

Multilingual Counter Narrative Type Classification

Yi-Ling Chung^{◇, ♣}, Marco Guerini[♣], and Rodrigo Agerri[♣]

[◇]University of Trento

[♣]Fondazione Bruno Kessler

[♣]HiTZ Center - Ixa, University of the Basque Country UPV/EHU

ychung@fbk.eu, guerini@fbk.eu, rodrigo.agerri@ehu.eus

Abstract

The growing interest in employing counter narratives for hatred intervention brings with it a focus on dataset creation and automation strategies. In this scenario, learning to recognize counter narrative types from natural text is expected to be useful for applications such as hate speech countering, where operators from non-governmental organizations are supposed to answer to hate with several and diverse arguments that can be mined from online sources. This paper presents the first multilingual work on counter narrative type classification, evaluating SoTA pre-trained language models in monolingual, multilingual and cross-lingual settings. When considering a fine-grained annotation of counter narrative classes, we report strong baseline classification results for the majority of the counter narrative types, especially if we translate every language to English before cross-lingual prediction. This suggests that knowledge about counter narratives can be successfully transferred across languages.

1 Introduction

The pervasive problem of online hate speech (HS) has motivated the research community to investigate methods for mitigating hatred, such as hate speech detection (Schmidt and Wiegand, 2017; Fortuna and Nunes, 2018) and, more recently, hate moderation through counter narratives (CNs). CNs are non-negative responses to hateful messages providing fact-bound arguments or alternative viewpoints. Distinct from standard approaches to hate intervention by content moderation, counter narratives are preferable as they preserve the right to freedom of speech and encourage peaceful conversations (Benesch, 2014; Schieb and Preuss, 2016).

In particular, the NLP community has started exploring CN generation (Qian et al., 2019; Tekiroğlu et al., 2020; Chung et al., 2021; Fanton et al., 2021; Zhu and Bhat, 2021), also in multilingual settings (Chung et al., 2020) with the aim of help-

ing non-governmental organizations (NGOs) fight HS. However, automatic generation and detection of counter narratives still face three important challenges. First, counter narratives vary considerably in manner of expression and strategies (Benesch et al., 2016), see Table 1, posing difficulties to automatic CN evaluation and CN classification. Second, few corpora are available since domain experts are required to obtain high-quality data. Third, the only existing work on classifying counter narratives targeted just English (Mathew et al., 2019).

In our view, automated mechanisms to distinguish CN characteristics are crucial for developing hate countering applications that can address personalization or for tasks such as counter narrative mining from user-generated content. Classifying counter narratives would also help to establish which counter narratives are more effective to target hate speech. Finally, it may also be useful for NGO operators to answer hate speech with diverse counter narratives (not just the most common ones) which could be mined from online sources.

HS: *Muslims and non-British people are raping, enslaving and murdering our women! How disgusting!*

CN1 (Denouncing): *This is not acceptable. Hatred cannot be tolerated and must be stopped.*

CN2 (Facts): *Please notice that, regardless of religions, huge numbers of girls and children experience domestic sexual abuse by a partner or family member. We should encourage people to help victims and prevent sexual crimes.*

Table 1: An example hate speech with counter narratives (CN1 and CN2) using different strategies.

To the best of our knowledge, this work is the first to tackle counter narrative classification from a multilingual perspective, predicting the types/strategies employed in a counter narrative message given a hateful message. Other contributions include: (1) the first empirical analysis of counter narrative characteristics on expert-based hate countering data targeting Islamophobia (Chung et al., 2019); (2) strong baseline re-

sults for multilingual and multi-class counter narrative classification by experimenting with the large pre-trained model XLM-R (Conneau et al., 2020); and (3) a set of cross-lingual zero-shot experiments showing that huge improvements can be obtained by translating the training data from other languages to English, which is coherent with previous work on the multilinguality of language models (Pires et al., 2019). This result also shows that common characteristics of some counter narrative types are carried across languages. Data splits and training scripts are available at <https://github.com/yilingchung/multilingualCN-classification>.

