



Hate versus politics: detection of hate against policy makers in Italian tweets

Armend Duzha¹ · Cristiano Casadei¹ · Michael Tosi¹ · Fabio Celli¹ 

Received: 23 December 2020 / Accepted: 21 July 2021 / Published online: 20 August 2021
© The Author(s), under exclusive licence to Springer Nature Switzerland AG 2021

Abstract

Accurate detection of hate speech against politicians, policy making and political ideas is crucial to maintain democracy and free speech. Unfortunately, the amount of labelled data necessary for training models to detect hate speech are limited and domain-dependent. In this paper, we address the issue of classification of hate speech against policy makers from Twitter in Italian, producing the first resource of this type in this language. We collected and annotated 1264 tweets, examined the cases of disagreements between annotators, and performed in-domain and cross-domain hate speech classifications with different features and algorithms. We achieved a performance of ROC AUC 0.83 and analyzed the most predictive attributes, also finding the different language features in the anti-policymakers and anti-immigration domains. Finally, we visualized networks of hashtags to capture the topics used in hateful and normal tweets.

Keywords Hate speech · Natural language processing · Social media · Policy making

Introduction and background

The rise of Natural Language Processing (NLP) tasks focused on hate speech Badjatiya et al. (2017) and the analysis of online debates Celli et al. (2014) have both highlight bad behaviors in social media, such as offensive language against vulnerable groups (e.g., immigrants, minorities, etc.) Poletto et al. (2017), as well as aggressive language against women Saha et al. (2018). An under-researched - yet important - area of investigation is anti-policy hate: the hate speech against politicians, policy making and laws at any level (national, regional and local). While anti-policy hate speech has been addressed in Arabic Guellil et al. (2020), most European languages have been under-researched.

✉ Fabio Celli
fabio.celli@maggioli.it

¹ Maggioli S.p.A, Via Bornaccino 101, 47822 Santarcangelo di Romagna, Italy

In recent years, scientific research contributed to the automatic detection of hate speech from text with datasets annotated with hate labels, aggressiveness, offensiveness, and other related dimensions Sanguinetti et al. (2018). Scholars have presented systems for the detection of hate speech in social media focused on specific targets, such as immigrants Del Vigna et al. (2017), and language domains, such as racism Kwok and Wang (2013), misogyny Frenda et al. (2019) or cyberbullying Menini et al. (2019). Each type of hate speech has its own vocabulary and its own dynamics, thus the selection of a specific domain is crucial to obtain clean data and to restrict the scope of experiments and learning tasks.

We have formulated three Research Questions:

- RQ1: How different are hate speech domains, such as anti-immigrants and anti-policy?
- RQ2: Is it possible to perform cross-domain training to exploit techniques and models trained in one domain (i.e. anti-immigration) to detect hate speech in another domain (i.e. against policy-makers)?
- RQ3: Is it possible to identify and track the topics of public debate involved/not involved in hate speech?

In order to address RQ1, we performed correlation and classification analysis. The former was carried out to measure how different language features are related to hate speech in different domains, the latter to test the performance of classifiers in different domains. To address RQ2, we performed cross-domain classification and applied hate speech models trained in an anti-immigration domain to a policy-making domain. Finally, to address RQ3, we extracted the hashtags from tweets labelled as hateful and non-hateful, visualized the network of co-occurrences with a Hifan Hu graph Yifan and Shi (2015).

With this research, we aim to provide actionable insights for evidence-based decision-making Kyriazis et al. (2020), as online hate is often a predictor of offline crime Williams et al. (2020). We selected Twitter as the source of data and Italian as the target language for two reasons:

- (1) There are datasets annotated with anti-immigrant hate speech labels in Italian, but no datasets annotated with anti policy making hate speech labels,
- (2) Italy has, at least since the elections in 2018, a large audience that pays attention to hyper-partisan sources on Twitter that are prone to produce and retweet messages of hate against policy making Giglietto et al. (2019).

This paper contributes to the scientific research in NLP and hate speech detection in two ways. First: the production of a new corpus, annotated with hate speech labels, in an under-resourced language (Italian). Second: the classification of hate speech tweets against policy making, and its comparison to the classification of hate speech against immigrants.

The paper is structured as follows: after a literature review ([‘Related work’](#)), we collect a stream of tweets in Italian using keywords (i.e., hashtags) related to

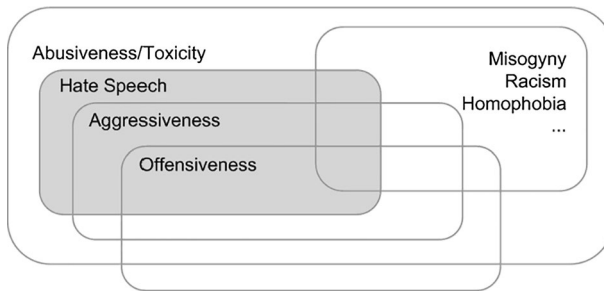


Fig. 1 Relation between hate speech and related concepts. *Source* Poletto et al. (2020)

laws and regulations ([‘Data collection and annotation’](#)). We then train, test, and evaluate models for hate speech from existing resources, analyze the predictive power of each feature, visualize the results ([‘Experiments and discussion’](#)), and draw conclusions ([‘Conclusion and future work’](#)).

