

Article

Multimodal Hate Speech Detection in Greek Social Media

Konstantinos Perifanos^{1,†} , Dionysis Goutsos^{1,†} 

¹ National and Kapodistrian University of Athens; kperifanos@phil.uoa.gr

¹ National and Kapodistrian University of Athens; dgoutsos@phil.uoa.gr

* Correspondence: kperifanos@phil.uoa.gr

Abstract: Hateful and abusive speech presents a major challenge for all online social media platforms. Recent advances in Natural Language Processing and Natural Language Understanding allow more accurate detection of hate speech in textual streams. This study presents a multimodal approach to hate speech detection by combining Computer Vision and Natural Language processing models for abusive context detection. Our study focuses on Twitter messages and, more specifically, on hateful, xenophobic and racist speech in Greek aimed at refugees and migrants. In our approach we combine transfer learning and fine-tuning of Bidirectional Encoder Representations from Transformers (BERT) and Residual Neural Networks (Resnet). Our contribution includes the development of a new dataset for hate speech classification, consisting of tweet ids, along with the code to obtain their visual appearance, as they would have been rendered in a web browser. We have also released a pre-trained Language Model trained on Greek tweets, which has been used in our experiments. We report a consistently high level of accuracy (accuracy score=0.970, f1-score=0.947 in our best model) in racist and xenophobic speech detection.

Keywords: Multimodal Machine Learning; Deep Learning; Hate Speech Detection

1. Introduction

Abusive language and behaviour in social media platforms is a significant problem that affects online platforms, as well as the communities in which they are employed. The problem of hate speech is rather serious, as it escapes the virtual boundaries of social media platforms. Williams et. al [1] report that increase in hate speech observed in social media platforms correlates with hate crimes. It has also been noted¹ that hate speech that originates in social media can transform and disseminate into mainstream media and politic narrative. More recently, a number of alt-right related accounts have been suspended in various social media platforms, including Twitter² and Facebook³, as well as Instagram and Snapchat.

In a report published by Facebook [2], the task of fighting hate speech violations in social networks is defined as *preservation of integrity*. While social media platforms have enforced strict policies about integrity and hateful content⁴, the problem remains a difficult one as it involves several layers of complexity: the computational complexity, since the volume of the content is huge, the subtleties and the cultural aspects of each language, the problem of low resource languages, as well as the inherent ambiguity in natural language. In addition, hateful content producers are quickly adapting to platform restrictions, aiming at bypassing stricter AI and NLP systems for hateful language detection. One way of doing so is by utilising visual streams and hateful memes, as well as by using contextual visual means to propagate hateful content. This is done, for example, by using text embedded in images or screenshots, aiming at preventing

¹ <https://www.businessinsider.com/how-online-hate-speech-moves-from-the-fringes-to-the-mainstream-2018-10>

² <https://www.bbc.com/news/technology-55638558>

³ <https://www.nbcnews.com/tech/tech-news/facebook-bans-qanon-across-its-platforms-n1242339>

⁴ <https://help.twitter.com/en/rules-and-policies/hateful-conduct-policy>, https://www.facebook.com/communitystandards/recentupdates/hate_speech/

Natural Language Processing systems that operate directly on text to immediately flag the content in question as hateful.

In this paper we follow and extend the work of [3–5] on anti-refugee and anti-migrant hate speech detection. We apply hate speech detection to Greek and enrich this with a multimodal approach, in order to take into account hateful content that does not necessarily carry textual streams. Although our work focuses on racist and xenophobic speech, it can be relatively easily extended to other forms of abusive content concerning sexual orientation, political views, religion, disabilities or misogyny.

2. Related Work

2.1. Hate speech detection as a text classification problem

Hate speech is defined in Cambridge Dictionary as "public speech that expresses hate or encourages violence towards a person or group based on something such as race, religion, sex, or sexual orientation".

