

INFOTEC-NLP at HOMO-LAT 2025: Testing a Novel Multi-region Spanish Model to Monitor Opinion in Latin American LGBTQI+ Social Media

Guillermo Ruiz¹, Rogelio Campos², Tania Ramirez-delreal^{3,4}, Daniela Moctezuma⁴, Mario Graff^{1,3} and Eric Sadit Tellez^{1,3,*}

¹INFOTEC Centro Público de Investigación en Tecnologías de la Información y Comunicación, 112 Circuito Tecnopolo Sur, Parque Industrial Tecnopolo 2, Aguascalientes, 20326, México.

²Independent researcher, Puebla, México.

³Secretaría de Ciencia, Humanidades, Tecnología e Innovación (SECIHTI), 1582 Insurgentes Sur 1582, Crédito Constructor, Ciudad de México, 03940 México

⁴Centro de Investigación en Ciencias de Información Geoespacial (CentroGeo), Aguascalientes, Ags., 20213, Mexico

Abstract

This paper presents our solution notebook for the HOMO-LAT@IberLEF2025 challenge, which consists of two tasks that require models that take advantage of subtle differences in language between Spanish-speaking countries to identify the polarity of LGBTQI+ discourse in Reddit posts coming from Latin American Spanish-speaking countries. Our approach uses our BILMALAT model, a novel region-aware Large Language Model based on RoBERTa that is designed to understand and process Spanish variants from several countries, making it uniquely suited for capturing the diverse linguistic and cultural nuances present in the HOMO-LAT dataset. We show the high performance of BILMALAT in this challenge, demonstrating its efficacy in *opinion mining* analysis. Our solution achieved the third position in both HOMO-LAT Tasks 1 and 2, solely fine-tuning the model for the task without extra strategies to balance the data and tackling large messages beyond the capacity of our model.

Keywords

hate and hope speech identification, multi-dialect Spanish language models, text classification models

1. Introduction

Sentiment analysis, also known as opinion mining, applied to the online discourse on the LGBTQI+ community in Latin America, provides an indispensable framework for comprehending societal perceptions and attitudes. Social media platforms, e.g., Reddit, which facilitate discussions, represent extensive repositories of public sentiment, encompassing a spectrum from supportive to explicitly harmful. Examining this content is essential to identify the predominant biases, monitor the dissemination of discriminatory language, and evaluate the effectiveness of inclusion initiatives in various Spanish-speaking contexts. This paper outlines our participation in the HOMO-LAT 2025 [1] shared task from IberLEF [2], specifically addressing this significant need by focusing on the polarity analysis of LGBTQI+-related terms in Latin American Spanish.

1.1. Related work

To promote safer online environments, Bel-Enguix et al. [3] introduced HOMOMEX, a shared task that seeks computational models for automatically identifying hate speech against LGBTQI+ people on social networks, particularly focusing on Mexican Spanish tweets. Subsequently, Gómez-Adorno et al.

IberLEF 2025, September 2025, Zaragoza, Spain

*Corresponding author.

✉ luis.ruiz@infotec.mx (G. Ruiz); rcamp.ser@gmail.com (R. Campos); tramirez@centrogeo.edu.mx (T. Ramirez-delreal); dmoctezuma@centrogeo.edu.mx (D. Moctezuma); mario.graff@infotec.edu.mx (M. Graff); eric.tellez@infotec.mx (E. S. Tellez)

🌐 <https://sadit.github.io/> (E. S. Tellez)

🆔 0000-0001-7422-7011 (G. Ruiz); 0000-0002-1638-5086 (T. Ramirez-delreal); 0000-0003-3038-4642 (D. Moctezuma); 0000-0001-6573-4142 (M. Graff); 0000-0001-5804-9868 (E. S. Tellez)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

[4] presented a second edition of the challenge that maintained its focus on the Mexican Spanish dialect while expanding the variety of message sources by incorporating song lyrics. The present iteration has evolved into HOMO-LAT, a task centered on identifying polarity across various Latin American countries using Reddit posts instead of tweets. This new version presents distinct challenges, regarding multi-region data, where the training and test sets are sourced from the same countries, and cross-region data, where the training and test sets comprise different countries. In HOMO-LAT, regular messages indicate the country of origin, which allows the prediction model to take advantage of this information.