2 Counter Arguments v.s. Counter Narratives

While counter arguments and counter narratives have a similar purpose - to present an alternative stance to a statement - a counter argument is not necessarily a counter narrative, and vice versa. The main differences between the two are briefly discussed below, and point for the development of specific classification approaches for CNs.

Presence of an argument. A counter argument is typically defined as an argument (equivalent to a claim or standpoint) contradicting the initial statement with supporting evidence or reasoning (Khishfe, 2012; Schiappa and Nordin, 2013; Stab et al., 2018). The presence of supported evidence or reasoning is essential for counter-argumentation and argumentative information retrieval, regardless of it being implicit or explicit (e.g., Statement: *Nuclear energy should be used to replace fossil fuels.* v.s. Counter argument: *Renewable energy is a better option for replacing fossil fuels than unsustainable and expensive energy like nuclear power.*). In contrast, counter narratives can be formulated without using arguments. For example, a counter narrative can be without supporting evidence (e.g., Hate speech: *Islam is a plague.* v.s. Counter narrative: *How can you use such language to describe a religion of 1.6 billion people?*).

Non-negative arguments. Following guidelines¹ generally adopted by NGOs, an appropriate counter narrative should not contain hostile language, prejudice, and unlawful content. Using negative tones as responses may risk escalating the conversation and engaging in hateful sentiment.

¹As an example: <https://getthetrollsout.org/stoppinghate>

On the contrary, counter arguments are not bound by such constraints.

3 Counter Narrative Type Classification

Classifying counter narrative types is a new and challenging task. We hypothesize that underlying linguistic features in one language can facilitate the classification performance in other languages. Accordingly, we approach this task in a multilingual setting focusing on English, French and Italian. We experimented in the following learning settings: (1) **monolingual**, in which the train and test set are in the same language; (2) **multilingual**, in which we train in 3 languages and test on each of them; (3) **zero-shot cross-lingual**, where we train in one language and evaluate on the other two unseen languages; and, (4) **zero-shot translated**, similar to the multilingual experiment but translating the Italian and French training data into English.

3.1 Dataset

	Train			Test		
	EN	IT	FR	EN	IT	FR
Facts	957	1080	1329	237	270	333
Question	258	285	342	66	72	84
Denouncing	405	153	705	102	36	174
Humor	162	96	213	42	24	51
Hypocrisy	270	111	348	66	27	87
Total	2052	1725	2937	513	429	729

Table 2: Number of instances by types.

In our experiments we use CONAN (Chung et al., 2019), the only multilingual hate counter-argument dataset currently available that has been niche-sourced to NGO operators. The dataset consists of 14k HS-CN pairs with counter narrative type annotation over English, French, and Italian. For each language, an original pair is augmented with two paraphrases of the original HS coupled with the original CN. In CONAN annotation, one CN could be assigned more than one CN type. As a first investigation into multilingual counter narrative classification, we select the pairs annotated with just one CN type for simplicity, discarding 25%, 17%, and 27% of samples for EN, FR and IT, respectively. Considering that the class balance is skewed, we narrow down the categories to majority classes over 3 languages: *facts* (35%), *denouncing* (13%), *question* (9%), *hypocrisy* (7%), and *humor* (5%). We then randomly sampled 80% and 20% of the dataset for the training and testing, ensuring that one HS and its two HS paraphrases paired with the same counter narrative are kept in the same split,

so that the same CN does not appear by chance in both the training and testing sets. While the class imbalance poses a challenge, it reflects the practical scenario where certain types are less frequent.

Since CONAN contains only positive examples of CN, we further created 200 instances for the classes *support* and *unrelated*.² *Support* is a response that endorses the hate speech and *unrelated* is a text not connected to the hate speech in any sense. Clearly these two classes do not fall into the categories of CN and in fact they do not exist in CONAN. We include them to ensure that the models are exposed to varied non-CN text. Such setting can avoid model overfitting and increase the applicability of a system to the real world full of noisy content. The amount of these instances (200) is set to be close to the less populated class (*humor*). For the *unrelated* pairs, we randomly sampled data from Wikilingua³ (Ladhak et al., 2020), featuring topics unrelated to islamophobia. The *support* pairs consist of HS paired with each of the two paraphrases. Table 2 reports the distribution of training and testing examples per class across the three languages.