Research related to hate and abusive speech detection in social media has gained significant momentum over recent years [6,7], as the problem of toxicity is becoming increasingly damaging [8]. Prior to 2012 most studies focused on traditional hand-crafted features and vocabularies, in approaches similar to author profiling techniques. More recently and with the advance of representation learning, efforts have shifted to dense representations produced by techniques such as word2vec, paragraph2vec and node2vec [9–11]. These techniques are producing dense continuous representations on word or sub-word level, which are then combined to produce sentence or author representations. [12]

Davidson et. al [13] aim at predicting if a text is offensive, hate speech or neutral. They approach feature engineering with collections of unigrams, bigrams and trigrams in texts, after applying porter stemmer and lowercasing texts. They apply a series of models and report best precision, recall and f1-score of 0.91, 0.90, 0.90 respectively. Mishra et.al [5] use graph convolutional networks to attack the problem, utilising the social graph information as part of the model. Waseem and Hovy focus on data collection [3] and on linguistic features that improve quality, while Waseem [4] provides a list of criteria for the annotation process of hate speech.

Jaki and Smedt [14] provide an extensive sociolinguistic analysis of right-wing hate speech on Twitter and move on to identify patterns of non-linguistic symbols and signs to denote the ideology of users engaging in hate speech. Similar behaviour is found in our own dataset and this is an additional motivation to explore the visual modality of the rendered tweets.

Among work on other languages than English, Poletto et al. [15] discuss the process and challenges of the development of a hate speech corpus for Italian, while Pereira et. al. [16] focus on Spanish and develop HaterNet, a system for the identification and classification of hate speech in Twitter, as well as the analysis of hate trends and other negative sentiments.

Studies on hate speech in Greek can be found in Baider [17], in which covert hate speech in the form of irony in Greek and Cypriot social media is discussed. Lekea and Karampelas [18] present a methodology for automatically detecting the presence of hate speech in relation to terrorism. Pitenis et. al. [19] focus on offensive speech in Greek social media and explore various classical and Deep Learning models. Neural-based methods for representation learning focused on Twitter users, that can be used for text classification are discussed in [12].

2.2. Multimodal Learning and hate speech detection

Multimodal machine learning aims at integrating and modelling multiple communicative modalities, including linguistic, acoustic and visual messages [20]. The rise of Deep Learning and Representation Learning enables the combination of multiple modalities in a single learning framework in an intuitive and efficient way [21].

In the context of hate speech detection, Facebook introduced the Hateful Memes Challenge and dataset [22], in which two separate non-toxic channels of information can be combined to form a potentially hateful meme. The challenge is to build multimodal machine learning frameworks to correctly classify memes.

There are two major architectures to tackle multimodality, the early-fusion and the late-fusion approach. In the early-fusion approach the different modalities are combined before attempting to classify the content. Late-fusion systems, instead, first operate on each modality and then attempt to combine results and report a final decision.

In this paper we follow an early-fusion multimodal system paradigm, in which we first combine text-based representations with image-based representations into an overall representation and we then attempt classification. Essentially, two different backbone networks operating in text and image are fine-tuned together to optimise prediction accuracy.

3. Dataset

While there have been some public available datasets of labelled hateful tweets, publicly available in the form of tweet ids ([3,4,19]), as Mishra observes [5], many of these tweets have been deleted and abusive users have been suspended due to violations of Twitter's policy. Thus, the available datasets are no longer valid for baselines and comparisons and can only be partially be used as additional data for model training.

As we are mostly interested in tweets in Greek, we followed steps from [3] to create a new dataset. Initially, we collected all the tweets from the hashtag #απέλαση (deportation), along with two months of tweets containing the racist slang term "λάθρο", which is used to refer to undocumented immigrants (λάθρο from λαθραίος, illegal), as prime instances of hateful tweets. Interestingly enough, during the same period of time two major events occurred, namely the conviction of the neo-Nazi party Golden Dawn as a criminal organisation and the expected beginning of the trial for the murder of LGBTQ activist Zak Kostopoulos⁵ (which was to be indefinitely postponed because of Covid-19 restrictions). Similarly to what Jaki and De Smedt observe [14], there is an overlap of neo-Nazi, far-right and alt-right social media accounts that systematically target refugees, LGBTQ activists, feminists and human right advocates and this is reflected in our dataset, especially with hashtag combinations.

We then extracted a set of 500 Twitter users from these tweets and further enriched the user base with accounts appearing in mentions and replies in the bootstrapped data. We also included known media and public figure accounts, resulting a set of 1263 users in total. We collected approximately 126,000 tweets in total out of which we sampled and labelled 4004, obtaining a dataset of 1040 toxic and 2964 non-toxic tweets, which were used for model training and evaluation. For the labelling process we asked for the help of three human right activists in accordance with the process described by [4]. It is also important to note here that not all tweets from the bootstrap are racist/xenophobic and toxic, as human rights users also used these hashtags in an attempt to mitigate xenophobic propaganda. For the annotation task we used the docanno tool [23].