When trained on extensive text datasets, it is acknowledged that transformer-based language models (LMs) [5] are highly effective for linguistic modeling; BERT [6] and RoBERTa [7] are notable examples of such models. Although the initial training of these models is expensive, they can be efficiently fine-tuned for specific applications once pre-trained. Original BERT focuses on the English language and provides a multilingual model that covers Spanish. Since BERT models are highly dependent on how they were trained, including the data presented, several Spanish models have been introduced, improving performance in most tasks. For example, BETO [8], BERTIN [9], and RoBERTuito [10]. These models are highly performant and work for tasks in generic Spanish. However, experiments like HOMO-LAT show that regional information is important for several domains, and we need more sophisticated models, as noted in [11], to be able to consider regional diversity.

In addition, when considering the study of specific language variants, academic investigations such as those cited in [12] and [13] have examined regional variations in the Spanish language among different countries. The conclusions of these studies underscore the need for further initiatives to improve the Spanish language resources. Tellez et al. [11] present several lexical and regional word embeddings and BERT-based models tailored to Spanish-speaking countries. The study emphasizes that training models in the same language variant as the dataset defining a task results in improved performance. Jorge Ortiz-Fuentes¹ has fine-tuned BETO models specifically for the Chilean Spanish, reinforcing the idea that models designed for specific regions perform better in regional tasks.

1.2. Task description

The HOMO-LAT of IberLEF 2025 asks systems to accurately predict the polarity of a Reddit post, which involves determining whether the post is *positive*, *negative*, or *neutral* regarding the LGBTQI+ community. These posts are composed in Spanish by users from various Latin American nations. Solutions should consider the regional context to account for local references and regionalisms that might offer different interpretations, making it crucial to identify the intent of these communications. The posts are categorized into three types: identification of hate speech, hopeful messages, and sentiment analysis. The task comprises two subtasks: Task 1 was assessed using a test dataset from the same countries as in the training phase. Task 2 used a test consisting of messages from Bolivia, Costa Rica, Cuba, the Dominican Republic, Ecuador, El Salvador, Guatemala, Honduras, Nicaragua, Panama, Paraguay, Peru, Puerto Rico, Uruguay, and Venezuela.

The training set contains 5,700 Reddit posts and the development set includes 1,400 messages. Each post has at least one keyword from the following: *trans*, *lgbt*, *gay/gays*, *lesbiana/as*, *bisexual/es*, *asexual/es*, *transexual/les*, *travesti/is*, *queer/s*, *transgénero*, *pansexual/es*, *intersexual/es*. The training has 3174 posts classified as *neutral*, 1960 as *negative*, and 633 as *positive*. For the development partition, 888 are *neutral*, 475 *negative*, and 80 *positive*.

Roadmap

This section introduces the underlying opinion mining problem behind HOMO-LAT@IberLEF2025 and presents a review of the related work. Section 2 briefly describes our BILMALAT model and describes how we use it to solve the HOMO-LAT tasks. Section 3 details our experimental methodology and lists our results. Finally, conclusions are given in Section 4.

¹Models and comparisons can be accessed at <https://huggingface.co/dccuchile/patana-chilean-spanish-bert> and <https://huggingface.co/dccuchile/tulio-chilean-spanish-bert>.

Country	Token
Argentina	_ar
Colombia	_co
Chile	_cl
Mexico	_mx
Uruguay	_uy
Panama	_pa
Peru	_pe
Paraguay	_py
Guatemala	_gt
Venezuela	_ve
Honduras	_hn
El Salvador	_sv
Republica Dominicana	_do
Bolivia	_bo
Ecuador	_ec
Nicaragua	_ni
Costa Rica	_cr
Cuba	_cu

Table 1

Tokens indicating the country of the message; the current model lacks representation for Puerto Rico due to some issues with our training procedure.

2. Our system solution

Our approach was to fine-tune an early bird version of our BILMALAT² model. BILMALAT uses a RoBERTa-based training strategy on 600 million tweets from Latin America countries, trained from scratch. The model was designed such that the where and when of a given text is captured, i.e. it learns about regions and periods of time. Our first real-world application is to use it in HOMO-LAT; the model can be accessed and tested with the Huggingface framework³.

