International Conference on Formal Methods
and Foundations of Artificial Intelligence
Eszterházy Károly Catholic University
Eger, Hungary, June 5–7, 2025

Fine-Tuning BiomedBERT with LoRA and Pseudo-Labeling for Robust DDI Classification

Ioan-Flaviu Gheorghita[a], **Vlad-Ioan Bocanet**[b], **Laszlo Barna Iantovics**[c]

[a]Doctoral School of Letters, Humanities and Applied Sciences, George Emil Palade University of Medicine, Pharmacy, Science, and Technology of Târgu Mureş, Târgu Mureş, România.
gheorghita.flaviu-ioan.21@stud.umfst.ro

[b]Dep. of Manufacturing Engineering, Technical University of Cluj-Napoca, Cluj-Napoca 400641, România.
vlad.bocanet@tcm.utcluj.ro

[c]Dep. of Electrical Engineering and Information Technology, George Emil Palade University of Medicine, Pharmacy, Science, and Technology of Târgu Mureş, Târgu Mureş, România.
barna.iantovics@umfst.ro

Abstract

In clinical decision support systems (CDSS), where precise classification of drug-drug interactions (DDIs) [10] can directly affect treatment safety and outcomes, understanding how drugs interact—-whether they enhance each other's effects or interfere with each other—is a major issue. We present a scalable and efficient method for classifying DDIs using a fine-tuned biomedical language model. Our method minimizes resource use during fine-tuning [1, 2] by using BiomedBERT [3, 5], a domain-specific variation of BERT [6, 9] pre-trained on biomedical literature, further adapted through Low-Rank Adaptation (LoRA) [4]. The goal is to categorize DDIs into two clinically relevant types: Synergistic and Antagonistic interactions. Initially, BiomedBERT is fine-tuned on a labeled subset of interaction descriptions sourced from DrugBank. To address the challenge of limited labeled data, we introduce a pseudo-labeling strategy: the fine-tuned model is used to

infer interaction types for the remaining unlabeled data. A checkpointing mechanism stores predictions and confidence scores incrementally, so guaranteeing the process is resumable and strong against system interruptions. Unlike traditional approaches depending on rigorous confidence thresholds, our framework logs all predictions so enabling manual or rule-based refining downstream. This enables the construction of large, high-quality labeled datasets without extensive human annotation. Advanced machine learning pipelines base their work on the resulting classification dataset. This architecture's modularity lets it be easily adapted to other biomedical classification chores or extra interaction categories. Furthermore, our approach is model-agnostic and can be extended to other transformer-based backbones fine-tuned for pharmacovigilance or electronic health record mining. These design decisions guarantee wide applicability over clinical natural language processing (NLP) pipelines and enable scalable deployment in actual settings. Above all, this study gets one ready for a more broad clinical application. The designated classified interactions are meant to enable next studies on DDI prediction models for integration into CDSS [8], systems [7] which can help doctors choose safe and efficient polypharmacy plans, maximize drug combination treatments, and avoid side effects by means of optimizations. Through the identification of synergistic and antagonistic drug effects, our method helps to create intelligent, data-driven tools for individualized medicine and therapeutic planning. Our pipeline—built on LoRA-based adaptation and iterative pseudo-labeling—offers a reproducible, low-resource strategy for biomedical NLP and DDI classification at scale and it lays the foundation for better, evidence-based decision-making in multi-drug treatment environments by bridging the gap between raw biomedical data and pragmatic clinical utility.

References

[1] J. Benjamin, A. Morales, D. Smith: *LoRA for Efficient Fine-Tuning of Foundation Models*, Journal of Machine Learning Research 24.1 (2023), pp. 1–18, doi: 10.5555/123456789.

[2] T. Dettmers, A. Pagnoni, A. Holtzman, et al.: *QLoRA: Efficient Fine-Tuning of Quantized LLMs*, arXiv preprint arXiv:2305.14314 (2023), doi: 10.48550/arXiv.2305.14314.

[3] Y. Gu, R. Tinn, H. Cheng, et al.: *BioBERT-KG: Biomedical Knowledge Graph Construction Based on Pretrained Language Model*, arXiv preprint arXiv:2201.00065 (2022), doi: 10.48550/arXiv.2201.00065.

[4] E. J. Hu, Y. Shen, P. Wallis, et al.: *LoRA: Low-Rank Adaptation of Large Language Models*, arXiv preprint arXiv:2106.09685 (2021), doi: 10.48550/arXiv.2106.09685.

[5] J. Lee, W. Yoon, S. Kim, et al.: *BioBERT: a pre-trained biomedical language representation model for biomedical text mining*, Bioinformatics 36.4 (2020), pp. 1234–1240, doi: 10.1093/bioinformatics/btz682.

[6] Z. Li, H. Liu, J. Wang: *DDI-BERT: A BERT-based Method for Drug–Drug Interaction Extraction from Biomedical Text*, BMC Bioinformatics 24.1 (2023), p. 12, doi: 10.1186/s12859-023-05106-1.

[7] B. Shickel, P. J. Tighe, A. Bihorac, P. Rashidi: *Deep learning in clinical decision support: A review*, Journal of Biomedical Informatics 84 (2018), pp. 405–419, doi: 10.1016/j.jbi.2018.11.002.

[8] R. T. Sutton, D. Pincock, D. C. Baumgart, et al.: *An overview of clinical decision support systems: benefits, risks, and strategies for success*, NPJ Digital Medicine 3.17 (2020), doi: 10.1038/s41746-020-0221-y.

[9] W. Yuan, J. Zhang, J. Zeng, et al.: *BioMedLM: Foundation models for biomedical natural language processing*, arXiv preprint arXiv:2210.10341 (2022), doi: 10.48550/arXiv.2210.10341.

[10] J. Zhu, C. Liu, et al.: *A Review of Drug–Drug Interaction Prediction Using Deep Learning*, Current Drug Metabolism 23.5 (2022), pp. 392–404, doi: 10.2174/1389200223666220203114035.