While this study mostly focuses on hateful and racist content produced in the Greek language, we collected tweets in English as well, both toxic and non-toxic, which we decided to keep in the corpus. As mentioned above, due to the strictness of social media platform policies, hateful content is expected to quite quickly disappear as toxic users are suspended and hateful tweets are removed, most probably in a much faster rate than in the past [5]. Account suspension and content removal happened during the collection of our data, making data collection and model evaluation and comparison even more challenging. In this sense it does not make too much sense anymore to collect and store tweet ids and urls in the form of a corpus, as most of these urls/ids are bound to become quickly obsolete.

⁵ <https://www.theguardian.com/world/2020/dec/20/long-fight-for-justice-over-death-of-greek-lgbt-activist-zak-kostopoulos>

For the computer vision models, we used the tweet urls to screen capture the visual appearance of tweets, as they are rendered in a web browser. For this task we used the Selenium Webdriver library⁶ with a headless Chrome browser of virtual window size of 600x600 pixels. Even during the labelling and collection process, several tweets have been deleted and toxic users have been suspended from Twitter and we had to manually remove this content for the tasks focusing on the visual component.

The dataset of the tweet ids we used for this work can be found in the link below⁷.

4. Methodology

In our approach we explore and combine text and image modalities to detect hate speech. There are multiple reasons to combine these modalities, as opposed to following the traditional, text-only approach. Specifically, users often use messages encoded in images to avoid NLP-based hate speech detection systems. Similarly, one can post a hate speech news item with or without text, while the rendering of the tweet, including the link preview, will still be hate speech; a NLP detection system thus needs to follow the link to determine if indeed the content of this particular url is toxic. Additionally, it is quite common among users engaging in hate speech to use visual elements to denote their ideology [14]. This is also common in the Greek context, in which users tend to include the Greek flag in both their usernames and their background images. This additional information can be leveraged by a machine learning model to improve classification quality. Following this observation, we also attempt a direct classification algorithm on user profile screenshots as a binary classification problem: users that have at least one hateful tweet vs. users that do not have hateful tweets in our dataset. The resulting model achieves a score of 75% accuracy by fine-tuning a resnet18 [24] backbone.

4.1. Text modality

We trained and evaluated our model using the Greek version of BERT⁸ and, more specifically, bert-base-greek-uncased-v1 (12-layer, 768-hidden, 12-heads, 110M parameters).

The tweets have been lower-cased and accents have been stripped before they were fed to the classifier. The overall architecture is a neural network which takes as input the text, generates the contextual representation of the tweet and then the output is fed to a linear layer with one 1 output. We use Cross Entropy Loss and fine-tune the network using Adam Optimiser with learning rate $lr = 10^{-5}$.

We train the network for 10 epochs in 80% of the data and validate it on the remaining 20%. The results are shown in table 1.

Note that we do not explicitly use author information, but we rather rely only on the text that is tweeted. Using user and social network information is expected to increase classification accuracy, given the fact that there are users that systematically use toxic language against refugees and LGBTQ members.

4.1.1. A RoBERTa LM for Greek tweets

Parallel to using the Greek version of BERT, we also trained our own language model, consisting mostly of Greek tweets. The training data set consists of 23 million tweets in Greek, of 5000 users in total, spanning from 2008 to 2018. The trained model is a small version of RoBERTa [25] and is publicly available in huggingface model zoo⁹.

4.2. Image modality

In this task, we completely omit tweet text and only use the visual representation of the tweet, as would have been rendered in a browser. Here, we explore two configura-

⁶ <https://selenium-python.readthedocs.io/>

⁷ <https://github.com/kperi/MultimodalHateSpeechDetection>

⁸ <https://github.com/nlpauieb/greek-bert>

⁹ <https://huggingface.co/Konstantinos/BERTaTweetGR>

Model	Modality	Accuracy	F1-score
BERTaTweetGR	text	0.894	0.891
nlpaueb/greek-bert	text	0.944	0.939
resnet18	image	0.915	0.849
resnet34	image	0.915	0.858
resnet50	image	0.916	0.863
resnet101	image	0.917	0.860
resnet18 + nlpaueb/BERTaTweetGR	text + image	0.94	0.931
resnet18 + nlpaueb/greek-bert	text + image	0.970	0.947
resnet34 + nlpaueb/greek-bert	text + image	0.964	0.939
resnet50 + nlpaueb/greek-bert	text + image	0.960	0.933
resnet101 + nlpaueb/greek-bert	text + image	0.960	0.930

Table 1: Summary of results

tions. In the first we crop the top part of the tweet to prevent the model from learning the user visual representation from the rendered twitter handle, while in the second we feed the entire screenshot into the deep learning model.

