論文内容の要旨

申請者氏名　LI XINBAI

The rapid advancement of language models (LMs) has prompted extensive research into enhancing their knowledge and reasoning capabilities. While pre-training language models (PLMs) acquire vast amounts of semantic knowledge from large-scale corpora, their ability to handle domain-specific understanding and reasoning tasks remains limited. Existing knowledge enhancement (KE) methods have been applied to both small-scale models and large language models (LLMs), leveraging external knowledge to improve prediction performance. However, challenges remain in effectively integrating structured knowledge into LLMs and in efficiently training smaller models under limited computational resources. This thesis explores KE strategies tailored for both LLMs and smaller models.
For LLMs, a key challenge is the natural integration of structured data, such as knowledge graphs (KGs), into text-based inputs suitable for LLM processing. To address this, we propose GenKP, a knowledge prompt generation framework that injects knowledge into LLMs via in-context learning (ICL). GenKP utilizes LLMs in combination with KGs to generate knowledge samples, refining them through weighted verification and BM25 ranking to reduce noise and enhance factual accuracy. Experimental results demonstrate that GenKP improves LLM performance, outperforming traditional triple and template-based knowledge injection approaches.
For smaller models, we investigate the incorporation of external knowledge through knowledge distillation (KD). While LLMs can serve as teacher models to transfer knowledge, they lack specialized domain expertise. To mitigate this, we introduce Retrieval-Enhanced KD (RED), a novel framework that expands knowledge sources in KD by incorporating retrieved external knowledge. This method enables LLMs to assess and integrate relevant external information while training student models. Furthermore, we propose compartment, traction, and filter mechanisms to address challenges in external knowledge insertion. Experimental evaluations across four datasets in two specialized domains indicate that our approach significantly enhances the performance of smaller models. Notably, even in scenarios where retrieved knowledge is noisy or mismatched, our approach ensures robustness and accurate predictions.
Overall, this thesis presents novel KE methodologies for both LLMs and smaller models, tackling challenges in knowledge integration, retrieval, and knowledge transfer. By bridging structured knowledge with LLMs and optimizing KD techniques for smaller models, this research advances the effectiveness of KE approaches in improving model reasoning and prediction capabilities.

論文審査結果の要旨

申請者氏名　　LI XINBAI

審査委員一同は、令和7年5月19日 本論文申請者に対して審査を行った. 関連研究の調査の不足と, 本研究との関連が希薄であった点が指摘されたものの軽微な指摘であり, 改善が行われたことが確認された.