

HateBR: A Large Expert Annotated Corpus of Brazilian Instagram Comments for Offensive Language and Hate Speech Detection

Francielle Vargas*†, Isabelle Carvalho*, Fabiana Góes*
Thiago A.S. Pardo*, Fabrício Benevenuto†

*Institute of Mathematical and Computer Sciences, University of São Paulo, Brazil

†Computer Science Department, Federal University of Minas Gerais, Brazil

{francielleavargas,isabelle.carvalho,fabianagoes}@usp.br, taspardo@icmc.usp.br, fabricio@dcc.ufmg.br

Abstract

Due to the severity of the social media offensive and hateful comments in Brazil, and the lack of research in Portuguese, this paper provides the first large-scale expert annotated corpus of Brazilian Instagram comments for hate speech and offensive language detection. The HateBR corpus was collected from the comment section of Brazilian politicians' accounts on Instagram and manually annotated by specialists, reaching a high inter-annotator agreement. The corpus consists of 7,000 documents annotated according to three different layers: a binary classification (offensive versus non-offensive comments), offensiveness-level classification (highly, moderately, and slightly offensive), and nine hate speech groups (xenophobia, racism, homophobia, sexism, religious intolerance, partyism, apology for the dictatorship, antisemitism, and fatphobia). We also implemented baseline experiments for offensive language and hate speech detection and compared them with a literature baseline. Results show that the baseline experiments on our corpus outperform the current state-of-the-art for the Portuguese language.

Keywords: hate speech and offensive language detection, corpus annotation, natural language processing

1. Introduction

Offensive language and hate speech detection has attracted interest from different institutions and has become an important research topic (Poletto et al., 2021; Pitenis et al., 2020; Zannettou et al., 2020; Çöltekin, 2020; Guest et al., 2021). While this challenging endeavour is undoubtedly a relevant research line, also has its implications for society concerning race, gender, religion, and origin. Furthermore, automated methods for hateful and offensive comments detection may bolster web security in revealing individuals harboring malicious intentions towards specific groups (Gao et al., 2017).

In Brazil, hate speech is prohibited, nevertheless the regulation is not effective due to the difficulty of identifying, quantifying and classifying this kind of online content. The data on hate crimes in Brazil is very worrisome: during the 2018 year's election period, denunciations with xenophobia content had an increase of 2,369%; apology and public incitement to violence and crimes against life, 630%; neo-nazism, 548%; homophobia, 350%; racism, 218%; and religious intolerance, 145%¹.

The state-of-the-art has focused on different tasks, such as automatically detecting hate speech groups, for example, racism (Hasanuzzaman et al., 2017), antisemitism (Zannettou et al., 2020), religious intolerance (Ghosh Chowdhury et al., 2019), misogyny and sexism (Guest et al., 2021; Jha and Mamidi, 2017), and cyberbullying (Safi Samghabadi et al., 2020); filtering pages

with hate and violence (Liu and Forss, 2015); offensive language detection (Zampieri et al., 2019; Steimel et al., 2019); and toxicity (Leite et al., 2020; Guimarães et al., 2020). Schmidt and Wiegand (2017) present a comprehensive survey of Natural Language Processing (NLP) techniques applied to hate speech detection, and Poletto et al. (2021) describe resources and benchmark corpora for hate speech detection. Multilingual hate speech detection is studied by ranasinghe-zampieri-2020-multilingual,steimeletal2019investigating,basileetal-2019-semeval.

Due to the relevance of this topic and the severity of online hate speech context in Brazil, the proposition of a reliable annotated corpus is fundamental to carry out experiments and to build automatic systems for offensive language and hate speech detection. Nevertheless, the annotation process of offensive content is intrinsically challenging, bearing in mind that what is considered offensive is influenced by pragmatic (contextual) factors, and people may have different perspectives on an offense. On account of that, Poletto et al. (2021) claim that authors in the field have discussed aspects related to the implications of an annotation process for offensive language and hate speech phenomena, which inspired a multi-layer annotation scheme (Zampieri et al., 2019), target-aware annotation (Basile et al., 2019), and the implicit-explicit distinction in the annotation (Caselli et al., 2020). Corroborating those authors, we claim that, as being particularly challenging the offensive language and hate speech detection, a well-defined annotation schema has a considerable impact among the consistency and quality of the data, and the performance of the derived machine learning classifiers.

¹<https://www.bbc.com/portuguese/brasil-46146756>

In this paper, we provide the first large-scale expert annotated corpus of Brazilian Instagram comments for hate speech and offensive language detection in Brazilian Portuguese. The HateBR corpus was collected from different accounts of Brazilian politicians from Instagram social media. The political context was chosen due to the identification of a wide variety of serious offensive and hateful attacks against different groups. The entire annotation schema was proposed and annotated by different specialists: a linguist, a hate speech expert, NLP and machine learning researchers, and handled by accurate guidelines and training steps, in order to ensure the same understanding of the tasks, and towards minimizing bias. Furthermore, baseline experiments were implemented, whose results (85% of F1-score) overcame the current state-of-the-art for the Portuguese language. More precisely, the main contributions of this paper are:

- The first large-scale expert annotated corpus for offensive language and hate speech on the web and social media in Brazilian Portuguese. The corpus titled “HateBR” consists of 7,000 Instagram comments annotated in three different layers (offensive versus non-offensive; offensive comments sorted into offensiveness-levels such as highly, moderately, and slightly; and nine hate speech groups: xenophobia, racism, homophobia, sexism, religious intolerance, partyism, an apology to dictatorship, antisemitism, and fatphobia).
- A new expert annotation schema for hate speech and offensive language detection, which is divided into three layers: offensive classification, offensiveness-classification and hate speech classification.