According to [14], the use of certain images and symbols that can be used to easily flag toxic behaviour is common among far right and alt-right groups. For example, flags, crosses and similar symbols are quite frequent in these groups. We therefore investigate if a deep learning model can detect the presence of these symbols and increase accuracy. We then fine-tune resnet{18,34,50,101} models, pre-trained on the ImageNet dataset [26], achieving a score of 0.91 accuracy with resnet101 and f1-score of 0.863 with resnet50, as shown in table 1.

4.3. Multimodal learning

Finally, we combine both modalities in a single model, in order to learn joint representations of text and tweet images. We follow the early-fusion architecture and combining the representations of the BERT and Resnet models into a single representation vector, followed by a feedforward network, in which we train for 20 epochs. The combination of the two modalities indeed increases the classification accuracy by approximately 2.5

5. Results

A detailed list of the results obtained by our models are presented in the table 1.

Initially, we observe that by using only the image modality the classification rates are high, compared to other systems that use text only approach.

It is interesting to note that we have a misclassified tweet that contains no text in our validation set. It is also interesting that the model sometimes confuses tweets that contain words that are non-toxic by definition but are quite frequently used in toxic tweets. Typical examples are words such as "Ισλάμ" (Islam) or "πρόσφυγας" (refugee). This is in agreement with [3]. Since our training dataset has been initially seeded with all the tweets of a specific hashtag, after aggregating all labelled and predicted positive hate speech tweets, we find that the top 3 most toxic accounts are in agreement with the analysis performed by Smyrniaios¹⁰.

6. Interpretability

One interesting question that can be raised is on what the model focuses in order to make a decision. This is very interesting for both model interpretability and model debugging: we would like to be confident that our model learns the correct features from the input space and we do not leak unwanted information to the model.