Related work

Hate speech is defined as any expression *that is abusive, insulting, intimidating, harassing, and/or incites, supports and facilitates violence, hatred, or discrimination. It is directed against people (individuals or groups) on the basis of their race, ethnic origin, religion, gender, age, physical condition, disability, sexual orientation, political conviction, and so forth* Karmen and Melita (2012). A recent study defined the relationships between hate speech and related concepts (see Fig. 1), highlighting the fact that involved phenomena make hate speech especially hard to model, with the risk of creating data that is biased and making the models prone to overfitting. In addition to this, literature also reports cases of annotators’ insensitivity to differences in dialects and offenses Sap et al. (2019) that make annotation difficult. For these reasons, one of the largest challenges in the field of hate speech is to investigate architectures that are explainable, stable and well-performing across different languages and domains Poletto et al. (2020).

Another key issue is that many recent approaches based on word embeddings Kenneth (2017), Deep Learning algorithms and BERT Pre-trained transformers Jacob et al. (2018) Tenney et al. (2019) Polignano et al. (2019), are vulnerable to undesirable bias in training data, especially in the political domain Wich et al. (2020), and suffer from poor interpretability MacAvaney et al. (2019). In other words, it can be difficult to understand how the systems based on Deep Learning techniques make their decisions about hateful/non-hateful messages. Moreover, the decisions taken by systems might be based on biased and unfair models. A method for explaining the decisions of transformer models is to look at the attention vectors Clark et al. (2019). Yet, studies show that learned attention weights are frequently uncorrelated with gradient-based measures of feature importance, thus, different attention distributions can nonetheless yield similar predictions Jain and Wallace

Table 1 State-of-the-art in hate speech classification

Algorithm	Features	F1
SVM, RNN	PolSbjLex, 2 Wemb	0.799
SVM+RNN+LR meta	ptWemb badWLex	0.793
SVM	EmoLex, Wemb, social	0.783
RNN	ptWemb	0.663
SVM+RNN+RF meta	ptWemb	0.649

The reported systems use as algorithms: *SVM* support vector machines, *RNN* recurrent neural networks, *LR* linear regression, ensemble learning (meta), *RF* random forests. As features: word-embedding extracted from training set (Wemb) or pre-trained (ptWemb), Polarity/Subjectivity (PolSbjLex), bad words (badWLex), emotion lexica (EmoLex) and social media metadata (social)

(2019). In a context of policy making, the transparency of the decisions and the possibility to interpret the results should be considered a priority.

Despite there being many studies about hate speech in Natural Language Processing (NLP) against various targets, such as anti-immigrants, there are few works in the field of hate speech detection against politicians and policy making. Previous approaches to this task exploited transparent Machine Learning (ML) algorithms, such as Gaussian Naïve Bayes, Random Forests or Support Vector Machines (SVM), as well as Deep Learning algorithms, such as convolutional neural network (CNN), Multi-Layer Perceptrons (MLP), Recurrent Neural Networks (RNN) with long-short-term memory (LSTM) or bi-directional long-short-term memory (Bi-LSTM) on top of word embeddings extracted from the training set or pre-trained from other resources with transfer learning. These studies show that good results can be obtained with Bi-LSTM, MLP and SVM Guellil et al. (2020).

Studies that provided useful datasets in the field of hate speech include SemEval 2019 who studies multilingual hate speech against immigrants and women in English and Spanish Basile et al. (2019). In Italian there are two main corpora, both about anti-immigrant hate: the Italian HS corpus Poletto et al. (2017) and HaSpeeDe-tw2018, the dataset released during the EVALITA campaign in 2018 Sanguinetti et al. (2020b). The former is a collection of more than 5,700 tweets manually annotated with hate speech, aggressiveness, irony and other forms of potentially harassing communication. The latter, is a dataset (3000 tweets for training and 1000 for testing) manually annotated with hate speech labels. The results of HaSpeeDe-tw2018, reported in Table 1, are the state-of-the-art in hate speech detection in Italian and show that lexical resources, such as polarity and emotion lexica, are useful to this task Bosco et al. (2018), Fersini et al. (2018).

Most hate speech recognition systems at HaSpeeDe-tw2018 exploit SVM, Recurrent Neural Networks with LSTM or ensemble learning (meta) Bai et al. (2018), Michele et al. (2018), De la Pena Sarracén et al. (2018), and word-embeddings as features Santucci et al. (2018), pre-trained or extracted from the training set. Some systems also use cross-platform data (i.e. Facebook and Twitter) and shows that this strategy yields similar results for Twitter Corazza et al. (2019). Crucially, the best performing systems make use of lexical resources for polarity, subjectivity and

emotions Cimino et al. (2018), showing that word embeddings are more effective when combined with lexical resources. The current state-of-the-art in the HaSpeeDe task from Twitter is 0.808 macro-F1, obtained using transformer-based models Sanguinetti et al. (2020a). Regarding the visualization, the heuristic power of network graphs has been known in computational social sciences for at least one decade. For example, network graphs of topics or Twitter hashtags can be used to analyze sentiment polarization of hyper-partisan topics Kiran and Weber (2017). Another example, networks of replies annotated with personality types can represent the conversational dynamics of neurotics and emotionally stable users Celli and Rossi (2013).

