

Fine-Tuned Large Language Models for Symptom Recognition from Spanish Clinical Text

Mai A. Shaaban¹, Abbas Akkasi^{2,*}, Adnan Khan², Majid Komeili², Mohammad Yaqub¹

¹Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE, ²School of Computer Science, Carleton University, Ottawa, CA.

*Corresponding author: +1- 613 222 8134, E-mail: abbasakkasi@cunet.carleton.ca

Abstract

The accurate recognition of symptoms in clinical reports is significantly important in the fields of healthcare and biomedical natural language processing. These entities serve as essential building blocks for clinical information extraction, enabling retrieval of critical medical insights from vast amounts of textual data. Furthermore, the ability to identify and categorize these entities is fundamental for developing advanced clinical decision support systems, aiding healthcare professionals in diagnosis and treatment planning. In this study, we participated in SympTEMIST – a shared task on detection of symptoms, signs and findings in Spanish medical documents. We combine a set of large language models finetuned with the data released by the task’s organizers.

Introduction

Biomedical Named Entity Recognition (BioNER) has emerged as a critical component in the realm of biomedical and clinical informatics, with the primary goal of extracting meaningful entities from vast repositories of textual clinical data. Among the diverse categories of entities in the biomedical domain, recognizing the mention of symptoms, including signs and findings from clinical reports, stands out as an indispensable and profoundly significant task (1). The recognition of symptoms, signs, and findings in clinical reports holds profound implications for various facets of healthcare and research, making it an essential area of focus in the BioNER domain. Clinical narratives, often recorded in the form of electronic health records (EHR), radiology reports, pathology reports, and medical literature, are a treasure trove of information regarding the health status of patients. These reports contain a wealth of information about observed clinical phenomena, the manifestations of diseases, and diagnostic clues essential for patient care, clinical decision-making, and medical research. Key drivers for the importance of recognizing these entities include symptoms, signs, and findings that are fundamental components of diagnosing and managing a patient's health condition (2).

Timely and accurate recognition of these entities can significantly enhance clinical decision support systems, enabling healthcare providers to make more informed decisions and improve patient outcomes. However, the early identification of symptoms and signs, especially in the context of infectious diseases and outbreaks, is critical for public health surveillance. In addition, recognizing symptoms and findings empowers researchers to systematically extract structured information, supporting large-scale analyses, epidemiological studies, and the development of novel medical insights (2, 3).

The SympTEMIST (3) track is coordinated by the natural language processing (NLP) for Biomedical Information Analysis team at the Barcelona Supercomputing Center and is endorsed by Spanish and European initiatives, including DataTools4Heart, AI4HF, BARITONE, and AI4ProfHealth. This event is scheduled to take place within the framework of the BioCreative 2023 conference. The SympTEMIST challenge encompasses three distinct subtasks, namely SymptomNER, SymptomNorm, and SymptomMultiNorm. In this study, we focused our participation exclusively on the initial subtask, which centers on the identification of symptoms within clinical reports. We addressed the challenge by combining the Language Model Models (LLMs) that had been finetuned using data specifically tailored to this task. The outcomes achieved may not be comparable to NER systems designed for general text domains. However, they offer valuable insights into the potential efficiency of ensemble models based on LLMs. This is especially true when these models are equipped with a diverse set of classifiers that have been finetuned for increased efficiency. The upcoming sections of this paper are organized as follows: we begin with a review of recent research related to the task. The subsequent section, "Materials and Methods" outlines our approach for addressing BioNER in Spanish clinical data along with the data splits, experimental setup, and hyperparameters. Following this, the "Results and Discussion" section is dedicated to analyzing the results obtained. Finally, the paper concludes with proposed future directions.

