods have reported promising results, we notice that their evaluations are largely biased towards detecting content that is non-hate, as opposed to detecting and classifying real hateful content. A limited number of studies [3, 30] have shown that, for example, state of the art methods that detect sexism messages can only obtain an accuracy (F1) of between 15 and 60 percentage points lower than detecting non-hate messages. These results suggest that it is much harder to detect hateful content and their types than non-hate¹. However, from a practical point of view, we argue that the ability to correctly (Precision) and thoroughly (Recall) detect and identify specific types of hate speech is more desirable. For example, social media companies need to flag up hateful content for moderation, while law en-

forcement need to identify hateful messages and their nature as forensic evidence.

Motivated by these observations, our work makes two major contributions to the research of online hate speech detection. **First**, we conduct a data analysis to quantify and qualify the linguistic characteristics of such content on the social media, in order to understand the challenging case of detecting hateful content compared to non-hate. By comparison, we show that hateful content exhibits a ‘long tail’ pattern compared to non-hate due to their lack of unique, discriminative linguistic features, and this makes them very difficult to identify using conventional features widely adopted in many language-based tasks. **Second**, we propose Deep Neural Network (DNN) structures that are empirically shown to be very effective feature extractors for identifying specific types of hate speech. These include two DNN models inspired and adapted from literature of other Machine Learning tasks: one that simulates skip-gram like feature extraction based on modified Convolutional Neural Networks (CNN), and another that extracts orderly information between features using Gated Recurrent Unit networks (GRU).

Evaluated on the largest collection of English Twitter datasets, we show that our proposed methods can outperform state of the art methods by up to 6 percentage points in macro-average F1, or 9 percentage points in the more challenging task of detecting and classifying hateful content. Our thorough evaluation on all currently available public Twitter datasets sets a new benchmark for future research in this area. And our findings encourage future work to take a renewed perspective, i.e., to consider the challenging case of long tail.

The remainder of this paper is structured as follows. Section 2 reviews related work on hate speech detection and other relevant fields; Section 3 describes our data analysis to understand the challenges of hate speech detection on Twitter; Section 4 introduces our methods; Section 5 presents experiments and results; and finally Section 6 concludes this work and discusses future work.

2. Related Work

2.1. Terminology and Scope

Recent years have seen an increasing number of research on hate speech detection as well as other related areas. As a result, the term ‘hate speech’ is

¹Even in a binary setting of the task (i.e., either a message is hate or not), the high accuracy obtainable on detecting non-hate does not automatically translate to high accuracy of the other task due to the highly imbalanced nature in such datasets, as we shall show later.

often seen to co-exist or become mixed with other terms such as ‘offensive’, ‘profane’, and ‘abusive languages’, and ‘cyberbullying’. To distinguish them, we identify that hate speech 1) targets individual or groups on the basis of their characteristics; 2) demonstrates a clear intention to incite harm, or to promote hatred; 3) may or may not use offensive or profane words. For example: ‘Assimilate? No they all need to go back to their own countries. #BanMuslims Sorry if someone disagrees too bad.’

In contrast, ‘All you perverts (other than me) who posted today, needs to leave the O Board. Dfasdfdasfads’ is an example of abusive language, which often bears the purpose of insulting individuals or groups, and can include hate speech, derogatory and offensive language [26]. ‘i spend my money how i want bitch its my business’ is an example of offensive or profane language, which is typically characterised by the use of swearing or curse words. ‘Our class prom night just got ruined because u showed up. Who invited u anyway?’ is an example of bullying, which has the purpose to harass, threaten or intimidate typically individuals rather than groups.

In the following, we cover state of the art in all these areas with a focus on hate speech². Our methods and experiments will only address hate speech, due to both dataset availability and the goal of this work.

2.2. Methods of Hate Speech Detection and Related Problems

Existing methods primarily cast the problem as a supervised document classification task [33]. These can be divided into two categories: one relies on manual feature engineering that are then consumed by algorithms such as SVM, Naive Bayes, and Logistic Regression [3, 9, 11, 15, 19, 23, 35–39] (**classic methods**); the other represents the more recent deep learning paradigm that employs neural networks to automatically learn multi-layers of abstract features from raw data [13, 26, 30, 34] (**deep learning methods**).

Classic methods require manually designing and encoding features of data instances into feature vectors, which are then directly used by classifiers.

²We will indicate explicitly where works address a related problem rather than hate speech.

Schmidt et al. [33] summarised several types of features used in the state of the art. *Simple surface features* such as bag of words, word and character n-grams have been used as fundamental features in hate speech detection [3, 4, 9, 15, 19, 34–37], as well as other related tasks such as the detection of offensive and abusive content [5, 23, 26], discrimination [39], and cyberbullying [41]. Other surface features can include URL mentions, hashtags, punctuations, word and document lengths, capitalisation, etc [5, 9, 26]. *Word generalisation* includes the use of low-dimensional, dense vectorial word representations usually learned by clustering [35], topic modelling [38, 41], and word embeddings [11, 26, 34, 39] from unlabelled corpora. Such word representations are then used to construct feature vectors of messages. *Sentiment analysis* makes use of the degree of polarity expressed in a message [3, 9, 14, 34]. *Lexical resources* are often used to look up specific negative words (such as slurs, insults, etc.) in messages [3, 14, 26, 38]. *Linguistic features* utilise syntactic information such as Part of Speech (PoS) and certain dependency relations as features [3, 5, 9, 14, 41]. *Meta-information* refers to data about messages, such as gender identity of a user associated with a message [36, 37], or high frequency of profane words in a user’s post history [8, 38]. In addition, *Knowledge-Based features* such as messages mapped to stereotypical concepts in a knowledge base [10] and *multimodal information* such as image captions and pixel features [41] were used in cyberbullying detection but only in very confined context [33].

In terms of classifiers, existing methods are predominantly supervised. Among these, Support Vector Machines (SVM) is the most popular algorithm [3, 5, 9, 15, 23, 35, 38, 39], while other algorithms such as Naive Bayes [5, 9, 19, 23, 39], Logistic Regression [9, 11, 23, 36, 37], and Random Forest [9, 38] are also used.

Deep learning based methods employ deep artificial neural networks to learn abstract feature representations from input data through its multiple stacked layers for the classification of hate speech. The input can take various forms of feature encoding, including any of those used in the classic methods. However, the key difference is that in such a model the input features are not directly used for classification. Instead, the multi-layer structure learns new abstract feature representations that prove to be more effective for learning. For this reason, deep learning based methods typically shift its focus from manual feature engineering to the

network structure, which is carefully designed to automatically extract useful features from a simple input feature representation. Indeed we notice a clear trend in the literature that shifts towards the adoption of deep learning based methods and studies have also shown them to perform better than classic methods on this task [13, 30]. Note that this categorisation excludes those methods [11, 23, 39] that used DNN to learn word or text embeddings and subsequently apply another classifier (e.g., SVM, logistic regression) to use such embeddings as features for classification. Instead, we focus on DNN methods that perform the classification task itself.

To the best of our knowledge, methods belong to this category include [1, 13, 30, 34, 40], all of which used simple word and/or character based one-hot encoding as input features to their models, while Vigna et al. [34] also used word polarity. The most popular network architectures are Convolutional Neural Network (CNN) and Recurrent Neural Network (RNN), typically Long Short-Term Memory network (LSTM). In the literature, CNN is well known as an effective network to act as ‘feature extractors’, whereas RNN is good for modelling orderly sequence learning problems [29]. In the context of hate speech classification, intuitively, CNN extracts word or character combinations [1, 13, 30] (e.g., phrases, n-grams), RNN learns word or character dependencies (orderly information) in Tweets [1, 34].

In our previous work [40], we showed benefits of combining both structures in such tasks by using a hybrid CNN and GRU structure. This work largely extends it in several ways. First, we adapt the model to multiple CNN layers; second, we propose a new DNN architecture based on the idea of extracting skip-gram like features for this task; third, we conduct data analysis to understand the challenges in hate speech detection due to the linguistic characteristics in the data; and finally, we perform an extended evaluation of our methods, particularly their capability on addressing these challenges.

2.3. Evaluation of Hate Speech Detection Methods

Evaluation of the performance of hate speech (and also other related content) detection typically adopts the classic Precision, Recall and F1 metrics. Precision measures the percentage of true positives among the set of hate speech messages identified by a system; Recall measures the percentage of true positives among the set of real hate speech messages we ex-

pect the system to capture (also called ‘**ground truth**’ or ‘**gold standard**’), and F1 calculates the harmonic mean of the two. The three metrics are usually applied to each class in a dataset, and often an aggregated figure is computed either using **micro-average** or **macro-average**. The first sums up the individual true positives, false positives, and false negatives identified by a system regardless of different classes to calculate overall Precision, Recall and F1 scores. The second takes the average of the Precision, Recall and F1 on different classes.

Existing studies on hate speech detection have primarily reported their results using micro-average Precision, Recall and F1 [1, 13, 30, 36, 37, 40]. The problem with this is that in an unbalanced dataset where instances of one class (to be called the ‘dominant class’) significantly out-number others (to be called ‘minority classes’), micro-averaging can mask the real performance on minority classes. Thus a significantly lower or higher F1 score on a minority class (when compared to the majority class) is unlikely to cause significant change in micro-F1 on the entire dataset. As we will show in Section 3, hate speech detection is a typical task dealing with extremely unbalanced datasets, where real hateful content only accounts for a very small percentage of the entire dataset, while the large majority is non-hate but exhibits similar linguistic characteristics to hateful content. As argued before, practical applications often need to focus on detecting hateful content and identifying their types. In this case, reporting micro F1 on the entire dataset will not properly reflect a system’s ability to deal with hateful content as opposed to non-hate. Unfortunately, only a very limited number of work has reported performance on a per-class basis [4, 30]. As an example, when compared to the micro F1 scores obtained on the entire dataset, the highest F1 score reported for detecting sexism messages is 47 percentage points lower in [4] while 11 points lower in [30]. This has largely motivated our study to understand what causes hate speech to be so difficult to classify from a linguistic point of view, and to evaluate hate speech detection methods by giving more focus on their capability of classifying real hateful content.