In the next section, we describe how we created the dataset and annotated it with hate speech labels.

Data collection and annotation

In order to monitor the reactions of society towards policy making, we retrieved a stream of tweets in Italian from March to May 2020, using snowball sampling. Starting from a set of seed hashtags, for instance: #dpcm (decree of the president of the council of ministers), #legge (law) and #leggedibilancio (budget law), we retrieved a sample of tweets and then added the new hashtags contained in this sample to extend the list of seed hashtags and retrieve new tweets. We called this dataset Policycorpus. We removed duplicates, retweets and tweets containing only hashtags and urls. In total we obtain a set of 1264 tweets (1000 for training and 264 for testing). The amount of hate labels in the Policycorpus is 11% (1124 normal and 140 hate tweets). It is strongly unbalanced, like in the it-HS corpus (17% of hate tweets), because it reflects the raw distribution of hate tweets in Twitter. The HaSpeeDe-tw corpus (32% of hate tweets) instead has a distribution that oversample hate tweets. At the end of the sampling process, the list of seeds included about 60 hashtags referring to

- Laws, such as #decretorilancio (#relaunchdecree), #leggelettorale (#elector-allaw), #decretosicurezza (#securitydecree)
- Politicians and policy makers, such as #Salvini, #decretoSalvini (#Salvinidecree), #Renzi, #Meloni, #DraghiPremier
- Political parties, such as #lega (#league), #pd (#Democratic Party)
- Political tv shows, such as #ottoemezzo, #nonelarena, #noneladurso, #Piaz-zapulita
- Topics of the public debate, such as #COVID, #precari (#precariousworkers), #sicurezza (#security), #giustizia (#justice), #ItalExit
- Hyper-partisan slogans, such as #vergognaConte (#shameonConte), #conted-imentiti (#ConteResign) or #noicontrosalvini (#WeareagainstSalvini)

This is the first corpus in Italian annotated with hate speech against policy makers. We plan to make this resource available under request¹.

To produce gold standard labels, we asked two Italian experts of communication, to manually label the tweets in the Polycyrcorpus, distinguishing between hate and normal tweets according to the following guidelines: By definition, hate speech is *any expression that is abusive, insulting, intimidating, harassing, and/or incites to violence, hatred, or discrimination. It is directed against people on the basis of their race, ethnic origin, religion, gender, age, physical condition, disability, sexual orientation, political conviction, and so forth.* Translated Examples:

- (1) “a clear #NO to #Netherlands that we would like users of the #MES economic resources but in exchange for Italy’s renunciation of its budgetary autonomy. To Netherlands we say: thank you and goodbye, WE ARE NOT INTERESTED !!” is normal because it does not contain hate, insults, intimidation, violence or discrimination.
- (2) “... There is a weekly catwalk of the #jackal #no #notAtAll! Listening to a Po #clown after a true PATRIOT a doctor from #Bergamo cannot be held, seen or heard. Giletti should stop inviting certain SLACKERS FROM THE PO VALLEY! #COVID-19 #NonelArena” contains hate speech, including insults like #clown and #jackal.
- (3) “I have my say ... #Draghi is a great economist but we don’t need a #Monti-style economist ... We don’t need another technical #government to obey the banking lobby! We need a political leader! We need a #ItalExit! We need the #Lira! #No to #DraghiPremier” is a normal case, despite the strong negative sentiment. It might be controversial for the presence of the term *lobby*, often used in abusive contexts, but in this case, it is not directed against people on the basis of their race, ethnic origin, religion, gender, age, physical condition, disability, sexual orientation or political conviction.

The Inter-Annotator Agreement is $k = 0.53$. Although the score is not high, it is in line with the score reported in the literature for hate speech against immigrants ($k = 0.54$)Poletto et al. (2017) and indicates that the detection of hate speech is a hard task for humans. All the examples of disagreement were discussed and an agreement was reached between the annotators. The cases of disagreements occurred more often when the sentiment of the tweet was negative, this was mainly due to:

The use of vulgar expressions not explicitly directed against specific people but generically against political choices.

The negative interpretation of hyper-partisan hashtags, such as #contedimetiti (#ConteResign) or #noicontrosalvini (#WeareagainstSalvini), in tweets without explicit insults or abusive language.

¹ For dataset requests please contact the corresponding author.

The substitution of explicit insults with derogatory words, such as the word “circus” instead of “clowns”. In the next section, we report and discuss the results of the experiments.

