

A Margin Theory for Contextual Stochastic Linear Optimization: From Generalization to Active Learning

ABSTRACT— In the setting of contextual stochastic linear optimization (CSLO), the cost vector parameters of a linear optimization problem are unknown at the time of decision-making but may be predicted based on contextual features. The Predict-then-Optimize framework and the associated Smart Predict-then-Optimize (SPO) loss function provide a natural way to evaluate and train machine learning models based on the cost of decisions induced by the predicted parameters. On the other hand, in classical machine learning settings like binary classification, the notion of margin provides an intuitive and elegant way to analyze the generalization properties and performance guarantees of models and informs algorithmic developments. In this talk, we develop a margin theory for CSLO that relies on the notion of distance to degeneracy to place a margin around problematic predictions that do not yield unique optimal solutions. We derive margin-based generalization bounds for the SPO loss function that hold for broad classes of CSLO problems and that improve upon combinatorial bounds in situations with a favorable margin structure. Finally, we utilize the distance to degeneracy to develop a margin-based method for an online active learning variant of CSLO that sequentially decides whether to request the "labels" of feature samples from an unlabeled data stream. When the data distribution again satisfies some favorable margin properties, we theoretically and numerically demonstrate an improved label complexity over the standard supervised approach that acquires all labels. This talk represents several joint works with Adam Elmachtoub, Ambuj Tewari, Othman El Balghiti, Mo Liu, Heyuan Liu, and Max Shen.



Dr. Paul Grigas

Assistant Professor
Dept. of Industrial Engineering
and Operations Research
University of California, Berkeley

SPEAKER BIO — Paul Grigas is an Assistant Professor of Industrial Engineering and Operations Research at the University of California, Berkeley. Paul's research interests are broadly in optimization, machine learning, and data-driven decision-making, with particular emphasis on contextual stochastic optimization and algorithms at the interface of machine learning and continuous optimization. Paul's research is funded by the National Science Foundation, including an NSF CRII Award. Paul was awarded 1st place in the 2020 INFORMS Junior Faculty Interest Group (JFIG) Paper Competition and the 2015 INFORMS Optimization Society Student Paper Prize. He received his B.S. in Operations Research and Information Engineering (ORIE) from Cornell University in 2011, and his Ph.D. in Operations Research from MIT in 2016.