In what follows, we briefly introduce the main related work. Section 3 describes the HateBR corpus development, as well as the proposed annotation schema and its evaluation. In Sections 4 and 5, HateBR corpus statistics and experiments are presented. At last, final remarks are discussed in Section 6.

2. Related Work

Most of hate speech and offensive language corpora are proposed for the English language (Zampieri et al., 2019; Fersini et al., 2018; Davidson et al., 2017; Gao and Huang, 2017; Jha and Mamidi, 2017; Golbeck et al., 2017). For the French language, a corpus of Facebook and Twitter annotated data for Islamophobia, sexism, homophobia, religion intolerance and disability detection was also proposed (Chung et al., 2019; Ousidhoum et al., 2019). For the German language, a new anti-foreigner prejudice corpus was proposed. This corpus is composed of 5,836 Facebook posts and hierarchically annotated with slightly and explicitly/substantially offensive language according to six targets: foreigners, government, press, community,

other, and unknown (Bretschneider and Peters, 2017). For the Greek language, an annotated corpus of Twitter and Gazeta posts for offensive content detection is also available (Pitenis et al., 2020; Pavlopoulos et al., 2017). For the Slovene and Croatian languages, a large-scale corpus composed of 17,000,000 posts, composed of 2% of abusive language on a leading media company website was built (Ljubešić et al., 2018). For the Arabic language, there is a corpus of 6,136 twitter posts, which is annotated according to religion intolerance subcategories (Albadi et al., 2018). For the Indonesian language, a hate speech annotated corpus from Twitter data was also proposed (Alfina et al., 2017).

For the Portuguese language, a corpus composed of 5,668 tweets in European and Brazilian Portuguese, and automated methods using a hierarchy of hate to identify social groups of discrimination was proposed by Fortuna et al. (2019). They used pre-trained Glove word embeddings with 300 dimensions for feature extraction and a LSTM architecture proposed in Badjatiya et al. (2017). The authors obtained 78% of F1-score using cross-validation. Furthermore, Fortuna et al. (2021) built a new specialized lexicon specifically for European Portuguese, which, according to the authors, may be useful to detect a broader spectrum of content referring to minorities. Also, for Brazilian Portuguese, a corpus composed of 1,250 comments collected from G1 Brazilian online newspaper ² was proposed by de Pelle and Moreira (2017). The authors report the annotation of a binary class: offensive and non-offensive comments and seven hate groups (racism, sexism, homophobia, xenophobia, religious intolerance, and cursing). The authors evaluated a set of features based in n-grams and Information Gain (InfoGain) algorithm (Witten et al., 2016). Classical machine learning methods such as Support Vector Machine (SVM) (Scholkopf and Smola, 2001) with linear kernel, and Multinomial Naive Bayes (NB) (Eyheramendy et al., 2003) were applied. The best model obtained 80% of F1-Score.

3. HateBR Corpus Development

In this section, we describe in detail the process for the building, annotation, and evaluation of the proposed corpus.

3.1. Approach Overview

The entire process for the corpus construction occurred for approximately six months, between August 2020 to January 2021. This project was performed by different specialists (e.g., a linguist, a hate speech expert, and NLP and machine learning researchers) and led by the linguist and hate speech expert in order to ensure the reliability and quality of the annotated data. Figure 1 exhibits an overview of the proposed approach for the construction of the HateBR corpus.

²<https://g1.globo.com/>




Figure 1: The proposed approach for the HateBR corpus construction.

As shown by Figure 1, in the first step - domain definition - the political domain was selected. In the second step - criteria for selection of accounts - the following criteria were defined: six different public accounts, being three liberal-party and three conservative-party accounts, from four women and two men. In the third step - data extraction - we implemented an Instagram API using the following parameters: post id, maximum number of 500 comments per post, and only public accounts were selected. Then, we extracted five hundred comments for each post published across six months from the second half of 2019 year were collected. For instance, five hundred comments were collected from the same account on an Instagram post published in August 2019. In the same settings, other five hundred comments were collected from a second post published in September 2019, and so on. In total, thirty posts were selected from six predefined Instagram accounts. Subsequently to the data extraction step, we proposed an approach for data cleaning. The data cleaning step basically consists of removing noise, such as links, characters without semantic value, and also comments that presented only emoticons, laughs (kkk, hahah, hshshs), or mentions (e.g., @namesomeone) without any textual content. Hashtags and emotions were kept. After these steps, we obtained the initial version of HateBR corpus without labels.

For the corpus annotation, we defined a set of criteria for selection of annotators, such as higher levels of education (e.g., Ph.D. and Ph.D. candidate); only specialists (e.g., linguists, hate speech experts and computer scientists); and diverse profiles, such as distinct political orientations and colors in order to minimize bias. Subsequently, we began the annotation process and proposed a new annotation schema, determining more precisely offensive language and hate speech classification. After all the previous steps are completed, the corpus was annotated using different levels of classification. The first level consists of a binary classification in offensive language versus non-offensive language; each of the 7,000 Instagram comments was annotated with an offensive (3,500 comments) or a non-offensive (3,500 comments) label. The second layer consists of offensiveness-level classification (highly, moderately, and weakly). Each of the 3,500 comments classified as offensive in the first layer was classified in offensiveness levels: highly offensive (778 comments), moderately offensive (1,044 comments), and

slightly offensive (1,678 comments). Lastly, in the third layer, offensive comments that incited violence or hate against groups, based on specific characteristics (e.g., physical appearance, religion) received a label of hate speech (727 comments) considering nine identified hate speech groups (xenophobia, racism, homophobia, sexism, religious intolerance, partyism, apology for the dictatorship, antisemitism, and fatphobia). Still in the third layer, offensive comments that did not present violence or hate against groups received a label of no hate speech (2,773 comments). Finally, we evaluated the proposed annotation process using annotation agreement metrics, as Kappa (McHugh, 2012; Sim and Wright, 2005) and Fleiss (Fleiss, 1971), which obtained a high inter-annotator agreement for offensive language classification (75% Kappa and 74% Fleiss), and moderate for offensiveness-level classification (47% Kappa and 46% Fleiss).

