

CODEOFCONDUCT at Multilingual Counterspeech Generation: A Context-Aware Model for Robust Counterspeech Generation in Low-Resource Languages

Michael Bennie¹, Bushi Xiao¹, Chryseis Xinyi Liu¹, Demi Zhang¹,
Jian Meng¹, Alayo Tripp¹

¹University of Florida, United States

Correspondence: michaelbennie@ufl.edu, xiaobushi@ufl.edu, liu.x1@ufl.edu, zhang.yidan@ufl.edu

Abstract

This paper introduces a context-aware model for robust counterspeech generation, which achieved significant success in the MCG-COLING-2025 shared task. Our approach particularly excelled in low-resource language settings. By leveraging a simulated annealing algorithm fine-tuned on multilingual datasets, the model generates factually accurate responses to hate speech.

We demonstrate state-of-the-art performance across four languages (Basque, English, Italian, and Spanish), with our system ranking first for Basque, second for Italian, and third for both English and Spanish. Notably, our model swept all three top positions for Basque, highlighting its effectiveness in low-resource scenarios.

Evaluation of the shared task employs both traditional metrics (BLEU, ROUGE, BERTScore, Novelty) and JudgeLM based on LLM. We present a detailed analysis of our results, including an empirical evaluation of the model performance and comprehensive score distributions across evaluation metrics.

This work contributes to the growing body of research on multilingual counterspeech generation, offering insights into developing robust models that can adapt to diverse linguistic and cultural contexts in the fight against online hate speech.

1 Introduction

Hate speech (HS) encompasses expressions that demean or target individuals or groups based on characteristics such as race, ethnicity, gender, sexual orientation, or religion (de Gibert et al., 2018). Although HS makes up only a small fraction of social media content, its impact is significant, affecting nearly one-third of people (Vidgen et al., 2019). The prevalence of HS on social media has become a critical societal concern. Traditional approaches, such as content removal and moderation, have been

widely implemented but are often criticized for infringing on free speech. As an alternative, counterspeech (CS) has emerged as a promising solution to mitigate HS while upholding the principle of free expression (Poudhar et al., 2024).

Counterspeech is defined as speech intended to counteract and neutralize harmful language. It has demonstrated effectiveness in real-world applications (Cepollaro et al., 2023), but its manual creation is labor-intensive and impractical at scale given the high volume of HS online. This challenge has driven interest in automating CS generation using NLP technologies. The growing need for effective counterspeech highlights the importance of strategies that foster healthier online environments. However, multilingual CS generation remains a significant challenge, particularly in low-resource settings where data scarcity limits model development. Research has focused on understanding HS, crafting effective CS, and addressing the unique challenges of multilingual and resource-constrained contexts.

The MCG-COLING-2025 shared task addressed these challenges by inviting researchers to generate respectful, specific, and truthful counterspeech across multiple languages, including Basque, English, Italian, and Spanish. In this paper, we present *CODEOFCONDUCT*, a context-aware model that achieved state-of-the-art performance in this shared task. Our model demonstrates exceptional effectiveness in low-resource scenarios, particularly for Basque, and offers valuable insights into leveraging multilingual datasets and advanced optimization techniques for counterspeech generation.

Our contributions

- A novel application of the simulated annealing algorithm in fine-tuning for counterspeech generation.
- Competitive performance across four lan-

guages, including top-ranking results in low-resource language Basque.

- Critical analysis of provided CS evaluation tools, especially Judge-LM.

2 Background

The languages chosen for the shared task cover a wide range of linguistic features, from Basque’s complex agglutinative morphology to the more straightforward syntax of English. This diversity allows for testing models’ adaptability to varying linguistic challenges. Furthermore, the inclusion of background knowledge and cultural differences in multiple languages adds an additional layer of complexity, requiring models to integrate contextual information effectively.