Experiments and discussion

Our goal is to create models of hate speech that automatically predict hateful tweets against policy makers in the Polycorpus. First, we describe the features extracted from text, then we perform in-domain and cross-domain classification, and finally, we conduct feature analysis and visualize the hashtag networks. As discussed in ‘[Related work](#)’, we aim to develop explainable Artificial Intelligence (AI) models, hence we also exploited ML algorithms based on lexical resource (Lex), such as SVM, Adaboost and Random Forests, in addition to more advanced techniques, for instance, neural networks based on the ALBERTo pretrained transformer model. We ran two different experiments:

- In experiment one, we tried to answer to RQ2, using different algorithms to train models on the existing corpora. We then perform a cross-domain classification, evaluating the predictions trained on HaSpeeDe-tw and it-HS to the Polycorpus test set (‘[In-domain and cross-domain classification](#)’);
- In experiment two, we tried to answer to RQ1, with a feature analysis to understand which features are best predictors of hate speech in the policy making domain with respect to anti-immigration domain (‘[Feature analysis](#)’);

Finally, to answer RQ3, we visualized the networks of hashtags in order to understand the relationships between topics used in normal and hateful tweets (‘[Hashtags network analysis](#)’). Before all, we described the features extracted from text.

Feature extraction

Building upon the previous work presented in the literature, we adopted linguistic resources for the extraction of features to use with ML algorithms. In particular, we used:

- LIWC Tausczik and Pennebaker (2010), a linguistic resource available in many languages, including Italian Alparone et al. (2004), that maps words to 68 psycholinguistics dimensions, such as linguistic dimensions (i.e. pronouns, articles, tense), psychological processes (i.e cognitive mechanisms, sensations, certainty, causation) human processes (i.e. sex, social life, family), personal concerns (i.e. leisure, money, religion, death) and spoken categories (i.e. assent, nonfluencies)
- NRC Mohammad et al. (2013), a linguistic resource that maps words to 10 emotion and polarity features, for instance positive words, negative words, anger, anticipation, fear, sadness, joy, surprise, trust and disgust.

Table 2 Results of the classification of hate speech in Italian on the Italian HS corpus (it-HS), HaSpeeDe-tw2018 (HaSpeeDe-tw) and Polycorpus (PC) with different algorithms, lexical features (Lex) and transformer embeddings (AlBERTo)

Dataset-algorithm	ROC	<i>F1</i>
it-HS baseline	0.5	0.7
it-HS Lex+random forest	0.69	0.78
it-HS Lex+support vector machines	0.51	0.76
it-HS Lex+adaboost bayesian net	0.67	0.79
it-HS AlBERTo+neural networks	0.85	0.76
HaSpeeDe-tw baseline	0.5	0.4
HaSpeeDe-tw Lex+random forest	0.73	0.71
HaSpeeDe-tw Lex+support vector machines	0.63	0.72
HaSpeeDe-tw Lex+adaboost bayes net	0.73	0.72
HaSpeeDe-tw AlBERTo+neural networks	0.88	0.82
PC baseline	0.5	0.8
PC Lex+random forest	0.69	0.87
PC Lex+Support Vector Machines	0.52	0.88
PC Lex+adaboost decision tree	0.57	0.88
PC AlBERTo+Neural Networks	0.83	0.89

The best performances are marked in bold

- Other 22 language-independent stylistic features Celli (2015), including positive/negative emoticons/emojis, ratio of punctuation, question and expression marks, numbers, operators, links, hashtags, mentions or emails addresses, parentheses, lowercase/uppercase and ratio of repeated bigrams.

These dictionaries extract a matrix of 100 features, less sparse than bag-of-words. In addition to this, we used a transformer model trained on Italian tweets: AlBERTo, that extracts a dense matrix of more than 700 embedding features Polignano et al. (2019).

In-domain and cross-domain classification

Hate speech labels are naturally unbalanced, as normal tweets are - fortunately - the large majority, especially in the Polycorpus and it-HS corpus. As this is a natural condition, we chose to keep the labels unbalanced and measure the performances with two metrics: ROC AUC curve, which is insensitive to class imbalance, and weighted-average F-measure that takes into account the difference of performance for the two classes. In this experiment, we trained and tested various algorithms, we used a training-test split as evaluation settings, which is 88–12% in the it-HS corpus, 75–25% in HaSpeeDe-tw2018 and 80–20% in the Polycorpus (Table 2).

A closer look at the per-class performance obtained with the best algorithm (AlBERTo + neural networks), reveals that in general the algorithm has a higher performance in the detection of normal tweets and lower performance in the recognition of hate tweets, which have a poor recall. The fact that recall is higher

Table 3 Per-class results of the classification on each corpus with the best algorithm (AlBERTo + neural networks)

Corpus and label	Precision	Recall	F1-score
PC normal	0.92	0.98	0.95
PC hate	0.76	0.41	0.54
it-HS normal	0.71	0.92	0.81
it-HS hate	0.86	0.56	0.68
HaSpeeDe-tw18	0.86	0.89	0.88
HaSpeeDe-tw18	0.75	0.71	0.73

Table 4 Results of the cross-domain classification of hate speech in Italian on the Polycorpus-test (PC-test) with the models trained on the HaSpeeDe-tw2018 and Italian HS corpora

