

Abstract: Penalized regression is an attractive methodology for dealing with high-dimensional data where classical likelihood approaches to modeling break down. However, its widespread adoption has been hindered by a lack of inferential tools. For example, penalized regression is very useful for variable selection, but how confident should one be about those selections? How many of those selections would likely have occurred by chance alone? What properties should confidence intervals even have in high dimensions? There have been many recent developments in this area; in this talk, I will present methods and techniques to address these questions that I have developed over the past few years.