3. Dataset Analysis - the Case of Long Tail

We first start with an analysis of typical datasets used in the studies of hate speech detection on Twitter. From this we show the very unbalanced nature of

such data, and compare the linguistic characteristics of hate speech against non-hate to discuss the challenge of detecting and classifying hateful content.

3.1. Public Twitter Datasets

We use the largest collection of publicly available English Twitter datasets previously compiled in our work [40]. This includes seven datasets published in previous research. **DT** consolidates the dataset by [9] into two types, ‘hate’ and ‘non-hate’. The Tweets do not have a focused topic but were collected using a controlled vocabulary of abusive words. **RM** contains ‘hate’ and ‘non-hate’ Tweets focused on refugee and muslim discussions. **WZ** is initially published by [36] and contains ‘sexism’, ‘racism’, and ‘non-hate’; the same authors created another smaller dataset annotated by domain experts and amateurs separately. The authors showed that the two sets of annotations were different as the supervised classifiers obtained different results on them. We will use **WZ-S.amt** to denote the dataset annotated by amateurs, and **WZ-S.exp** to denote the dataset annotated by experts. **WZ-S.gb** merges the **WZ-S.amt** and **WZ-S.exp** datasets by taking the majority vote from both amateur and expert annotations where the expert was given double weights [13]. **WZ.pj** combines the **WZ** and the **WZ-S.exp** datasets [30]. All of the **WZ-S.amt**, **WZ-S.exp**, **WZ-S.gb**, and **WZ.pj** datasets contain ‘sexism’, ‘racism’, and ‘non-hate’ Tweets, but also added a ‘both’ class that includes Tweets considered to be both ‘sexism’ and ‘racism’. However, there are only several handful of instances of this class and they were found to be insufficient for model learning. Therefore, following [30] we exclude this class from these datasets. Table 1 summarises the statistics of these datasets.

3.2. Dataset Analysis

As shown in Table 1, all datasets are significantly biased towards non-hate, as hate Tweets account between only 5.8% (DT) and 31.6% (WZ). When we inspect specific types of hate, some can be even more scarce, such as ‘racism’ and as mentioned before, the extreme case of ‘both’. This has two implications. First, an evaluation measure such as the micro F1 that looks at a system’s performance on the entire dataset regardless of class difference can be biased to the system’s ability of detecting ‘non-hate’. In other words, a hypothetical system that achieves almost perfect F1 in identifying ‘racism’ Tweets can still be overshadowed

Dataset	#Tweets	Classes (%)
WZ	16,093	racism (12%) sexism (19.6%) neither (68.4%)
WZ-S.amt	6,579	racism (1.9%) sexism (16.3%) neither (81.8%)
WZ-S.exp	6,559	racism (1.3%) sexism (11.8%) neither (86.9%)
WZ-S.gb	6,567	racism (1.4%) sexism (13.9%) neither (84.7%)
WZ.pj	18,593	racism (10.8%) sexism (20.3%) neither (68.9%)
DT	24,783	hate (5.8%) non-hate (94.2%)
RM	2,435	hate (17%) non-hate (83%)

Table 1

Statistics of datasets used in the experiment

owed by its poor F1 in identifying ‘non-hate’, and vice versa. Second, compared to non-hate, the training data for hate Tweets are very scarce. This may not be an issue that is easy to address as it seems, since the datasets are collected from Twitter and reflect the real nature of data imbalance in this domain. Thus to annotate more training data for hateful content we will almost certainly have to spend significantly more effort annotating non-hate. Also, as we shall show in the following, this problem may not be easily mitigated by conventional methods of over- or under-sampling. Because the real challenge is the lack of unique, discriminative linguistic characteristics in hate Tweets compared to non-hate.

As a proxy to quantify and compare the linguistic characteristics of hate and non-hate Tweets, we propose to study the ‘uniqueness’ of the vocabulary for each class. We argue that this can be a reasonable reflection of the features used for classifying each class. On the one hand, most types of features are derived from words; on the other hand, our previous work already showed that the most effective features in such tasks are based on surface words [32].

Specifically, we start with applying a state of the art Tweet normalisation tool³ to tokenise and transform each Tweet into a sequence of words. This is done to mitigate the noise due to the colloquial nature of the data. The process involves, for example, spelling correction, elongated word normalisation (‘yaaaaa’ becomes ‘yay’), word segmentation on hashtags (‘#ban-refugees’ becomes ‘ban refugees’), and unpacking

³<https://github.com/cbaziotis/ekphrasis>

contractions (e.g., ‘can’t’ becomes ‘can not’). Then we lemmatise each word to return its dictionary form. We refer to this process as ‘pre-processing’ and the output as pre-processed Tweets.

Next, given a Tweet t_i , let c_n be the class label of t_i , $words(t_i)$ returns the set of different words from t_i , and $uwords(c_n)$ returns the set of class-unique words that are found only for c_n (i.e., they do not appear in any other classes), then for each dataset, we measure for each Tweet a ‘uniqueness’ score u as:

$$u(t_i) = \frac{|words(t_i) \cap uwords(c_n)|}{|words(t_i)|} \quad (1)$$

This measures the fraction of class-unique words in a Tweet, depending on the class of this Tweet. Intuitively, the score can be considered as an indication of ‘uniqueness’ of the features found in a Tweet. A high value indicates that the Tweet can potentially contain more features that are unique to its class, and as a result, we can expect the Tweet to be relatively easy to classify. On the other hand, a low value indicates that many features of this Tweet are potentially non-discriminative as they may also be found across multiple classes, and therefore, we can expect the Tweet to be relatively difficult to classify.

We then compute this score for every Tweet in a dataset, and compare the number of Tweets with different uniqueness scores within each class. To better visualise this distribution, we bin the scores into 11 ranges as $u \in [0, 0]$, $u \in (0, 0.1]$, $u \in (0.1, 0.2]$, ..., $u \in (0.9, 1.0]$. In other words, the first range includes only Tweets with a uniqueness score of 0, then the remaining 10 ranges are defined with a 0.1 increment in the score. In Figure 1 we plot for each dataset, 1) the distribution of Tweets over these ranges regardless of their class (as indicated by the length of the dark horizontal bar, measured against the x axis that shows accumulative percentage of the dataset); and 2) the distribution of Tweets belong to each class (as indicated by the call-out boxes). For simplicity, we label each range using its higher bound on the y axis. As an example, $u \in [0, 0]$ is labelled as 0, and $u \in (0, 0.1]$ as 0.1.

Using the WZ-S.amt dataset as an example, the figure shows that almost 30% of Tweets (the bottom horizontal bars in the figure) in this dataset have a uniqueness score of 0. In other words, these Tweets contain no class-unique words. This can cause difficulty in extracting class-unique features from these Tweets, making them very difficult to classify. The call-out box for

this part of data shows that it contains 52% of sexism and 48% of racism Tweets. In fact, on this dataset, 76% of sexism and 81% of racism Tweets (adding up figures from the call-out boxes for the bottom three horizontal bars) only have a uniqueness score of 0.2 or lower. On those Tweets that have a uniqueness score of 0.4 or higher (the top six horizontal bars), i.e., those that may be deemed as relatively ‘easier’ to classify, we find only 2% of sexism and 3% of racism Tweets. In contrast, it is 17% for non-hate Tweets.

We can notice very similar patterns on all the datasets in this analysis. Overall, it shows that the majority of hate Tweets potentially lack discriminative features and as a result, they ‘sit in the long tail’ of the dataset as ranked by the uniqueness of Tweets. Note also that comparing the larger datasets WZ.pj and WZ against the smaller ones (i.e., the WZ-S ones), although both the absolute number and the percentage of the racism and sexism Tweets are increased significantly in the two larger datasets (see Table 3.1), this does not improve the long tail situation. Indeed, one can hate or not using the same words. And as a result, increasing the dataset size and improving class balance may not always guarantee a solution.

4. Methodology

In this section, we describe our DNN based methods that implement the intuition of extracting dependency between words or phrases as features from Tweets. To illustrate this idea, consider the example Tweet ‘These muslim refugees are troublemakers and parasites, they should be deported from my country’. Each of the words such as ‘muslim’, ‘refugee’, ‘troublemakers’, ‘parasites’, and ‘deported’ alone are not always indicative features of hate speech, as they can be used in any context. However, combinations such as ‘muslim refugees, troublemakers’, ‘refugees, troublemakers’, ‘refugees, parasites’, ‘refugees, deported’, and ‘they, deported’ can be more indicative features. Clearly, in these examples, the pair of words or phrases form certain dependence on each other, and such sequences cannot be captured by n-gram like features.

