Detecting changes in slope with an L_0 penalty

Prof. Dr. Paul Fearnhead
(Lancaster University)

December 6, 2017

Whilst there are many approaches to detecting changes in mean for a univariate time-series, the problem of detecting multiple changes in slope has comparatively been ignored. Part of the reason for this is that detecting changes in slope is much more challenging. For example, simple binary segmentation procedures do not work for this problem, whilst efficient dynamic programming methods that work well for the change in mean problem cannot be directly used for detecting changes in slope. We present a novel dynamic programming approach, CPOP, for finding the "best" continuous piecewise-linear fit to data. We define best based on a criterion that measures fit to data using the residual sum of squares, but penalises complexity based on an L_0 penalty on changes in slope. We show that using such a criterion is more reliable at estimating changepoint locations than approaches that penalise complexity using an L_1 penalty. Empirically CPOP has good computational properties, and can analyse a time-series with over 10,000 observations and over 100 changes in a few minutes. Our method is used to analyse data on the motion of bacteria, and provides fits to the data that both have substantially smaller residual sum of squares and are more parsimonious than two competing approaches.
Joint work with Robert Maidstone and Adam Letchford.