Ph.D. Dissertation Defense

Candidate: Xiangyu Tian
Degree: Doctor of Philosophy
School/Department: Interdisciplinary / Data Science
Date: Thursday, April 30, 2026
Time: 11:00 am – 12:30 pm
Location: North Building 316
Title: Risk Measures for Multi-Class Classification: Robustness, Fairness, and Asymptotic Guarantees

Chairperson: Dr. Darinka Dentcheva, Department of Mathematics
Committee Members: Dr. Violet Chen, School of Business
Dr. Kathrin Smetana, Department of Mathematics
Dr. Philippos Mordohai, Department of Computer Science
Dr. Andrzej Ruszczynski, Department of Management Science & Information Systems, Rutgers University

Abstract

The first part of the thesis develops a new framework for multi-class classification based on the theory of coherent risk measures for systemic risk. In this framework, the loss functions for each class are evaluated using coherent contextual risk measures, and a systemic risk measure determines the overall classification risk. We construct risk-averse counterparts to a popular multi-class classification method. To this end, we formulate a two-stage stochastic programming problem to construct the classifier. We design a novel risk-averse regularized decomposition method whose computational effort grows linearly with the number of data points. We further extend the risk-averse approach to kernel-based multi-class problems. Numerical experiments demonstrate that the proposed framework is particularly effective when the data is noisy, corrupted, or scarce relative to the problem dimension. The risk-averse classifiers exhibit greater robustness and better generalization than their risk-neutral counterparts, with the performance gap widening as the number of classes increases. The proposed approach can also be used to construct risk-averse counterparts to other classification methods.

The second part of the thesis addresses fairness in classification using contextual risk measures within classes. The sensitive attributes (e.g., demographic groups) are modeled as contexts. The conditional (group-wise) classification risks are aggregated using specific nonlinear class aggregators, which are then aggregated again to obtain an overall risk evaluation. We develop a three-level regularized decomposition algorithm based on the previous chapter to solve the resulting optimization problem. We demonstrate on a drug-use dataset that contextual-risk-based classifiers achieve comparable or improved fairness relative to a Wasserstein distributionally robust fair classifier, while maintaining higher predictive performance and lower variability than the existing robust–fair baseline. We further discuss fairness in the more general multi-class and multi-group setting and propose a chi-square test for evaluating group fairness. Experiments on the Adult income dataset with race as a multi-group sensitive attribute confirm the effectiveness of our approach.

The third part of the thesis investigates the asymptotic optimality and convergence rates of risk-averse multi-class classifiers. We represent the risk measures as composite functionals and prove the consistency of both the total risk estimator and the optimal classifier as the sample size increases. Furthermore, we establish a central limit theorem for the systemic risk estimator. This result characterizes the limiting distribution of the estimation error, enabling the construction of confidence intervals and statistical hypothesis tests for the classification risk. We illustrate the theoretical findings through numerical experiments on both real image datasets and synthetic data, confirming the predicted convergence to normality.