We propose two DNN structures that may capture such features. Our previous work [40] combines traditional a CNN with a GRU layer [40] and for the sake of completeness, we also include its details below. Our other method combines traditional CNN with

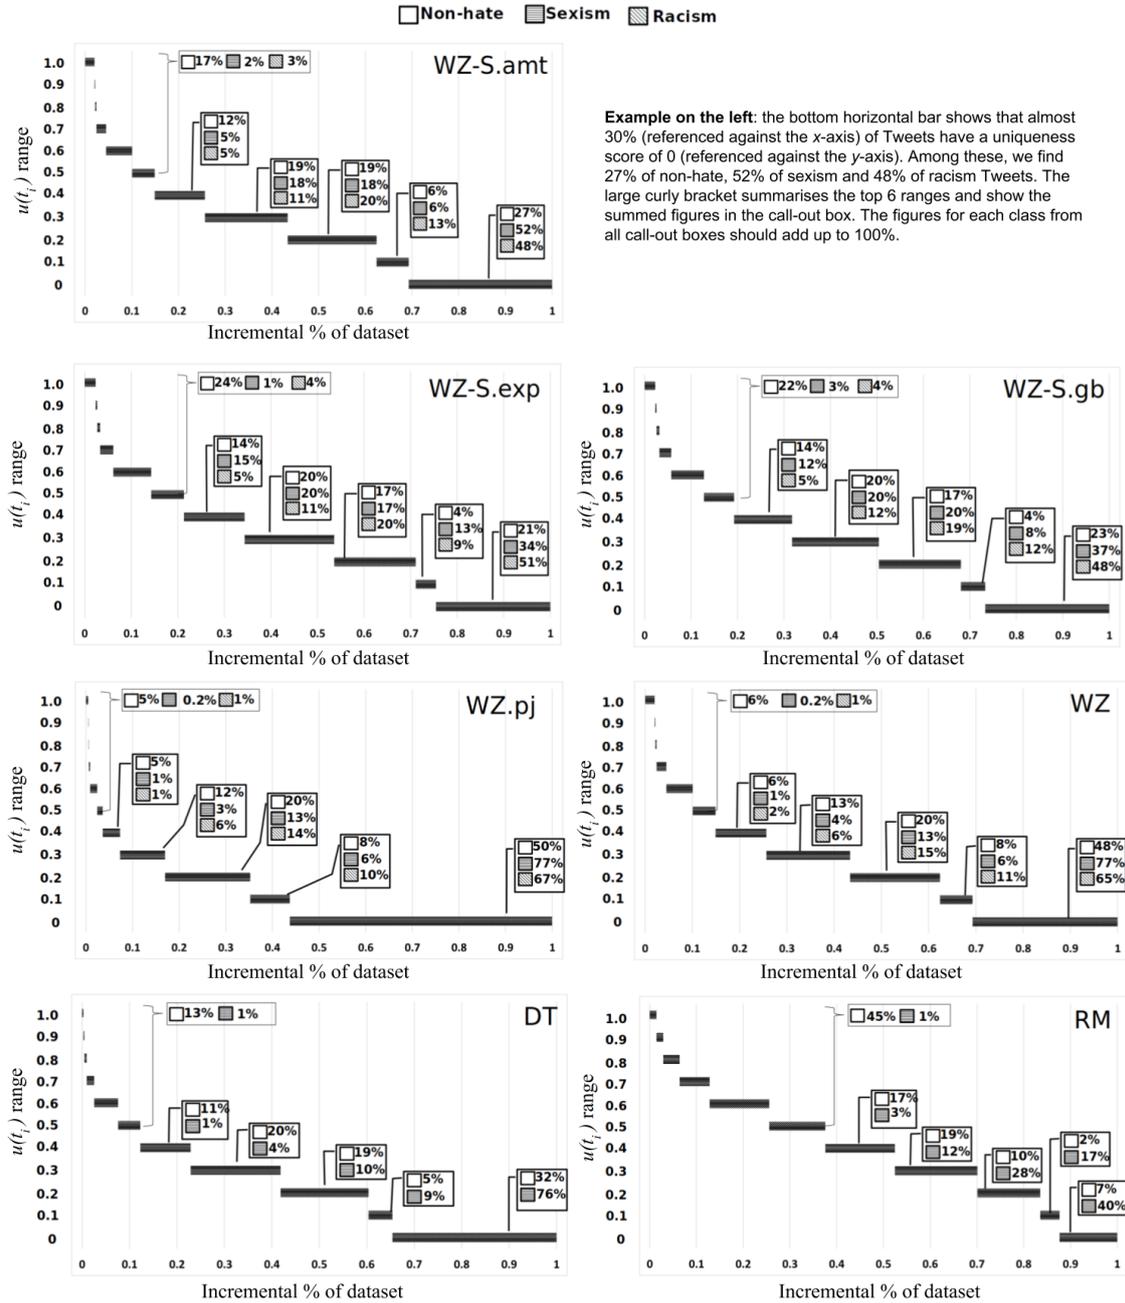


Fig. 1. Distribution of Tweets in each dataset over the 11 ranges of the uniqueness scores. The x-axis shows accumulative percentage of the dataset; the y-axis shows the labels of these ranges. The call-out boxes show for each class, the fraction of Tweets fall under that range (in case a class is not present it has a fraction of 0%).

some modified CNN layers serving as skip-gram extractors - to be called ‘skipped CNN’. Both structures modify a common, base CNN model (Section 4.1) that acts as the n-gram feature extractor, while the added GRU (Section 4.1.1) and the skipped CNN (Section 4.1.2) components are expected to extract the dependent sequences of such n-grams, as illustrated above.

4.1. The Base CNN Model

The Base CNN model is illustrated in Figure 2. Given a Tweet, we firstly apply the pre-processing described in Section 3.2 to normalise and transform the Tweet into a sequence of words. This sequence is then passed to a word embedding layer, which maps the sequence into a real vector domain (word embeddings). Specifically, each word is mapped onto a fixed dimensional real valued vector, where each element is the weight for that dimension for that word. Word embeddings are often trained on very large unlabelled corpus, and comparatively, the datasets used in this study are much smaller. Therefore in this work, we use pre-trained word embeddings that are publicly available (to be detailed in Section 6). One potential issue with pre-trained embeddings is Out-Of-Vocabulary (OOV) words, particularly on Twitter data due to its colloquial nature. Thus the pre-processing also helps to reduce the noise in the language and hence the scale of OOV. For example, by hashtag segmentation we transform an OOV ‘#BanIslam’ into ‘Ban’ and ‘Islam’ that are more likely to be included in the pre-trained embedding models.

The embedding layer passes an input feature space with a shape of 100×300 to three 1D convolutional layers, each uses 100 filters but different window sizes of 2, 3, and 4 respectively. Intuitively, each CNN layer can be considered as extractors of bi-gram, tri-gram and quad-gram features. The rectified linear unit function is used for activation in these CNNs. The output of these CNNs are concatenated to a single layer, which is then further down-sampled by a 1D max pooling layer with a pool size of 4. The output of max pooling is fed into the final softmax layer to predict probability distribution over all possible classes (n), which will depend on individual datasets.

One of the recent trends in text processing tasks on Twitter is the use of character based n-grams and embeddings instead of word based, such as in [23, 30]. The main reason for this is to cope with the noisy and informal nature of the language in Tweets. We do not use character based models, mainly because the liter-

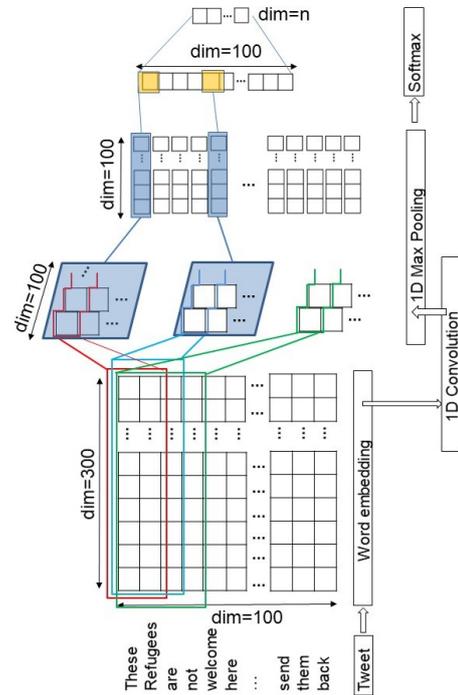


Fig. 2. The Base CNN model uses three different window sizes to extract features. This diagram is best viewed in colour.

atures that compared word based and character based models are rather inconclusive. Although Mehdad et al. [23] obtained better results using character based models, Park et al. [30] and Gamback et al. [13] reported the opposite. Further, our pre-processing already reduces the noise in the language to some extent. And word-based models also better fit our intuitions explained before.

4.1.1. CNN + GRU

With this model, we extend the Base CNN model by adding a GRU layer that takes input from the max pooling layer. This treats the features as timesteps and outputs 100 hidden units per timestep. Compared to LSTM, which is a popular type of RNN, the key difference in a GRU is that it has two gates (reset and update gates) whereas an LSTM has three gates (namely input, output and forget gates). Thus GRU is a simpler structure with fewer parameters to train. In theory, this makes it faster to train and generalise better on small data; while empirically it is shown to achieve comparable results to LSTM [7]. Next, a global max pooling layer ‘flattens’ the output space by taking the highest value in each timestep dimension, producing a feature vector that is finally fed into the softmax layer.

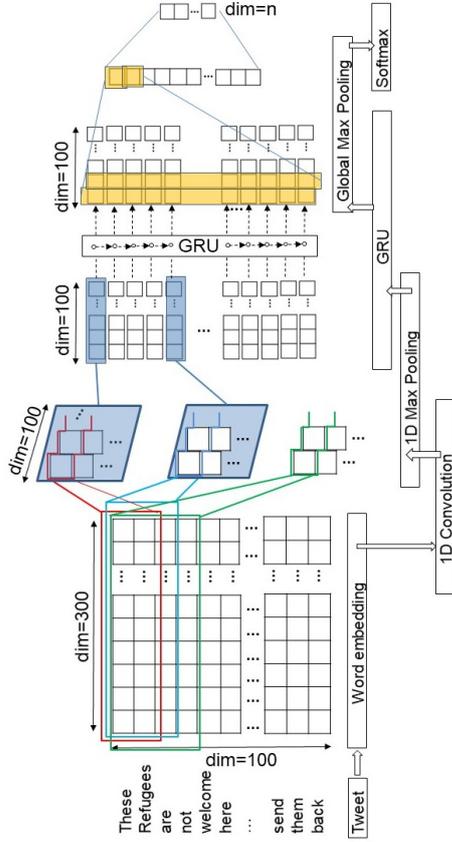


Fig. 3. The CNN+GRU architecture. This diagram is best viewed in colour.

