

Proposal for a 3-year PhD or 2-year postdoctoral contract in AI for healthcare

Title: Step-by-step guidance to clinical decision,
by combining data- and knowledge-driven approaches

1 Supervision

Co-supervisor: Adrien Coulet
Research group: HeKA research team (Inria, Inserm, Université Paris Cité)
Contact: adrien.coulet@inria.fr

Co-supervisor: Bastien Rance
Research group: HeKA (Inria, Inserm, UP Cité), Centre de Rech. des Cordeliers
Department of medical informatics, Georges Pompidou Hospital, AP-HP
Contact: bastien.rance@aphp.fr

2 Motivation and background

In the last years, deep-learning-based clinical decision support has mostly been over-simplified to a single classification task, mostly for demonstrating the feasibility of reproducing a final medical decision on the basis of previously acquired healthcare data [ROC⁺18, MLKD16]. This led to the development and publication of many models, with very little perfusion in clinics, due to their lack of explainability [SAD24], fairness [BHN23], usability, and need [CSF⁺20]. For example, a single class prediction based on patient history appear to be of limited use to practitioners that first consider their last observations, and wait for decision support to integrate medical history to their own medical reasoning [AMCD21].

In previous works, we proposed to learn decision pathways from Electronic Health Records (EHRs) to propose a step-by-step guidance to the medical decision [MNC24a]. These pathways are build on the basis of data collected along healthcare activity, (i.e., traces of real-word clinical practice) and for this reason follow to some extent medical reasoning. This data-driven approach can be opposed to the classical decision processes in medicine that follow Clinical Practice Guidelines (CPGs), which compile in narratives the state-of-the-art knowledge authored by a college of specialists of the domain. We consider the latter as a knowledge-driven approach.

Both data- or knowledge- driven approaches have been shown to be associated to unfair decisions for various reasons, including sampling biases when constituting training datasets, or study cohorts used to build CPGs [RCB⁺19].

We think that Large Language Models (LLMs) are promising tools to support knowledge-driven decision support [CMC24, BJS24]. But by default, they do not consider the amount of observational data acquired in EHRs. In addition, we think that hybrid systems that combines data- and knowledge-driven approaches offer the best promises for the design for explainable and fair decision supports in biomedicine. Consequently, the following research project proposes to study

the combination of (i) LLM, as a provider to the state-of-the-art knowledge, and (ii) data-driven approaches such as the one developed by Muyama *et al.* for decision support [MNC24a].

3 Objective

LLM as a provider of state-of-the-art knowledge – Despite their undeniable role in clinical decision-making, CPGs have limitations. First, their development requires input from a diverse panel of experts, leading to a time-consuming and costly process. This implies that updates to CPG take several years, which makes them unsuitable for rapidly evolving medical practices, such as the introduction of new diagnostic tests or the emergence of a new disease. Second, due to the substantial resources needed to develop these guidelines, it is not practical to create CPGs for every known medical condition. Third, guidelines are generally designed to address the needs of the majority of the population, potentially overlooking rare conditions or uncommon populations, potentially leading to unfair recommendations.

The first objective of this research project is to study how LLMs may embed state-of-the-art knowledge represented within CPGs, and thus support medical decision. Several challenges are associated with this objective such as the fact that CPGs include semi-structured information as decision trees or score tables, which are key to guide recommendations [Mil19, ACD⁺19], or the fact that important element of knowledge may be spread within over several documents, including potential inconsistencies.

Combine data-driven approach with LLM – Several methods can be leveraged to learn clinical pathways from data including sequence and process mining, to discover patterns within hospital processes [MNC24b]. Reinforcement Learning (RL) methods are natural candidates to learning sequential decisions, Deep Reinforcement Learning (DRL) adopts deep neural networks to learn the optimal policy that is central to RL and facilitate handling cases where the number of states is potentially large [LNF19]. To identify optimal treatment regimens for patients, [LWL⁺22] proposed a deep Q network-based model for EHR data. Patient records are used to model the state, action and transition probability of the model, used, in turn, to determine the individual optimal dose. Resulting constructs can in turn be compared with state-of-the-art CPGs, but both sources are not combined.

The second objective of the project is to enable LLM and a data-driven approach, such as a DRL one, to be combined. One potential direction is to find inspiration in the reinforcement learning (RL) from human feedback paradigm that is used to adapt LLM from human preferences [ZSW⁺20]. Here LLM model could be adapted through a number of RL iterations where the feedback is not provided by a human, but through healthcare data that trace historical clinical decision made by clinicians.

Evaluate various approaches – A transversal objective of the project is to enable the evaluation of the various developed systems, with the following objectives in mind : correctness, optimality, fairness and utility. This relies, through others, on enabling the comparison of various clinical decision pathways in terms of precision, recall (for instance in regards to a ground truth diagnosis), number of steps (*i.e.*, tests, questions), time before diagnosis, cost and others. To this matter, the candidate could propose metrics that will enable comparing pathways one-to-one and many-to-many. One potential direction would be to rely on a common representation framework of clinical pathways by leveraging the experience on scientific workflows of the ShareFAIR PEPR project consortium. In particular, we could aim at designing a shareable, traceable, understandable format that supports reasoning mechanisms, to facilitate further comparison.

