Doctoral School of Information and Biomedical Technologies

Polish Academy of Sciences (TIB PAN)

SUBJECT (EN): Neuro-symbolic integration of Large Language Models and Knowledge Graphs for clinical reasoning in acute and rare conditions.

SUBJECT (PL): Integracja neuro-symboliczna dużych modeli językowych i grafów wiedzy we wnioskowaniu klinicznym w stanach ostrych i rzadkich

SUPERVISORS: Prof. Jan Mielniczuk, PhD, Michał J. Dąbrowski – Institute of Computer Science, Polish Academy of Sciences

DESCRIPTION: The rapid advancement of generative Artificial Intelligence in medicine has highlighted a profound methodological divide. Large Language Models (LLMs) excel at soft-reasoning, identifying latent patterns, and synthesizing heterogeneous data. However, their clinical utility is severely limited by non-transparent reasoning pathways, hallucinations, and a lack of deterministic verification. Conversely, clinical Knowledge Graphs (KGs) and structured ontologies offer explicit, verifiable representations of medical knowledge, but they are rigid, brittle under data missingness, and fail to process unstructured, streaming Patient Health Records effectively.

This doctoral project focuses on fundamental methodological research at the intersection of neuro-symbolic AI, aiming to develop novel frameworks that fundamentally integrate KGs and LLMs. Moving away from standard, static Retrieval-Augmented Generation (RAG), the research will investigate dynamic, real-time clinical reasoning loops. The core technical focus will involve designing self-correcting agentic architectures utilizing the Reflexion mechanism—where the LLM’s probabilistic clinical hypotheses are iteratively cross-examined, grounded, and corrected against deterministic constraints derived from a biomedical Knowledge Graph.

A key scientific challenge of this thesis is ensuring high-fidelity reasoning over incomplete, heterogeneous, and asynchronous Electronic Health Record (EHR) data streams, specifically tailored for acute and rare medical conditions (where time-criticality and data sparsity overlap). Furthermore, to address the strict constraints of the EU AI Act and GDPR, the project will theoretically and experimentally investigate the trade-offs of deploying these complex agentic frameworks within strictly isolated, On-Premise environments, bypassing reliance on commercial cloud APIs.

The primary output of the PhD thesis will be new domain-independent algorithms, formal neuro-symbolic alignment methods, and uncertainty-quantification models verified against real-world, high-fidelity benchmarks.

PROFILE OF THE CANDIDATE: The ideal candidate should possess a strong foundational background enabling independent research at the intersection of computer science, graph theory, and artificial intelligence.

Requirements:

- M.Sc. in Computer Science, Bioinformatics, Data Science, Mathematics, or a closely related quantitative field.
- Strong knowledge of Machine Learning and Natural Language Processing, including the mathematical foundations of transformer architectures and LLMs.
- Proficiency in Python and deep learning frameworks (e.g., PyTorch, Hugging Face ecosystem).
- Solid understanding of data processing pipelines (data cleaning, integration of heterogeneous data sources).
- Excellent analytical skills and an eager, proactive approach to solving complex theoretical and computational problems.
- Fluency in written and spoken English, with the ability to write scientific papers and present research results.

Desirable assets (nice-to-have):

- Experience with Knowledge Graphs, Graph Neural Networks (GNNs), or graph databases (e.g., Neo4j).
- Familiarity with medical ontologies, clinical data standards (e.g., SNOMED-CT, UMLS, HPO), and data privacy regulations (GDPR, AI Act).
- Prior research experience, participation in scientific projects, or peer-reviewed publications.

Candidate should contact the author of the proposal before formal submission of documents (michal.dabrowski@ipipan.waw.pl)

BIBLIOGRAPHY:

[1] *Zhao, Weike, et al. "An agentic system for rare disease diagnosis with traceable reasoning." Nature (2026): 1-10.*

[2] *Tseng, Rachel Marjorie Wei Wen, et al. "Prospective real-world implementation of deep learning systems in healthcare: a systematic review guided by implementation science." npj Digital Medicine (2026).*

[3] *Zhang, Ming, et al. "LLMEval-Med: a real-world clinical benchmark for medical LLMs with physician validation." arXiv preprint arXiv:2506.04078 (2025).*