

Large Scale Crowdsourcing and Characterization of Twitter Abusive Behavior

Antigoni-Maria Founta[‡], Constantinos Djouvas^{*}, Despoina Chatzakou[‡], Ilias Leontiadis[#],
Jeremy Blackburn[§], Gianluca Stringhini[†], Athena Vakali[‡], Michael Sirivianos^{*}, Nicolas Kourtellis[#]

[‡]Aristotle University of Thessaloniki ^{*}Cyprus University of Technology [#]Telefonica Research

[§]University of Alabama at Birmingham [†]University College London

{founanti,deppych,avakali}@csd.auth.gr, {costas.tziouvas,michael.sirivianos}@cut.ac.cy,
{nicolas.kourtellis,iliass.leontiadis}@telefonica.com, blackburn@uab.edu,g.stringhini@ucl.ac.uk

Abstract

In recent years, offensive, abusive and hateful language, sexism, racism and other types of aggressive and cyberbullying behavior have been manifesting with increased frequency, and in many online social media platforms. In fact, past scientific work focused on studying these forms in popular media, such as Facebook and Twitter.

Building on such work, we present an 8-month study of the various forms of abusive behavior on Twitter, in a holistic fashion. Departing from past work, we examine a wide variety of labeling schemes, which cover different forms of abusive behavior, at the same time. We propose an incremental and iterative methodology, that utilizes the power of crowdsourcing to annotate a large scale collection of tweets with a set of abuse-related labels. In fact, by applying our methodology including statistical analysis for label merging or elimination, we identify a reduced but robust set of labels. Finally, we offer a first overview and findings of our collected and annotated dataset of 100 thousand tweets, which we make publicly available for further scientific exploration.

Introduction

The rise of hateful behavior online has recently become a topic of interest. The research community has studied hate speech (Djuric et al. 2015), cyberbullying (Hosseinmardi et al. 2015), and semi-organized online harassment campaigns (Hine et al. 2017), while also proposing systems to automatically detect and block abusive behavior (Ribeiro et al. 2017), (Davidson et al. 2017), (Serra et al. 2017). To their credit, social network platforms are also taking steps to mitigate damage, e.g., providing users with tools to flag abusive behavior (Kayes et al. 2015), (Twitter 2017).

Unfortunately, abusive content poses some unique challenges to researchers and practitioners. First and foremost, even defining what qualifies as abuse is not straightforward, which in turn makes it difficult to extract ground truth on which studies can be then based on. Unlike other types of malicious activity, e.g., spam or malware, the accounts carrying out this type of behavior are usually *controlled by humans*, not bots. This makes techniques based on grouping

together similar messages or searching for automated activity ineffective. Second, this activity typically manifests in a minority of situations and not a lot of examples can be gathered in a random collection of posts. Elaborate techniques must be employed to boost the collection of such cases for the automated techniques to work.

To deal with these issues, crowdsourcing proved to be a promising direction towards developing labeled datasets. However, human labeling poses a number of challenges. One of the main issues is the existence of different types of abuse and labels to describe them (e.g., offensive language, hate speech, aggressive behavior, cyberbullying, etc.). It is often difficult, even for a human, to distinguish (Chatzakou et al. 2017) what qualifies as each. For example, certain types of language, such as sarcasm, can be misinterpreted by annotators if the messages are not presented in context. Another challenge is striking a good balance between the number of annotators employed per task, their payment, and how much time the crowdsourcing process takes to complete.

This paper tackles three challenges faced when trying to collect large scale ground truth on abusive behavior: (i) difficulty of the crowdsourced workers to distinguish between different abusive categories (e.g., hate speech vs. offensive language vs. abusive language), (ii) different occurrence rates for different categories of abuse, and (iii) scaling the multi-labeled annotation process to thousand tweets, while maintaining quality of annotation and time-budget constraints.

Into the direction of addressing the aforementioned challenges we proceed with the following contributions:

- A methodology to detect and cut through the confusion of crowdworkers when they are asked to distinguish between nuanced labels.
- A boosted sampling approach that maintains an unbiased dataset, while ensuring more annotations for the minority class.
- The design and development of a data collection platform that optimizes for costs.
- An annotated dataset of abusive behavior on 100k tweets, created using the methodologies and platform we developed. [available upon request]

Related Work

In the recent past, a few datasets have been collected, annotated, and released by other researchers around abusive behavior on Twitter, using various labels and different methodologies for annotation.

A widely used and available dataset in this field, is the one in (Waseem and Hovy 2016), employed by several studies working on hate speech detection such as in (Badjatiya et al. 2017), (Gambäck and Sikdar 2017), (Park and Fung 2017), and (Jha and Mamidi 2017). That work focused on disambiguating different types of hate speech, and more specifically between *racism* and *sexism*, at the level of tweets (i.e., whether a tweet is racist or sexist). The authors collected tweets based on a set of hate-related terms and users, and manually annotated a subset of their dataset using an outside annotator for reviewing. The final dataset consists of almost 17k tweets, 12% of which are labeled as racist, 20% sexist, and the rest are considered normal. After the annotation process, they investigated which features better assist the detection of hate speech by building upon this dataset. Their findings show that only gender plays an important role, while geographic and word-length distributions are almost completely ineffective.

Another dataset that has already been used in different studies (e.g., (Malmasi and Zampieri 2017), (Olteanu, Talamadupula, and Varshney 2017)) is the one presented in (Davidson et al. 2017). The focus of that work was mainly to distinguish between *hateful* and *offensive* language. According to the authors, offensive language contains offensive terms which are not necessarily inappropriate, while hate speech intends to be derogatory, humiliating, or insulting. Since the automatic distinction between the two is important, this dataset is focusing on the differences. They started their annotation process by identifying and collecting a set of possible hateful users, and extracted their tweets. Then, they sampled this collection for 25k tweets containing terms from a hate speech lexicon. Finally, this sampled dataset was placed in CrowdFlower to be annotated by workers. With an intercoder-agreement score of 92%, the vast majority of the final annotations was offensive (77%), and only a very small percentage was hateful (6%), while the rest were normal.