Biomedical usecases – The project will be applied on three distinct biomedical use cases : the differential diagnosis of anemia, hypertension and the management of intracranial aneurysms [EdST⁺22]

4 Context

The position will take place in the HeKA team (Inria, Inserm, Université Paris Cité), physically located at PariSanté Campus, 2-10 rue d'Oradour-sur-Glane, 75015 Paris. The work is a collaboration between the HeKA team, and the Pitié-Salpêtrière hospitals (HEGP) of the AP-HP (Assistance Publique - Hôpitaux de Paris).

This project is funded by the ShareFAIR PEPR Digital Health project, funded by the French Research National Agency.

5 Application

Only via <https://jobs.inria.fr/public/classic/en/offres/2024-08428>

Références

- [ACD⁺19] Martin Aringer, Karen Costenbader, David Daikh, Ralph Brinks, Marta Mosca, Rosalind Ramsey-Goldman, Josef S Smolen, David Wofsy, Dimitrios T Boumpas, Diane L Kamen, et al. 2019 european league against rheumatism/american college of rheumatology classification criteria for systemic lupus erythematosus. *Arthritis & rheumatology*, 71(9) :1400–1412, 2019.
- [AMCD21] Julia Adler-Milstein, Jonathan H Chen, and Gurpreet Dhaliwal. Next-generation artificial intelligence for diagnosis : from predicting diagnostic labels to “wayfinding”. *Jama*, 326(24) :2467–2468, 2021.
- [BHN23] Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning : Limitations and Opportunities*. MIT Press, 2023.
- [BJS24] Suhana Bedi, Sneha S Jain, and Nigam H Shah. Evaluating the clinical benefits of llms. *Nature Medicine*, 30(9) :2409–2410, 2024.
- [CMC24] Elisa Castagnari, Lillian Muyama, and Adrien Coulet. Prompting Large Language Models for Supporting the Differential Diagnosis of Anemia. In *LLMs4MI 2024 @FLLM 2024 - First International Workshop on Large Language Models Applications in Medical Informatics In conjunction with the The Second International Conference on Foundation and Large Language Models*, Dubai, United Arab Emirates, November 2024. IEEE.
- [CSF⁺20] Christine M Cutillo, Karlie R Sharma, Luca Foschini, Shinjini Kundu, Maxine Mackintosh, Kenneth D Mandl, and MI in Healthcare Workshop Working Group Beck Tyler 1 Collier Elaine 1 Colvis Christine 1 Gersing Kenneth 1 Gordon Valery 1 Jensen Roxanne 8 Shabestari Behrouz 9 Southall Noel 1. Machine intelligence in healthcare—perspectives on trustworthiness, explainability, usability, and transparency. *NPJ digital medicine*, 3(1) :47, 2020.
- [EdST⁺22] Nima Etminan, Diana Aguiar de Sousa, Cindy Tiseo, Romain Bourcier, Hubert Desal, Antti Lindgren, Timo Koivisto, David Netuka, Simone Peschillo, Sabrina Lémeret, et al. European stroke organisation (eso) guidelines on management of unruptured intracranial aneurysms. *European Stroke Journal*, 7(3) :LXXXI–CVI, 2022.
- [LNF19] Siqi Liu, Kee Yuan Ngiam, and Mengling Feng. Deep reinforcement learning for clinical decision support : a brief survey. *arXiv preprint arXiv :1907.09475*, 2019.
- [LWL⁺22] Tianhao Li, Zhishun Wang, Wei Lu, Qian Zhang, and Dengfeng Li. Electronic health records based reinforcement learning for treatment optimizing. *Information Systems*, 104 :101878, 2022.

- [Mil19] Michael Miller. Acc/aha lipids & ascvd guidelines : 2018 update. *Metabolism-Clinical and Experimental*, 99 :116–118, 2019.
- [MLKD16] Riccardo Miotto, Li Li, Brian A Kidd, and Joel T Dudley. Deep patient : an unsupervised representation to predict the future of patients from the electronic health records. *Scientific reports*, 6(1) :1–10, 2016.
- [MNC24a] Lillian Muyama, Antoine Neuraz, and Adrien Coulet. Deep reinforcement learning for personalized diagnostic decision pathways using electronic health records : A comparative study on anemia and systemic lupus erythematosus. *Artificial Intelligence in Medicine*, 157 :102994, 2024.
- [MNC24b] Lillian Muyama, Antoine Neuraz, and Adrien Coulet. Machine learning approaches for the discovery of clinical pathways from patient data : A systematic review. *Journal of Biomedical Informatics*, page 104746, 2024.
- [RCB⁺19] Fatima Rodriguez, Sukyung Chung, Manuel R Blum, Adrien Coulet, Sanjay Basu, and Latha P Palaniappan. Atherosclerotic cardiovascular disease risk prediction in disaggregated asian and hispanic subgroups using electronic health records. *Journal of the American Heart Association*, 8(14) :e011874, 2019.
- [ROC⁺18] Alvin Rajkomar, Eyal Oren, Kai Chen, Andrew M Dai, Nissan Hajaj, Michaela Hardt, Peter J Liu, Xiaobing Liu, Jake Marcus, Mimi Sun, et al. Scalable and accurate deep learning with electronic health records. *NPJ digital medicine*, 1(1) :1–10, 2018.
- [SAD24] A review of explainable artificial intelligence in healthcare. *Computers and Electrical Engineering*, 118 :109370, 2024.
- [ZSW⁺20] Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences, 2020.