Dataset-algorithm	ROC	<i>F1</i>
Pc-test baseline	0.5	0.8
HaSpeeDe-tw-train to PC-test Lex+random forest	0.71	0.77
HaSpeeDe-tw-train to PC-test Lex+SVM	0.56	0.73
HaSpeeDe-tw-train to PC-test Lex+adaboost bayesian net	0.71	0.76
HaSpeeDe-tw-train to PC-test AlBERTo+Neural Networks	0.61	0.72
it-HS-train to PC-test Lex+random forest	0.81	0.88
it-HS-train to PC-test Lex+support vector machines	0.5	0.8
it-HS-train to PC-test Lex+adaboost bayesian net	0.66	0.89
it-HS-train to PC-test AlBERTo+neural networks	0.53	0.82

in the HaSpeeDe-tw corpus than in the Polycorpus suggests that balancing the number of hate examples with the normal ones has a positive effect on recall. Precision is similar in these two datasets (0.75): the it-HS corpus has a higher precision on the hate class, but the recall follows the same pattern of the other two corpora. We present these results in Table 3.

In attempt to address RQ2, we used the models trained on the HaSpeeDe-tw and it-HS corpora in the previous experiment to automatically produce predictions on the Polycorpus test set, thus performing a cross-domain backtest. Given the differences between domains we expect poor results in the next experiment, the results of which are presented in Table 4.

As expected, the results of cross-domain classification show that the domain shift had a huge impact on the performance of the classifiers, particularly from HaSpeeDe-tw to Polycorpus, where the results measured with weighted-average F1 are below the majority baseline, suggesting that the features are so different that the model cannot use them in the correct way. Surprisingly, the models trained on the it-HS corpus produced good results, but only the ones trained with ML algorithms, particularly random forests and adaboost, that are more capable of using weak features. AlBERTo and Neural Networks in this case performed only slightly better than the majority baseline. We believe that the large training size of it-HS corpus had a positive effect for the cross-domain adaptation.

Table 5 Results of the correlation ranking between different lexical features and hate speech

Lex feature	best corr. to HS
<i>HaSpeeDe-tw</i>	
Lower case (style)	+ 0.25*
Punctuation to word ratio (style)	+ 0.18
Expression marks (style)	+ 0.17*
Uppercase char ratio (style)	- 0.15**
Numbers (style)	- 0.21*
Non-ending punctuation (style)	- 0.3*
<i>it-HS</i>	
Expression marks (style)	+ 0.11**
Certainty (LIWC)	+ 0.1*
Anger (LIWC)	+ 0.1**
Prepositions (LIWC)	- 0.05*
Articles (LIWC)	- 0.05*
Ending punctuation (style)	- 0.05**
<i>Policycorpus</i>	
Swears (LIWC)	+ 0.15*
Initial uppercase (style)	+ 0.15**
Sexual (LIWC)	+ 0.13**
Upper case (style)	+ 0.12**
Positive emotions (LIWC)	- 0.05**
Numbers (style)	- 0.06*
Lower case (style)	- 0.07*

***P* value lower than 0.01, **P* value lower than 0.05

Feature analysis

The cross-domain classification highlighted the difference in the features between the corpora. To measure this difference, and answer RQ1, we computed the Pearson correlation between the lexical features and the hate speech scores. In Table 5 we present the best lexical features correlated to hate speech in each dataset. Positive correlation indicates the best features to classify hate messages and negative correlations indicate the best feature to classify normal messages. All these features were used in the classification experiments.

The analysis revealed that Stylographic features, such as the ratio of lowercase and uppercase characters, have a strong predictive power in the HaSpeeDe-tw2018 dataset, but not in the it-HS corpus, where there is more variety. LIWC features, such as sexual, anger and swear word ratios, are among the best predictors of hate speech against politicians. This experiment clearly shows that the most useful features for the detection of hate speech in the domain of anti-immigration are punctuation (the more there is punctuation, the more a message is non-hateful) and expression marks (the more exclamations, the more a message is likely to be hateful). In Policycorpus there are sexual and swear words as markers of hateful messages and lower case,

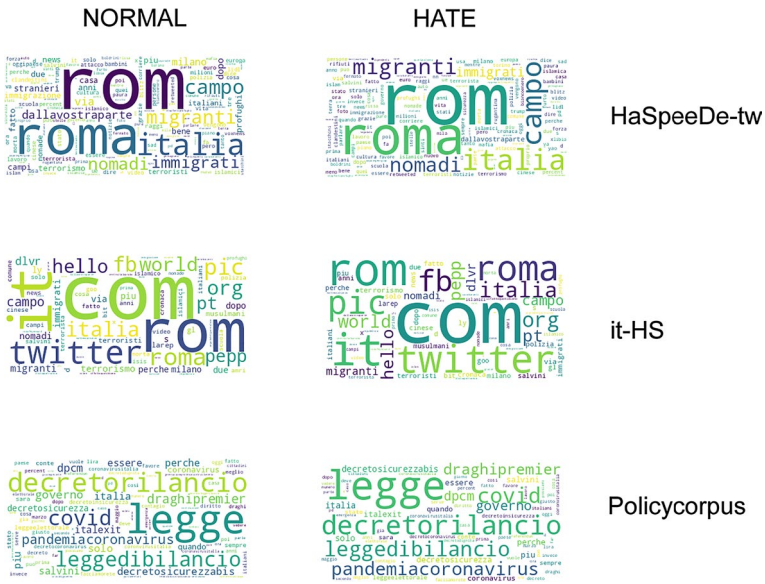


Fig. 2 Visualization of the average activations of each token in the attention vectors, associated to hate and normal labels, for each corpus

