

Title: **Optimal Kernel Shapes for Local Linear Regression**

Author(s): **Dirk Ormoneit and Trevor Hastie**

Technical Report number (Dept. of Statistics, Stanford Univ.): **1999-11**

Date: **June 1999**

Abstract:

Local linear regression performs very well in many low-dimensional forecasting problems. In high-dimensional spaces, its performance typically decays due to the well-known “curse-of-dimensionality”. Specifically, the volume of a weighting kernel that contains a fixed number of samples increases exponentially with the number of dimensions. The bias of a local linear estimate may thus become unacceptable for many real-world data sets. A possible way to control the bias is by varying the “shape” of the weighting kernel. In this work we suggest a new, data-driven method to estimating the optimal kernel shape. Experiments using two artificially generated data sets and data from the UC Irvine repository show the benefits of kernel shaping.