

ReAD: Reinforcement-Guided Capability Distillation



Xueqi Cheng



Presentation Abstract

Knowledge distillation (KD) compresses a large model into a smaller one that preserves the capabilities needed for a downstream task, yet existing methods assume that capabilities can be optimized independently and often overlook how distillation reshapes a model’s broader capability profile. Our analysis shows that improving one capability frequently shifts others in unintended ways and that a substantial portion of the distillation budget is spent on capabilities that offer little benefit to the target task, leading to inefficient and unpredictable outcomes. We reframe knowledge distillation through the lens of compressed intelligence, studying how intelligence decomposes into interacting capabilities and how limited capacity should be optimally allocated among them. Building on these insights, we propose ReAD, a reinforcement-guided capability distillation framework that learns to identify essential capabilities, generate targeted synthetic data, and allocate distillation effort across capabilities under explicit budget constraints. Experimental results demonstrate that ReAD produces higher-quality distilled models for the given task with significantly lower computation cost.

Speaker’s Profile



I’m Xueqi Cheng, a Ph.D. student in Computer Science at Florida State University, advised by Dr. Yushun Dong in the Responsible AI (RAI) Lab. My research aims to enhance the utility, security, and efficiency of Machine Learning as a Service (MLaaS), with applications ranging from graph-based learning to natural language processing and computer vision. I am also broadly interested in social network analysis and AI for social good, focusing on how AI can help address key societal challenges.

Seminar Time: 2026-02-24 Tuesday, 11:45 AM - 12:45 PM

Location (In-Person Only): Room 353 LOVE Building

Contact FSU Student Seminar Organizers

Dr. Yushun Dong

Email: yd24f@fsu.edu