3.2 Models

The XLM-R language model (Conneau et al., 2020), pre-trained on CommonCrawl for 100 languages, reports strong performance on several cross-lingual downstream tasks such as natural language inference or named entity recognition, and also on multilingual stance detection, close related to CN classification (Zotova et al., 2021). In this paper we leverage XLM-R to provide a strong baseline in the task of CN type classification in both monolingual and multilingual settings, using the Transformers library (Wolf et al., 2020). For every experiment, we fine-tuned the base version of XLM-R for 10 epochs with batch size 32 and 2e-5 learning rate.

4 Experimental Results

Similarly to stance detection and argument mining tasks (Mohammad et al., 2016; Stab et al., 2018), we report average macro F1 score over the CN types to avoid obtaining very high scores simply

²This allows, for instance, to address cases where we need to identify counter narratives in a conversation containing both abusive and non-abusive language.

³Wikilingua is a multilingual dataset for summarization covering how-to guides on various topics written by human from WikiHow.

deriving from the dominant classes in the dataset. Furthermore, we also follow previous literature on stance detection (Mohammad et al., 2016) and report only the performance over the relevant classes.

Monolingual results. As shown by Table 3, monolingual models consistently yield the best performance in predicting the majority classes *facts* and *question*. As for *denouncing*, we obtain results above average for English, French and Italian (0.65, 0.58, and 0.51 respectively). Worst results are obtained for *humor* and *hypocrisy*. For Italian and French the model completely fails to identify *humor*, the most difficult and under-represented class; for English, the prediction is moderate (0.45). Lastly, the low results for the *hypocrisy* class seemed to be caused by the difficulty for the model in discriminating *hypocrisy* from *facts*. For example, 38% of the prediction errors for the *hypocrisy* class in English are caused by wrongly classifying *hypocrisy* instances as *facts* (more details are provided in the Appendix, Figure 1).

In a post-hoc manual analysis, annotators expressed difficulties in differentiating these two classes, difficulties illustrated by the examples provided in Table 4. This issue seems to be further confirmed by annotation statistics obtained from CONAN: among all the instances that contains the *hypocrisy* label, 45% were annotated with *hypocrisy* alone, 27% with *hypocrisy* and *facts*, while other multi-label cases were much lower. We hypothesize that *hypocrisy* could be a subclass of *facts*. After all, pointing out the contradiction in a hate statement may imply correcting misstatements via facts. Future work can try to merge two classes together to improve classification performance.

Multilingual results. As it is reported in Table 3, multilingual learning obtains better overall results than monolingual variants. In general the model trained on all languages improve the per-type performance although we observe decreases in *denouncing* and *hypocrisy* for Italian. Also, *humor* is still challenging for each language with more than 50% of miss-categorization.

Zero-shot cross-lingual results. For each target language, we provide the average F1 of zero-shot cross-lingual models trained on the other 2 source languages in Table 3. Although the results are lower than in the monolingual setting, results show that cross-lingual transfer can be a feasible strategy for classifying counter narratives for languages for

Target language	Monolingual			Multilingual			Zero-shot			Zero-shot transl _{en}		
	EN	IT	FR	EN	IT	FR	EN	IT	FR	EN	IT	FR
Average	0.60	0.59	0.49	0.65	0.56	0.49	0.45	0.46	0.40	0.65	0.69	0.72
Facts	0.84	0.91	0.71	0.82	0.94	0.74	0.71	0.86	0.67	0.85	0.95	0.83
Question	0.73	0.91	0.68	0.81	0.91	0.69	0.65	0.85	0.52	0.75	0.93	0.87
Denouncing	0.65	0.51	0.58	0.68	0.46	0.61	0.28	0.36	0.39	0.70	0.69	0.75
Humor	0.45	0.08	0.15	0.44	0.00	0.05	0.34	0.00	0.20	0.56	0.42	0.51
Hypocrisy	0.35	0.52	0.36	0.49	0.48	0.35	0.30	0.25	0.23	0.40	0.46	0.62

Table 3: Macro-F1 results under various configurations. For each target language in zero-shot setting, the scores are averaged over the models trained on the other two source languages.