¹⁰ <http://ephemeron.eu/2220>

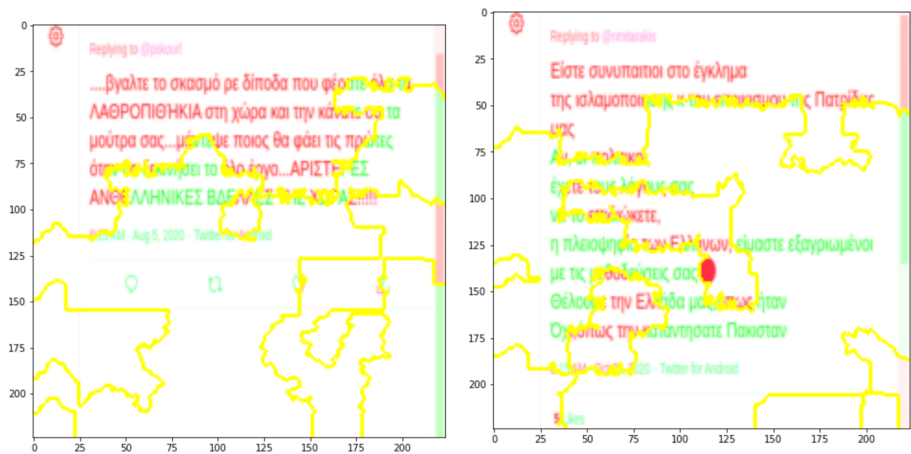


Figure 1. Regions of the input image contributing to toxicity - toxic prediction

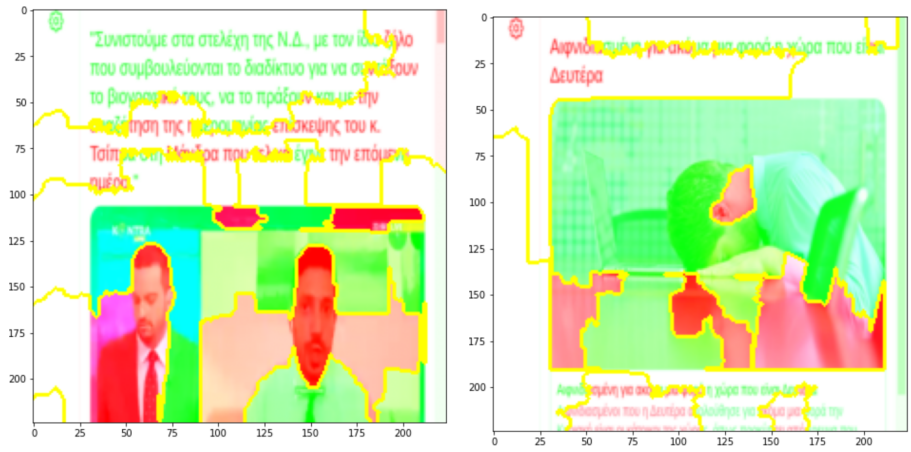


Figure 2. Regions of the input image contributing to non toxicity - non toxic prediction

We use the Lime library (LIME - Local Interpretable Model-Agnostic Explanations [27]) for model interpretability.

Figure 1 presents examples of instances correctly classified as toxic by the model. The red regions of the examples refer to boundaries of the input image that are contributing to toxicity and the green ones to non-toxicity. Similarly figure 2 shows input regions that are contributing to instances classified as non-toxic.

The same process can be applied in text classifiers, resulting to contribution in word level per class, as shown in figures 3, 4.

7. Scaling and applications

Using our best model we have been able to generate predictions for more than 30K tweets per day, using a small virtual machine (4GB RAM, no GPU for test time).




Figure 3. Words contributing to toxic prediction




Figure 4. Words contributing to non toxic prediction

Prediction time using GPU is significantly faster and our benchmarks using a Tesla GPU show that we can automatically label more than 100000 tweets in approximately 10 minutes. The main bottleneck for this process is the screen capture task to obtain the visual modality as it is subject to rate limit and several passes are required to ensure that the screen has been correctly captured.

The framework we recommend in this work can be relatively easily adapted to more classes such as hate speech, offensive speech or other forms of abusive speech including sexism.

8. Conclusion and further research

In this paper we investigate the problem of hate speech in social networks and, more specifically, racist speech against refugees in Greek. Our approach follows the most recent trends in Natural Language Processing and suggests that visual and textual modalities combined in a late-fusion multimodal learning setting can improve overall detection accuracy. As part of this work we have made publicly available a RoBERTa-based Language Model trained on Greek tweets. Our NLP models have been developed using the transformers python library [28] and our Computer Vision models using pytorch [29]. Our research focuses only on single tweets and does not take into account social graph information. One interesting research direction to follow would be to combine the multimodal learning approach with social graph information in a single framework, essentially combining Graph Neural Networks with multimodal representations.

References

- Williams, M.L.; Burnap, P.; Javed, A.; Liu, H.; Ozalp, S. Hate in the machine: anti-Black and anti-Muslim social media posts as predictors of offline racially and religiously aggravated crime. *The British Journal of Criminology* **2020**, *60*, 93–117.
- Halevy, A.; Ferrer, C.C.; Ma, H.; Ozertem, U.; Pantel, P.; Saeidi, M.; Silvestri, F.; Stoyanov, V. Preserving Integrity in Online Social Networks, 2020, [arXiv:cs.SI/2009.10311].
- Waseem, Z.; Hovy, D. Hateful Symbols or Hateful People? Predictive Features for Hate Speech Detection on Twitter. Proceedings of the NAACL Student Research Workshop; Association for Computational Linguistics: San Diego, California, 2016; pp. 88–93.
- Waseem, Z. Are You a Racist or Am I Seeing Things? Annotator Influence on Hate Speech Detection on Twitter. Proceedings of the First Workshop on NLP and Computational Social Science; Association for Computational Linguistics: Austin, Texas, 2016; pp. 138–142.
- Mishra, P.; Del Tredici, M.; Yannakoudakis, H.; Shutova, E. Abusive language detection with graph convolutional networks. *arXiv preprint arXiv:1904.04073* **2019**.
- Guberman, J.; Schmitz, C.; Hemphill, L. Quantifying toxicity and verbal violence on Twitter. Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion, 2016, pp. 277–280.
- Gunasekara, I.; Nejadgholi, I. A review of standard text classification practices for multi-label toxicity identification of online content. Proceedings of the 2nd workshop on abusive language online (ALW2), 2018, pp. 21–25.
- Parent, M.C.; Gobble, T.D.; Rochlen, A. Social media behavior, toxic masculinity, and depression. *Psychology of Men & Masculinities* **2019**, *20*, 277.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G.; Dean, J. Distributed representations of words and phrases and their compositionality. *arXiv preprint arXiv:1310.4546* **2013**.
- Le, Q.; Mikolov, T. Distributed representations of sentences and documents. International conference on machine learning. PMLR, 2014, pp. 1188–1196.

