University of Pittsburgh | School of Computing and Information

Proposal Defense

Doctor of Philosophy in Computer Science

"**Improving Text Revision with Large Language Model**" by **Zhexiong Liu**

Date: April 4, 2025
Time: 1:00 – 3:00 pm
Place: Room 6106 Sennott Square, 3810 Forbes Ave, Pittsburgh, PA 15213

Committee:

- Diane Litman, Department of Computer Science, School of Computing and Information, University of Pittsburgh
- Adriana Kovashka, Department of Computer Science, School of Computing and Information, University of Pittsburgh
- Lorraine Li, Department of Computer Science, School of Computing and Information, University of Pittsburgh
- Richard Correnti, Learning Sciences and Policy, School of Education, University of Pittsburgh

Abstract:

Text revision is a fundamental process in natural language processing (NLP) that involves reviewing and refining written work to improve revision quality. For example, effective revision addresses revision issues at surface levels (e.g., grammar, punctuation, style) and content levels (e.g., coherence, clarity, reasoning). Despite its importance, developing automated text revision systems that provide immediate feedback on revision improvement remains a challenge. A key bottleneck lies in the complexity of revision intentions, for instance, writers often revise text for multiple or overlapping purposes (e.g., clarification, specification, meaning-change), making it difficult for existing systems to identify these intentions effectively. This highlights the need for advanced approaches to disentangle revision intentions and leverage them to improve overall text quality.

This proposal explores the potential of large language models (LLMs) to solve multiple revision tasks, including identifying revision intentions, generating optimized revisions, and assessing revision quality. Specifically, the proposal introduces a novel importance-redundancy LLM fine-tuning (IR-tuning) framework for identifying revision intentions, which dynamically and efficiently updates model weights on certain LLM layers while freezing the others during the fine-tuning process. Furthermore, it proposes an intention-aware fine-tuning (Intention-tuning) framework to generate optimized text revisions by simply fine-tuning specific LLM layers with additional revision-intention objectives. This method enforces LLMs to learn revision intention information on selected LLM layers and ensures the same selected LLM layers are used to generate optimized revisions. The proposed frameworks enhance LLMs' ability to capture and apply intention information to efficiently solve revision tasks, while also providing insights into layer-wise LLM fine-tuning approaches for solving general classification and generation tasks. Broadly, this proposal could advance LLM-based automated systems for assessing text writing and revision in NLP and educational domains.