For BILMALAT to be region- and period-aware, we use additional tokens to indicate the region. The first token corresponds to the country, the exact list is shown in Table 1. The second and third tokens correspond to the year and month of publication. The tokens can be *_2015, _2016, ..., _2023* and *_01, ..., _12* for the year and month, respectively. The rest of the tokens constitute the actual post.

Here, we describe our solution. We first used the BETO⁴ model as a baseline. We fine-tuned the model in the training set for 6 epochs with a learning rate of 10^{-5} . We conducted various experiments to choose these hyperparameters.

We then fine-tuned the BILMALAT with the messages provided. Since the training data included the country of origin, we could include the first token, for the second and third we simply used a generic *x* and *_01* to force the model to predict without temporal bias⁵. We fine-tuned the model using different hyperparameters.

3. Results

Our experiments were carried out on Google Colab⁶. We employed the Huggingface framework with a Pytorch back-end. Our BETO baseline obtained a best accuracy of 0.5385 and an F1-score of 0.4487.

²<https://huggingface.co/guillermoruiz/bilmaLAT>

³<https://huggingface.co/>

⁴<https://huggingface.co/dccuchile/bert-base-spanish-wwm-cased>

⁵For instance, our internal tests show that our model is aware of different meanings for terms like *infección* in COVID times and other periods; we need much more research to take advantage of this extra-information.

⁶Google Colab site <https://colab.research.google.com/>.

Test	Learning rate	Epoch	Best epoch	Validation loss	Acc.	F1-Macro
1	10^{-4}	9	4	1.2860	0.5350	0.4113
2	10^{-4}	9	3	1.1139	0.5613	0.4194
3	10^{-5}	9	7	1.1222	0.5482	0.4403
4	10^{-5}	9	5	1.0866	0.5738	0.4493
5	10^{-5}	9	4	1.0467	0.5398	0.4516
6	10^{-5}	9	3	0.9242	0.5849	0.4643
7	10^{-6}	10	9	0.8220	0.6091	0.3748

Table 2

Results on the development set.

Our model selection procedure was quite simple, focused on fine-tuning the HOMO-LAT base-model; using the given training and development sets, we varied the learning rate in 10^{-4} , 10^{-5} , and 10^{-6} ; and looking for the best model among the computed epochs. This process was not systematic and was performed manually due to the lack of time. Table 2 displays these results for Task 1. We found that models with 10^{-5} show a higher performance, notoriously a learning rate of 10^{-6} produced very low scores, which means that we should give more epochs, but we were unable to do it due to our late joining the challenge. We did not fine-tune for task 2 and use the same model.

Official Results

Our BILMALAT achieved a third place of seven in the gold standard with an F1 score of 0.5137 vs. 0.5296 of the best system. Our system also achieved third place (of seven) with an F1 score of 0.4639, where an F1 score 0.5086 was the best achieved.

4. Conclusions

This challenge was a perfect scenario to test our BILMALAT and compare with the popular alternative BETO model. Our approach, a language model enriched with geographical and cultural context, demonstrated a clear advantage in navigating the subtle expressions of sentiment prevalent in the dataset due to the message’s origin. Our model supports up to 128 input tokens, which might be enough for tweets, but the Reddit posts used in this task were much longer, which gives them greater complexity.

This result underscores the need for fine-grained regional awareness in natural language processing, moving beyond generic linguistic models to include local information in the training data. The insights gained from this competition provide a valuable contribution to the field of sentiment analysis and open the way for more sophisticated and culturally sensitive AI applications.

Future work

Due to our delayed awareness of the challenge, we had limited time to explore advanced strategies for task management. Consequently, we could not test sophisticated techniques like data augmentation or other methods to address the dataset’s inherent imbalance and the uneven distribution between training and development data. Additionally, similar to other BERT-based approaches, our model is constrained by sentence size, meaning it cannot handle huge messages without additional strategies, i.e., BILMALAT requires additional strategies to process large messages. With further focus on addressing the unique features of the challenge and overcoming the limitations of BERT-based models, our results can be enhanced.