There is also the recent work by (Golbeck et al. 2017), which is focused on online trolling and harassment on Twitter. The authors in that work first used various online sources, such as blocklists, to produce a list of keywords that can be used for collecting harassing tweets with high probability. Subsequently, they created code guidelines on the annotation task and trained coders to label the tweets using labels such as “the very worst,” “threats,” “hate speech,” “direct harassment,” “potentially offensive,” and “non-harassment.” Their aim was to label tweets as harassing, only if they really were “the worst of the worst content.” Their final dataset includes 35k of tweets annotated by 2 or 3 coders.

The previous datasets focus on the used language by annotating text content, i.e., tweets. However, there are other works focusing on user characteristics. That is, they provide annotation of Twitter users based on their exhibited behav-

ior. As (Ribeiro et al. 2017) emphasize, identifying content as hateful raises major issues, while “characterizing and detecting hateful users [...] presents plenty of opportunities to explore a richer feature space.” Therefore, detecting inappropriate user behavior is a closely related task, although under a different, but equally important scope. One such work is (Chatzakou et al. 2017), where the authors detected cyberbullying and cyber-aggression by collecting and annotating a dataset from Twitter. Their annotation methodology is very close to ours, by employing crowdsourcing in CrowdFlower for the annotation task. Their final dataset is consisted of 9,484 tweets, labeled into one of four categories: 1) bullying, 2) aggressive, 3) spam, or 4) normal. The aggressive and bullying labels make up about 8% of the dataset, spam makes up about 33%, and the remainder of the annotations are normal.

While all aforementioned works fall under the same domain, i.e., annotating inappropriate speech, one crucial challenge still remains unaddressed. Specifically, and while studying the existing literature, we noticed that there is an important gap regarding the principled selection of the most appropriate labels for annotating aggressive online behavior. In fact, in past literature, types or labels of inappropriate speech are usually used interchangeably, or selected randomly among available ones, for use in the annotation task. Indeed, past studies do not explain or justify the selection of types of inappropriate speech which they employ for their annotations. Furthermore, they either use only a subset of popular labels, and consider the others covered (without however establishing why this is so), or combine them together under the same umbrella label. In this work, we take a step back and propose a principled methodology to narrow down the list of possible labels used in this space. Our methodology is iterative, to account for limited time and budget, as well as allow for controlled statistical analysis of label selection by annotators. We use the final set of selected labels for a large scale crowdsourcing study, to annotate 100k tweets with appropriate labels on abusive behavior. Table 1 summarizes the past works relevant to the topic, which have released their datasets for scientific exploration.

Overview of Methodology

Our goal: The overall goal of this work is to create a large and highly accurate annotated dataset of tweets (100k) via a crowdsourcing platform like CrowdFlower (CF). Unlike previous work, we are interested in workers selecting from more than one potential category of abusive behavior (i.e., two or more labels), in order to study the correlation between them and make adjustments on the final labels used. Annotating such a large dataset exhibits some unique challenges since we must *minimize the cost without compromising the quality* of the annotation. This can only be achieved by carefully tuning the task and the platform, appropriately selecting the samples to annotate, and striking a balance between worker payment and quality.

Dataset	# Tweets	Labels	Annotators
(Chatzakou et al. 2017)	9,484	aggressive, bullying, spam, normal	5
(Waseem and Hovy 2016)	16,914	racist, sexist, normal	1
(Davidson et al. 2017)	24,802	hateful, offensive (but not hateful), neither	3 or more
(Golbeck et al. 2017)	35,000	the worst, threats, hate speech, direct harassment, potentially offensive, non-harassment	2-3
Present study	100,000	offensive, abusive, hateful speech, aggressive, cyberbullying, spam, normal	5-20

Table 1: Summary of related work datasets.

Challenges with Crowdsourced Platforms

There are several challenges that need to be resolved to best balance high-quality annotations and minimal cost.

Labels: We want to build a dataset that distinguishes between various expressions of online abuse, e.g., abusive and aggressive, hateful, offensive, cyberbullying. However, even if detailed definitions and examples are provided, crowdsourced workers might still find it challenging to consistently label examples. Thus, the overall design of the task (e.g., how to phrase the questions) and the selection of labels to choose from, is an important challenge. To address this, we use a number of preliminary annotation rounds that aim to identify the exact nature of any confusion. This facilitates the elimination of any ambiguity on the labels used during the main annotation and, thus, achieving high accuracy and consistency.

Sampling: In the grand scheme of things, abusive tweets are quite rare (between 0.1% and 3%, depending on the label). Therefore, even large-scale datasets might contain just a few samples. For typical machine learning algorithms which can take benefit from such a dataset, few samples means less opportunity to train on the specific behaviors, and overall, worse classification performance. One way to deal with this extreme imbalance is to pre-select tweets that are likely to be abusive (e.g., those that contain hate words), however, this approach also biases the dataset.

To address this sampling issue, we designed a boosted random sampling technique. A large part of the dataset is randomly sampled, but then boosted with tweets that are likely to belong into one or more of the minority classes. We use text analysis and preliminary crowdsourcing rounds to design a model that can pre-select the tweets of the boosted set. Both sets are then mixed together and given to the crowdsourcing platform for the final annotation.

Judgments: The next challenge to address is determining the proper number of crowdworker decisions that are necessary for a high confidence annotation. As expected, this is largely dependent on a combination of factors: complexity of the task, worker reward, quality of annotators, etc. The solution we settled in is to employ a large number of annotators during the exploratory rounds (up to 20 annotators per tweet) to establish the general level of agreement we should expect, given the number of annotators.

Payment: The payment to crowdworkers also plays an important role in their annotations. Studies have shown that




Figure 1: Data Preparation Pipeline (Step 1). Pre-filtering and spam removal to clean tweets. (A) random set of unboosted tweets. (B) boosted sampling to produce a set of tweets biased towards abusive behavior. Sub-datasets D_1 and D_2 are used in the subsequent Steps 2 and 3.

when participants are payed fairly, it positively affects their results, but always depending on the type of task to be performed (Ye, You, and Jr. 2017). In our case, and similarly to (Chatzakou et al. 2017), we started with a default payment schema (5¢ for a batch of 10 tweets), and used the preliminary rounds to adjust as needed.