2.1 Opportunities and Gaps

Multilingual counterspeech generation is challenging due to variations in linguistic structure, cultural norms, and resource availability. English dominates this field due to its resource-rich environment, while languages like Basque lack sufficient annotated data and pretrained models (Faisal et al. (2021)). Approaches to address these challenges often leverage transfer learning and multilingual pre-training. Models like mBERT (Devlin et al. (2019)) and XLM-RoBERTa (Conneau et al. (2020)) have demonstrated robust cross-lingual transfer capabilities, enabling better performance in low-resource settings. Fine-tuning these models on task-specific data, such as the HS-CN pairs in this shared task, has shown promise in generating meaningful and contextually relevant counterspeech.

Notably, the integration of background knowledge to enhance counterspeech quality has been explored in previous works (Qian et al. (2019)), which introduced external knowledge to improve the informativeness of responses. This concept is particularly relevant in the MCG-COLING-2025 task that utilizes the CONAN dataset, where models must use hate speech, background knowledge, and linguistic nuances to synthesize counter speech.

2.2 Low-Resource NLP: Basque as a Case Study

To bridge these gaps, which are due to their complex morphology and limited linguistic resources, techniques such as data augmentation (Feng et al., 2021), transfer learning, and multilingual pretraining have been explored. A Basque BERT model

(Agerri and et al., 2020) has been developed using a dataset drawn from the Basque edition of Wikipedia and news articles from various Basque media sources, illustrating the potential of specialized models in low-resource contexts.

For counterspeech generation, low-resource languages require innovative solutions to overcome data scarcity. Judge-EUS¹ has been utilized to enhance response quality, as demonstrated below in our approach.

2.3 Multilingual Hate Speech-Counter Narrative Dataset

For this shared task, the training dataset involves 596 Hate Speech-Counter Narrative (HS-CN) pairs curated by the task organizers. The hate speech instances were sourced from the Multitarget-CONAN dataset (Fanton et al., 2021), and the counterspeech responses were newly generated by the organizers. Each HS-CN pair is accompanied by five background knowledge sentences, some of which are relevant for crafting effective counterspeech.

The dataset spans four languages: Basque, English, Italian, and Spanish, representing a diverse typological spectrum:

- **Basque:** An agglutinative language isolate with unique grammatical structure.
- **Italian and Spanish:** Two Romance languages with high lexical and syntactic similarity.
- **English:** A Germanic language with relatively simpler morphological structures.

The dataset is divided into the following splits for each language:

- **Development:** 100 HS-CN pairs.
- **Training:** 396 HS-CN pairs.
- **Testing:** 100 HS-CN pairs (counter-narratives held out as blind test data).

2.4 Current Evaluation Metrics

- **JudgeLM:** An LLM-based ranking method for evaluating automatic counter-narrative generation (Zubiaga, 2024).
- **BLEU:** Measures token overlap between predictions and references (Papineni et al., 2002).

¹<https://huggingface.co/HiTZ/judge-eus>

- ROUGE-L: Computes sentence-level structure similarity and longest co-occurring n-grams (Lin, 2004).
- BERTScore: Calculates token-level similarity using contextual embeddings (Zhang et al., 2019).
- Novelty: Measures the proportion of non-singleton n-grams in generated text that do not appear in the training data (Wang and Wan, 2018).
- Genlen: The average length of generated predictions.

3 Methodology

We employed a three-stage approach for generating effective CS across multiple languages. The first stage utilizes a simulated annealing approach combined with LLMs to generate and select diverse responses. A word sampling mechanism extracts vocabularies from both predefined word lists and input text (HS and CN) to enrich the response generation. Each candidate response is evaluated using JudgeLM, with scores exponentially weighted to guide the sampling process. The second stage implements a Round-robin tournament evaluation system to rank and select the most effective responses, ensuring high-quality output even in low-resource language settings. In the last stage, we combined each of the first-ranked, second-ranked, third-ranked, and fourth-ranked answers into their own CSVs and then ran the evaluation script given by MCG-COLING² to find the top 3 runs.