HS	CN (Gold CN Type)	Prediction
<i>Not all Muslims are terrorists, but all terrorists are Muslims.</i>	<i>Only because when a non-muslim commits a terrorist act it is not called terrorism. Instead they're a 'lone wolf' with a 'traumatic past'. (Hypocrisy)</i>	Facts
<i>Women of our culture that decide to become Islamic are so selfish. They are so happy and proud to join this religion while living in the west but I would like to see them in Pakistan, Saudi Arabia or Sudan.</i>	<i>Maybe that 'sisterhood of tolerance and understanding' is challenging sexism everywhere including those men who feel they can tell women what they should believe in and how to live their lives. (Facts)</i>	Hypocrisy

Table 4: Examples of misclassified counter narratives.

which no training data is readily available, especially for the majority classes. Best results in this scenario are obtained when English is the source language, obtaining lower but close results with respect to the monolingual results in Italian and French (for more details about the results per language pair see Table 5 in the Appendix).

Zero-shot translated results. Data augmentation through translation has been widely employed to improve classification performance (Toledo-Ronen et al., 2020), also in cross-lingual settings (Zotova et al., 2021). At the same time, the cross-lingual capabilities of Transformer models such as XLM-R and mBERT are being actively investigated (Muller et al., 2021; Pires et al., 2019; Wu and Dredze, 2019). Thus, we conduct an additional experiment adding, to the English training set, the manually translated Italian and French training data before testing on the target languages (Italian and French). The aim is to investigate if XLM-R benefits from fine-tuning on a high resource language (English) instead of combining English with other languages, such as Italian and French, which are not so well represented in pre-trained multilingual models (Martin et al., 2020; Agerri et al., 2020; Espinosa et al., 2020). By doing so, in this ‘zero-shot transl_{en}’ setting we aim to expose the model with semantic knowledge from some target languages without actually seeing those languages (Italian and French). Table 3 shows that we obtain a huge performance jump with respect to the multilingual results, with 13 points improvement in macro-F1

score for Italian and 23 points for French. Interestingly, the improvement is more impressive for the *humor* class, the most challenging of them all.

5 Related Work

Broadly speaking, counter narrative type classification is related to stance detection (Toledo-Ronen et al., 2020; Schiller et al., 2021), which is crucial for argument search (Stab et al., 2018). In contrast to stance detection, that concentrates on binary or relatively simple classification – e.g., determine if an argument supports or contests a given topic (Sridhar et al., 2015; Rosenthal and McKeown, 2015; Stab et al., 2018) – we present a multi-class approach to counter narrative classification.

Counter narratives have been adopted as a direct and effective response to online hatred in several campaigns and on social media platforms including Twitter (Munger, 2017; Wright et al., 2017), Facebook (Schieb and Preuss, 2016), and Youtube (Ernst et al., 2017; Mathew et al., 2019). Although it has been argued that hate speech detection can benefit from CN classification, there are very few studies on this regard, with only one previous work on classifying counter narrative types (Mathew et al., 2019). However, unlike our present work, they consider hostile language as one of the main types of counter narratives, which is explicitly discouraged by NGOs working on hatred intervention. Furthermore, we investigate multilingual and cross-lingual CN classification leveraging a SoTA pre-trained multilingual language model.