3.2. Data Collection

Brazil occupies the third position in the worldwide ranking of Instagram’s audience with 110 million active Brazilian users with an audience of 93 million users³. Taking into consideration that Instagram is a powerful platform for mass media, we automatically collected Instagram comments to build our corpus. Tables 1 and 2 show the data collection statistics.

Table 1: Data collection statistics.

Data	Total
Amount of extracted comments	15,000
Amount of removed comments	8,000
Final corpus	7,000

Table 2: Accounts and posts information.

Profile	Total	Description
Gender	6 accounts	4 women and 2 men
Political	6 accounts	3 liberals and 3 conservative
Posts	30 posts	500 comments per post

As shown in Tables 1 and 2, corroborating our proposal of balancing the variables, such as gender and political party, we collected fifteen thousand comments from six public Instagram accounts of Brazilian politicians divided in three politicians from the liberal party, and three politicians from the conservative party, being four women and two men. We decided to select the most popular posts for each account during the second half of 2019, being five posts for each account and five hundred comments for each post. Thereafter, we removed eight thousand comments that presented only emoticons, laughs or mentions. In addition, labeled comments that were surplus aiming at balancing the classes of binary classification were also removed. Therefore, in these eight thousand removed items, there are both noises and surplus labeled comments.

³<https://www.statista.com/>

3.3. Annotation Process

A detailed description of our annotation process approach is presented in this Section.

3.3.1. Selection of Annotators

The first step of the annotation process consists of selection of annotators. Due to the degree of complexity of the offensive language and hate speech detection tasks, mainly because it involves a highly politicized domain, we decided to select only specialists at higher levels of education. In addition, towards minimizing bias and their negative impact on the results, we diversify the annotators’ profile, as shown in Table 3.

Table 3: Annotators’ profile.

Profile	Description
Education	PhD or PhD candidate
Gender	feminine
Political	liberal and conservative
Color	black and white
Brazilian region	north and southeast

As shown in Table 3, the annotators are from North and Southeast Brazilian regions, and have at least a PhD study. Furthermore, they are white and black women, and are aligned with liberal or conservative parties.

3.3.2. Annotation Schema

Matter of ongoing debate, offensive language and hate speech detection tackles a conceptual difficulty distinguishing hateful and offensive expressions from expressions that merely denote dislike or disagreement (Post, 2009). In spite of the enormous difficulty of these tasks, this paper provides a new annotation schema for automatic detection of offensive language and hate speech in Brazilian Portuguese, as shown in Figure 2. In our annotation schema, we accurately discriminate each one of these definitions - offensive language and hate speech - which will be described in the following paragraphs.

According to Zampieri et al. (2019), offensive posts include insults, threats, and messages containing any form of untargeted profanity. Accordingly, in this paper we assume that offensive language consists of a kind of language containing terms or expressions with any pejorative connotation, including swear words⁴, which may be explicit or implicit. Furthermore, as defined by Fortuna and Nunes (2018), in this paper we assume that hate speech is a kind of language that attacks or diminishes, that incites violence or hate against groups, based on specific characteristics such as physical appearance, religion, or others, and it may occur with different linguistic styles, even in subtle forms or when humor is used. Therefore, hate speech is a type of of-

⁴Swear words express the speaker’s emotional state tied to impoliteness and rudeness speech. They are a type of opinion that is highly confrontational, rude, or aggressive (Jay and Janschewitz, 2008; Culpeper et al., 2017)

fensive language used against groups target of discrimination (e.g., sexism, racism, homophobia).

Table 4 shows examples of offensive language and hate speech, which may be explicit or implicit, extracted from the HateBR corpus. Note that **bold** indicates terms or expressions with explicit pejorative connotation, and underline indicates “clues” of terms or expressions with implicit pejorative connotation. We also describe the terms originally written in Portuguese and their translation to English.

Table 4: Examples of comments classified as offensive language and hate speech extracted from the HateBR corpus.

Type	Instagram Comments	Translation
Offensive Language	Essa besta humana é o câncer do País, tem q voltar para jaula, <u>urgentemente!</u> E viva o Presidente Bolsonaro.	This human beast is the cancer of the country, it has to <u>go back to the cage</u> , <u>urgently!</u> And long live President Bolsonaro.
Non-Offensive Language	Quem falou isso pra vc deputada? O sergio moro ta aprovado pela maioria dos brasileiros.	Who said that to you deputy? Sergio Moro is approved by the majority of Brazilians.
Hate Speech	Vagabunda. Comunista. Mentirosa. O povo chileno nao merece uma desgraça desta	Bitch. Communist. Liar. The people from Chile do not deserve such a disgrace .
No-Hate Speech	Pois é, deveria <u>devolver o dinheiro</u> aos cofres públicos do Brasil. Canalha.	It should <u>refund money</u> to the public Brazilian coffers. Jerk.

