Exploiting Prior Knowledge and Contextual Awareness for Adaptive and Interpretable Healthcare Analytics

Chin Wang Cheong

Hong Kong Baptist University, Hong Kong
cwcheong@comp.hkbu.edu.hk

With the increasing adoption of electronic health records (EHRs), clinical data has become more accessible, offering new opportunities to enable predictive analytics for personalized clinical decision support. The EHR data are characterized by its high dimensionality of clinical features (diagnosis, medication, lab test results, etc.) where their interactions are complex and patient-specific. That makes application of existing machine learning methods to achieve highly accurate analytics results non-trivial. In addition, model interpretability is another important concern in the context of healthcare.

My thesis research work focuses on developing deep learning methods to achieve more personalized and reliable healthcare predictive analytics. Two key research questions are being addressed:

– How can *prior medical knowledge* be effectively integrated into predictive models to enhance interpretability and prediction accuracy?
– How can *context awareness* be enabled in predictive models to improve their reliability and personalization?

1 Proposed Methodologies and Contributions

Traditional deep learning often lacks mechanisms to incorporate domain knowledge and adapt to patient-specific contexts, which can result in reduced interpretability and suboptimal performance. In my thesis research, methodologies are developed to leverage: i) *medical knowledge* obtained *explicitly* from medical ontologies and *implicitly* from large language models via knowledge distillation techniques, and ii) *patient context-aware* analytics including *input-aware* feature selection and imputation, to enhance EHR-based predictive analytics.

1.1 Integration of Explicit and Implicit Medical Knowledge

Existing research of integrating medical ontologies into predictive models often assumes a direct mapping between ontologies and learned representations. However, ontologies categorize medical concepts based on structured taxonomies, which may not align consistently well with the data-driven embeddings. Mismatches between ontological structures and EHR-derived embeddings can degrade model performance. We propose a novel learning framework called **ADORE**

(Adaptive Ontology-Driven Representation Learning) which can adaptively restructure multi-relational medical ontologies (e.g., SNOMED-CT) to gain better alignment with the EHR in the embedding space, improving model interpretability and downstream predictive performance.

Large language models (LLMs) provides another rich source of knowledge, which can be distilled via prompt. While they have been explored as well for healthcare analytics applications, the most effective way of integrating them for predictive model learning remains open. We propose a novel knowledge distillation method called **HyperKD** (Hypergraph Knowledge Distillation) which can adaptively identify uncertain patient records and make queries to LLMs to extract need-to-know higher-order relationships of clinical features (represented as hypergraphs). HyperKD can effectively improve patient representation learning via data augmentation and semi-supervised hypergraph prediction.

1.2 Predictive Analytics with Patient Context Awareness

Context-aware analytics such as instance-wise feature selection and imputation methods have been studied. Existing approaches often suffer from high variance, poor interpretability, and limited adaptability to missing data, which are common in real-world EHR data. We propose a novel instance-wise feature selection method named **FlexGPC** (Flexible Grouped Predictive Coding) which can dynamically compose instance-specific feature masks by learning a set of shared feature groups to enhance interpretability and stability in patient data analysis.

To further extend from instance-wise feature selection, we explore the possibility of learning instance-wise feature imputation model end-to-end. We propose a novel hypernetwork-based imputation model called **HynetImpute** that can capture missing patterns using a hypernetwork which can generate feature-subset-specific encoders within a variational autoencoder (VAE) framework for effective patient context specific imputation. The distributional consistency between original and imputed data can also be well preserved.

2 Current Progress and Future Work

Rigorous experiments have been conducted using real-world EHR datasets MIMIC-III and eICU to demonstrate that the proposed methodologies, including ADORE, HyperKD, FlexGPC, and HynetImpute, can outperform the state-of-the-art methods proposed in the literature. We have also demonstrated via different case studies to demonstrated how better personalization and intepretability can be supported in the context of patient health condition.

My future research plan includes: i) extending the proposed methods to support multi-modal EHR data, ii) further strengthening the interpretability and clinical validation of the proposed methods, iii) designing active feature acquisition methods for more cost-effective clinical decision support, and iv) jointly integrating both medical ontologies and LLMs knowledge for enhancing model reliability.