The intuition is to pick the highest scoring features to represent a Tweet, which empirically works better than the normal configuration. The structure of this model is shown in Figure 3.

The GRU layer captures sequence orders that can be useful for this task. In an analogy, it learns dependency relationships between n-grams extracted by the CNN layer before. And as a result, it may capture co-occurring word n-grams as useful patterns for classification, such as the pairs of words and phrases illustrated before.

4.1.2. CNN + skipped CNN (sCNN)

With this model, we propose to extend the Base CNN model by adding CNNs that use ‘gapped window’ to extract features from its input, and we call these CNN layers ‘skipped CNNs’. A gapped window is one where inputs at certain (consecutive) positions of the window are ignored, such as those shown in Figure 4. We say that these positions within the window

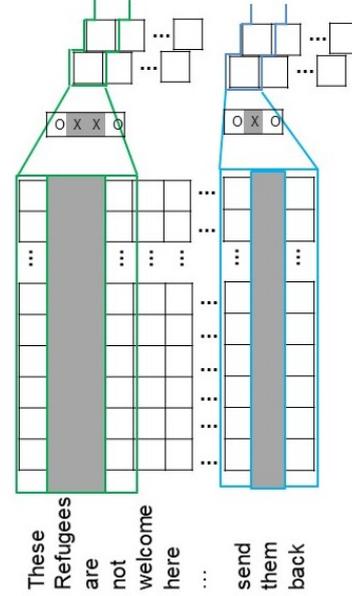


Fig. 4. Example of a 2 gapped size 4 window and a one gapped size 3 window. The ‘X’ indicates that input for the corresponding position in the window is ignored.

are ‘deactivated’ while other positions are ‘activated’. Specifically, given a window of size j , applying a gap of i consecutive positions will produce multiple shapes of size j windows, as illustrated in Algorithm 1.

Algorithm 1 Creation of i gapped size j windows. A sequence $[O, X, O, O]$ represents one possible shape of a 1 gapped size 4 window, where the first and the last two positions are activated (‘O’) and the second position is deactivated (‘X’).

```

1: Input:  $i : 0 < i < j, j : j > 0, w \leftarrow [p_1, \dots, p_j]$ 
2: Output:  $W \leftarrow \emptyset$  a set of  $j$  sized window shapes
3: for all  $k \in [2, j)$  and  $k \in \mathbb{N}_+$  do
4:   Set  $p_1$  in  $w$  to O
5:   Set  $p_j$  in  $w$  to O
6:   for all  $x \in [k, k + i)$  and  $x \in \mathbb{N}_+$  do
7:     Set  $p_x$  to X
8:     for all  $y \in [k + i + 1, j)$  and  $y \in \mathbb{N}_+$  do
9:       Set  $p_y$  in  $w$  to O
10:    end for
11:     $W \leftarrow W \cup \{w\}$ 
12:  end for
13: end for

```

As an example, applying a 1-gap to a size 4 window will produce two shapes: $[O, X, O, O], [O, O, X, O]$, where

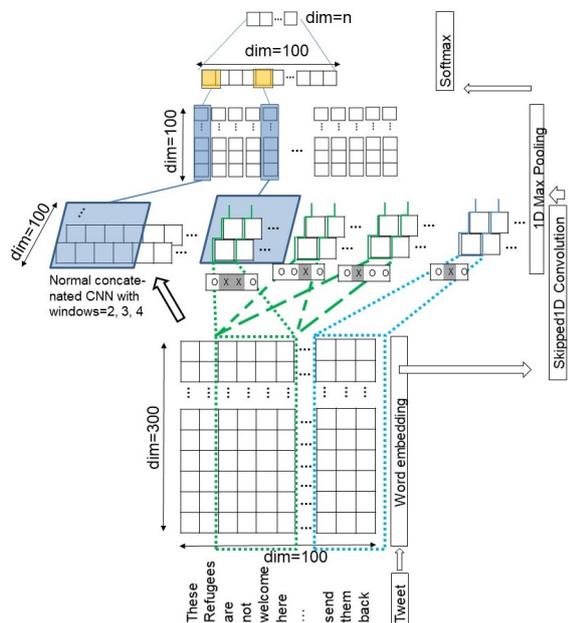


Fig. 5. The CNN+sCNN model concatenates features extracted by the normal CNN layers with window sizes of 2, 3, and 4, with features extracted by the four skipped CNN layers. This diagram is best viewed in colour.

‘O’ indicates an activated position and ‘X’ indicates a deactivated position in the window; while applying a 2-gap to a size 4 window will produce a single shape of [O,X,X,O].

To extend the Base CNN model, we add CNNs using 1-gapped size 3 windows, 1-gapped size 4 windows and 2-gapped size 4 windows, keeping the remaining parts of the structure the same. This results in a model illustrated in Figure 5.

Intuitively, the skipped CNNs can be considered as extractors of ‘skip-gram’ like features. In an analogy, we expect it to extract useful features such as ‘muslim refugees ? troublemakers’, ‘muslim ? ? troublemakers’, ‘refugees ? troublemakers’, and ‘they ? ? deported’ from the example sentence before, where ‘?’ is a wildcard representing any word token in a sequence.

To the best of our knowledge, the work by Nguyen et al. [25] is the only one that uses DNN models to extract skip-gram features that are used directly in NLP tasks. However, our method is different in two ways. First, Nguyen et al. addressed a task of mention detection from sentences, i.e., classifying word tokens in a sentence into sequences of particular entities or not. Our work deals with sentence classification. This means

that our modelling of the task input and their features are essentially different. Second, the authors used skip-gram features only, while our method adds skip-grams to conventional n-grams, as we concatenate the output from the skipped CNNs and the conventional CNNs. While the concept of skip-grams has become popular in training word embeddings with neural network models since Mikolov et al. [24], the task is different from directly using skip-gram as features for NLP.

4.1.3. Model Parameters

We use the categorical cross entropy loss function and the Adam optimiser to train the models, as the first is empirically found to be more effective on classification tasks than other commonly used loss functions including classification error and mean squared error [22], and the second is designed to improve the classic stochastic gradient descent (SGD) optimiser and in theory combines the advantages of two other common extensions of SGD (AdaGrad and RMSProp) [18].

Our choice of parameters described above are largely based on empirical findings reported previously, default values or anecdotal evidence. Arguably, these may not be the best settings for optimal results, which are always data-dependent. However, we show later in experiments that the models already obtain promising results even without extensive data-driven parameter tuning.

5. Experiment

In this section, we present our experiments for evaluation and discuss the results. We compare our CNN+GRU and CNN+sCNN methods against three re-implemented state of the art methods (Section 5.1), and discuss the results in Section 5.2. This is followed by an analysis to show how our methods have managed to effectively capture hate Tweets in the long tail (Section 5.3), and to discover the typical errors made by all methods compared (Section 5.4).

Word embeddings. We experimented with three different choices of pre-trained word embeddings: the Word2Vec embeddings trained on the 3-billion-word Google News corpus with a skip-gram model⁴, the ‘GloVe’ embeddings trained on a corpus of 840 billion tokens using a Web crawler [31], and the Twitter embeddings trained on 200 million Tweets with spam re-

⁴<https://github.com/mnihaltz/word2vec-GoogleNews-vectors>

moved [20]⁵. However, our results did not find any embeddings that can consistently outperform others on all tasks and datasets. In fact, this is rather unsurprising, as previous studies [6] suggested similar patterns: the superiority of one word embeddings model on intrinsic tasks (e.g., measuring similarity) is generally non-transferable to downstream applications, across tasks, domains, or even datasets. For this reason, below we only focus on results obtained using the Word2Vec embeddings, to be consistent with previous work such as [30]. Our full results obtained with the different embeddings are available in the Appendix.

Performance evaluation metrics. We use the standard Precision (P), Recall (R) and F1 measures for evaluation. For the sake of readability, we only present F1 scores in the following sections unless otherwise stated. Again full results can be found in the Appendix. Due to the significant class imbalance in the data, we show F1 obtained on both hate and non-hate Tweets separately.

Implementation. For all methods discussed in this work, we used the Python Keras⁶ with Theano backend⁷ and the scikit-learn⁸ library for implementation⁹. For DNN based methods, we fix the *epochs* to 10 and use a *mini-batch* of 100 on all datasets. These parameters are rather arbitrary and fixed for consistency. We run all experiments in a 5-fold cross validation setting and report the average across all folds.

5.1. State of the art

We re-implemented three state of the art methods covering both the classic and deep learning based methods. **First**, we use the SVM based method described in Davidson et al. [9]. A number of different types of features are used as below. Unless otherwise stated, these features are extracted from the pre-processed Tweets:

- Surface features: word unigram, bigram and trigram each weighted by TF-IDF; number of mentions, and hashtags¹⁰; number of characters, and words;

- Linguistic features: Part-of-Speech (PoS)¹¹ tag unigrams, bigrams, and trigrams, weighted by their TF-IDF and removing any candidates with a document frequency lower than 5; number of syllables; Flesch-Kincaid Grade Level and Flesch Reading Ease scores that to measure the ‘readability’ of a document
- Sentiment features: sentiment polarity scores of the tweet, calculated using a public API¹².