Related Work

Recent progress in NLP has led to significant breakthroughs in BioNER. LLMs including ChatGPT (4), BERT (5) and its derivatives (7– 10), have significantly impacted NER tasks. In (10), authors introduced an innovative deep language model that incorporates both syntactic and semantic analysis to extract patient-reported symptoms from clinical text. Transformers based NER model was devised to discern neurological signs and symptoms within text, subsequently mapping them to clinical concepts (11). LLMs such as RoBERTa (12) and DeBERTa (13), have pushed the state-of-the-art in BioNER.

LLMs encounter challenges in handling non-English languages like Spanish, the fourth most spoken language, due to limited NLP resources (14). Recent years have witnessed a surge in publications related to BioNER in Spanish. Significant advancements have primarily manifested through collaborative initiatives such as the IberLEF (15) and BioNLP open shared tasks (16). However, despite this growing interest, handling the Spanish, especially in clinical settings poses numerous challenges. Spanish, with its rich inflectional morphology indicating various syntactic, semantic, and grammatical aspects (such as gender, number, etc.), presents unique complexities.

Material and Methods

Our approach entails finetuning the LLMs, which include XLM-RL¹ and XLM-RB² (17), BBES³ and BBS⁴ (18), E5-B⁵ and E5-L⁶ (19). We also employ a simple ensemble technique called

¹ XLM-RL: XLM-RoBERTa Large

² XLM-RB: XLM-RoBERTa Base

³ BBES: bsc-bio-ehr-es

⁴ BBS: bsc-bio-es

⁵ E5-B: multilingual-e5-base

majority voting (MV). Opting for a solution among several alternatives can result in either an enhanced or deteriorated outcome of the MV technique, particularly if the majority of solutions are of lower quality.

XLM-RL (561M parameters) and XLM-RB (279M parameters) are multilingual masked language models, trained on a filtered CommonCrawl dataset of over two terabytes. We leverage BBS and BBES both pretrained on Spanish clinical data. BBS exclusively leverages biomedical resources and BBES incorporates both biomedical and EHR data. Extensive finetuning and evaluation on clinical NER tasks demonstrated their superior performance over other general-domain and domain-specific Spanish models. Lastly, we leverage E5-B (278M parameters) and E5-L (560M parameters) which are advanced text embedding models, trained through a contrastive approach using weak supervision. To mitigate risks associated with relying on a single model, we employ the MV technique, combining predictions of the six models to favor entity labels with the most frequent consensus across the ensemble.

Experimental Setup

The section provides an overview of the SympTEMIST corpus, preprocessing steps, dataset size, and model training configuration. The SympTEMIST corpus encompasses Spanish medical records, each uniquely identified by a filename. The dataset comprises a total of 744 examples. Each record is meticulously annotated, featuring an annotation ID, target label (SINTOMA), start and end positions within the text, and the corresponding text snippet. For preprocessing, we use the SpaCy tool for tokenization following the inside–outside–beginning (IOB) tagging scheme. We split the data into 95% training and 5% validation for our experiments. For testing, we received 250 examples from the SympTEMIST organizers. To finetune the models for submission, we utilized all 744 training examples, incorporating the 5% validation data. The hyperparameters for all six models were set as follows: a batch size of four, 70 training epochs in total, and a learning rate of $5e-05$, employing a linear scheduler. In the context of our experiments, we utilized PyTorch and Transformers as the primary software modules.

Results and Discussion

This section thoroughly analyzes validation and test data outcomes. As depicted in Figure 1, a consistent upward trend in performance was observed for the validation set as the number of epochs increased for all six models. This signifies that all models are learning and improving the ability to make accurate predictions. Notably, the XLM-RL model stood out by achieving the highest F1-score of 0.70. For evaluating on test data, we selected only three individual models out of the six models. The selection was based on the order of best-performing models on validation data, although validation data has a relatively smaller size than test data and may not fully capture the diversity inherent in the test data.

