
SEBASTIAN POKUTTA, Georgia Institute of Technology
Lazifying Conditional Gradients

Both, discrete optimization techniques as well as machine learning approaches have made significant progress over the last decades and can be considered important tools that are regularly used in applications. However, very little has been done at the intersection of the two disciplines despite their unprecedented real-world significance: as soon as the underlying set of decisions that we want to optimize over is of a combinatorial or discrete nature standard learning approaches fail due to unacceptable running times or real-world irrelevant guarantees. At the same time many strategic, tactical, and operational challenges that we face (e.g., in dynamic routing, ad allocation, pick path optimization, dynamic pricing, demand prediction, etc.) require a tight integration of data, learning, and decision making. In this talk I will provide a general method to significantly speed-up convex optimization and learning by modifying conditional gradient algorithms. This new approach is particularly effective in the context of combinatorial problems leading to several orders of magnitude in speed-up.

(joint work with Gábor Braun and Daniel Zink)