Our **second** and **third** methods are re-implementations of Gambäck et al. [13] (**GB**) and Park et al. [30] (**PK**) using word embeddings, which obtained better results than their character based counterparts in their original work. Both are based on concatenation of multiple CNNs and similar to the base CNN model described in Section 4.1, and are therefore good reference for comparison to analyse the benefits of our added GRU and skipped CNN layers. We kept the same hyper-parameters for these as well as our proposed methods for comparative analysis¹³.

5.2. Overall Results

Our experiments could not identify a best performing candidate among the three state of the art methods on all datasets, by all measures. Therefore, in the following discussion, unless otherwise stated, we compare our methods against *the best results achieved by any* of the three state of the art methods

Overall micro and macro F1. Firstly, we compare micro and macro F1 obtained by each method in Table 2. In terms of micro F1, both our CNN+sCNN and CNN+GRU methods obtained consistently the best results on all datasets. Nevertheless, the improvement over the best results from any of the three state of the art methods is rather incremental. The situation with macro F1 is however, quite different. Again compared against the state of the art, both our methods obtained consistently better results. The improvements in many cases are much higher: 2~6 and 1~4 percentage points (or ‘percent’ in short) for CNN+sCNN and CNN+GRU respectively. When only the categories of hate Tweets are considered (i.e., excluding non-hate, see ‘macro F1, hate’), the improvements are signifi-

⁵‘Set1’ in [20]

⁶<https://keras.io/>, version 2.0.2

⁷<http://deeplearning.net/software/theano/>, version 0.9.0

⁸<http://scikit-learn.org/>, version 0.19.1

⁹Code available at <https://github.com/ziqizhang/chase>

¹⁰Extracted from the original tweet before pre-processing.

¹¹The NLTK library is used.

¹²<https://github.com/cjhutto/vaderSentiment>

¹³In fact both papers did not detail their hyper-parameter settings, which is another reason for us to use consistent configurations as our methods.

cant in some cases: a maximum of 9 and 7 percent for CNN+sCNN and CNN+GRU respectively.

However, notice that both the overall and hate speech-only macro F1 scores are significantly lower than micro F1, for any methods, on any datasets. This further supports our earlier data analysis findings that classifying hate Tweets is a much harder task than non-hate, and micro F1 scores will overshadow a method's true performance on a per-class basis, due to the imbalanced nature of such data.

F1 per-class. To clearly compare each method's performance on classifying hate speech, we show the per-class results in Table 3. This demonstrates the benefits of our methods compared to the state of the art when focusing on any categories of hate Tweets rather than non-hate: the CNN+sCNN model was able to outperform the best results by any of the three comparison methods, achieving a maximum of 16 percent increase in F1; the CNN+GRU model was less good, but still obtained better results on four datasets, achieving a maximum of 12 percent increase in F1.

It is also worth to note that the best improvement obtained by our methods were found on the racism class in the three WZ-S datasets, and the hate class in the DT dataset. As shown in Table 1, these are all minority classes representing a very small population in the dataset (between 1 and 6%). This suggests that our proposed methods can be potentially very effective when there is a lack of training data.

sCNN v.s. GRU. Comparing CNN+sCNN with CNN+GRU using Table 3, we see that the first performed much better when only hate Tweets are considered, suggesting that the skipped CNNs may be more effective feature extractors than GRU for hate speech detection in very short texts such as Tweets. CNN+sCNN consistently outperformed the highest F1 by any of the three state of the art methods on any dataset, and it also achieved higher improvement. CNN+GRU on the other hand, obtained better F1 on four datasets and the same best F1 (as the state of the art) on two datasets. This also translates to overall better macro F1 by CNN+sCNN compared to CNN+GRU (Table 2).

Patterns observed with other word embeddings. While we show detailed results with different word embeddings in the Appendix, it is worth to mention here that the same patterns were noticed with these

different embeddings, i.e., both CNN+sCNN and CNN+GRU have outperformed the state of the art methods in the most cases but the first model obtained much better results. The best results however, were not always obtained with the Word2Vec embeddings on all datasets. For example, on the DT dataset, when using the Twitter embeddings, both methods made a further improvement of 11 (CNN+sCNN) and 6 (CNN+GRU) percent F1 when classifying hate Tweets compared to the Word2Vec embeddings (see Table 7). Compared against the best results by state of the art methods that also use Twitter embeddings (Park et al. in this case, see Table 6) and considering hate Tweets only, CNN+sCNN obtained an improvement of 11 percent and CNN+GRU 6 percent. Slightly better results were noted with the GloVe embedding on the same dataset.

Against previously reported results. For many reasons such as the difference in the re-generated datasets¹⁴, possibly different pre-processing methods, and unknown hyper-parameter settings from the previous work¹⁵, we cannot guarantee an identical replication of the previous methods in our re-implementation. Therefore in Table 4, we compare results by our methods against previously reported results (micro F1 is used as this is the case in all the previous work) on each dataset on an 'as-is' basis.

Table 4 shows that both our methods have achieved the best results on all datasets, outperforming state of the art on six and in some cases, quite significantly. Note that on the WZ.pj dataset where our methods did not obtain further improvement, the best reported state of the art result was obtained using a hybrid character-and-word embeddings CNN model [30]. Our methods in fact, outperformed both the word-only and character-only embeddings models in that same work.

5.3. Effectiveness on Identifying the Long Tail

While the results so far have shown that our proposed methods can obtain better performance in the task, it is unclear whether they are particularly ef-

¹⁴As noted in [40], we had to re-download Tweets using previously published Tweet IDs in the shared datasets. But some Tweets have become no longer available.

¹⁵The details of the pre-processing, network structures and many hyper-parameter settings are not reported in nearly all of the previous work. For comparability, as mentioned before, we kept the same configurations for our methods as well as the re-implemented state of the art methods.

Dataset	Measure	SVM	GB	PK	CNN+sCNN	CNN+GRU
WZ-S.amt	micro F1	.81	.92	.92	.92 (-)	.92 (-)
	macro F1	.59	.63	.64	.68 (+.04)	.65 (+.01)
	macro F1, hate	.45	.48	.49	.55 (+.06)	.51 (+.02)
WZ-S.exp	micro F1	.82	.91	.91	.92 (+.01)	.92 (+.01)
	macro F1	.58	.69	.69	.74 (+.05)	.70 (+.01)
	macro F1, hate	.42	.57	.56	.63 (+.06)	.58 (+.01)
WZ-S.gb	micro F1	.83	.92	.92	.93 (+.01)	.93 (+.01)
	macro F1	.64	.70	.70	.76 (+.06)	.74 (+.04)
	macro F1, hate	.51	.57	.57	.66 (+.09)	.64 (+.07)
WZ.pj	micro F1	.72	.82	.82	.83 (+.01)	.83 (+.01)
	macro F1	.66	.74	.75	.77 (+.02)	.75 (-)
	macro F1, hate	.59	.68	.69	.71 (+.02)	.69 (-)
WZ	micro F1	.72	.82	.82	.83 (+.01)	.83 (+.01)
	macro F1	.65	.75	.75	.77 (+.02)	.76 (+.01)
	macro F1, hate	.58	.70	.69	.71 (+.01)	.71 (+.01)
DT	micro F1	.79	.94	.94	.94 (-)	.94 (-)
	macro F1	.56	.61	.60	.64 (+.03)	.63 (+.02)
	macro F1, hate	.23	.25	.23	.30 (+.05)	.29 (+.04)
RM	micro F1	.79	.90	.90	.91 (+.01)	.90 (-)
	macro F1	.72	.80	.81	.83 (+.02)	.81 (+.01)
	macro F1, hate	.58	.67	.68	.71 (+.03)	.68 (-)

Table 2

Micro v.s. macro F1 results of different methods (using the Word2Vec embeddings). The best results on each row are highlighted in **bold**. Numbers within (brackets) indicate the improvement in F1 compared to the *best result* by any of the three state of the art methods.

Dataset	Class	SVM	GB	PK	CNN+sCNN	CNN+GRU
WZ-S.amt	racism	.22	.15	.18	.29 (+.07)	.21 (-.01)
	sexism	.68	.80	.80	.81 (+.01)	.80 (-)
	non-hate	.88	.95	.95	.95 (-)	.95 (-)
WZ-S.exp	racism	.26	.49	.46	.58 (+.09)	.50 (+.01)
	sexism	.58	.65	.65	.68 (+.03)	.65 (-)
	non-hate	.89	.95	.97	.96 (-.01)	.95 (-.02)
WZ-S.gb	racism	.38	.38	.37	.54 (+.16)	.50 (+.12)
	sexism	.64	.75	.76	.77 (+.02)	.77 (+.01)
	non-hate	.90	.96	.96	.96 (-)	.96 (-)
WZ.pj	racism	.60	.69	.70	.73 (+.03)	.70 (-)
	sexism	.58	.67	.68	.69 (+.01)	.68 (-)
	non-hate	.80	.88	.88	.88 (-)	.88 (-)
WZ	racism	.59	.72	.71	.74 (+.02)	.73 (+.01)
	sexism	.57	.67	.67	.68 (+.01)	.68 (+.01)
	non-hate	.79	.88	.88	.88 (-)	.88 (-)
DT	hate	.23	.25	.23	.30 (+.05)	.29 (+.04)
	non-hate	.88	.97	.97	.97 (-)	.97 (-)
RM	hate	.58	.67	.68	.71 (+.03)	.68 (-)
	non-hate	.86	.94	.94	.94 (-)	.94 (-)

Table 3

F1 results of different models for each class (using the Word2Vec embeddings). The best results on each row are highlighted in **bold**. Numbers within (brackets) indicate the improvement in F1 compared to the *best results* from any of the three state of the art methods.

