



Department of Mathematics and Statistics

Colloquium

Tuesday October 2

AMB 164 4:00 pm

Why does functional pruning yield such fast algorithms for optimal changepoint detection?

Toby Hocking

School of Informatics, Computing and Cyber Systems at NAU

Abstract

Detecting abrupt changes is an important problem in analyzing N data points gathered over time or space. In this setting, maximum likelihood inference amounts to minimizing a loss function (which encourages data fitting) plus a penalty on the number of changes in the model parameters (which discourages overfitting). Computing the optimal solution to this non-convex problem is possible using classical dynamic programming algorithms, but their $O(N^2)$ complexity is too slow for large data sequences. The functional pruning technique of Rigaiil involves storing the optimal cost using functions rather than scalar values. Empirical results from several recent papers show that the functional pruning technique consistently yields optimal algorithms of $O(N \log N)$ complexity, which is computationally tractable for very large N . The theoretical results of Rigaiil prove that functional pruning is $O(N^2)$ in the worst case and $O(N \log N)$ on average (for a special loss function). For future work it would be interesting to further study the average complexity of these algorithms, in order to provide more theoretical justification for these very fast empirical results.

Refreshments at 3:45