Crowdsourcing Methodology

Our methodology addresses the above challenges via a three-step process. We visualize these steps in three figures: Figure 1 illustrates the preliminary data preparation process. Figures 2 and 3 visualize the next two steps which involve the iterative annotation rounds.

Step 1: data collection and sampling. The first step of the process (Figure 1) is to collect a random set of tweets. To do so, we utilize the Twitter Stream API. We store the data in elastic search and we apply basic pre-filtering to exclude spam, tweets that have no content, tweets that are not in English, etc. Furthermore, we apply simple text analysis and machine learning to create the boosted set of tweets, that will be used to improve coverage over the minority classes. Finally, we randomly sample a small dataset (D_1), that is used for the exploratory analysis, and the remainder (D_2) for the large scale annotation.

Step 2: exploratory analysis. Considering the various trade-offs among the different parameters of the annotation task, we first introduce an iterative process that allows the researchers to properly adjust these parameters (Figure 2). This analysis is performed on a small sample (300 tweets in our case), as a means to enable quick and affordable testing among all the different design choices and parameters. These parameters include the payment, the type and number of labels, the presentation of these labels, the number




Figure 2: Exploratory Analysis (Step 2). Dataset D_1 is inputted in the platform for annotation under label set L , and in consecutive rounds. In each round, statistical analysis performed can narrow down the set of labels to L' . Final set of labels L'' can be inputted in Step 3.

of judgments required, trustworthiness of users, annotation process, etc. Furthermore, this process can reveal possible points of confusion (e.g., identify if two labels are frequently mixed).

During these iterations, we fine-tune the filters used to better boost the dataset, in order to contain more samples of the minority classes. After each iteration, an analysis of the results can reveal if a satisfying quality is reached and whether a given parameter has contributed to this, akin to an A/B testing. In an ideal scenario, a researcher can execute many such iterations to optimize better for the set of labels to be used, the money paid to workers, etc. In practice, and always due to limited budget and time, these iterations can only be a handful.

In our case, the process converged after three iterations, allowing us to identify the influence each one of the aforementioned parameters has in the annotated dataset. The main outcome of this step was to determine the most representative and clearly understood set of labels that should be used for the large scale annotation task. It also enabled us to assess the ideal number of judgments to strike a balance between cost and quality. Details about each round are given in the following sections.

Step 3: full annotation. The third step of the process is when we actually annotate the larger dataset (D_2), using the previously established settings and labels (L''). As shown in Figure 3, we built our own custom platform to host the annotation task. To accommodate such a large-scale task, we also created a database schema to store the data and the results, and to calculate the statistics.

In the next sections we examine these steps in detail.

Step 1: Data collection and sampling

In this section, we present the data preparation procedure. We detail how we collect and filter the data, the preprocessing part, and finally the necessary sampling. The pipeline is shown in Figure 1.

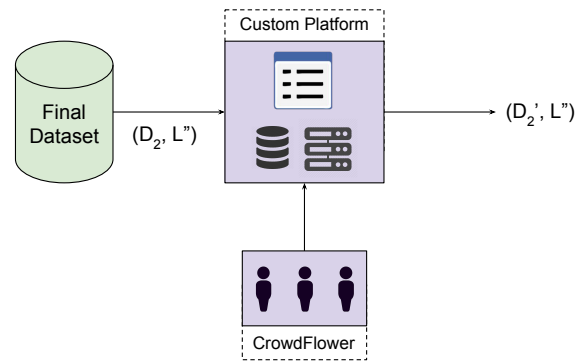


Figure 3: Final Annotation Round (Step 3). A larger dataset D_2 , with the final label set L'' can be used for large scale annotation. A custom-built platform used allows for better control of the annotation flow, and reduce dependencies on CrowdFlower specific design limitations.

Collection

The first step of the process is to collect a random set of tweets. To do so, we utilize the Twitter Stream API and we collect all the tweets provided by the API (1% of the entire traffic) during the period of 30th March 2017 - 9th April 2017, consisting of 32 million tweets in total.

Metadata Extraction

To facilitate filtering and sampling (the next two steps), each tweet is enriched with metadata (Figure 1). First of all, from the tweet's content we extract the number of URLs, hashtags, mentions, emojis/smilies, and numerals. Furthermore, we tag retweets and mentions. Finally, we extract metadata from Twitter, such as the detected language, the account age, etc. Additionally, we apply sentiment analysis, such as polarity and subjectivity of the tweet, using the TextBlob Python library. Finally, we count the number of offensive terms found using two dictionaries (HateBase¹ and an offensive words dictionary²).

Filtering

For the entire dataset we apply some basic preprocessing in order to filter out tweets that should not be annotated. Firstly, we remove tweets that are considered spam. There are numerous techniques for tweet spam detection (Dhingra and Mittal 2015), (Santos et al. 2014), (Wang et al. 2015), (Zhou and Sun), (Wang 2010). Inspired by these, we apply filtering criteria for the elimination of such spam-related tweets. Furthermore, we only keep original tweets (i.e., drop retweets without new content), while also remove those that have small text content (e.g., only URLs, images, etc.). Finally, we remove any tweets that are not written in English, using Twitter's language detection.

¹<https://www.hatebase.org>

²<https://www.noswearing.com/dictionary>

Boosted Sampling

As shown in Figure 1, after we collect and clean the data, the next step is to create the final dataset that will be used in the various rounds. One major issue that needs to be addressed when considering such datasets is the class imbalance of the behavior under study. More particularly, in the case of abuse detection, even though inappropriate content is very frequent in Online Social Networks, it is still a minority compared to the tremendous amount of “normal” data produced. Therefore, when selecting the data to create a sample that will be annotated, it is necessary to ensure there will be plenty of inappropriate annotations to work with, otherwise the dataset is not very useful for the research community. Therefore, we follow a sampling procedure and inject the selected data in the randomly sampled ones.