3.1 Comparison to other methods

Our approach of using JudgeLM to rank candidate counter-narratives (CNs) parallels prior methods proposed by Zubiaga et al. (2024), where Large Language Models (LLMs) generate CNs and an LLM-based evaluator selects the best response via pairwise comparisons. While both systems rely on tournament-style evaluation using an LLM judge, our framework fundamentally diverges by incorporating a simulated annealing stage before the round-robin tournament to generate a set of iteratively refined answers to compare. Instead of generating each CN in a single pass, we repeatedly mutate, expand, and re-score candidate responses,

²https://github.com/hitz-zentroa/eval-MCG-COLING-2025/blob/master/evaluation/bash/judge_full_pipeline.sh

enabling a broader exploration of the CN space and reducing the risk of local optima. By adjusting a temperature parameter, even lower-scoring CNs can remain viable candidates at early stages, fostering globally stronger outputs. Furthermore, we augment vocabulary sampling with tokens from both predefined lists and the original hate speech to try to create contextually grounded answers.

3.2 Data Used

As our model did not require training or any other outside data, we elected to only use the testing set (100 HS-CN pairs for each language) of the data provided to generate our answers. It was determined that the use of the development and training subsets of data was not necessary as we could directly test the quality of generated answers using JudgeLM.

3.3 Stage 1: Counterspeech Generation

In this stage, we implemented a simulated annealing algorithm (Algorithm 1 from the appendix) to generate effective counterspeech (CS) responses to hate speech (HS) instances across multiple languages. The algorithm iteratively refines CS candidates by exploring the search space in a manner inspired by thermodynamic annealing processes.

The algorithm begins with an initial CS candidate c_0 , which can be the HS instance h itself or another string. In our case, we used the background knowledge provided with each HS-CS pair from MCG-COLING as the initial string. At each iteration, we update the temperature parameter T by an increment ΔT , controlling the exploration-exploitation trade-off.

For each candidate CS c in the current set C , we generate a set of new candidates S by appending randomly sampled words from a language-specific word list. Notably, part of the vocabulary that we sample from is from the tokenized HS. This sampling enriches the vocabulary and introduces relevant words from the original HS in the candidate responses.

We evaluate each candidate $c' \in S$ using an LLM-based judge to obtain a score $E(c')$, reflecting its relevance, fluency, and effectiveness as a counterspeech. To prioritize higher-quality candidates while still allowing exploration of the search space, we compute selection probabilities using a Boltzmann-like distribution (Algorithm 2 from the appendix):

$$P(c') = \frac{T^{E(c')}}{\sum_{c'' \in S} T^{E(c'')}}.$$

This probability distribution ensures that candidates with higher scores are more likely to be selected, but candidates with lower scores still have a chance of being chosen, especially at lower temperatures. This mechanism allows the algorithm to avoid local optima by occasionally exploring less promising candidates.

We select k candidates from S based on the computed probabilities $P(c')$. For each selected candidate c' , we generate new CS responses \tilde{S} using Language Models (LLMs). These LLM-generated responses further diversify the candidate pool and introduce potentially high-quality CS that may not be reachable through simple word appending. The exact LLMs used for counterspeech generation can be found in the appendix in Table 2.

The new candidates $\tilde{c} \in \tilde{S}$ are evaluated, and their probabilities are computed in the same manner. We update the candidate set C with the top candidates from \tilde{S} based on their probabilities. This process is repeated for a predefined number of iterations or until a candidate reaches the target score S_{target} .

We optimized these hyper-parameters by experimenting on a small subset of 4 HS instances (one from each language) and measuring the average high score achieved. The results of this hyper-parameter tuning are presented in Table A.2 in the Appendix. We observed that increasing the number of iterations and candidates per loop improved the average high score, with the combination of 8 iterations and 6 candidates per loop achieving an average high score of 10, which meets our target score threshold.

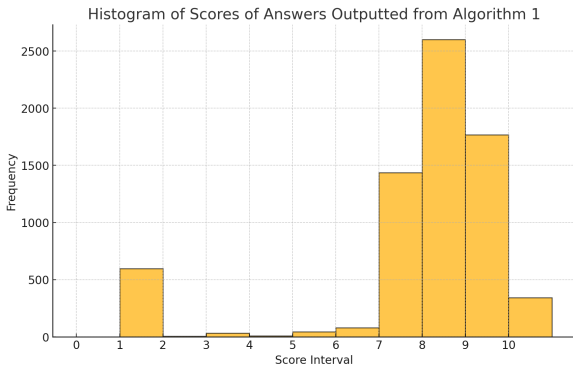


Figure 1: A histogram of scores $E(c)$ for every counterspeech generated from Algorithm 1