⁶ E5-L: multilingual-e5-large




Figure 1: The f1-score trend over train epochs for validation data

Table 1 indicates that the finetuned models predict a few false positives, but they miss more true positive entities on test data. Additionally, if the ensemble of models in the majority voting method generates diverse predictions, it can lead to some level of vote dilution. In other words, the majority voting may average out correct predictions with incorrect ones, resulting in a less accurate overall prediction. To mitigate this problem, a weighted majority voting Click or tap here to enter text. could be considered, where higher weights are assigned to top-performing models and lower weights to lower-performing models. Moreover, the observed results on test data with domain-specific models on Spanish clinical data highlight the pressing necessity for ongoing research and advancement in this area, as they outperform general-domain models.

Model Variant	Validation data			Test data		
	Precision	Recall	F1-score	Precision	Recall	F1-score
BBES	0.66	0.70	0.68	0.70	0.57	0.63
BBS	0.63	0.70	0.66	0.69	0.62	0.65
XLM-RL	0.70	0.71	0.70	0.62	0.50	0.56
MV	0.65	0.62	0.64	0.68	0.60	0.64

Table 2: Models performance on the validation and test data

Conclusion

In summary, the accurate identification of symptoms in clinical reports is crucial for healthcare and biomedical natural language processing. These entities are foundational for extracting vital medical insights from extensive text data. Recognizing and categorizing these entities is key for advanced clinical decision support systems to assist healthcare professionals. Our study, using domain-specific models fine-tuned with SympTEMIST data, demonstrates the urgency for continued research and progress in this field, as these models outperform general-domain ones when applied to Spanish clinical data. One potential future direction in this field involves augmenting the pool of LLMs with diverse architectures that have undergone extensive fine-tuning using substantial, similar datasets. The combination of their outcomes, taking into account their respective performance levels, is also worth considering.

References

1. Liang T, Xia C, Zhao Z, Jiang Y, Yin Y, Yu PS. Transferring From Textual Entailment to Biomedical Named Entity Recognition. *IEEE/ACM Trans Comput Biol Bioinform*. 2023 Jul 1;20(04):2577–86.
2. Li J, Wei Q, Ghiasvand O, Chen M, Lobanov V, Weng C, et al. A comparative study of pre-trained language models for named entity recognition in clinical trial eligibility criteria from multiple corpora. *BMC Med Inform Decis Mak* [Internet]. 2022 Sep 1 [cited 2023 Oct 16];22(3):1–10. Available from: <https://bmcmedinformdecismak.biomedcentral.com/articles/10.1186/s12911-022-01967-7>
3. Lima-López S, Farré-Maduell E, Gasco-Sánchez L, Rodríguez-Miret J, Krallinger M. Overview of SympTEMIST at BioCreative VIII: corpus, guidelines and evaluation of systems for the detection and normalization of symptoms, signs and findings from text. In: *Proceedings of the BioCreative VIII Challenge and Workshop: Curation and Evaluation in the era of Generative Models*. 2023.
4. Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, et al. Language Models are Few-Shot Learners. *Adv Neural Inf Process Syst* [Internet]. 2020 May 28 [cited 2023 Oct 18];2020-December. Available from: <https://arxiv.org/abs/2005.14165v4>
5. Devlin J, Chang MW, Lee K, Toutanova K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference* [Internet]. 2018 Oct 11 [cited 2023 Oct 19];1:4171–86. Available from: <https://arxiv.org/abs/1810.04805v2>
6. Wang G, Yang G, Du Z, Fan L, Li X. ClinicalGPT: Large Language Models Finetuned with Diverse Medical Data and Comprehensive Evaluation. 2023 Jun 16 [cited 2023 Oct 20]; Available from: <https://arxiv.org/abs/2306.09968v1>
7. Luo R, Sun L, Xia Y, Qin T, Zhang S, Poon H, et al. BioGPT: Generative Pre-trained Transformer for Biomedical Text Generation and Mining. *Brief Bioinform* [Internet]. 2022 Oct 19 [cited 2023 Oct 20];23(6). Available from: <https://arxiv.org/abs/2210.10341v3>
8. Gu Y, Tinn R, Cheng H, Lucas M, Usuyama N, Liu X, et al. Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing. *ACM Trans Comput Healthc* [Internet]. 2020 Jul 30 [cited 2023 Oct 20];3(1):24. Available from: <http://arxiv.org/abs/2007.15779>
9. Labrak Y, Bazoge A, Dufour R, Rouvier M, Morin E, Daille B, et al. DrBERT: A Robust Pre-trained Model in French for Biomedical and Clinical domains. 2023 Apr 3 [cited 2023 Oct 20];16207–21. Available from: <https://arxiv.org/abs/2304.00958v2>
10. Luo X, Gandhi P, Storey S, Huang K. A Deep Language Model for Symptom Extraction from Clinical Text and Its Application to Extract COVID-19 symptoms from Social Media. *IEEE J Biomed Health*