numbers and positive emotions as markers of non-hateful messages. It is interesting to note that lower case letters are correlated to hate speech in the anti-immigration domain, while in the anti-policy domain they are correlated to non-hateful messages. The similarity between the best features in it-HS and Polycyrcorus explained the good result obtained in the cross-domain classification with ML algorithms. We also exploited the attention vectors of AIBERTo to try to explain the poor performance in the cross-domain classification. Using the average activations of each token in the attention vectors, we computed the strongest predictors in the model. The results, represented as wordclouds in Fig. 2, show the most frequent tokens activated to detect hate and normal labels for each corpus. The clear difference from the tokens used in the anti-immigrant and anti-policy domains is a clue of the poor performance in cross-domain classification.

Hashtags network analysis

To address RQ3, we treated the ‘normal’ and ‘hate’ classes as nodes in a network that we plotted with Yifan Hu trees Yifan and Shi (2015). In this way we were able to visualize the network of hashtags connected to the ‘normal’ or ‘hate’ nodes in the Polycyrcorus. In other words, we were able to visualize hashtags appearing only in hate speech context, hashtags appearing only in normal contexts, and hashtags appearing in both networks. The results, depicted in Fig. 3, show a pattern with a blue cloud (above), that represents the network of hashtags in normal tweets and a red cloud (below), that represents the hashtags in hate messages. Between the two,

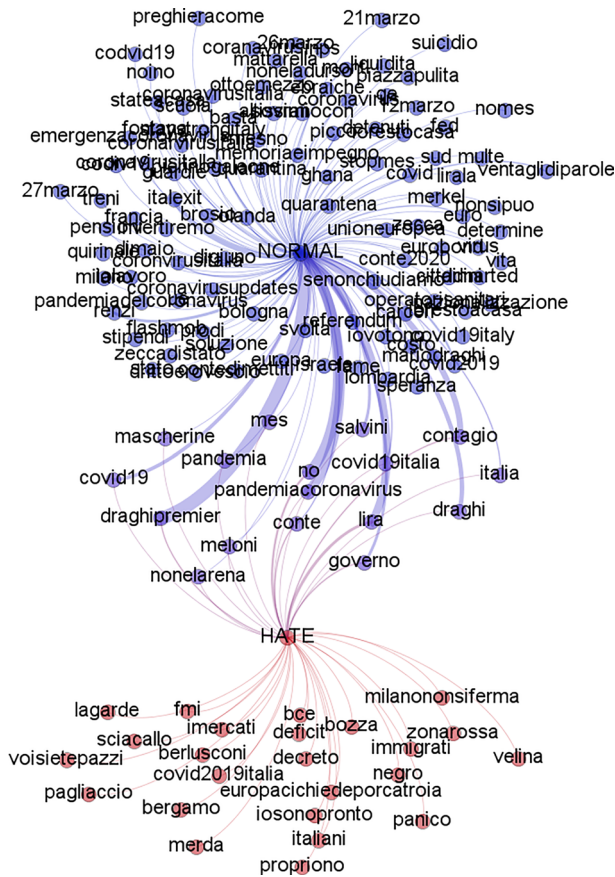


Fig. 3 Detailed visualization of the hashtag network of a small portion of the Polycorpus with Yifan Hu trees. The blue cloud (above) contains hashtags of normal messages, the red cloud (below) contains the hashtags from hate tweets and the smaller cloud between the two contains the hashtags in both messages

there is a smaller cloud of hashtags used in both contexts. A closer look to these hashtags reveals the topics of the public debate that are more controversial.

These topics include politicians (#Salvini, #Meloni, #Conte, #Draghi), economic issues (#lira, #MES), keywords related to the pandemic (#covid, #pandemia, #mascherine) and to political tv shows (#nonelarena).

Conclusion and future work

In this paper, we presented a new resource for the analysis of hate speech against policy makers on Twitter. The dataset, named Polycorpus, is the first of this type in Italian, an under-resourced language. We confirmed that the annotation of hate

speech is difficult, and detailed the cases of disagreements between annotators. Using this resource, we demonstrated that:

- Deep Learning algorithms and transformer-based models achieve state-of-the-art performances in both domains.
- Machine Learning algorithms are suitable for cross-domain classification from hate speech against immigrants to hate speech against policy makers.
- Hate speech against immigrants can be detected by looking at the style of the written text (i.e. punctuation and exclamation), while hate speech towards policy makers is based more on the vocabulary and psycholinguistic aspects (i.e. swear words).

