DISSERTATION PROPOSAL

Aidin Niaparast

"Decision-Making under Uncertainty: Online Joint Replenishment and Learning-Augmented Data Structures"

Wednesday, November 13, 2024
12:30pm
Tepper 4242 / Zoom link: Please contact lclee@andrew.cmu.edu

This dissertation investigates algorithms to cope with uncertainty in the input instance in both online worst-case and learning-augmented frameworks.

In the first chapter, we develop a new model for Online Joint Replenishment Problem, a key problem in supply chain and inventory management. Our model is the first to incorporate both holding and backlog costs, and generalizes many online and offline variants of this problem. We design a natural greedy algorithm for this problem and show that it is constant competitive against adversarial inputs using a novel dual-fitting analysis.

In chapters two and three, we investigate two fundamental data structures, namely Online List Labeling and Incremental Topological Ordering, in the learning-augmented setting and improve the running time of the classical algorithms for these problems. By leveraging predictions, we design algorithms that guarantee consistency, robustness, and smoothness with respect to predictions—that is, they have the best possible running time under perfect predictions, never perform worse than the best-known worst-case methods, and their running time degrades smoothly with the prediction error. We develop the general framework of Prediction Decomposition that allows us to warm start worst-case algorithms using predictions.

While almost all existing work in algorithms with predictions focuses on point predictions, in the fourth chapter, we explore the Binary Search problem using distributional predictions. We show that this is a richer setting: there are simple distributions where using the classical prediction-based algorithm with any single prediction does poorly. Motivated by this, the main contribution of this chapter is designing an algorithm whose performance is parameterized by a natural distance measure between the predicted and actual input distribution.

For the last three chapters, we demonstrate empirically that predictions, learned from a very small training dataset, are sufficient to provide significant speed-ups on real and synthetic datasets.

Proposed Committee:
R. Ravi and Benjamin Moseley (Co-Chairs), Ravi Kumar (Google), and Samuel McCauley (Williams College)

Proposal Documents:
Chapter 1: https://arxiv.org/abs/2410.18535
Chapter 2:
https://proceedings.neurips.cc/paper_files/paper/2023/hash/bd8284e53b6d177cbede82def77d4951-Abstract-Conference.html
Chapter 3: https://icml.cc/virtual/2024/poster/32743
Chapter 4: https://cmu.box.com/s/beixherkbort7afaag5e1r4n96jrz9vw