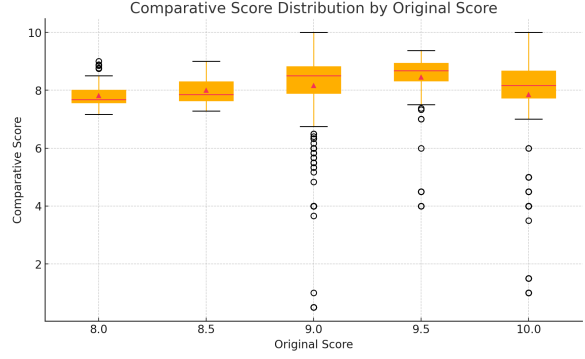


Figure 2: Box and whisker charts that compare the original scored value of a CS answer from stage 1 to the re-scored vales from stage 2.

We used these optimized parameters to generate sentences for MCG-COLING 2025 where we set the target scores of each CS answer (S_{target}) to be 9. After generating answers for each of the 400 total instances of hate speech, we were left with a set of a 6,915 answers. The distribution of the initial scores for each answer can be found in Figure 1. As can be seen, most answers achieved scores of 8 or higher from JudgeLM at this stage.

Assuming each iteration has k CS options, each of which makes k LLM generation calls that produce n answers, the worst-case time complexity of this algorithm would be $O(N_{\text{max}} \cdot n \cdot k^2)$ calls to JudgeLM.

3.4 Stage 2: Round Robin Ranking

While the simulated annealing algorithm can consistently generate answers that score a 10, these answers are not necessarily the best possible responses. Since JudgeLM comparatively ranks answers, the score of an answer can be affected by the quality of the other candidates' answers. This phenomenon is exemplified in Figure 2, which shows that there is a large variance (especially among high-scoring answers) between the score a specific counterspeech (c) received in the original algorithm ($E(c)$) and the score it received when recalculated in comparison to other high-scoring answers (RoundRobin(c, C)) where C is a set of possible CS answers for a given HS.

Ideally, we would compare every answer with every other answer at the end of each loop of the simulated annealing algorithm to identify the best responses. However, repeatedly performing these comprehensive comparisons in each iteration would cause the time complexity to skyrocket to $O(N_{\text{max}} \cdot n^2 \cdot k^4)$, where N_{max} is the maximum

number of iterations, n is the number of candidates per loop, and k is the number of candidates selected for further exploration. Such computational demands are impractical for our application.

Therefore, we implement a post-hoc algorithm to compare all the generated high-scoring answers and determine which ones are the best. The round-robin algorithm (Algorithm 3 from the appendix) addresses this by taking the final set of high-scoring answers from Algorithm 1 and assigning each an average score based on pairwise comparisons. Each response in the answer pool is evaluated against every other response using JudgeLM, creating a tournament-style evaluation where each pair of responses is compared twice by swapping their positions to reduce position bias. As this algorithm has a quadratic time complexity, we only used the top 6 answers in the calculation. These comparison scores are accumulated and averaged across all matches and then divided by the total number of matches to generate an average scoring. This was used to provide a robust overall ranking for each counterspeech response. The top-four ranked answers were then stored.

3.5 Stage 3: Generation of Submission Files

Upon completion of the ranking process, four files (one for each rank of answer) were generated for each language’s answers. A final sanity check was completed for each language by comparing each of the 4 files with a round-robin based scoring algorithm from MCG-COLING. For English, Spanish, and Italian, results were consistent with expectations; files containing higher ranking answers received a higher score. As seen in Figure 7, run 4, despite being made of 4th place answers, scored 3.5 points higher than the file for the second rank answers. As such, the final submission for Basque in the 2025 MCG-COLING task contained the first, second, and fourth rank answers.

4 Conclusion

This study showcases the effectiveness of the CODEOFCONDUCT model in crafting context-aware counterspeech, achieving exceptional results in both high- and low-resource language settings. By integrating a simulated annealing-based generation framework with a robust round-robin ranking mechanism, our approach secured leading positions in the MCG-COLING-2025 shared task across four diverse languages.