Declaration on Generative AI

During the preparation of this work, we applied the Writefull's model for grammar and spelling checks. After using these services, we reviewed and edited the content as needed and assume full responsibility for the content of the publication.

References

- [1] G. Bel-Enguix, H. Gómez-Adorno, S. Ojeda-Trueba, G. Sierra, J. Barco, E. Lee, J. Dunstan, R. Manrique, Overview of HOMO-LAT at IberLEF 2025: Human-centric polarity detection in Online Messages Oriented to the Latin American-speaking lgbtq+ populaTion, *Procesamiento del lenguaje natural* 75 (2025) –.
- [2] J. Á. González-Barba, L. Chiruzzo, S. M. Jiménez-Zafra, Overview of IberLEF 2025: Natural Language Processing Challenges for Spanish and other Iberian Languages, in: *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2025)*, co-located with the 41st Conference of the Spanish Society for Natural Language Processing (SEPLN 2025), CEUR-WS.org, 2025.
- [3] G. Bel-Enguix, H. Gómez-Adorno, G. Sierra, J. Vázquez, S. T. Andersen, S. Ojeda-Trueba, Overview of homo-mex at iberlef 2023: Hate speech detection in online messages directed towards the mexican spanish speaking lgbtq+ population, *Procesamiento del lenguaje natural* 71 (2023) 361–370.
- [4] H. Gómez-Adorno, G. Bel-Enguix, H. Calvo, S. Ojeda-Trueba, S. T. Andersen, J. Vázquez, T. Alcántara, M. Soto, C. Macias, Overview of homo-mex at iberlef 2024: Hate speech detection towards the mexican spanish speaking lgbt+ population, *Procesamiento del Lenguaje Natural* 73 (2024) 393–405.
- [5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, *Advances in neural information processing systems* 30 (2017).
- [6] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805* (2018).
- [7] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, 2019. URL: <https://arxiv.org/abs/1907.11692>. arXiv:1907.11692.
- [8] J. Cañete, G. Chaperon, R. Fuentes, J.-H. Ho, H. Kang, J. Pérez, Spanish pre-trained bert model and evaluation data, in: *PML4DC at ICLR 2020*, 2020.
- [9] J. D. la Rosa y Eduardo G. Ponferrada y Manu Romero y Paulo Villegas y Pablo González de Prado Salas y María Grandury, Bertin: Efficient pre-training of a spanish language model using perplexity sampling, *Procesamiento del Lenguaje Natural* 68 (2022) 13–23. URL: <http://journal.sepln.org/sepln/ojs/ojs/index.php/pln/article/view/6403>.
- [10] J. M. Pérez, D. A. Furman, L. Alonso Alemany, F. M. Luque, RoBERTuito: a pre-trained language model for social media text in Spanish, in: *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, European Language Resources Association, Marseille, France, 2022, pp. 7235–7243. URL: <https://aclanthology.org/2022.lrec-1.785>.
- [11] E. S. Tellez, D. Moctezuma, S. Miranda, M. Graff, G. Ruiz, Regionalized models for spanish language variations based on twitter, *Language Resources and Evaluation* 57 (2023). URL: <https://doi.org/10.1007/s10579-023-09640-9>. doi:10.1007/s10579-023-09640-9.
- [12] S. Jimenez, G. Dueñas, A. Gelbukh, C. A. Rodriguez-Diaz, S. Mancera, Automatic detection of regional words for pan-hispanic spanish on twitter, in: G. R. Simari, E. Fermé, F. Gutiérrez Segura, J. A. Rodríguez Melquiades (Eds.), *Advances in Artificial Intelligence – IBERAMIA 2018*, Springer International Publishing, Cham, 2018, pp. 404–416.
- [13] C. España-Bonet, A. Barrón-Cedeño, Elote, choclo and mazorca: on the varieties of Spanish, in: K. Duh, H. Gomez, S. Bethard (Eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Vol-*

ume 1: Long Papers), Association for Computational Linguistics, Mexico City, Mexico, 2024, pp. 3689–3711. URL: <https://aclanthology.org/2024.naacl-long.204>. doi:10.18653/v1/2024.naacl-long.204.