As it is shown in Table 4, there are explicit and implicit terms or expressions with pejorative connotation in offensive and hate speech comments. For example, in the comments classified as *offensive language* and *no-hate speech*, although the term “câncer” (cancer) may be found in non-pejorative contexts of use (e.g., he has cancer), in this comment context, it is used with a pejorative connotation. In contrast, the expression “besta humana” (human beast) and the term “canalha” (jerk) also present pejorative connotations even though they would be mainly found in pejorative contexts. Moving forward, it should be noted that both offensive and hate speech comments include implicit terms or expressions. For example, the expressions “voltar a jaula” (go back to the cage) and “devolver o dinheiro” (refund money) are clues that indicate the implicit pejorative terms “criminal” (criminoso) and “assaltante” (thief), respectively. Furthermore, *hate speech* comments consists of attacks against groups (e.g., sexism and partyism); and *non-offensive* comments do not present any terms or expressions with pejorative connotation.

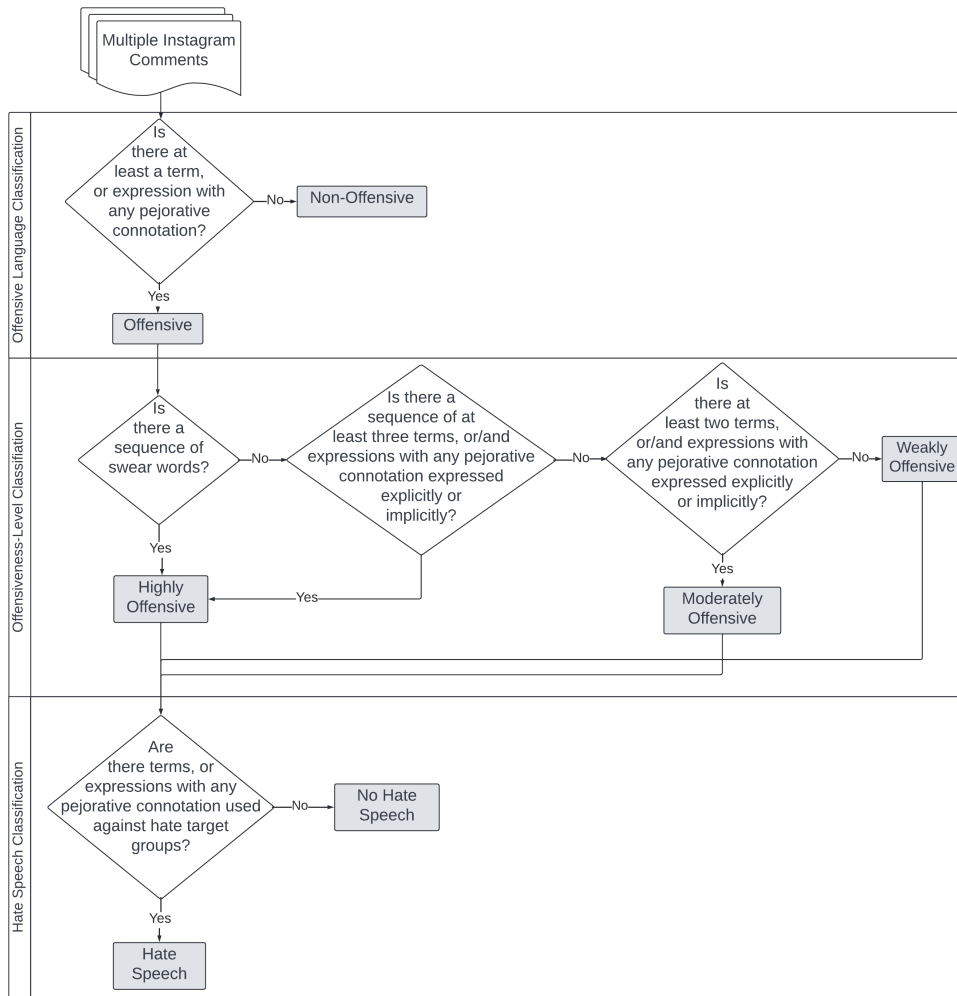


Figure 2: HateBR annotation schema.

Corroborating these offensive language and hate speech definitions, and taking advantage of our initial premise that a well-defined and suitable annotation schema is a great determining factor to improve the machine learning classifier’s performance, we introduce in this paper a new annotation schema for hate speech and offensive language detection on social media. The proposed annotation schema is shown in Figure 2. Note that our annotation schema is divided into three layers. In the first layer, we annotated the corpus using a binary classification (offensive or non-offensive comments). Subsequently, we selected only offensive comments obtained from the previous annotation layer, and classified them into offensiveness levels. The offensiveness-level classification consists of three classes: highly, moderately, and slightly. The third layer provides annotation of offensive comments with hate speech (one of the nine hate groups that we already introduced), and offensive comments without hate speech. We further describe in detail how the classification is performed in each annotation layer to following.

- **Offensive language classification:** We initially assume that comments that present at least one term or expression with any pejorative connotation should be classified as offensive, and comments that have no terms or expressions with any pejorative connotation should be classified as non-offensive comments.
- **Offensiveness-level classification:** In this paper, we introduce a fine-grained offensive annotation, which we called offensiveness-level classification. In this layer of annotation, comments classified as offensive were also annotated according to three offensiveness levels: highly, moderately, and slightly. We assume that offensive comments that present a sequence of swear words should immediately be classified as highly offensive. In the same setting, offensive comments containing a sequence of at least three terms or/and expressions with any pejorative connotation, which may be explicit or implicit, should also be classified as highly offensive. Moving forward, comments that do not meet these last two criteria, and present

at least two terms or expressions with any pejorative connotation, which may be explicit or implicit, should be classified as moderately offensive. Lastly, offensive comments that do not meet the previous criteria should be classified as slightly offensive.