As evidenced by the quantitative results in Figure 7, our model’s success with Basque—a language with limited NLP resources such as annotated datasets, pretrained models, and linguistic tools—stands out as a key achievement. This success stems from three key factors: (1) our simulated annealing approach’s effectiveness in handling Basque’s complex agglutinative morphology, (2) a word sampling strategy specifically enhanced for low-resource scenarios by incorporating domain-specific terminology, and (3) an evaluation system well-suited to Basque’s unique linguistic features. All three of the runs submitted were able to outclass all the other runs submitted by other groups. By working with only 100 HS-CN pairs for each language and leveraging multilingual pretrained models alongside innovative optimization techniques, we demonstrated how thoughtful methodologies can overcome resource constraints. Our sweeping success in Basque, coupled with strong rankings in Italian, English, and Spanish, highlights the versatility of our approach in navigating diverse linguistic and cultural challenges, with particular effectiveness in tackling the unique demands of low-resource languages.

Through the use of comprehensive evaluation metrics, including JudgeLM and traditional measures such as BLEU and ROUGE, we ensured that the generated counterspeech was both linguistically accurate and contextually appropriate. The combination of simulated annealing and round-robin evaluation is particularly well-suited for this task for several reasons: simulated annealing enables exploration of culturally appropriate responses through temperature-controlled sampling, while round-robin evaluation captures the nuanced effectiveness of counterspeech across different cultural contexts that simple metrics might miss. The final results highlighted the model’s strengths while also revealing opportunities for improvement, particularly in adapting evaluation frameworks to better reflect the nuances of multilingual and culturally specific outputs.

5 Limitations

Static Response Generation

Our CS generation system, while effective in current contexts, faces inherent temporal limitations. The greedy annealing algorithm’s reliance on static word lists and predefined evaluation metrics constrains its ability to adapt to rapidly evolving hate