For the boosted sample, we use the metadata extracted earlier. We choose tweets that, based on the sentiment analysis, show strong negative polarity (< -0.7) and contain at least one offensive word. Finally, we create two datasets: D_1 is a sample of just 300 tweets that is used for the exploratory analysis, and D_2 that contains 100K tweets that will be used for the final annotation.

Datasets

In total we work with two datasets, D_1 and D_2 , one for each step. In Table 2 we present the datasets used per round with some extra information regarding the annotations.

Step 2: Exploratory Rounds

The goal of this step is to tune the crowdsourcing parameters on a smaller dataset, in order to quickly get some insights but minimize the cost while doing so.

We focused our exploration on identifying the most representative labels related to the types of abusive content. We begin with the most extensively used labels found in literature, and at each round we looked into the results and merge/remove labels that were frequently confused by the annotators. Furthermore, these rounds helped us to further filter spam and get a more representative boosted sampling.

We begin with a first round that includes 300 tweets and 5 judgments per tweet. We collect annotations and assess if the plurality of labels is confusing, how the spamming annotation works, etc. Afterwards, we continue to a second round, where we focus only on the tweets that were marked as inappropriate in the first round, but requesting a much larger number of annotations to assess better the overlap of used labels. Finally, we conclude with a third round to validate the selected labels and confirm annotation agreement, before moving on to the Step 3 and the large scale annotation.

Definitions

Before starting, annotators are provided with definitions for each label which they have to acknowledge reading. The definitions are constructed based on all the descriptions we found in the related literature, as cited on each category, as well as Cambridge³ and Black’s Law⁴ dictionaries. In total,

³<http://dictionary.cambridge.org>

⁴<http://thelawdictionary.org>

the following definitions were displayed to the annotators:

- **Offensive Language:** Profanity, strongly impolite, rude or vulgar language expressed with fighting or hurtful words in order to insult a targeted individual or group. (Chen et al. 2012), (Razavi et al. 2010).
- **Abusive Language:** Any strongly impolite, rude or hurtful language using profanity, that can show a debasement of someone or something, or show intense emotion. (Papegnies et al. 2017), (Park and Fung 2017), (Nobata et al. 2016).
- **Hate Speech:** Language used to express hatred towards a targeted individual or group, or is intended to be derogatory, to humiliate, or to insult the members of the group, on the basis of attributes such as race, religion, ethnic origin, sexual orientation, disability, or gender. (Davidson et al. 2017), (Badjatiya et al. 2017), (Warner and Hirschberg 2012), (Schmidt and Wiegand 2017), (Djuric et al. 2015).
- **Aggressive Behavior:** Overt, angry and often violent social interaction delivered via electronic means, with the intention of inflicting damage or other unpleasantness upon another individual or group of people, who perceive such acts as derogatory, harmful, or unwanted. (Chatzakou et al. 2017), (Hosseinmardi et al. 2015).
- **Cyberbullying Behavior:** It’s the use of force, threat, or coercion to abuse, embarrass, intimidate, or aggressively dominate others, using electronic forms of contact. It typically denotes repeated and hostile behavior performed by a group or an individual. (Chatzakou et al. 2017), (Hosseinmardi et al. 2015), (Dinakar, Reichart, and Lieberman 2011), (Riadi and others 2017), (Kansara and Shekokar 2015).
- **Spam:** Posts that are consisted of related or unrelated advertising / marketing, selling products of adult nature, linking to malicious websites, phishing attempts and other kinds of unwanted information, usually executed repeatedly.
- **Normal:** all tweets that do not fall in any of the prior categories.

Therefore, our starting set of labels L is:

$$L = \{Offensive, Abusive, Hateful, Aggressive, Cyberbullying, Spam, Normal\}$$

First Round

On the first round, annotators were asked in a primary selection, to first classify tweets into three general categories: normal, spam, and inappropriate. In the case that inappropriate was selected, then a secondary panel offered them the five aforementioned inappropriate speech categories. This way, users could define more explicitly the type exhibited by the tweet. Furthermore, they had the option to suggest a new subcategory utilizing the “other” option and a text box. Finally, the participants were encouraged to select multiple subcategories whenever appropriate. The dataset described

	Dataset	Tweets	Judgments	Sampling Percentage
Step 2				
Round 1	D_1	300	5	33% boosted - 67% random
Round 2	<i>Sample of D_1</i>	88	10-20	92% boosted - 8% random
Validation Round	D_1	300	5	33% boosted - 67% random
Step 3				
Final Round	D_2	100k	5	10% boosted - 90% random

Table 2: Datasets per Round




Figure 4: Distributions of judgments per inappropriate label for the two exploratory rounds in Step 2.

above was sent to CrowdFlower for annotation, asking for five judgments per tweet.

In Figures 4 and 5 we see (in blue color) some of the results of this round. Figure 4 shows the distributions of judgments per inappropriate label. It is important to note here that the percentages of both rounds take under consideration only the total amount of inappropriate labels, since these are the ones we want to observe. We notice that Offensive and Abusive are the most popular categories, followed by Hateful and Aggressive. Cyberbullying is rarely used. Normal and Spam are not presented on the figure, but are very frequently used, with a percentage of 53% for Normal, and 15% for Spam, overall.

In Figure 5 we observe the agreement of the annotators, when there is majority voted, grouped in three majority categories, for convenience. The three categories are: i) Overwhelming majority, when at least 80% of the annotators agree, ii) Strong Majority, when at least 50% of the annotators agree and iii) Simple majority, for the rest of the cases. When two or more labels have equal number of judgments, the tweet is not included in any of the three categories, since it is not assigned with a majority vote. The results of the first round show a clear “win” of the Overwhelming category, which means that most participants agreed on their votes. On the other hand, most of the judgments of this round are Normal or Spam, therefore we can not be positive that the majority results refer to the inappropriate labels. The confusion becomes more clear when we run the second round, the results of which are presented below.