- **Hate speech classification:** We assume that offensive comments targeted against groups based on specific characteristics (e.g., physical appearance, religion, etc.) should be classified as hate speech. On the other hand, offensive comments not targeted against groups should not be classified as hate speech. The annotation of hate speech comments was accomplished according to nine hate speech groups (partyism, sexism, religion intolerance, apology for the dictatorship, fatphobia, homophobia, racism, antisemitism, and xenophobia).

Therefore, the annotators followed three main steps. In the first step, they classified each of the collected Instagram comments in offensive or non-offensive comments. In the second step, for the offensiveness-level classification, each one of 3,500 comments labeled as offensive in the previous step received one of the three following labels: highly, moderately and slightly offensive. Finally, in the third step, offensive comments were classified by each annotator into nine hate speech groups.

3.3.3. Annotation Evaluation

As already mentioned, our corpus was annotated by three different specialists. Each comment was annotated by each one to guarantee the reliability of the process. Besides that, the linguist and hate speech expert served as judges when a tie occurred. We also computed inter-annotator agreement using two different evaluation metrics: Cohen’s kappa (McHugh, 2012; Sim and Wright, 2005) and Fleiss’ kappa (Fleiss, 1971) - which we describe in detail below. A further evaluation of hate speech groups was also carried out. Firstly, the offensive comments annotated with any hate speech groups by at least two annotators were immediately validated. Then, offensive comments annotated with hate speech groups by only one annotator were submitted to a new checking step, in which the linguist decided whether that label should be validated or discarded.

Cohen’s kappa

This measure is described by equation 1, where po is the relative agreement observed between raters and pe is the hypothetical probability of random agreement. It shows the degree of agreement between two or more judges beyond what would be expected by chance (McHugh, 2012; Sim and Wright, 2005). Kappa values range from 0 to 1, and there are possible interpretations of these values (Landis and Koch, 1977). Each stratum represents the final value of the Kappa score and

the level of agreement among annotators. Note that a value from 0.0 to 0.20 is a slight agreement, from 0.21 to 0.40 is fair, from 0.41 to 0.60 is moderate, from 0.61 to 0.80 is substantial, and above 0.80 refers to almost perfect agreement.

$$k = \frac{po - pe}{1 - pe} \quad (1)$$

Considering the tasks presented in this paper, we would agree that NLP highly subjective tasks encompass considerable negative impact on the inter-annotation agreement results. In other words, the more subjective the task is, the more difficult it is to obtain good inter-annotation agreement score. However, based on the obtained results shown in Table 5, our annotation process presents substantial results according to strength of agreement from Cohen’s kappa. Accordingly, high inter-annotator agreement for offensive language classification (75%), and moderate inter-annotator agreement for offensiveness-level classification (47%) were reached. It should be noted that, although the moderate performance obtained for the offensiveness-level classification would be a further investigation issue, we must point out that this task is highly subjective and ambiguous, consequently presenting a wide range of challenges such as high disagreement. Note that “AB”, “BC”, and “CA” consists of obtained score agreement between two different human annotators.

Table 5: Cohen’s kappa.

Peer Agreement	AB	BC	CA	AVG
Offensive language	0.76	0.72	0.76	0.75
Offensiveness-level	0.46	0.44	0.50	0.47

Fleiss’ kappa

The Fleiss evaluation measure (Fleiss, 1971) is an extension of Cohen’s kappa for cases where there are more than two annotators (or methods). That being said, Fleiss’ kappa is applied when there is wide range of annotators that provide categorical ratings, such as binary or nominal scale, for a fixed number of items. The interpretation for the values of Fleiss’ kappa also follows the values proposed by Cohen’s kappa. In this paper, we also evaluated our annotation process using the Fleiss metric, as shown in Table 6.

Table 6: Fleiss’ kappa.

Fleiss’ kappa	ABC
Offensive language	0.74
Offensiveness-level	0.46

As it is shown in Table 6, a high inter-annotator agreement was reached for offensive language classification (74%), and a moderate inter-annotator agreement score (46%) was obtained for offensiveness-levels classification. Once again, the fine-grained offensive classification is a ambitious and challenging task due to high disagreement among annotators.

4. HateBR Corpus Statistics

As a result of this paper, we present statistics of the HateBR corpus. As shown in Tables 7 and 8, the corpus is composed of 7,000 document-level annotations. Firstly, the corpus was annotated into a binary class. Each of the 7,000 comments received a offensive label (3,500 comments) or non-offensive (3,500 comments) label. Additionally, the 3,500 comments identified as offensive were also classified according to offensiveness-level, being 1,678 slightly offensive, 1,044 moderately offensive, and 778 highly offensive. As shown in Tables 9 and 10, offensive comments were also categorized according to the nine hate speech groups (partyism, sexism, religion intolerance, apology for the dictatorship, fatphobia, homophobia, racism, antisemitism, and xenophobia). Further, over the Instagram subjects of posts, in which the comments were extracted, 70% are related to government issues, 6,6% fake news, 6,6% sexism, 6,6% racism, 6,6 % environment, and 3,3% economy.

Table 7: Offensive language.

Labels	Total
Non-offensive	3,500
Offensive	3,500
Total	7,000

Table 8: Offensiveness-level.

Labels	Total
Slightly offensive	1,678
Moderately offensive	1,044
Highly offensive	778
Total	3,500

Table 9: Hate speech groups.