We also visualized the spread of hate speech in Twitter against policy makers Hagen et al. (2019) and identified clusters of tweets that appear only in hate tweets and in both normal and hate tweets. We suggest that this method can be exploited to track which topics convey hatred towards policy-makers. Combining hate speech detection algorithms and visualizations, one can build a dashboard for monitoring hate speech on Twitter. The final aspect that we want to highlight is that the amount of data available, and its balance between classes, can help to improve the performance of the classifiers. In the future we plan to run experiments on domain-adaptation and collect more data for hate speech detection against policy makers.

Acknowledgements The research leading to the results presented in this paper has received funding from the PolicyCLOUD project, supported by the European Union's Horizon 2020 research and innovation programme under Grant Agreement no 870675.

Data Availability The authors plan to make Polycorpus, the dataset presented in this manuscript, available upon request, under the conditions set by the PolicyCLOUD project. For more information, please contact the corresponding author.

Declarations

Conflict of interest Three of authors of this paper declare to be employed by Maggioli S.p.A., a private Company with a financial interest in the field of Public Administration. They also declare their involvement in the PolicyCLOUD project, which has received funding from the European Union's Horizon 2020 research and innovation programme (Grant Agreement no 870675), and has financial interests in the subject and materials discussed in this manuscript.

References

- Alparone FR, Caso S, Agosti A, Rellini A (2004) The Italian liwc2001 dictionary. LIWC.net, Austin
- Badjatiya P, Gupta S, Gupta M, Varma V (2017) Deep learning for hate speech detection in tweets. In: Proceedings of the 26th international conference on world wide web companion, pp 759–760
- Bai X, Merenda F, Zaghi C, Caselli T, Nissim M (2018) Rug@ evalita 2018: hate speech detection in Italian social media. In EVALITA 2018. CEUR workshop proceedings (CEUR-WS.org)
- Basile V, Bosco C, Fersini E, Debora N, Patti V, Manuel RPF, Rosso P, Sanguinetti M, et al (2019) Semeval-2019 task 5: multilingual detection of hate speech against immigrants and women in twitter. In: 13th international workshop on semantic evaluation, pp 54–63. Assoc Comput Linguist

- Bosco C, Dell'Orletta F, Poletto F, Sanguinetti M, Tesconi M (2018) Overview of the evalita 2018 hate speech detection task. In: EVALITA 2018-sixth evaluation campaign of natural language processing and speech tools for Italian, vol 2263, pp 1–9. CEUR
- Celli F (2015) Lato: language-independent analysis tool. Technical report, University of Trento
- Celli F, Rossi L (2013) Long chains or stable communities? The role of emotional stability in twitter conversations. *Comput Intell*
- Celli F, Riccardi G, Ghosh A (2014) Corea: Italian news corpus with emotions and agreement. *Proc CLIC-it 2014*:98–102
- Cimino A, De Mattei L, Dell'Orletta F (2018) Multi-task learning in deep neural networks at evalita 2018. In: Proceedings of the 6th evaluation campaign of natural language processing and speech tools for Italian (EVALITA'18), pp 86–95
- Clark K, Khandelwal U, Levy O, Manning CD (2019) What does bert look at? An analysis of bert's attention. [arXiv:1906.04341](https://arxiv.org/abs/1906.04341)
- Corazza M, Menini S, Cabrio E, Tonelli S, Villata S (2019) Cross-platform evaluation for italian hate speech detection. In: *CLiC-it*
- De la Pena SGL, Pons RG, Enrique MCC, Rosso P (2018) Hate speech detection using attention-based lstm. In: EVALITA evaluation of NLP and speech tools for Italian, vol 12, p 235
- Del Vigna Fabio, Cimino Andrea, Dell'Orletta Felice, Petrocchi Marinella, and Tesconi Maurizio (2017) Hate me, hate me not: Hate speech detection on facebook. In *Proceedings of the First Italian Conference on Cybersecurity (ITASEC17)*, pp 86–95
- Fersini E, Nozza D, and Rosso P (2018) Overview of the evalita 2018 task on automatic misogyny identification (ami). In: Evaluation campaign of natural language processing and speech tools for Italian. Final workshop, EVALITA 2018, vol 2263. CEUR
- Frenda S, Ghanem B, Montes-y GM, Rosso P (2019) Online hate speech against women: automatic identification of misogyny and sexism on twitter. *J Intell Fuzzy Syst* 36(5):4743–4752
- Giglietto F, Righetti N, Marino G, Rossi L (2019) Multi-party media partisanship attention score. estimating partisan attention of news media sources using twitter data in the lead-up to 2018 Italian election. *Comun Polit* 20(1):85–108
- Guellil I, Adeel A, Azouaou F, Chennoufi S, Maafi H, Hamitouche T (2020) Detecting hate speech against politicians in arabic community on social media. *Int J Web Inf Syst*
- Hagen L, Keller TE, Yerden X, Luna-Reyes LF (2019) Open data visualizations and analytics as tools for policy-making. *Gov Inf Q* 36(4):101387
- Jacob D, Ming-Wei C, Kenton L, Kristina T. Bert (2018) Pre-training of deep bidirectional transformers for language understanding. [arXiv:1810.04805](https://arxiv.org/abs/1810.04805), 2018
- Jain S, Wallace BC (2019) Attention is not explanation. [arXiv:1902.10186](https://arxiv.org/abs/1902.10186)
- Karmen E, Melita PK (2012) “You don understand, this is a new war!” analysis of hate speech in news web sites comments. *Mass Commun Soc* 15(6):899–920
- Kiran GVR, Ingmar (2017) A long-term analysis of polarization on twitter. In: Eleventh international AAAI conference on web and social media
- Kenneth WC (2017) Word2vec. *Nat Lang Eng* 23(1):155–162
- Kwok I, Wang Y (2013) Locate the hate: detecting tweets against blacks. In Proceedings of the twenty-seventh AAAI conference on artificial intelligence, pp 1621–1622
- Kyriazis D, Biran O, Bouras T, Brisch K, Duzha A, del Hoyo R, Kiourtis A, Kranas P, Maglogiannis I, Manias G, et al (2020) Policycloud: analytics as a service facilitating efficient data-driven public policy management. In: IFIP international conference on artificial intelligence applications and innovations, pp 141–150. Springer
- MacAvaney S, Yao H-R, Yang E, Russell K, Goharian N, Frieder O (2019) Hate speech detection: challenges and solutions. *PLoS ONE* 14(8):e0221152
- Menini S, Moretti G, Corazza M, Cabrio E, Tonelli S, Villata S (2019) A system to monitor cyberbullying based on message classification and social network analysis. In: Proceedings of the third workshop on abusive language online, pp 105–110
- Michele C, Menini S, Pinar A, Sprugnoli R, Cabrio E, Tonelli S, Serena V (2018) Comparing different supervised approaches to hate speech detection. In: Sixth evaluation campaign of natural language processing and speech tools for Italian (Evalita 2018)
- Mohammad SM, Turney PD (2013) Nrc emotion lexicon. Technical report, National Research Council, Canada