6 Conclusion

We present the first work on multilingual CN type classification. Our results show that: (i) the performance is promising for the majority classes (*facts, question, denouncing*); (ii) classifying *humor* and *hypocrisy* CNs is still challenging; (iii) combining training data from the three source languages improves performance over the monolingual evaluation; and (iv), the best overall results are obtained in the ‘zero-shot trans_{en}’ approach where the training data for Italian and French is translated to English. This shows that some knowledge about CNs is transferred across languages. While this is coherent with previous literature about multilingual language models, the exact source of such successful transfer across languages remains an open topic.

Acknowledgements

Rodrigo Agerri is funded by the RYC-2017–23647 fellowship and by the DeepReading (RTI2018-096846-B-C21, MCIU/AEI/FEDER, UE, Spanish Ministry of Science, Innovation and Universities) and DeepText (KK-2020/00088, Basque Government) projects.

References

- Rodrigo Agerri, Iñaki San Vicente, Jon Ander Campos, Ander Barrena, Xabier Saralegi, Aitor Soroa, and Eneko Agirre. 2020. [Give your text representation models some love: the case for Basque](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4781–4788, Marseille, France. European Language Resources Association.
- Susan Benesch. 2014. [Countering dangerous speech: New ideas for genocide prevention](#). Washington, DC: United States Holocaust Memorial Museum.
- Susan Benesch, Derek Ruths, Kelly P Dillon, Haji Mohammad Saleem, and Lucas Wright. 2016. [Counter-speech on twitter: A field study](#). *Dangerous Speech Project*.
- Yi-Ling Chung, Elizaveta Kuzmenko, Serra Sinem Tekiroglu, and Marco Guerini. 2019. [CONAN - COUNTER NARRATIVES THROUGH NICHE SOURCING: 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. Association for Computational Linguistics.
- Yi-Ling Chung, Serra Sinem Tekiroglu, and Marco Guerini. 2020. [Italian counter narrative generation to fight online hate speech](#). In *Proceedings of the Seventh Italian Conference on Computational Linguistics*, Bologna, Italy.
- Yi-Ling Chung, Serra Sinem Tekiroglu, and Marco Guerini. 2021. [Towards knowledge-grounded counter narrative generation for hate speech](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 899–914, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Julian Ernst, Josephine B Schmitt, Diana Rieger, Ann Kristin Beier, Peter Vorderer, Gary Bente, and Hans-Joachim Roth. 2017. [Hate beneath the counter speech? a qualitative content analysis of user comments on youtube related to counter speech videos](#). *Journal for Deradicalization*, 10:1–49.
- María S. Espinosa, Rodrigo Agerri, Álvaro Rodrigo, and Roberto Centeno. 2020. [Deepreading @ sardistance 2020: Combining textual, social and emotional features](#). In *EVALITA*.
- Margherita Fanton, Helena Bonaldi, Serra Sinem Tekiroglu, and Marco Guerini. 2021. [Human-in-the-loop for data collection: a multi-target counter narrative dataset to fight online hate speech](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3226–3240, Online. Association for Computational Linguistics.
- Paula Fortuna and Sérgio Nunes. 2018. [A survey on automatic detection of hate speech in text](#). *ACM Computing Surveys (CSUR)*, 51(4):85.
- Rola Khishfe. 2012. [Relationship between nature of science understandings and argumentation skills: A role for counterargument and contextual factors](#). *Journal of Research in Science Teaching*, 49(4):489–514.
- Faisal Ladhak, Esin Durmus, Claire Cardie, and Kathleen McKeown. 2020. [WikiLingua: A new benchmark dataset for cross-lingual abstractive summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4034–4048, Online. Association for Computational Linguistics.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamé Seddah, and Benoît Sagot. 2020. [CamemBERT: a tasty French language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.

