Research Statement

LIM Ee Peng
School of Computing and Information Systems, Singapore Management University
Tel: (65) 6828-0781; Email: aseplim@smu.edu.sg
10 (Day) 12 (Month) 2023 (Year)

Background

My research interests cover data mining, machine learning, and information retrieval. In my research, I develop computational methods that model and analyse user generated structured and unstructured data for the purpose of extracting knowledge embedded in the data and improving social and urban experience through recommendation and prediction. In my research, I venture beyond standard datasets and study novel problems in emerging applications that contribute to social good. I also actively seek research grants to support my work and collaboration with others.

In the following, I summarize my recent works in a few research areas.

Research Areas

I recently started to work on **conversational mental health agent research** as part of a new funded project by NMRC. A key challenge in this research is to develop role-playing conversational agents personalized to the characteristics and profiles of clients in mental health counselling. These agents simulating different clients allow safe assessment of counsellor trainees as well as safe evaluation of agents simulating counsellors. Nevertheless, the ability to personalize the utterances to counselling clients, whether conducted by humans or LLMs, has not been well studied. To bridge this gap, our study introduces a novel evaluation challenge: user verification in agent-generated conversations, which aimed to verify whether two sets of utterances originate from the same client. To this end, we assemble a large dataset collection encompassing thousands of clients and their utterances. We also develop and evaluate user verification models under experiment setups. We further utilize the client verification models to evaluate the personalization abilities of LLM-based role-playing models. Comprehensive experiments suggest that the current role-playing models fail in accurately mimicking clients, primarily due to their inherent linguistic characteristics. This work has been published as a long paper in ACL2024. In the next step of our research, we will thus investigate into the development of AI agents that could perform much better in user verification and behave consistently according to the assigned client profiles.

In **recommender systems research**, we investigate into using LLMs as strong sequential recommenders. We explore the in-context learning (ICL) approach to sequential recommendation and investigate the effects of instruction format, task consistency, demonstration selection, and number of demonstrations. As increasing the number of demonstrations in ICL does not improve accuracy despite using a long prompt, we propose a novel method called LLMSRec-Syn that incorporates multiple demonstration users into one aggregated demonstration. Our experiments on three recommendation datasets show that LLMSRec-Syn outperforms state-of-the-art LLM-based sequential recommendation methods. In some cases, LLMSRec-Syn can

perform on par with or even better than supervised learning methods. This work has been published at NAACL-HLT2024.

In **financial forecasting research**, my ex PhD student, Gary Ang, studied dynamic multimodal networks, such as dynamic relationship networks between companies that evolve across time due to changes in business strategies and alliances, which are associated with dynamic company attributes from multiple modalities such as textual online news, categorical events, and numerical financial-related data. Such information can be useful in predictive tasks involving companies. Environmental, social, and governance (ESG) ratings of companies are important for assessing the sustainability risks of companies. The process of generating ESG ratings by expert analysts is, however, laborious and time-intensive. We thus explore the use of dynamic multimodal networks extracted from the web for forecasting ESG ratings. Learning such dynamic multimodal networks from the web for forecasting ESG ratings is, however, challenging due to its heterogeneity and the low signal-to-noise ratios and non-stationary distributions of web information. In this research, we propose the Dynamic Multimodal Slot Concept Attention-based Network (DynScan) model. DynScan utilizes slot attention mechanisms together with slot concept alignment and disentanglement loss functions to learn latent slot concepts from dynamic multimodal networks to improve performance on ESG rating forecasting tasks. DynScan is evaluated on forecasting tasks on six datasets, comprising three ESG ratings across two sets of companies. Our experiments show that DynScan outperforms other state-of-the-art models on these forecasting tasks. We also visualize the slot concepts learned by DynScan on five synthetic datasets and three real-world datasets and observe distinct and meaningful slot concepts being learned by DynScan across both synthetic and real-world datasets. This work has been published at TWeb.

Future Research. I shall continue using my research expertise in data mining, information retrieval and AI to explore more varied research problems in novel application domains. I will also expand my collaboration with domain experts to identify and work on important social and urban research problems.

Selected Publications and Outputs

Yizhe Yang, Palakorn Achananuparp, Heyan Huang, Jing Jiang, Ee-Peng Lim: Speaker Verification in Agent-generated Conversations. ACL (1) 2024: 5655-5676.

Lei Wang, Ee-Peng Lim. The Whole is Better than the Sum: Using Aggregated Demonstrations in In-Context Learning for Sequential Recommendation. NAACL-HLT (Findings) 2024: 876-895.

Gary Ang, Ee-Peng Lim. Learning Dynamic Multimodal Network Slot Concepts from the Web for Forecasting Environmental, Social and Governance Ratings. ACM Trans. Web 18(3): 38:1-38:32 (2024)