Dataset	State of the art	CNN+sCNN	CNN+GRU
WZ-S.amt	.84 [36]	.92	.92
WZ-S.exp	.91 [36]	.92	.92
WZ-S.gb	.78 [13]	.93	.93
WZ.pj	.83 [30]	.83	.83
WZ	.74 [37]	.83	.83
DT	.90 [9]	.94	.94

Table 4

Comparing micro F1 on each dataset (using the Word2Vec embeddings) against previously reported results on an ‘as-is’ basis. The best performing result on each dataset is highlighted in **bold**. For [37] and [36], we used the result reported under their ‘Best Feature’ setting.

fective on classifying Tweets in the long tail of such datasets as we have shown before in Figure 1. To understand this, we undertake a further analysis below.

On each dataset, we compare the output from our proposed methods against that from Gamback et al. as a reference. We identify the additional Tweets that were correctly classified by either of our CNN+sCNN or CNN+GRU methods. We refer to these Tweets as **additional true positives**. Next, following the same process as that for Figure 1, we 1) compute the uniqueness score of each Tweet (Equation 1) as indicator of the fraction of class-unique words in each of them; 2) bin the scores into 11 ranges; and 3) show the distribution of additional true positives found on each dataset by our methods over these ranges in Figure 6 (i.e., each column in the figure corresponds to a method-dataset pair). Again for simplicity, we label each range using its higher bound on the y axis (e.g., $u \in [0, 0]$ is labelled as 0, and $u \in (0, 0.1]$ as 0.1). As an example, the leftmost column shows that comparing the output from our CNN+GRU model against Gamback et al. on the WZ-S.amt dataset, 38% of the additional true positives have a uniqueness score of 0.

The figure shows that in general, the vast majority of additional true positives have low uniqueness scores. The pattern is particularly strong on WZ and WZ.pj datasets, where the majority of additional true positives have very low uniqueness scores and in a very large number of cases, a substantial fraction of them (between 50 and 60%) have $u(t_i) = 0$, suggesting that these Tweets have no class-unique words at all and therefore, we expect them to potentially have fewer class-unique features. However, our methods still managed to classify them correctly, while the method by Gamback et al. could not. We believe these results are convincing evidence that our methods of using skipped CNN or GRU structures on such tasks can significantly

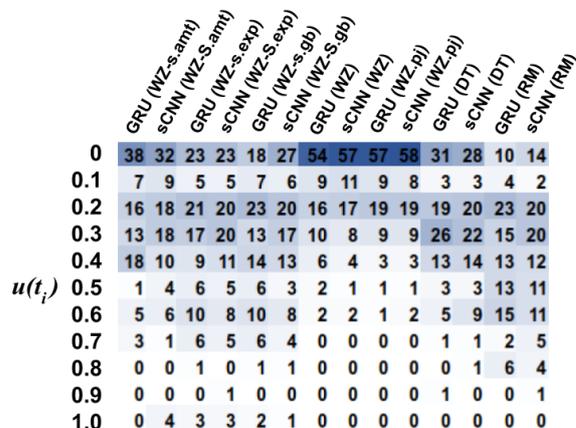


Fig. 6. (Best viewed in colour) Distribution of additional true positives (compared against Gamback et al.) identified by CNN+sCNN (sCNN for shorthand) and CNN+GRU (GRU) over different ranges of uniqueness scores (Equation 1) when using the Word2Vec embeddings. Each row in the heatmap corresponds to a uniqueness score range. Each column corresponds to a method-dataset pair. The numbers in each column sum up to 100% while the colour scale within each cell is determined by the number in that cell.

improve our capability of classifying Tweets that lack discriminative features. Again the results shown in Figure 6 are based on the Word2Vec embeddings but we noticed the same patterns with other embeddings, as shown in Figure 7 in the Appendix.

5.4. Error Analysis

To understand further the challenges of the task, we manually analysed a sample of 200 Tweets covering all classes to identify ones that are incorrectly classified by all methods. We generally split these errors into four categories.

Implicitness (46%) represents the largest group of errors, where the Tweet does not contain explicit lexical or syntactic patterns as useful classification features. Interpretation of such Tweets almost certainly requires complicated reasoning and cultural and social background knowledge. For example, subtle metaphors such as ‘expecting gender equality is the same as genocide’, stereotypical views such as in ‘... these same girls ... didn’t cook that well and aren’t very nice’ are often found in false negative hate Tweets.

Non-discriminative features (24%) is the second

majority case, where the classifiers were confused by certain features that are frequent, seemingly indicative of hate speech but in fact, can be found in both hate and non-hate Tweets. For example, one would assume that the presence of the phrase ‘white trash’ or pattern ‘* trash’ is more likely to be a strong indicator of hate speech than not, such as in ‘White bus drivers are all white trash...’. However, our analysis shows that such seemingly ‘obvious’ features are also prevalent in non-hate tweets such as ‘... I’m a piece of white trash I say it proudly’. The second example does not qualify as hate speech since it does not ‘target individual or groups’ or ‘has the intention to incite harm’.

There is also a large group of Tweets that require interpretation of **contextual information** (18%) such as the threads of conversation that they are part of, or from the included URLs in the Tweets to fully understand their nature. In these cases, the language alone often has no connotation of hate. For example, in ‘what they tell you is their intention is not their intention. <https://t.co/8cmfoOZwxz>’, the language itself does not imply hate. However, when it is combined with the video content referenced by the link, the Tweet incites hatred towards particular religious groups. The content referenced by URLs can be videos, images, websites, or even other Tweets. Another example is ‘@XXX Doing nothing does require an inordinate amount of skill’ (where ‘XXX’ is anonymised Twitter user handle) that is part of a conversation that makes derogatory remarks towards feminists.

Finally, we also identify a fair amount of **disputable annotations** (12%) that we could not completely agree with, even if the context as discussed above has been taken into account. For example, the Tweet ‘@XXX @XXX He got one serve, not two. Had to defend the doubles lines also’ is part of a conversation discussing a sports event and is annotated as sexism in the WS.pj dataset. However, we did not find anything of a hateful nature in the conversation. Another example in the same dataset is ‘@XXX Picwhatting? And you have quoted none of the tweets. What are you trying to say ...?’ is questioning a point raised in another Tweet which we consider as

sexism, but this Tweet itself has been annotated as sexism.

Taking into all such examples into consideration, we see that completely detecting hateful Tweets purely based on their linguistic content still remains extremely challenging, if not impossible.

6. Conclusion and Future Work

The propagation of hate speech on social media has been increasing significantly in recent years, both due to the anonymity and mobility of such platforms, as well as the changing political climate from many places in the world. Despite substantial effort from law enforcement departments, legislative bodies as well as millions of investment from social media companies, it is widely recognised that effective counter-measures rely on automated semantic analysis of such content. A crucial task in this direction is the detection and classification of hate speech based on its targeting characteristics.

This work makes several contributions to state of the art in this research area. Firstly, we undertook a thorough data analysis to understand the extremely unbalanced nature and the lack of discriminative features of hateful content in the typical datasets one has to deal with in such tasks. Secondly, we proposed new DNN based methods for such tasks, particularly designed to capture implicit features that are potentially useful for classification. Finally, our methods were thoroughly evaluated on the largest collection of Twitter datasets for hate speech, to show that they can be particularly effective on detecting and classifying hateful content (as opposed to non-hate), which we have shown to be more challenging and arguably more important in practice. Our results set a new benchmarking reference in this area of research.

Lessons learned. *First*, we showed that the very challenging nature of identifying hate speech from short text such as Tweets is due to the fact that hate Tweets are found in the long tail of a dataset due to their lack of unique, discriminative features. We further showed in experiments that for this very reason, the practice of ‘micro-averaging’ over both hate and non-hate classes in a dataset adopted for reporting results by most previous work can be questionable. It can significantly bias the evaluation towards the dominant non-hate class in a dataset, overshadowing a method’s ability to identify real hate speech.

Second, our proposed ‘skipped’ CNN or GRU structures are able to discover implicit features that can be potentially useful for identifying hate Tweets in the long tail. Interestingly, this may suggest that both structures can be potentially effective in the case where there is a lack of training data, and we plan to further evaluate this in the future. Among the two, the skipped CNNs performs much better.

Future work. We aim to explore the following directions of research in the future.

First, we will explore other branches of methods that aim at compensating the lack of training data in supervised learning tasks. Methods such as transfer learning could be potentially promising, as they study the problem of adapting supervised models trained in a resource-rich context to a resource-scare context. We will investigate, for example, whether features discovered from one hate class can be transferred to another, thus enhancing the training of each other.

Second, as shown in our data analysis as well as error analysis, the presence of abstract concepts such as ‘sexism’, ‘racism’ or even ‘hate’ in general is very difficult to detect if solely based on linguistic content. Therefore, we see the need to go beyond pure text classification and explore possibilities to model and integrate knowledge about users, social groups, mutual communication and even background knowledge (e.g., concepts expressed from Tweets) encoded in existing semantic knowledge bases.

Finally, our methods prove to be effective for classifying Tweets, a type of short texts. We aim to investigate whether the benefits of such DNN structures can generalise to other short text classification tasks, such as in the context of sentences.

