Optimal Inference After Model Selection

Will Fithian, Stanford University

To perform inference after model selection, we propose controlling the selective type I error; i.e., the error rate of a test given that it was performed. By doing so, we recover long-run frequency properties among selected hypotheses analogous to those that apply in the classical (non-adaptive) context. Our proposal is closely related to data splitting and has a similar intuitive justification, but is more powerful. Exploiting the classical theory of Lehmann and Scheffe (1955), we derive most powerful unbiased selective tests and confidence intervals for inference in exponential family models after arbitrary selection procedures. For linear regression, we derive new selective z-tests that generalize recent proposals for inference after model selection and improve on their power, and new selective t-tests that do not require knowledge of the error variance.

Based on joint work with Dennis Sun and Jonathan Taylor.