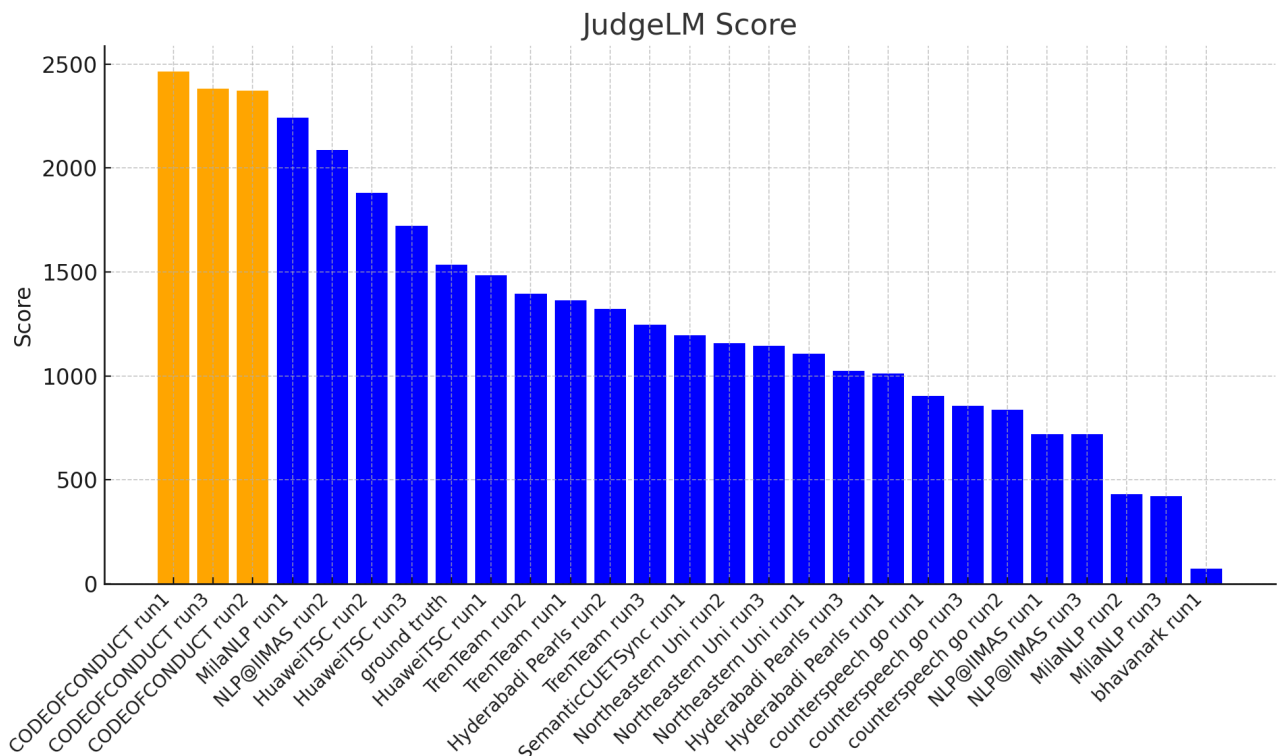


Figure 3: Chart depicting the JudgeLM scores for each Basque run. Bars drawn in yellow represent the results from the CODEOFCONDUCT submission.

speech patterns. As Lupu et al. (2023) highlighted, hate speech can shift dramatically with offline events, challenging our model’s static approach. Vidgen and Derczynski (2020) emphasized how the "garbage in, garbage out" principle affects such systems, highlighting the need for more dynamic algorithmic solutions that can adapt to emerging patterns in real-time.

Local Optimality

While the simulated annealing approach ensures rapid initial solution generation and efficient task distribution, it may not achieve global optimality. This occurs when the temperature decreases too rapidly or when the initial sampling conditions restrict exploration of the full solution space. Although we mitigate these issues through parameter tuning, the fundamental trade-off between exploration and exploitation remains a challenge.

Computational Cost

The method of combining simulated annealing and round-robin evaluation introduces significant computational cost. Using an NVIDIA A100 GPU, processing one language’s test set requires approximately 10 hours of computation time, primarily

due to two factors: (1) the quadratic number of calculations needed for each simulated annealing iteration, and (2) the quadratic time complexity of round-robin evaluation required for generating high-quality responses. High computational cost makes it challenging to meet the demands of real-world applications.

Evaluation Metrics

While JudgeLM demonstrates competence in English response evaluation, its lack of fine-tuning on multilingual counterspeech data affects its reliability. Basque responses require a separate JudgeEUS model, leading to potential inconsistencies in evaluation standards across languages. These models may overemphasize lexical similarities while missing cultural nuances and language-specific expressions, potentially leading to responses that score well numerically but fail to resonate with speakers of non-English languages.

References

Rodrigo Agerri and et al. 2020. Give your text representation models some love: the case for Basque. In *Proceedings of the Twelfth Language Resources and*

- Evaluation Conference*, pages 4781–4788, Marseille, France. European Language Resources Association.
- Bianca Cepollaro, Maxime Lepoutre, and Robert Mark Simpson. 2023. [Counterspeech](#). *Philosophy Compass*, 18(1):e12890.
- 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 (ACL)*, pages 8440–8451.
- Ona de Gibert, Naiara Perez, Aitor García-Pablos, and Montse Cuadros. 2018. [Hate speech dataset from a white supremacy forum](#). In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, pages 11–20, Brussels, Belgium. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186.
- Aamir Faisal, Muhammad Bilal Zafar, and Muhammad Imran Malik. 2021. Low-resource multilingual hate speech detection: A comparative study of classical and neural approaches. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)*, 20(6):1–20.
- Margherita Fanton, Helena Bonaldi, Serra Sinem Tekiroğlu, 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.
- Siqi Feng, Junjie Hu, Graham Neubig, and Yiming Yang. 2021. A survey of data augmentation approaches for nlp. In *Proceedings of the 2021 Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 45–55.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81.
- Yonatan Lupu, Richard Sear, Nicolas Velásquez, Rhys Leahy, Nicholas Johnson Restrepo, Beth Goldberg, and Neil F Johnson. 2023. [Offline events and online hate](#). *PloS one*, 18(1):e0278511.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pages 311–318.
- Aashima Poudhar, Ioannis Konstas, and Gavin Abercrombie. 2024. [A strategy labelled dataset of counterspeech](#). In *Proceedings of the 8th Workshop on Online Abuse and Harms (WOAH 2024)*, pages 256–265, Mexico City, Mexico. 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 (EMNLP)*, pages 4757–4766.
- Bertie Vidgen and Leon Derczynski. 2020. [Directions in abusive language training data: Garbage in, garbage out](#). *CoRR*, abs/2004.01670.
- Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019. [Challenges and frontiers in abusive content detection](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 80–93, Florence, Italy. Association for Computational Linguistics.
- Ke Wang and Xiaojun Wan. 2018. [Sentigan: Generating sentimental texts via mixture adversarial networks](#). In *International Joint Conference on Artificial Intelligence*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019. [Bertscore: Evaluating text generation with BERT](#). *CoRR*, abs/1904.09675.
- Arkaitz Zubiaga. 2024. Judgelm: A llm-based ranking method for evaluating automatic counter-narrative generation. In *Proceedings of COLING 2024*.
- Irune Zubiaga, Aitor Soroa, and Rodrigo Agerri. 2024. [Ixa at refutes 2024: Leveraging language models for counter narrative generation](#). In *IberLEF@SEPLN*.