References

- [1] Pinkesh Badjatiya, Shashank Gupta, Manish Gupta, and Vasudeva Varma. Deep learning for hate speech detection in tweets. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 759–760, 2017.
- [2] BBCNews. Countering hate speech online, Last accessed: July 2017, <http://eeagrants.org/News/2012/>.
- [3] Pete Burnap and Matthew L. Williams. Cyber hate speech on twitter: An application of machine classification and statistical modeling for policy and decision making. *Policy and Internet*, 7(2):223–242, 2015.
- [4] Pete Burnap and Matthew L. Williams. Us and them: identifying cyber hate on twitter across multiple protected characteristics. *EPJ Data Science*, 5(11):1–15, 2016.
- [5] Ying Chen, Yilu Zhou, Sencun Zhu, and Heng Xu. Detecting offensive language in social media to protect adolescent online safety. In *Proceedings of the 2012 ASE/IEEE International Conference on Social Computing and 2012 ASE/IEEE International Conference on Privacy, Security, Risk and Trust, SOCIALCOM-PASSAT ’12*, pages 71–80, Washington, DC, USA, 2012. IEEE Computer Society.
- [6] Billy Chiu, Anna Korhonen, and Sampo Pyysalo. Intrinsic evaluation of word vectors fails to predict extrinsic performance. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP, RepEval@ACL 2016, Berlin, Germany, August 2016*, pages 1–6, 2016.
- [7] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Deep Learning and Representation Learning Workshop*, 2014.
- [8] Maral Dadvar, Dolf Trieschnigg, Roeland Ordelman, and Franciska de Jong. Improving cyberbullying detection with user context. In *Proceedings of the 35th European Conference on Advances in Information Retrieval, ECIR’13*, pages 693–696, Berlin, Heidelberg, 2013. Springer-Verlag.
- [9] Thoams Davidson, Dana Warmesley, Michael Macy, and Ingmar Weber. Automated hate speech detection and the problem of offensive language. In *Proceedings of the 11th Conference on Web and Social Media. AAAI*, 2017.
- [10] Karthik Dinakar, Birago Jones, Catherine Havasi, Henry Lieberman, and Rosalind Picard. Common sense reasoning for detection, prevention, and mitigation of cyberbullying. *ACM Trans. Interact. Syst.*, 2(3):18:1–18:30, September 2012.
- [11] Nemanja Djuric, Jing Zhou, Robin Morris, Mihajlo Grbovic, Vladan Radosavljevic, and Narayan Bhamidipati. Hate speech detection with comment embeddings. In *Proceedings of the 24th International Conference on World Wide Web*, pages 29–30. ACM, 2015.
- [12] Igin Galiardone, Danit Gal, Thiago Alves, and Gabriela Martinez. Countering online hate speech. *UNESCO Series on Internet Freedom*, pages 1–73, 2015.
- [13] Björn Gambäck and Utpal Kumar Sikdar. Using convolutional neural networks to classify hate-speech. In *Proceedings of the First Workshop on Abusive Language Online*, pages 85–90. Association for Computational Linguistics, 2017.
- [14] Njagi Dennis Gitari, Zhang Zuping, Hanyurwimfura Damien, and Jun Long. A lexicon-based approach for hate speech detection. *International Journal of Multimedia and Ubiquitous Engineering*, 10(10):215–230, 2015.
- [15] Edel Greevy and Alan F. Smeaton. Classifying racist texts using a support vector machine. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’04*, pages 468–469, New York, NY, USA, 2004. ACM.
- [16] Guardian. Anti-muslim hate crime surges after manchester and london bridge attacks, Last accessed: July 2017, <https://www.theguardian.com>.
- [17] Guardian. Zuckerberg on refugee crisis: ‘hate speech has no place on facebook’, Last accessed: July 2017, <https://www.theguardian.com>.
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference for Learning Representations*, 2015.

- [19] Irene Kwok and Yuzhou Wang. Locate the hate: Detecting tweets against blacks. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, AAAI'13, pages 1621–1622. AAAI Press, 2013.
- [20] Quanzhi Li, Sameena Shah, Xiaomo Liu, and Armineh Nourbakhsh. Data sets: Word embeddings learned from tweets and general data. In *Proceedings of the Eleventh International Conference on Web and Social Media, ICWSM 2017, Montréal, Québec, Canada, May 15-18, 2017.*, pages 428–436, 2017.
- [21] Natasha Lomas. Facebook, google, twitter commit to hate speech action in germany, Last accessed: July 2017.
- [22] James D. McCaffrey. Why you should use cross-entropy error instead of classification error or mean squared error for neural network classifier training, Last accessed: Jan 2018, <https://jamesmccaffrey.wordpress.com>.
- [23] Yashar Mehdad and Joel Tetreault. Do characters abuse more than words? In *Proceedings of the SIGDIAL 2016 Conference*, pages 299–303, Los Angeles, USA, 2016. Association for Computational Linguistics.
- [24] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781, 2013.
- [25] Thien Huu Nguyen and Ralph Grishman. Modeling skip-grams for event detection with convolutional neural networks. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 886–891, 2016.
- [26] Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. Abusive language detection in online user content. In *Proceedings of the 25th International Conference on World Wide Web*, pages 145–153, 2016.
- [27] John T. Nockleby. *Hate Speech*, pages 1277–1279. Macmillan, New York, 2000.
- [28] A. Okeowo. Hate on the rise after trump's election, Last accessed: July 2017, <http://www.newyorker.com/>.
- [29] Francisco Javier Ordóñez and Daniel Roggen. Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition. *Sensors*, 16(1), 2016.
- [30] Jo Ho Park and Pascale Fung. One-step and two-step classification for abusive language detection on twitter. In *ALWI: 1st Workshop on Abusive Language Online*, Vancouver, Canada, 2017. Association for Computational Linguistics.
- [31] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- [32] David Robinson, Ziqi Zhang, and John Tepper. Hate speech detection on twitter: Feature engineering v.s. feature selection. In *Proceedings of the 15th Extended Semantic Web Conference, Poster Volume, ESWC'18*. CEUR-WS.org, 2018.
- [33] Anna Schmidt and Michael Wiegand. A survey on hate speech detection using natural language processing. In *International Workshop on Natural Language Processing for Social Media*, pages 1–10. Association for Computational Linguistics, 2017.
- [34] Fabio Del Vigna, Andrea Cimino, Felice Dell'Orletta, Marinella Petrocchi, and Maurizio Tesconi. Hate me, hate me not: Hate speech detection on facebook. In *Proceedings of the First Italian Conference on Cybersecurity*, pages 86–95, 2017.
- [35] William Warner and Julia Hirschberg. Detecting hate speech on the world wide web. In *Proceedings of the Second Workshop on Language in Social Media, LSM '12*, pages 19–26. Association for Computational Linguistics, 2012.
- [36] Zeerak Waseem. Are you a racist or am i seeing things? annotator influence on hate speech detection on twitter. In *Proc. of the Workshop on NLP and Computational Social Science*, pages 138–142. Association for Computational Linguistics, 2016.
- [37] 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*, pages 88–93. Association for Computational Linguistics, 2016.
- [38] Guang Xiang, Bin Fan, Ling Wang, Jason Hong, and Carolyn Rose. Detecting offensive tweets via topical feature discovery over a large scale twitter corpus. In *Conference on Information and Knowledge Management*, pages 1980–1984. ACM, 2012.
- [39] Shuhan Yuan, Xintao Wu, and Yang Xiang. A two phase deep learning model for identifying discrimination from tweets. In *Proceedings of 19th International Conference on Extending Database Technology*, pages 696–697, 2016.
- [40] Ziqi Zhang, David Robinson, and John Tepper. Detecting hate speech on twitter using a convolution-gru based deep neural network. In *Proceedings of the 15th Extended Semantic Web Conference, ESWC'18*, pages 745–760. Springer, 2018.
- [41] Haoti Zhong, Hao Li, Anna Squicciarini, Sarah Rajtmajer, Christopher Griffin, David Miller, and Cornelia Caragea. Content-driven detection of cyberbullying on the instagram social network. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI'16*, pages 3952–3958. AAAI Press, 2016.

Appendix A. Full results

Our full results obtained with different word embeddings are shown in Table 6 for the three re-implemented state of the art methods, and Table 7 for the proposed CNN+sCNN and CNN+GRU methods. **e.w2v**, **e.twt**, and **e.glv** each denotes the Word2Vec, Twitter, and GloVe embeddings respectively. We also analyse the percentage of OOV in each pre-trained word embeddings models, and show the statistics in Table 5.

From the results, we cannot identify any word embeddings that consistently outperform others on all tasks and datasets, and there is also no strong correlation between the percentage of OOV in each word embeddings model and the obtained F1 with that model. Using the Gambäck et al. baseline as an example (Table 6), despite being the least complete embeddings model, e.w2v still obtained the best F1 when classifying racism Tweets on 6 datasets. On the contrary, despite being the most complete embedding model, e.twt only obtained the best F1 when classifying sexism Tweets on four datasets.

Although counter-intuitive, this may not be very much a surprise, considering previous findings from [6]. The authors showed that the performance of word embeddings on intrinsic evaluation tasks (e.g., word similarity) does not always correlate to that on extrinsic, or downstream tasks. In details, they showed that the context window size used to train word embeddings can affect the trade-off between capturing the domain relevance and functions of words. A large context window not only reduces sparsity by introducing more contexts for each word, it also better captures the topics of words. A small window on the other hand, captures word function. In sequence labelling tasks, they showed that word embeddings trained using a context window of just 1 performed the best.

Arguably in our task, the topical relevance of words may be more important for the classification of hate speech. Although the Twitter based embeddings model has the best coverage, it may have suffered from insufficient context during training, since Tweets are often very short compared to other corpora used to train

Word2Vec and GloVe embeddings. As a result, the topical relevance of words captured by these embeddings may have lower quality than we expect, and therefore empirically, they do not lead to consistently better performance than Word2Vec or GloVe embeddings that are trained on very large, context rich, general purpose corpus.

Dataset	e.twt	e.w2v	e.glv
WZ	4%	13%	6%
WZ-S.amt	3%	10%	5%
WZ-S.exp	3%	10%	5%
WZ-S.amt	3%	10%	5%
WZ.pj	4%	14%	7%
DT	6%	25%	11%
RM	4%	11%	6%

Table 5

Percentage of OOV in each pre-trained embedding model across all datasets.

