

Classification and regression with structured outputs

Recently there has been considerable interest in methods of discriminative training for classification problems with structured outputs. Examples include hidden state reconstruction in a hidden Markov model, natural language parsing, and node classification in a graph (such as the a graph of web pages and links). In all of these examples the input is a structured object such as a string or graph and the output is also a structured objects, such as a parse tree or a labeling of the nodes in the input. This talk will discuss recently developed generalizations of support vector machines and present some PAC-Bayesian margin bounds for the structured case.