11. Grover, A.; Leskovec, J. node2vec: Scalable feature learning for networks. *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 855–864.
12. Perifanos, K.; Florou, E.; Goutsos, D. Neural Embeddings for Idiolect Identification. *2018 9th International Conference on Information, Intelligence, Systems and Applications (IISA)*, 2018, pp. 1–3.
13. Davidson, T.; Warmley, D.; Macy, M.; Weber, I. Automated Hate Speech Detection and the Problem of Offensive Language, 2017, [\[arXiv:cs.CL/1703.04009\]](https://arxiv.org/abs/1703.04009).
14. Jaki, S.; Smedt, T.D. Right-wing German Hate Speech on Twitter: Analysis and Automatic Detection. *CoRR* **2019**, *abs/1910.07518*, [\[1910.07518\]](https://arxiv.org/abs/1910.07518).
15. Poletto, F.; Stranisci, M.; Sanguinetti, M.; Patti, V.; Bosco, C. Hate speech annotation: Analysis of an italian twitter corpus. *4th Italian Conference on Computational Linguistics, CLiC-it 2017. CEUR-WS*, 2017, Vol. 2006, pp. 1–6.
16. Pereira-Kohatsu, J.C.; Quijano-Sánchez, L.; Liberatore, F.; Camacho-Collados, M. Detecting and monitoring hate speech in Twitter. *Sensors* **2019**, *19*, 4654.
17. Baider, F.; Constantinou, M. Covert hate speech: A contrastive study of Greek and Greek Cypriot online discussions with an emphasis on irony. *Journal of Language Aggression and Conflict* **2020**, *8*, 262–287.
18. Lekea, I.K.; Karampelas, P. Detecting hate speech within the terrorist argument: a Greek case. *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 2018, pp. 1084–1091.
19. Pitenis, Z.; Zampieri, M.; Ranasinghe, T. Offensive language identification in greek. *arXiv preprint arXiv:2003.07459* **2020**.
20. Morency, L.P.; Baltrušaitis, T. Multimodal Machine Learning: Integrating Language, Vision and Speech. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*; Association for Computational Linguistics: Vancouver, Canada, 2017; pp. 3–5.
21. Lippe, P.; Holla, N.; Chandra, S.; Rajamanickam, S.; Antoniou, G.; Shutova, E.; Yannakoudakis, H. A Multimodal Framework for the Detection of Hateful Memes. *arXiv preprint arXiv:2012.12871* **2020**.
22. Kiela, D.; Firooz, H.; Mohan, A.; Goswami, V.; Singh, A.; Ringshia, P.; Testuggine, D. The Hateful Memes Challenge: Detecting Hate Speech in Multimodal Memes, 2020, [\[arXiv:cs.AI/2005.04790\]](https://arxiv.org/abs/2005.04790).
23. Nakayama, H.; Kubo, T.; Kamura, J.; Taniguchi, Y.; Liang, X. doccano: Text Annotation Tool for Human, 2018. Software available from <https://github.com/doccano/doccano>.
24. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition, 2015, [\[arXiv:cs.CV/1512.03385\]](https://arxiv.org/abs/1512.03385).
25. Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; Stoyanov, V. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* **2019**.
26. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Fei-Fei, L. Imagenet: A large-scale hierarchical image database. *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
27. Ribeiro, M.T.; Singh, S.; Guestrin, C. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, 2016, pp. 1135–1144.
28. Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; Davison, J.; Shleifer, S.; von Platen, P.; Ma, C.; Jernite, Y.; Plu, J.; Xu, C.; Scao, T.L.; Gugger, S.; Drame, M.; Lhoest, Q.; Rush, A.M. Transformers: State-of-the-Art Natural Language Processing. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*; Association for Computational Linguistics: Online, 2020; pp. 38–45.
29. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; Chintala, S. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*; Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; Garnett, R., Eds.; Curran Associates, Inc., 2019; pp. 8024–8035.