Figure 5: Categories of majority distributions for all preliminary rounds in Step 2.

Second Round

Results presented on the first round provide some insights on the correlation among inappropriate speech categories. However, our confidence on these results is low, mainly because of the low amount of annotations per tweet. For this particular reason, we decide to proceed in a new annotation round. Here, we use only the tweets that were previously annotated as Inappropriate, with a high agreement score. In total, these tweets are 88 out of the initial 300. Each tweet was consequently annotated by at least 10 workers, but usually around 20. We kept the same setup regarding labels and instructions, as we want to be able to compare the two rounds afterwards.

The results of Figure 4 show a similar, yet not identical, distribution of the five labels. Offensive and Abusive are still the “leading” labels, although Abusive is slightly more popular in this round. Hateful and Aggressive follow again, in the same order. Finally, Cyberbullying is again very low. On the other hand, majorities in Figure 5 have completely changed. As it appears, in most cases annotators disagree about the labels, and only very few have a high amount of agreement. This clearly shows that the task of choosing between our set of labels is not trivial.

Comparison of the Exploratory Rounds

In order to study how labels are related, we compare the results of the previous two rounds. More specifically, we calculate correlations of the various labels, measure their similarities and report on co-occurrences. In this section, we study these statistics and reach the final decision over which labels will be kept for the validation round.

We begin by measuring the correlation and similarities among the inappropriate category labels. Such correlations and similarities will allow us to measure how closely each label appears to have been selected with another label, given a set of tweets. We calculate correlations and their significance using Pearson, Spearman, and Kendall Tau Correlation Coefficients. The similarity between labels is measured using Cosine Similarity. For each pair of labels, we calculate their similarity vectors, in order to gain some insight on the correspondence of pairs in accordance with their ranking. That is, for each label, we construct a vector of votes, with each cell representing a tweet annotated. Then, we compute the similarity of these vectors, in all-pairwise fashion between the available labels. Similarly, we compute the correlation of labels using these vectors. In Table 3 we present these results for both rounds, in order to compare.

On the first round, all three correlation coefficient metrics show low correlation between most of the labels. The only correlation that seems consistent statistically significant in all cases is between Abusive and Aggressive ($p < 0.05$). Moreover, Aggressive and Cyberbullying seem to be somewhat correlated, but the significance is not consistent. Finally, Kendall Tau also shows a significant relationship between Offensive and Aggressive, which does not appear in the other two cases. The rest of the combinations do not exhibit any important correlations.

When it comes to the second round, there are generally low correlations with no statistical significance between the labels, with some exceptions. Offensive and Abusive are correlated in statistically significant fashion ($p < 0.05$), and this is consistent across all metrics. Furthermore, in most of the cases, both are also significantly correlated to Cyberbullying. Spearman and Kendall also show some correlation between Offensive and Abusive with Aggressive. Hateful, never seems to be correlated with the rest of the labels.

Regarding the cosine similarities, we see in Table 3 that in the first round, the values of similarities are not very high. Nevertheless, the most highly similar pairs are Offensive and Aggressive with Abusive. On the other hand, during the second round, we notice that the values of the similarities are much higher than before. Again the most similar pair of labels is Offensive - Abusive, followed by Abusive - Hateful. We notice in general, in this annotation round, that hateful seems to be more related than it was on the previous round, but the correlation results still don't indicate a strong correspondence.

To support the previous results, we also calculate the co-occurrences of the various labels for each one of the three majority agreement groups (Overwhelming, Strong and Simple). Due to space limitations, we can not fully present the co-occurrences results here. However, we briefly state what we observed and how they support our final learnings. The results show that users seem to be very confused about selecting a label, resulting in low levels of agreement for most of the inappropriate tweets. On the second annotation round, for example, Abusive seems to be used a lot of times and is often the majority label, but it's always confused with many of the other labels (especially Offensive). Offensive is also confused with Abusive and Aggressive, and fre-

quently also with Hateful. Finally, Cyberbullying never becomes a majority-label. As expected, this becomes more intense with harder to categorize tweets, i.e., tweets that most annotators disagreed in their judgments.

Insights learned: The results of the first two annotation rounds allowed us to draw the following conclusions regarding the use of the five inappropriate speech labels:

1. We can form three groups of labels according to their popularity: Abusive and Offensive are the two most popular labels, Hateful and Aggressive are somewhat popular, and Cyberbullying is rarely used. Therefore Cyberbullying can be safely eliminated from the list of inappropriate labels, mainly due to the very few times it was selected on both the annotation rounds. This decision, though, is also supported by the very nature of Cyberbullying, which according to its definition should be repetitive. Yet, in our case we have no sense of time or repetition, since we work with individual tweets.
2. Abusive, Offensive and Aggressive seem to be significantly correlated, highly coexisting in the annotations and very similar (according to the similarity results). Abusive is the most popular among the three and the most central (i.e. the other two labels are much more related with this than with each other).
3. While Hateful is frequently coexisting with other labels, indicating a confusion among users in the use of this label, it does not appear to be significantly correlated with any other label. This is also supported by the definition of Hateful, since there is a well-defined description of the target groups of this category, compared to the rest.

Validation Round

Given the above insights, we proceed with extra validation rounds, before the large-scale annotation. To do so, we first remove Cyberbullying (due to point 1). Then, utilizing the insights from point 2, we merge Abusive, Offensive and Aggressive into a single category. To choose which label to use, we run one annotation campaign for each keyword, and we achieved similar results. Therefore, we kept Abusive as the keyword for this category of tweets. Finally, we decided to keep Hateful separately, as explained in point 3. Thus, in these validation rounds, users were presented with four labels:

$$L' = \{Abusive, Hateful, Normal, Spam\}$$

The dataset used is again D_1 (containing 300 tweets) and we required again 5 judgments per tweet.