- Binny Mathew, Punyajoy Saha, Hardik Tharad, Subham Rajgaria, Prajwal Singhania, Suman Kalyan Maity, Pawan Goyal, and Animesh Mukherjee. 2019. Thou shalt not hate: Countering online hate speech. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 13, pages 369–380.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. [SemEval-2016 task 6: Detecting stance in tweets](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Association for Computational Linguistics.
- Benjamin Muller, Yanai Elazar, Benoît Sagot, and Djamel Seddah. 2021. [First align, then predict: Understanding the cross-lingual ability of multilingual BERT](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2214–2231, Online. Association for Computational Linguistics.
- Kevin Munger. 2017. [Tweetment effects on the tweeted: Experimentally reducing racist harassment](#). *Political Behavior*, 39(3):629–649.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Jing Qian, Anna Bethke, Yinyin Liu, Elizabeth Belding, and William Yang Wang. 2019. [A benchmark dataset for learning to intervene in online hate speech](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 4755–4764, Hong Kong, China. Association for Computational Linguistics.
- Sara Rosenthal and Kathy McKeown. 2015. [I couldn't agree more: The role of conversational structure in agreement and disagreement detection in online discussions](#). In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 168–177, Prague, Czech Republic. Association for Computational Linguistics.
- Edward Schiappa and John P. Nordin. 2013. *Argumentation: Keeping Faith with Reason*. Pearson UK.
- Carla Schieb and Mike Preuss. 2016. Governing hate speech by means of counterspeech on facebook. In *66th ICA Annual Conference*, pages 1–23, Fukuoka, Japan.
- Benjamin Schiller, Johannes Daxenberger, and Iryna Gurevych. 2021. Stance detection benchmark: How robust is your stance detection? *KI-Künstliche Intelligenz*, pages 1–13.
- Anna Schmidt and Michael Wiegand. 2017. A survey on hate speech detection using natural language processing. In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, pages 1–10.
- Dhanya Sridhar, James Foulds, Bert Huang, Lise Getoor, and Marilyn Walker. 2015. [Joint models of disagreement and stance in online debate](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 116–125, Beijing, China. Association for Computational Linguistics.
- Christian Stab, Tristan Miller, Benjamin Schiller, Pranav Rai, and Iryna Gurevych. 2018. [Cross-topic argument mining from heterogeneous sources](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3664–3674, Brussels, Belgium. Association for Computational Linguistics.
- Serra Sinem Tekiroğlu, Yi-Ling Chung, and Marco Guerini. 2020. [Generating counter narratives against online hate speech: Data and strategies](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1177–1190, Online. Association for Computational Linguistics.
- Orith Toledo-Ronen, Matan Orbach, Yonatan Bilu, Artem Spector, and Noam Slonim. 2020. [Multilingual argument mining: Datasets and analysis](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 303–317, Online. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Lucas Wright, Derek Ruths, Kelly P Dillon, Haji Mohammad Saleem, and Susan Benesch. 2017. [Vectors for counterspeech on Twitter](#). In *Proceedings of the First Workshop on Abusive Language Online*, pages 57–62, Vancouver, BC, Canada. Association for Computational Linguistics.
- Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages

833–844, Hong Kong, China. Association for Computational Linguistics.

Wanzheng Zhu and Suma Bhat. 2021. [Generate, prune, select: A pipeline for counterspeech generation against online hate speech](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 134–149, Online. Association for Computational Linguistics.

Elena Zotova, Rodrigo Agerri, and German Rigau. 2021. [Semi-automatic generation of multilingual datasets for stance detection in twitter](#). *Expert Syst. Appl.*, 170:114547.

A Appendices

	IT -> EN	FR -> EN	EN -> IT	FR -> IT	EN -> FR	IT -> FR
Average	0.41	0.50	0.48	0.45	0.43	0.37
Facts	0.71	0.71	0.88	0.84	0.73	0.60
Question	0.56	0.74	0.84	0.87	0.48	0.55
Denouncing	0.30	0.25	0.46	0.26	0.44	0.33
Humor	0.24	0.43	0.00	0.00	0.22	0.18
Hypocrisy	0.22	0.38	0.24	0.27	0.28	0.17

Table 5: Zero-shot cross-lingual results in terms of macro-F1 per type.




Figure 1: Confusion matrix on monolingual training for EN, IT, and FR from left to right (upper part); multilingual model tested on EN, IT, and FR from left to right (down part). The predictions are represented by columns and gold class is represented in rows.