A Appendix

A.1 Algorithms

Algorithm 1: Simulated Annealing for Counterspeech Generation

Data: Hate speech h , initial counterspeech c_0 , target score S_{target} , max iterations N_{max} , candidates per loop k , temp increment ΔT , initial temp T_0

Result: Optimal counterspeech c^*

```

1 Initialize  $C \leftarrow \{c_0\}$ ,  $T \leftarrow T_0$ ;
2 for  $i = 1$  to  $N_{\text{max}}$  do
3    $T \leftarrow T + \Delta T$ ;
4    $C_{\text{new}} \leftarrow \emptyset$ ;
5   foreach  $c \in C$  do
6     Generate candidates  $S$  by appending
       random words to  $c$ ;
7     Evaluate scores  $E(S)$  using LLM
       judge;
8     Compute probabilities  $P(S)$  using
       Algorithm 2;
9     Select  $k$  candidates from  $S$  based on
        $P(S)$ ;
10    Generate new counterspeeches  $\tilde{S}$ 
       using LLMs on selected
       candidates;
11    Evaluate scores  $E(\tilde{S})$ ;
12    Compute probabilities  $P(\tilde{S})$ ;
13    Add top  $k$  candidates from  $\tilde{S}$  to
        $C_{\text{new}}$  based on  $P(\tilde{S})$ ;
14  end
15  Update  $C \leftarrow C_{\text{new}}$ ;
16  if  $\exists c \in C$  such that  $E(c) \geq S_{\text{target}}$  then
17    return  $c^*$  with highest  $E(c)$  in  $C$ ;
18  end
19 end
20 return  $c^*$  with highest  $E(c)$  in  $C$ ;

```

Algorithm 2: Compute Probabilities

Data: Candidates S , scores $E(S)$, temperature T

Result: Probabilities $P(S)$

```

1 foreach  $c' \in S$  do
2   Compute probability:

```

$$P(c') = \frac{T^{E(c')}}{\sum_{c'' \in S} T^{E(c'')}}$$

```

3 end

```

Algorithm 3: Round Robin Ranking for Counterspeech Evaluation

Data: Counterspeech c , set of other counterspeeches C

Result: Average score of c

```

1 Initialize  $total\_score \leftarrow 0$ ;
2 foreach  $a \in C$  do
3   Create question comparing  $c$  and  $a$ ;
4    $normal\_results \leftarrow$  Evaluate  $c$  vs.  $a$ 
       (normal order) using JudgeLM;
5    $reversed\_results \leftarrow$  Evaluate  $a$  vs.  $c$ 
       (reversed order) using JudgeLM;
6    $score_c \leftarrow$ 
        $normal\_results["output1"] +$ 
        $reversed\_results["output2"]$ ;
7    $total\_score \leftarrow total\_score + score_c$ ;
8 end
9 Compute
    $average\_score \leftarrow total\_score / (2 \times |C|)$ ;
10 return  $average\_score$ ;