We begin with Figure 6, where we notice the distribution of judgments towards labels. We see that the final inappropriate labels are much more frequently used now than before, something that was expected as we merged most of them. Hateful is still not as frequent as Abusive, but it still appears in almost 7% of the judgments, therefore can not be eliminated. In the agreement graph of Figure 5, we clearly notice a vast improvement from the previous rounds. Almost 70% of the tweets reach an Overwhelming Agreement (more than 3 out of 5 annotators agree), while annotators disagree highly only in very few tweets. This result is of course

	PCC	p-PCC	SCC	p-SCC	KTCC	p-KTCC	CosSim
First Round (300 tweets)							
Offensive - Abusive	0.057672	0.4863	0.109460	0.1854	0.095695	0.0844	0.536908
Offensive - Hateful	-0.064017	0.4395	-0.008290	0.9203	-0.007995	0.8854	0.410749
Offensive - Aggressive	-0.122807	0.1370	-0.124149	0.1327	-0.110099	0.0471	0.348367
Offensive - Cyberbullying	-0.083271	0.3143	-0.055111	0.5059	-0.049595	0.3711	0.209020
Abusive - Hateful	-0.096501	0.2433	-0.036050	0.6636	-0.033111	0.5504	0.320653
Abusive - Aggressive	0.195979	0.0170	0.324244	0.0001	0.285359	0.0000	0.478639
Abusive - Cyberbullying	0.042049	0.6118	0.065070	0.4320	0.060229	0.2774	0.251285
Hateful - Aggressive	-0.042224	0.6104	0.024063	0.7716	0.020994	0.7050	0.279881
Hateful - Cyberbullying	-0.076778	0.3537	-0.080666	0.3298	-0.076305	0.1688	0.133986
Aggressive - Cyberbullying	-0.142836	0.0833	-0.169740	0.0392	-0.161572	0.0036	0.066667
Second Round (88 tweets)							
Offensive - Abusive	0.322597	0.0022	0.408552	0.0001	0.294481	0.0000	0.741228
Offensive - Hateful	-0.076287	0.4799	-0.130442	0.2258	-0.095233	0.1889	0.544227
Offensive - Aggressive	0.056567	0.6006	0.245213	0.0213	0.186104	0.0102	0.482113
Offensive - Cyberbullying	0.230017	0.0311	0.191246	0.0743	0.157496	0.0298	0.497397
Abusive - Hateful	0.126504	0.2402	0.118195	0.2727	0.079851	0.2706	0.619584
Abusive - Aggressive	0.011576	0.9148	0.270948	0.0107	0.199341	0.0060	0.451047
Abusive - Cyberbullying	0.243139	0.0225	0.241344	0.0235	0.202128	0.0053	0.501380
Hateful - Aggressive	-0.072054	0.5047	-0.039991	0.7114	-0.030565	0.6733	0.374166
Hateful - Cyberbullying	0.013788	0.8985	0.003095	0.9772	0.003917	0.9569	0.350230
Aggressive - Cyberbullying	-0.001821	0.9866	0.125888	0.2425	0.109380	0.1313	0.268009

Table 3: Correlation Coefficients, p-values and Cosine Similarity values for each pair of inappropriate labels in the Exploratory Rounds.




Figure 6: Distributions of judgments per inappropriate label for the validation round.

highly connected with the fact that we presented a simpler task for the workers to complete. The correlations coefficient and similarities table (Table 4) also depicts some consistent results. More specifically, we have statistically significant negative correlations when observing the interactions of Abusive with all other labels (i.e., the appearance of the Abusive label is negatively correlated with the appearance of the other labels) and no important correlation in any other combination of labels. In the last column, we see the cosine similarities between the pairs. We notice that the vectors are much less similar, while also Abusive is clearly different than Normal and Spam and not very similar with Hateful either. Finally, from the co-occurrences results we saw that Abusive and Hateful are still sometimes confused (showing that our task is still not trivial), but this is not very frequent across tweets.

Step 3: Large Scale Annotation

Based on the decisions drawn upon the previous results, we launched the final, large-scale annotation task of 100k tweets, with 5 judgments per tweet. The setup for this round is kept the same with the last validation round, since it's already tested. Thus, the final labels we decided upon, as analyzed earlier, are:

$$L'' = \{Abusive, Hateful, Normal, Spam\}$$

Annotator Profiles

In order for CrowdFlower workers to participate in the annotation task, we first require some basic demographic information, such as gender, age, annual income, level of education and nationality. This is to have a better understanding of the annotators' profiles. Here, we analyze these demographics. First of all, we start with gender. Almost two thirds of the participants are male (67.5%) and one third is female (32.5%). Moreover, even though we provided an "other" option, almost no workers selected it. Regarding their educational level, most of them have a Bachelor Degree (52.2%), followed by Secondary Education (26.3%) and Master's Degree (20.1%), while very few have a PhD (1.5%).

The age of the participants ranges from 18 to over 87; 30.9% are between 18 and 24, 29.2% are 25-31, 19.8% are 32-38 and the remainder above 39 years old. More than half of them claim to have an income level below €10k (58.7%), 14.5% are between €10k and €20k and the rest spread across €20k and €100k. Finally, their nationalities vary a lot, coming from 87 different countries in total. However, by far the most frequent country of origin is Venezuela (44.8%), followed by USA (7.9%), Egypt (7.6%) and India (5.6%). Overall, the annotators from the top 10 countries

	PCC	p-PCC	SCC	p-SCC	KTCC	p-KTCC	CosSim
Abusive - Hateful	-0.378388	0.0000	-0.436322	0.0000	-0.375022	0.0000	0.382150
Abusive - Normal	-0.839857	0.0000	-0.843253	0.0000	-0.738567	0.0000	0.195936
Abusive - Spam	-0.314023	0.0001	-0.350580	0.0000	-0.304493	0.0000	0.136757
Hateful - Normal	-0.068685	0.4101	0.024373	0.7703	0.014200	0.7992	0.406726
Hateful - Spam	-0.072026	0.3876	-0.027309	0.7435	-0.025878	0.6430	0.194662
Normal - Spam	0.028358	0.7340	0.130268	0.1171	0.111400	0.0460	0.267608

Table 4: Correlation Coefficients, p-values and Cosine Similarity values for the labels of Validation Round.