Labels	Total
Partyism	496
Sexism	97
Religion Intolerance	47
Apology for the Dictatorship	32
Fatphobia	27
Homophobia	17
Racism	8
Antisemitism	2
Xenophobia	1
Total	727

Table 10: Post subjects.

Subjects	Total
Political-government	21
Political-fake news	2
Political-sexism	2
Political-racism	2
Political-environment	2
Political-economy	1
Total	30

5. Experiments

Towards investigation and validation related to suitability of the proposed expert annotated corpus for online offensive language and hate speech detection, we implemented baseline experiments using two different representations and four machine learning methods. The representations implemented were *n-grams*, more specifically the *unigram language model*, and the bag-of-ngrams with tf-idf⁵ preprocessing. The machine learning methods applied were Naive Bayes (NB) (Eyheramendy et al., 2003), Support Vector Machine (SVM) with linear kernel (Scholkopf and Smola, 2001), Multilayer Perceptron (MLP) with only one hidden layer (Haykin, 2009), and Logistic Regression (LR) (Ayyadevara, 2018). In our experiments, we used the Python 3.6, scikit-learn and pandas libraries, and sliced our data in 80% train, 10% test, and 10% validation. Results are shown in Table 11.

As shown in Table 11, we evaluated two different tasks: offensive language and hate speech detection. For the offensive language detection task, we implemented both baseline representations (unigram and tf-idf) over the 7,500 comments from HateBR, being 3,500 offensive labels and 3,500 non-offensive labels. As result, a high performance was achieved. The best model for this task obtained 85% of F1-score. In the same setting, for the hate speech detection task, we also implemented both baseline representation (unigram and tf-idf) over the 3,500 offensive comments from the HateBR corpus, being 727 hate speech labels and 2,773 non hate speech labels. In this experiment, we applied a class balancing technique called undersampling (Witten et al., 2016), aiming at unbalanced classes of hate speech. In particular, we adopted this method due to the fact that it makes overfitting unlikely. As result, a high performance was also obtained for hate speech detection task. The best model obtained 78% of F1-score. Furthermore, although the main focus of this paper is to provide a well-structured, suitable and expert annotation schema for offensive language and hate speech detection, these baseline experiments were presented to corroborate our initial premise that a well-defined and structured annotation schema leads to a good classification performance for highly complex and subjective tasks.

Despite the fact that dataset comparison is a challenging task in NLP, we propose a comparison among annotated corpora for the Portuguese language with our HateBR expert annotated corpus. Additionally, a comparison for the European and Brazilian Portuguese is also presented. Tables 12 and 13 show the results.

Note that the proposed corpus is the first large-scale manually annotated corpus for Portuguese, which consists of 7,000 Instagram comments annotated with three different classes. Besides that, as shown in Table 12, corpora proposed by literature for offensive lan-

⁵Term Frequency(TF) — Inverse Dense Frequency(IDF)

Table 11: NB, SVM, MLP and LR Evaluation.

Tasks	Features set	Class	Precision				Recall				F1-Score			
			NB	SVM	MLP	LR	NB	SVM	MLP	LR	NB	SVM	MLP	LR
Offensive Language Detection	unigram	0	0.72	0.82	0.83	0.83	0.89	0.79	0.87	0.87	0.79	0.81	0.85	0.85
		1	0.84	0.78	0.85	0.85	0.62	0.81	0.81	0.81	0.71	0.79	0.83	0.83
		Avg	0.78	0.80	0.84	0.84	0.75	0.80	0.84	0.84	0.75	0.80	0.84	0.84
	tf-idf	0	0.75	0.85	0.85	0.85	0.85	0.85	0.85	0.85	0.80	0.85	0.85	0.85
		1	0.81	0.84	0.83	0.83	0.69	0.84	0.84	0.84	0.75	0.84	0.84	0.84
		Avg	0.78	0.85	0.84	0.84	0.77	0.85	0.84	0.84	0.77	0.85	0.84	0.84
Hate Speech Detection	unigram	0	0.71	0.61	0.69	0.69	0.76	0.79	0.89	0.89	0.73	0.69	0.77	0.77
		1	0.80	0.79	0.89	0.89	0.76	0.61	0.68	0.68	0.78	0.69	0.77	0.77
		Avg	0.76	0.70	0.79	0.79	0.76	0.70	0.79	0.79	0.76	0.69	0.77	0.77
	tf-idf	0	0.74	0.64	0.69	0.69	0.77	0.82	0.85	0.85	0.76	0.75	0.76	0.76
		1	0.82	0.84	0.86	0.86	0.78	0.71	0.70	0.70	0.80	0.77	0.77	0.77
		Avg	0.78	0.76	0.77	0.77	0.78	0.77	0.78	0.78	0.78	0.76	0.77	0.77

Table 12: Hate speech and offensive language detection in Portuguese: datasets.

Authors	Total	Type	Classes	Hate Groups	Agreement	Balanced
Fortuna et al. (2019)	5,668	tweets	hate x no-hate	sexism, body, origin, homophobia, racism, ideology, religion, health, other-lifestyle	72%	no
de Pelle and Moreira (2017)	1,250	website comments	offensive x non-offensive	racism, sexism, homophobia, xenophobia, religious intolerance, and cursing	71%	no
HateBR corpus	7,000	instagram	(offensive x non-offensive); (offensiveness: slightly x moderately x highly); (hate x no-hate)	xenophobia, racism, homophobia, sexism, religious intolerance, partyism, apology to the dictatorship, antisemitism, and fatphobia	75%	yes

Table 13: Hate speech and offensive language detection in Portuguese: models and methods.

