



Proposal Defense
Doctor of Philosophy in Information Science

“A Mutually Evolving Knowledge Distillation Framework for Enhancing LLM Reasoning Capabilities” by Zhuochun Li

Date: October 9, 2026

Time: 1 to 3 p.m.

Place: Room 828, Information Sciences Building, 135 N. Bellefield Ave, Pittsburgh PA 15260

Committee:

- Daqing He, Advisor and Professor, Department of Informatics and Networked Systems, School of Computing and Information
- Konstantinos Pelechrinis, Associate Professor, Department of Informatics and Networked Systems, School of Computing and Information
- Pengfei Zhou, Assistant Professor, Department of Informatics and Networked Systems, School of Computing and Information
- Lorraine (Xiang) Li, Assistant Professor, Department of Computer Science, School of Computing and Information

Abstract:

Large language models (LLMs) have demonstrated strong multi-step reasoning capabilities, but their computational and memory requirements limit deployment in low-cost and resource-constrained settings. Knowledge distillation offers a promising way to transfer these capabilities to small language models (SLMs). However, current methods focus predominantly on verifiable tasks based on outcomes, such as mathematics and code generation, which may also transfer unfaithful reasoning steps produced by the teacher. Moreover, they treat distillation as a static, one-directional process that cannot adapt to the student's evolving needs. In this work, we aim to develop a Mutually Evolving Knowledge Distillation (MEKD) framework to improve the broad and reliable reasoning capabilities of SLMs. The framework has three interconnected components. Faithfulness verification identifies unreliable teacher reasoning before distillation. Process-aligned training combines outcome supervision with intermediate reasoning structure to guide student learning. Adaptive co-evolution adapts teacher supervision to the student's capabilities and failure patterns. The proposed research will first examine the effectiveness of reasoning distillation across mathematical, scientific, and commonsense tasks and assess the faithfulness of teacher rationales within each task regime. Building on this foundation, the research will proceed in two complementary directions. The first will develop process-oriented training methods that preserve logical progression while allowing the student to explore alternative reasoning paths. The second will promote adaptive teacher--student interaction, in which student errors inform the content and difficulty of subsequent guidance. Together, these efforts seek to advance knowledge distillation from static behavioral imitation toward a diverse, faithful, and adaptive learning process. This framework will enable SLMs to acquire robust reasoning capabilities more efficiently and effectively.