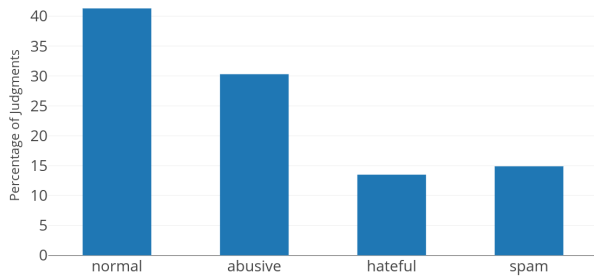


Figure 7: Label distributions of the large-scale annotated dataset

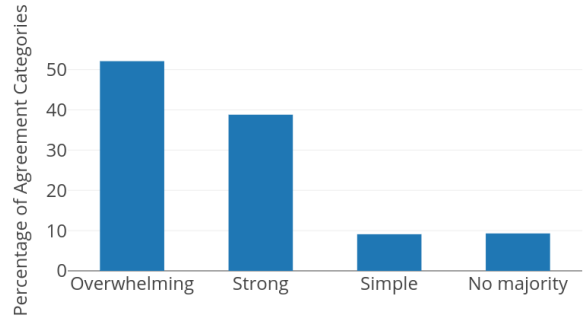


Figure 8: Distributions of categorized agreement for the large-scale annotated dataset

contribute 81.3% of all annotations.

Results

The distributions of the judgments (Figure 7) are smoother than earlier rounds, however still very similar. The two inappropriate labels cover a little less than half of all the judgments. Abusive is again very popular (30.3%) compared to hateful (13.5%), however normal is still the most popular label (41.3%) and spam is again not as frequent (almost 15%).

Regarding the three categories of agreement scores (Overwhelming, Strong and Simple) shown in Figure 8, we observe that more than half of the tweets (~52%) achieve an overwhelming agreement, which means that at least 4 out of 5 annotators agreed on the label. This percentage is slightly lower than the validation round, but we also have fewer “Normal” tweets which are the most agreed upon. Nevertheless, the remaining ~38.8% of tweets still reach an agreement of more than 3 out of 5 votes and only very few (~9.2%) achieve majority with only two annotators.

Finally, we compare the two sampled categories (Random sample vs Boosted sample) in Figure 9. We clearly see that, as expected, the boosted sample is by far richer in Abusive and Hateful content than the random sample. In fact, there are almost as many Abusive tweets in boosted, as Normal in random (~35% in both cases), and a very low total of ~4% Abusive and Hateful tweets in the random sample. This means that our decision to use a boosted sample in our methodology proved crucial: the amount of inappropriate tweets would have been too low to produce any important results, without the boosted sample.




Figure 9: Label distributions for Random vs Boosted-Sampled categories

Conclusion

In this work, we provided a methodology for annotating a large-scale dataset of inappropriate speech and the resulting labeled dataset. This annotation focused on various facets of abusive or hateful language in Twitter. We selected these two types, out of several inappropriate speech categories, based on an empirical analysis of the relationships between the corresponding labels. More specifically, we selected the most popularly used types of inappropriate speech in literature, and conduct a series of annotation rounds to understand how crowdworkers use these labels.

We analyzed these annotations in terms of correlations and similarities between the labels, and calculated their co-occurrences. After statistical analysis of these similarities between labels, we merged some of them and eliminate oth-

ers, to conclude to the most representative set. In this case, it was Abusive - and eliminated some less important ones such as Cyberbullying. When we obtained the final structure of the annotation task, we annotated the large-scale dataset.

With this present work, we make available 1) our followed methodology, 2) our code used for the custom-built platform for annotation, and 3) our final annotated dataset. We hope these three items will be useful to other researchers performing such annotations, or building machine learning models on top of such datasets.

Acknowledgements

The authors acknowledge research funding from the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No 691025.