```

A.2 Data

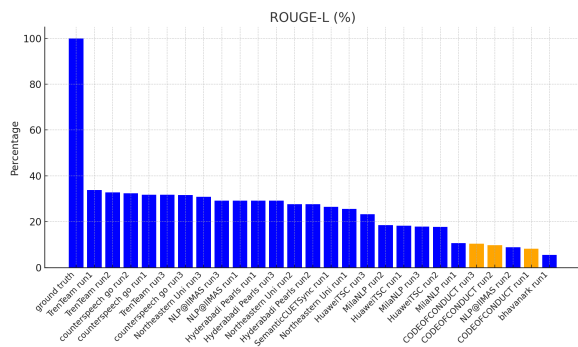


Figure 4: Chart depicting the ROUGE-L scores for each Basque run. Bars drawn in yellow represent the results from the CODEOFCONDUCT submission.

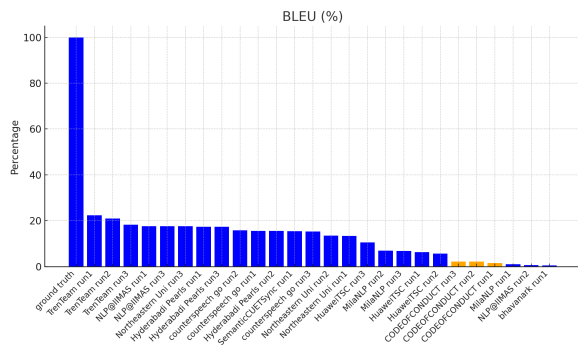


Figure 5: Chart depicting the BLEU scores for each Basque run. Bars drawn in yellow represent the results from the CODEOFCONDUCT submission.

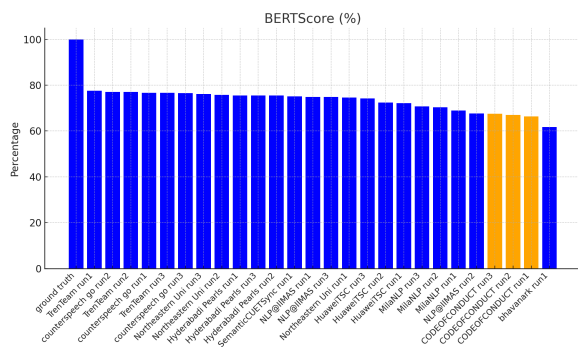


Figure 6: Chart depicting the BERT scores for each Basque run. Bars drawn in yellow represent the results from the CODEOFCONDUCT submission.

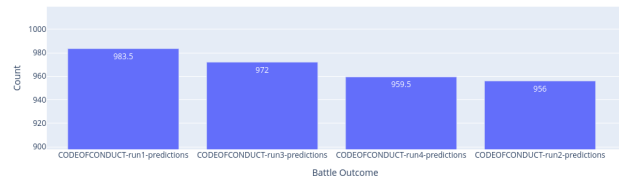


Figure 7: Scoring that compares the 4 runs of generated HS-CS files for Basque

Iterations	Candidates per Loop	Average High Score
2	2	8.875
4	2	9.375
6	2	9.000
8	2	9.500
2	4	9.125
4	4	9.875
6	4	9.875
8	4	9.500
2	6	9.625
4	6	9.625
6	6	9.750
8	6	10.000

Table 1: Hyper-parameter Tuning Results: Average High Scores

Model Name	Version	Parameters	HuggingFace Link
Hermes	3	8B	NousResearch/Hermes-3-Llama-3.1-8B
Zephyr	Beta	7B	HuggingFaceH4/zephyr-7b-beta
Meta-Llamaz	3	8B	NousResearch/Hermes-3-Llama-3.1-8B
Llama	3 Instruct	8B	meta-llama/Meta-Llama-3-8B-Instruct
Nous Hermes	2	7B	NousResearch/Nous-Hermes-2-Mixtral-8x7B-DPO
Llama	3.1	70B	meta-llama/Llama-3.1-70B-Instruct
Qwen	2.5 Instruct	72B	Qwen/Qwen2.5-72B-Instruct

Table 2: Model used for CS generation.