Authors	Set of Features	Learning Method	F1-Score
Fortuna et al. (2019)	Embeddings	LSTM	0.78
de Pelle and Moreira (2017)	N-grams, InfoGain	SVM, NB	0.80
HateBR corpus	N-grams, Tf-idf	NB, SVM, MLP, LR	0.85

guage and hate speech detection in Portuguese present a considerably smaller size compared to our corpus. The inter-human agreement score obtained in HateBR also overcame the other proposals. The annotation process proposed by de Pelle and Moreira (2017) was carried out over a corpus composed of only 1,250 comments, annotated with binary classification (offensive and non-offensive). Our corpus also presents balanced classes for offensive language classification (3,500 offensive comments x 3,500 non-offensive comments).

Last but certainly not least, as shown in Table 13, the results obtained with baseline experiments on our corpus clearly overcame the current ML models proposed by literature for the Portuguese language. Notice that Fortuna et al. (2019) proposed a sophisticated set of features and ML methods, which obtained an inferior performance compared to our baseline experiments. de Pelle and Moreira (2017) also implemented models for offensive comments detection, and, even though their models have been trained over an unrepresentative corpus, the baseline experiments performed on our corpus still overcame their performance.

6. Final Remarks

This paper provides the first large-scale expert annotated corpus of Brazilian Portuguese Instagram comments for offensive language and hate speech detection. The HateBR corpus was annotated by different specialists, and consists of 7,000 documents annotated with three different layers. The first layer consists of 3,500 comments annotated as offensive and 3,500 comments annotated as non-offensive. In the second layer, offensive comments were annotated according to offensiveness-level: slightly, moderately, and highly. In the third layer, offensive comments were also annotated considering nine hate speech groups. We evaluated the proposed annotation schema and a high human annotation agreement was obtained. Finally, baseline experiments were implemented, obtaining relevant results (85% of F1-score) that overcame the current literature baselines for the Portuguese language.

Acknowledgements

The authors are grateful to CNPq, FAPEMIG, and FAPESP for partially funding this project.