References

- [Badjatiya et al. 2017] Badjatiya, P.; Gupta, S.; Gupta, M.; and Varma, V. 2017. Deep learning for hate speech detection in tweets. In *Proceedings of the 26th International Conference on World Wide Web Companion*, 759–760. International World Wide Web Conferences Steering Committee.
- [Chatzakou et al. 2017] Chatzakou, D.; Kourtellis, N.; Blackburn, J.; De Cristofaro, E.; Stringhini, G.; and Vakali, A. 2017. Mean birds: Detecting aggression and bullying on twitter. In *9th ACM WebScience*.
- [Chen et al. 2012] Chen, Y.; Zhou, Y.; Zhu, S.; and Xu, H. 2012. Detecting offensive language in social media to protect adolescent online safety. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Conference on Social Computing (SocialCom)*, 71–80. IEEE.
- [Davidson et al. 2017] Davidson, T.; Warmlesley, D.; Macy, M.; and Weber, I. 2017. Automated hate speech detection and the problem of offensive language. In *AAAI ICWSM*.
- [Dhingra and Mittal 2015] Dhingra, A., and Mittal, S. 2015. Content based spam classification in twitter using multi-layer perceptron learning. *Int. J. Latest Trends Eng. Technol* 5(4).
- [Dinakar, Reichart, and Lieberman 2011] Dinakar, K.; Reichart, R.; and Lieberman, H. 2011. Modeling the detection of textual cyberbullying. *The Social Mobile Web* 11(02).
- [Djuric et al. 2015] Djuric, N.; Zhou, J.; Morris, R.; Grbovic, M.; Radosavljevic, V.; and Bhamidipati, N. 2015. Hate speech detection with comment embeddings. In *Proceedings of the 24th International Conference on World Wide Web*, 29–30. ACM.
- [Gambäck and Sikdar 2017] Gambäck, B., and Sikdar, U. K. 2017. Using convolutional neural networks to classify hate-speech. In *Proceedings of the First Workshop on Abusive Language Online*, 85–90.
- [Golbeck et al. 2017] Golbeck, J.; Ashktorab, Z.; Banjo, R. O.; Berlinger, A.; Bhagwan, S.; Buntain, C.; Cheakalos, P.; Geller, A. A.; Gergory, Q.; Gnanasekaran, R. K.; Gunasekaran, R. R.; Hoffman, K. M.; Hottle, J.; Jienjiltert, V.; Khare, S.; Lau, R.; Martindale, M. J.; Naik, S.; Nixon, H. L.; Ramachandran, P.; Rogers, K. M.; Rogers, L.; Sarin, M. S.; Shahane, G.; Thanki, J.; Vengataraman, P.; Wan, Z.; and Wu, D. M. 2017. A large labeled corpus for online harassment research. In *9th ACM Web Science*, 229–233.
- [Hine et al. 2017] Hine, G. E.; Onalapo, J.; De Cristofaro, E.; Kourtellis, N.; Leontiadis, I.; Samaras, R.; Stringhini, G.; and Blackburn, J. 2017. Kek, cucks, and god emperor trump: A measurement study of 4chan’s politically incorrect forum and its effects on the web. In *ICWSM*.
- [Hosseinmardi et al. 2015] Hosseinmardi, H.; Mattson, S. A.; Rafiq, R. I.; Han, R.; Lv, Q.; and Mishra, S. 2015. Analyzing labeled cyberbullying incidents on the instagram social network. In *International Conference on Social Informatics*, 49–66. Springer.
- [Jha and Mamidi 2017] Jha, A., and Mamidi, R. 2017. When does a compliment become sexist? analysis and classification of ambivalent sexism using twitter data. In *Proceedings of the Second Workshop on NLP and Computational Social Science*, 7–16.
- [Kansara and Shekokar 2015] Kansara, K. B., and Shekokar, N. M. 2015. A framework for cyberbullying detection in social network. *International Journal of Current Engineering and Technology* 5.
- [Kayes et al. 2015] Kayes, I.; Kourtellis, N.; Quercia, D.; Iamnitich, A.; and Bonchi, F. 2015. The Social World of Content Abusers in Community Question Answering. In *WWW*.
- [Malmasi and Zampieri 2017] Malmasi, S., and Zampieri, M. 2017. Detecting hate speech in social media. *arXiv preprint arXiv:1712.06427*.
- [Nobata et al. 2016] Nobata, C.; Tetreault, J.; Thomas, A.; Mehdad, Y.; and Chang, Y. 2016. Abusive language detection in online user content. In *Proceedings of the 25th International Conference on World Wide Web*, 145–153. International World Wide Web Conferences Steering Committee.
- [Olteanu, Talamadupula, and Varshney 2017] Olteanu, A.; Talamadupula, K.; and Varshney, K. R. 2017. The limits of abstract evaluation metrics: The case of hate speech detection. In *Proceedings of the 2017 ACM on Web Science Conference*, 405–406. ACM.
- [Papegnies et al. 2017] Papegnies, E.; Labatut, V.; Dufour, R.; and Linarès, G. 2017. Detection of abusive messages in an on-line community. In *Conférence en Recherche d’Information et Applications*.
- [Park and Fung 2017] Park, J. H., and Fung, P. 2017. One-step and two-step classification for abusive language detection on twitter. *arXiv preprint arXiv:1706.01206*.
- [Razavi et al. 2010] Razavi, A.; Inkpen, D.; Uritsky, S.; and Matwin, S. 2010. Offensive language detection using multi-level classification. *Advances in Artificial Intelligence* 16–27.
- [Riadi and others 2017] Riadi, I., et al. 2017. Detection of cyberbullying on social media using data mining techniques. *International Journal of Computer Science and Information Security* 15(3):244.
- [Ribeiro et al. 2017] Ribeiro, M. H.; Calais, P. H.; Santos, Y. A.; Almeida, V. A.; and Meira Jr, W. 2017. ” like sheep among wolves”: Characterizing hateful users on twitter. *arXiv preprint arXiv:1801.00317*.
- [Santos et al. 2014] Santos, I.; Minambres-Marcos, I.; Laorden, C.; Galán-García, P.; Santamaría-Ibirika, A.; and Bringas, P. G. 2014. Twitter content-based spam filtering. In *International Joint Conference SOCO13-CISIS13-ICEUTE13*, 449–458. Springer.
- [Schmidt and Wiegand 2017] Schmidt, A., and Wiegand, M. 2017. A survey on hate speech detection using natural language processing. In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media. Association for Computational Linguistics, Valencia, Spain*, 1–10.
- [Serra et al. 2017] Serra, J.; Leontiadis, I.; Spathis, D.; Blackburn, J.; Stringhini, G.; and Vakali, A. 2017. Class-based prediction errors to detect hate speech with out-of-vocabulary words. In *Abusive Language Workshop*.
- [Twitter 2017] Twitter. 2017. A calendar of our safety work. <https://goo.gl/sX2Apb>.

- [Wang et al. 2015] Wang, B.; Zubiaga, A.; Liakata, M.; and Procter, R. 2015. Making the most of tweet-inherent features for social spam detection on twitter. *arXiv preprint arXiv:1503.07405*.
- [Wang 2010] Wang, A. H. 2010. Don't follow me: Spam detection in twitter. In *Security and Cryptography (SECRYPT), Proceedings of the 2010 International Conference on*, 1–10. IEEE.
- [Warner and Hirschberg 2012] Warner, W., and Hirschberg, J. 2012. Detecting hate speech on the world wide web. In *Proceedings of the Second Workshop on Language in Social Media*, 19–26. Association for Computational Linguistics.
- [Waseem and Hovy 2016] Waseem, Z., and Hovy, D. 2016. Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *SRW@ HLT-NAACL*, 88–93.
- [Ye, You, and Jr. 2017] Ye, T.; You, S.; and Jr., L. P. R. 2017. When Does More Money Work? Examining the Role of Perceived Fairness in Pay on the Performance Quality of Crowdworkers. In *ICWSM*.
- [Zhou and Sun] Zhou, Z., and Sun, L. Network-based spam filter on twitter. <http://snap.stanford.edu/class/cs224w-2014/projects2014/cs224w-81-final.pdf>.