Dataset and classes		SVM			Gambäck et al. [13]									Park et al. [30]								
					e.w2v			e.twt			e.glv			e.w2v			e.twt			e.glv		
		P	R	F	P	R	F	P	R	F	P	R	F	P	R	F	P	R	F	P	R	F
WZ-S.amt	racism	.17	.44	.22	.51	.09	.15	.37	.05	.09	.42	.03	.06	.53	.11	.18	.47	.07	.13	.56	.04	.07
	sexism	.59	.79	.68	.85	.75	.80	.87	.71	.78	.87	.75	.81	.85	.75	.80	.86	.72	.78	.86	.75	.80
	non-hate	.95	.82	.88	.93	.97	.95	.92	.98	.95	.93	.98	.95	.93	.97	.95	.93	.98	.95	.93	.98	.95
WZ-S.exp	racism	.21	.51	.26	.61	.41	.49	.58	.31	.40	.53	.18	.26	.59	.38	.46	.56	.26	.35	.57	.21	.31
	sexism	.46	.76	.58	.74	.58	.65	.75	.60	.67	.74	.61	.67	.74	.58	.65	.75	.59	.66	.73	.61	.67
	non-hate	.96	.83	.89	.93	.97	.95	.94	.97	.95	.93	.97	.95	.93	.97	.97	.93	.97	.95	.94	.97	.95
WZ-S.gb	racism	.28	.70	.38	.53	.30	.38	.53	.22	.31	.61	.21	.31	.48	.30	.37	.57	.22	.31	.64	.22	.33
	sexism	.54	.79	.64	.81	.70	.75	.84	.69	.76	.81	.72	.76	.81	.71	.76	.84	.70	.76	.81	.72	.76
	non-hate	.96	.85	.90	.94	.97	.96	.94	.98	.96	.94	.97	.96	.94	.97	.96	.94	.98	.96	.94	.97	.96
WZ.pj	racism	.54	.68	.60	.75	.65	.69	.74	.68	.71	.75	.65	.69	.75	.65	.70	.75	.65	.70	.74	.67	.70
	sexism	.51	.66	.58	.77	.60	.67	.77	.59	.67	.78	.57	.66	.77	.60	.68	.78	.57	.66	.77	.57	.68
	non-hate	.86	.75	.80	.84	.92	.88	.84	.92	.88	.83	.92	.88	.84	.92	.88	.84	.92	.88	.84	.92	.88
WZ	racism	.54	.65	.59	.75	.68	.72	.74	.71	.72	.76	.65	.70	.75	.68	.71	.76	.69	.72	.76	.67	.71
	sexism	.51	.66	.57	.79	.58	.67	.79	.57	.66	.81	.53	.64	.80	.57	.67	.80	.57	.66	.81	.53	.64
	non-hate	.85	.75	.79	.84	.92	.88	.84	.91	.88	.83	.93	.88	.84	.92	.88	.84	.92	.88	.83	.93	.88
DT	hate	.15	.52	.23	.51	.16	.25	.52	.19	.27	.54	.20	.29	.54	.15	.23	.54	.20	.29	.54	.21	.30
	non-hate	.97	.81	.88	.95	.99	.97	.95	.99	.97	.95	.99	.97	.95	.99	.97	.95	.99	.97	.95	.99	.97
RM	hate	.45	.81	.58	.74	.62	.67	.75	.57	.65	.76	.59	.66	.75	.62	.68	.76	.56	.64	.74	.59	.66
	non-hate	.96	.78	.86	.92	.96	.94	.92	.96	.94	.92	.96	.94	.93	.96	.94	.91	.96	.94	.92	.96	.94

Table 6

Full results obtained by all baseline models.

Dataset and classes		CNN+sCNN									CNN+GRU								
		e.w2v			e.twt			e.glv			e.w2v			e.twt			e.glv		
		P	R	F	P	R	F	P	R	F	P	R	F	P	R	F	P	R	F
WZ-S.amt	racism	.45	.22	.29	.43	.26	.32	.36	.21	.26	.44	.14	.21	.44	.15	.22	.48	.18	.26
	sexism	.84	.78	.81	.80	.79	.79	.81	.81	.81	.84	.75	.80	.86	.74	.79	.83	.77	.80
	non-hate	.94	.97	.95	.94	.96	.95	.95	.96	.95	.93	.97	.95	.93	.97	.95	.94	.97	.95
WZ-S.exp	racism	.60	.57	.58	.59	.67	.63	.52	.64	.57	.52	.48	.50	.58	.26	.36	.50	.26	.34
	sexism	.72	.65	.68	.68	.76	.72	.67	.74	.70	.67	.64	.65	.75	.60	.67	.74	.61	.67
	non-hate	.93	.95	.96	.96	.95	.95	.96	.95	.95	.94	.96	.95	.93	.97	.95	.94	.97	.95
WZ-S.gb	racism	.55	.53	.54	.53	.63	.58	.50	.52	.51	.54	.47	.50	.62	.38	.47	.66	.34	.45
	sexism	.82	.73	.77	.78	.79	.79	.78	.80	.79	.79	.76	.77	.82	.74	.78	.81	.75	.78
	non-hate	.94	.97	.96	.96	.96	.96	.96	.96	.96	.95	.96	.96	.95	.97	.96	.95	.97	.96
WZ.pj	racism	.72	.73	.73	.64	.86	.74	.69	.80	.74	.73	.67	.70	.73	.72	.72	.72	.76	.74
	sexism	.75	.64	.69	.69	.71	.70	.70	.67	.68	.76	.61	.68	.75	.62	.68	.75	.62	.68
	non-hate	.86	.90	.88	.89	.84	.86	.87	.87	.87	.85	.91	.88	.85	.90	.88	.86	.90	.88
WZ	racism	.73	.74	.74	.64	.87	.74	.69	.84	.76	.75	.70	.73	.74	.74	.74	.74	.72	.73
	sexism	.79	.59	.68	.73	.63	.68	.73	.63	.68	.76	.62	.68	.76	.59	.67	.78	.57	.66
	non-hate	.85	.91	.88	.87	.86	.87	.87	.87	.87	.85	.91	.88	.85	.90	.88	.84	.91	.88
DT	hate	.51	.21	.30	.49	.35	.41	.50	.38	.43	.45	.21	.29	.48	.28	.35	.52	.32	.39
	non-hate	.95	.99	.97	.96	.98	.97	.96	.98	.97	.95	.98	.97	.96	.98	.97	.96	.98	.97
RM	hate	.73	.69	.71	.64	.75	.69	.74	.65	.69	.71	.65	.68	.72	.60	.66	.71	.64	.67
	non-hate	.94	.95	.94	.95	.91	.93	.93	.95	.94	.93	.95	.94	.92	.95	.94	.93	.95	.94

Table 7

Full results obtained by the CNN+GRU and CNN+sCNN models.

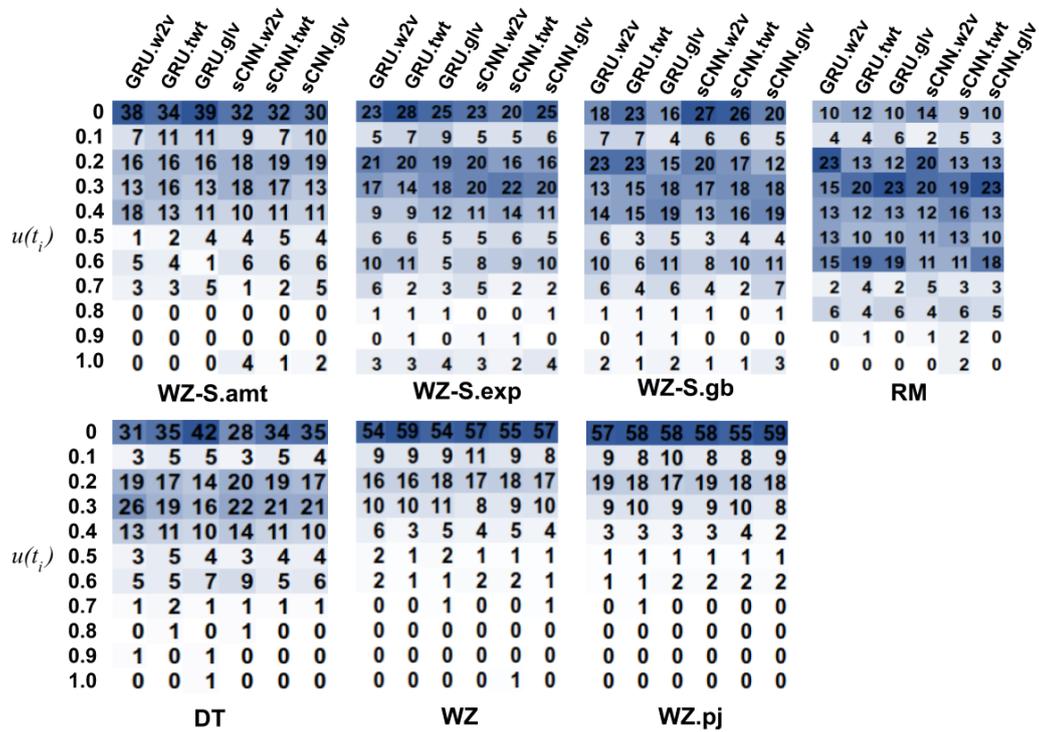


Fig. 7. (Best viewed in colour) Distribution of additional true positives (compared against Gambäck et al.) identified by CNN+sCNN (sCNN for shorthand) and CNN+GRU (GRU) over different ranges of uniqueness scores (see Equation 1) on each dataset. Each row in a heatmap corresponds to a uniqueness score range. The numbers in each column sum up to 100% while the colour scale within each cell is determined by the number label for that cell.