7. Bibliographical References

- Albadi, N., Kurdi, M., and Mishra, S. (2018). Are they our brothers? Analysis and detection of religious hate speech in the Arabic twittersphere. In *Proceedings of the 10th International Conference on Advances in Social Networks Analysis and Mining*, pages 69–76, Barcelona, Spain.
- Alfina, I., Mulia, R., Fanany, M. I., and Ekanata, Y. (2017). Hate speech detection in the Indonesian language: A dataset and preliminary study. In *Proceedings of the 9th International Conference on Advanced Computer Science and Information*, pages 233–238, Bali, Indonesia.
- Ayyadevara, V. K. (2018). Logistic regression. In *Pro Machine Learning Algorithms*, pages 49–69. Springer.
- 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*, page 759–760, Canton of Geneva, Switzerland.
- Basile, V., Bosco, C., Fersini, E., Nozza, D., Patti, V., Rangel Pardo, F. M., Rosso, P., and Sanguinetti, M. (2019). SemEval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 54–63, Minneapolis, USA.
- Bretschneider, U. and Peters, R. (2017). Detecting offensive statements towards foreigners in social media. In *Proceedings of the 50th Hawaii International Conference on System Sciences*, pages 2213–2222, Hawaii, USA.
- Caselli, T., Basile, V., Mitrović, J., Kartoziya, I., and Granitzer, M. (2020). I feel offended, don't be abusive! implicit/explicit messages in offensive and abusive language. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6193–6202, Marseille, France.
- Chung, Y.-L., Kuzmenko, E., Tekiroglu, S. S., and Guerini, M. (2019). CONAN—Counter Narratives through Nichesourcing: A multilingual dataset of responses to fight online hate speech. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2819–2829, Florence, Italy.
- Çöltekin, Ç. (2020). A corpus of Turkish offensive language on social media. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6174–6184, Marseille, France.
- Culpeper, J., Iganski, P., and Sweiry, A. (2017). Linguistic impoliteness and religiously aggravated hate crime in England and Wales. *Journal of Language Aggression and Conflict*, 5(1):1–29.
- Davidson, T., Warmesley, D., Macy, M. W., and Weber, I. (2017). Automated hate speech detection and the problem of offensive language. In *Proceedings of the 11th International Conference on Web and Social Media*, pages 512–515, Quebec, Canada.
- de Pelle, R. and Moreira, V. (2017). Offensive comments in the Brazilian web: A dataset and baseline results. In *Anais do VI Brazilian Workshop on Social Network Analysis and Mining*, pages 510–519, Rio Grande do Sul, Brazil.
- Eyheramendy, S., Lewis, D. D., and Madigan, D. (2003). On the naive bayes model for text categorization. In *Proceedings of the 9th International Workshop on Artificial Intelligence and Statistics*, pages 93–100, Florida, USA.
- Fersini, E., Rosso, P., and Anzovino, M. (2018). Overview of the task on automatic misogyny identification. In *Proceedings of the 3rd Workshop on Evaluation of Human Language Technologies for Iberian Languages*, pages 214–228, Sevilla, Spain.
- Fleiss, J. L. (1971). Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Fortuna, P. and Nunes, S. (2018). A survey on automatic detection of hate speech in text. *ACM Computing Surveys*, 51(4):1–30.
- Fortuna, P., Rocha da Silva, J., Soler-Company, J., Wanner, L., and Nunes, S. (2019). A hierarchically-labeled Portuguese hate speech dataset. In *Proceedings of the 3rd Workshop on Abusive Language Online*, pages 94–104, Florence, Italy.
- Fortuna, P., Cortez, V., Sozinho Ramalho, M., and Pérez-Mayos, L. (2021). MIN_PT: An European Portuguese lexicon for minorities related terms. In *Proceedings of the 5th Workshop on Online Abuse and Harms*, pages 76–80, Held Online.
- Gao, L. and Huang, R. (2017). Detecting online hate speech using context aware models. In *Proceedings of the International Conference Recent Advances in Natural Language Processing*, pages 260–266, Varna, Bulgaria.
- Gao, L., Kuppessmith, A., and Huang, R. (2017). Recognizing explicit and implicit hate speech using a weakly supervised two-path bootstrapping approach. In *Proceedings of the 8th International Joint Conference on Natural Language Processing*, pages 774–782, Taipei, Taiwan.
- Ghosh Chowdhury, A., Didolkar, A., Sawhney, R., and Shah, R. R. (2019). ARHNet - leveraging community interaction for detection of religious hate speech in Arabic. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 273–280, Florence, Italy.
- 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., Jienjittlert, 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 *Proceedings of the 9th ACM Web Science Conference*, pages 229–233, New York, USA.
- Guest, E., Vidgen, B., Mittos, A., Sastry, N., Tyson, G., and Margetts, H. (2021). An expert annotated dataset for the detection of online misogyny. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1336–1350, Held Online.
- Guimarães, S. S., Reis, J. C. S., Ribeiro, F. N., and Benvenuto, F. (2020). Characterizing toxicity on facebook comments in Brazil. In *Proceedings of the 26th Brazilian Symposium on Multimedia and the Web*, pages 253–260, Maranhão, Brazil.
- Hasanuzzaman, M., Dias, G., and Way, A. (2017). Demographic word embeddings for racism detection on Twitter. In *Proceedings of the 8th International Joint Conference on Natural Language Processing*, pages 926–936, Taipei, Taiwan.
- Haykin, S. (2009). *Neural networks and learning machines*. Pearson Upper Saddle River, 3 edition.
- Jay, T. and Janschewitz, K. (2008). The pragmatics of swearing. *Journal of Politeness Research - language Behaviour Culture*, 4(2):267–288.
- 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 2nd Workshop on NLP and Computational Social Science*, pages 7–16, Vancouver, Canada.
- Landis, J. R. and Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, pages 159–174.
- Leite, J. A., Silva, D., Bontcheva, K., and Scarton, C. (2020). Toxic language detection in social media for Brazilian Portuguese: New dataset and multilingual analysis. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 914–924, Suzhou, China.
- Liu, S. and Forss, T. (2015). Text classification models for web content filtering and online safety. In *Proceedings of the 15th IEEE International Conference on Data Mining Workshop*, pages 961–968, New Jersey, USA.
- Ljubešić, N., Erjavec, T., and Fišer, D. (2018). Datasets of Slovene and Croatian moderated news comments. In *Proceedings of the 2nd Workshop on Abusive Language Online*, pages 124–131, Brussels, Belgium.
- McHugh, M. L. (2012). Interrater reliability: the kappa statistic. *Biochemia medica*, 22(3):276–282.
- Ousidhoum, N., Lin, Z., Zhang, H., Song, Y., and Yeung, D.-Y. (2019). Multilingual and multi-aspect hate speech analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 4675–4684, Hong Kong, China.
- Pavlopoulos, J., Malakasiotis, P., and Androutsopoulos, I. (2017). Deep learning for user comment moderation. In *Proceedings of the 1st Workshop on Abusive Language Online*, pages 25–35, British Columbia, Canada.
- Pitenis, Z., Zampieri, M., and Ranasinghe, T. (2020). Offensive language identification in Greek. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5113–5119, Marseille, France.
- Poletto, F., Basile, V., Sanguinetti, M., Bosco, C., and Patti, V. (2021). Resources and benchmark corpora for hate speech detection: a systematic review. *Language Resources and Evaluation*, 55(3):477–523.
- Post, R. (2009). Hate speech. In *Extreme Speech and Democracy*, pages 123–138. Oxford Scholarship Online.
- Safi Samghabadi, N., López Monroy, A. P., and Solorio, T. (2020). Detecting early signs of cyberbullying in social media. In *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*, pages 144–149, Marseille, France.
- Schmidt, A. and Wiegand, M. (2017). A survey on hate speech detection using natural language processing. In *Proceedings of the 5th International Workshop on Natural Language Processing for Social Media*, pages 1–10, Valencia, Spain.
- Scholkopf, B. and Smola, A. J. (2001). *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, Cambridge.
- Sim, J. and Wright, C. C. (2005). The kappa statistic in reliability studies: Use, interpretation, and sample size requirements. *Physical therapy*, 85(3):257–268.
- Steimel, K., Dakota, D., Chen, Y., and Kübler, S. (2019). Investigating multilingual abusive language detection: A cautionary tale. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*, pages 1151–1160, Varna, Bulgaria.
- Witten, I. H., Frank, E., Hall, M. A., and Pal, C. J. (2016). *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann.
- Zampieri, M., Malmasi, S., Nakov, P., Rosenthal, S., Farra, N., and Kumar, R. (2019). Predicting the type and target of offensive posts in social media. In *Proceedings of the 17th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1415–1420, Minnesota, USA.
- Zannettou, S., Finkelstein, J., Bradlyn, B., and Blackburn, J. (2020). A quantitative approach to understanding online antisemitism. In *Proceedings of the 14th International AAAI Conference on Web and Social Media*, pages 786–797, Georgia, USA.