- Inform [Internet]. 2022 Apr 1 [cited 2023 Oct 16];26(4):1737. Available from: [/pmc/articles/PMC9074854/](https://pubmed.ncbi.nlm.nih.gov/36569804/)
11. Azizi S, Hier DB, Wunsch DC. Enhanced neurologic concept recognition using a named entity recognition model based on transformers. *Front Digit Health* [Internet]. 2022 Dec 8 [cited 2023 Oct 16];4. Available from: <https://pubmed.ncbi.nlm.nih.gov/36569804/>
 12. Liu Y, Ott M, Goyal N, Du J, Joshi M, Chen D, et al. RoBERTa: A Robustly Optimized BERT Pretraining Approach. 2019 Jul 26 [cited 2023 Oct 18]; Available from: <https://arxiv.org/abs/1907.11692v1>
 13. He P, Liu X, Gao J, Chen W. DeBERTa: Decoding-enhanced BERT with Disentangled Attention. *ICLR 2021 - 9th International Conference on Learning Representations* [Internet]. 2020 Jun 5 [cited 2023 Oct 18]; Available from: <https://arxiv.org/abs/2006.03654v6>
 14. Lewis MPaul, Summer Institute of Linguistics. *Ethnologue : languages of the world*. 2009 [cited 2023 Oct 19];1248. Available from: <https://books.google.com/books/about/Ethnologue.html?id=FVVPgAACAAJ>
 15. Gonzalo J, Montes-Y-Gómez M, Rangel F. Overview of IberLEF 2022: Natural Language Processing Challenges for Spanish and other Iberian Languages. 2022 [cited 2023 Oct 18]; Available from: <https://nlp.uned.es/>
 16. Akhtyamova L, Martínez P, Verspoor K, Cardiff J. Testing contextualized word embeddings to improve NER in Spanish clinical case narratives. *IEEE Access*. 2020;8:164717–26.
 17. Conneau A, Khandelwal K, Goyal N, Chaudhary V, Wenzek G, Guzmán F, et al. Unsupervised Cross-lingual Representation Learning at Scale. *Proceedings of the Annual Meeting of the Association for Computational Linguistics* [Internet]. 2019 Nov 5 [cited 2023 Oct 15];8440–51. Available from: <https://arxiv.org/abs/1911.02116v2>
 18. Carrino CP, Llop J, Pàmies M, Gutiérrez-Fandiño A, Armengol-Estapé J, Silveira-Ocampo J, et al. Pretrained Biomedical Language Models for Clinical NLP in Spanish. *Proceedings of the Annual Meeting of the Association for Computational Linguistics* [Internet]. 2022 [cited 2023 Oct 15];193–9. Available from: <https://aclanthology.org/2022.bionlp-1.19>
 19. Wang L, Yang N, Huang X, Jiao B, Yang L, Jiang D, et al. Text Embeddings by Weakly-Supervised Contrastive Pre-training. 2022 Dec 7 [cited 2023 Oct 15]; Available from: <https://arxiv.org/abs/2212.03533v1>
 20. Dogan A, Birant D. A Weighted Majority Voting Ensemble Approach for Classification. *UBMK 2019 - Proceedings, 4th International Conference on Computer Science and Engineering*. 2019 Sep 1;366–71.