- Poletto F, Stranisci M, Sanguinetti M, Patti V, Bosco C (2017) Hate speech annotation: Analysis of an Italian twitter corpus. In: 4th Italian conference on computational linguistics, CLiC-it 2017, vol 2006, pp 1–6. CEUR-WS
- Poletto F, Basile V, Sanguinetti M, Bosco C, Patti V (2020) Resources and benchmark corpora for hate speech detection: a systematic review. *Language Resources and Evaluation*, pp 1–47
- Polignano M, Basile P, de Gemmis M, Semeraro G, Basile V (2019) Alberto: Italian bert language understanding model for nlp challenging tasks based on tweets. In: CLiC-it
- Saha P, Mathew B, Goyal P, Mukherjee A (2018) Hateminers: detecting hate speech against women. [arXiv:1812.06700](https://arxiv.org/abs/1812.06700)
- Sanguinetti M, Poletto F, Bosco C, Patti V, Stranisci M (2018) An Italian twitter corpus of hate speech against immigrants. In: Proceedings of the eleventh international conference on language resources and evaluation (LREC 2018)
- Sanguinetti M, Comandini G, Di Nuovo E, Frenda S, Stranisci M, Bosco C, Caselli T, Patti V, Russo I (2020a) Haspeede 2@ evalita2020: overview of the evalita 2020 hate speech detection task. In: Proceedings of seventh evaluation campaign of natural language processing and speech tools for Italian. Final workshop (EVALITA 2020), Online. CEUR.org
- Sanguinetti M, Comandini G, Di Nuovo E, Frenda S, Stranisci M, Bosco C, Caselli T, Patti V, Russo I (2020b) Overview of the evalita 2020 second hate speech detection task (haspeede 2). In: Basile V, Croce D, Di Maro M, Passaro L (eds) Proceedings of the 7th evaluation campaign of natural language processing and speech tools for Italian (EVALITA 2020), Online. CEUR.org
- Santucci V, Spina S, Milani A, Biondi G, Di Bari G (2018) Detecting hate speech for italian language in social media. In: EVALITA 2018, co-located with the fifth Italian conference on computational linguistics (CLiC-it 2018), vol 2263
- Sap M, Card D, Gabriel S, Choi Y, Smith NA (2019) The risk of racial bias in hate speech detection. In: Proceedings of the 57th annual meeting of the association for computational linguistics, pp 1668–1678
- Tausczik YR, Pennebaker JW (2010) The psychological meaning of words: Liwc and computerized text analysis methods. *J Lang Soc Psychol* 29(1):24–54
- Tenney I, Das D, Pavlick E (2019) Bert rediscovers the classical nlp pipeline. [arXiv:1905.05950](https://arxiv.org/abs/1905.05950)
- Wich M, Bauer J, Groh G (2020) Impact of politically biased data on hate speech classification. In: Proceedings of the fourth workshop on online abuse and harms, pp 54–64
- Williams ML, Burnap P, Javed A, Liu H, Ozalp S (2020) Hate in the machine: anti-black and anti-muslim social media posts as predictors of offline racially and religiously aggravated crime. *Br J Criminol* 60(1):93–117
- Yifan H, Shi L (2015) Visualizing large graphs. *Wiley interdisciplinary reviews: Comput Stat* 7(2):115–136