

# Local polynomial regression and variable selection

Hugh Miller<sup>1</sup> and Peter Hall<sup>1</sup>

*The University of Melbourne*

**Abstract:** We propose a method for incorporating variable selection into local polynomial regression. This can improve the accuracy of the regression by extending the bandwidth in directions corresponding to those variables judged to be unimportant. It also increases our understanding of the dataset by highlighting areas where these variables are redundant. The approach has the potential to effect complete variable removal as well as perform partial removal when a variable redundancy applies only to particular regions of the data. We define a nonparametric oracle property and show that this is more than satisfied by our approach under asymptotic analysis. The usefulness of the method is demonstrated through simulated and real data numerical examples.

## 1. Introduction

The classical regression problem is concerned with predicting a noisy continuous response using a  $d$ -dimensional predictor vector with support on some  $d$ -dimensional subspace. This functional relationship is often taken to be smooth, and methods for estimating it range from parametric models, which specify the form of the relationship between predictors and response, through to nonparametric models, which have fewer prior assumptions about the shape of the fit. An important consideration for fitting such a regression model is whether all  $d$  predictors are in fact necessary. If a particular predictor has no relationship to the response, the model can be made both simpler and more accurate by removing it. In recent years there has been strong interest in techniques that automatically generate such “sparse” models. Most attention has been given to parametric forms, and in particular the linear model, where the response is assumed to vary linearly with the predictors. There has also been some investigation into variable selection for nonlinear models, notably through the use of smoothing splines and local regression.

One common feature of the existing sparse methods is that the variable selection is “global” in nature, attempting to universally include or exclude a predictor. Such an approach does not naturally reconcile well with some nonparametric techniques, such as local polynomial regression, which focus on a “local” subset of the data to estimate the response. In this local context it would be more helpful to understand local variable influence, since predictors that are irrelevant in some regions may in fact be important elsewhere in the subspace. Just as in the global setting, such information would allow us to improve the accuracy and parsimony of a model, but at a local level.

---

<sup>1</sup>Department of Mathematics and Statistics, The University of Melbourne, Melbourne, VIC 3052, Australia, e-mail: h.miller@ms.unimelb.edu.au; halpstat@ms.unimelb.edu.au

AMS 2000 subject classifications: Primary 62G08; secondary 62G20.

Keywords and phrases: variable selection, local regression, adaptive bandwidth, local variable significance.

However, this approach to variable selection can be problematic. Most notably, variable significance affects the definition of “local”. To illustrate concretely, suppose that two data points were close in every dimension except one. In typical local regression these points would not be considered close, and so the response at one point would not impact the other. If, however, we establish that the one predictor they differ by is not influential over a range that includes both these points, then they should actually be treated as neighbouring, and be treated as such in the model. Any methodology seeking to incorporate local variable influence needs to accommodate such potential situations.

Understanding local variable significance can also give additional insight into a dataset. If a variable is not important in certain regions of the support, knowledge of this allows us to discount it in certain circumstances, simplifying our understanding of the problem. For example, if none of the variables are relevant in a region, we may treat the response as locally constant and so know that we can ignore predictor effects when an observation lies in this region.

A final consideration is theoretical performance. In particular we shall present an approach that is “oracle”; that is, its performance is comparable to that of a particularly well-informed statistician, who has been provided in advance with the correct variables. It is interesting to note that variable interactions often cause sparse parametric approaches to fail to be oracle, but in the local nonparametric setting this is not an issue, because such interactions vanish as the neighbourhood of consideration shrinks.

In this paper we propose a flexible and adaptive approach to local variable selection using local polynomial regression. The key technique is careful adjustment of the local regression bandwidths to allow for variable redundancy. The method has been named LABAVS, standing for “locally adaptive bandwidth and variable selection”. Section 2 will introduce the LABAVS algorithm, including a motivating example and possible variations. Section 3 deals with theoretical properties and in particular presents a result showing that the performance of LABAVS is better than oracle when the dimension remains fixed. Section 4 presents numerical results for both real and simulated data, showing that the algorithm can improve prediction accuracy and is also a useful tool in arriving at an intuitive understanding of the data. Proofs are available in a longer version of this paper [13].

LABAVS is perhaps best viewed as an improvement to local polynomial regression, and will retain some of the advantages and disadvantages associated with this approach. In particular, it still suffers the “curse of dimensionality,” in that it struggles to detect local patterns when the dimension of genuine variables increases beyond a few. It is not the first attempt at incorporating variable selection into local polynomial regression; the papers by Lafferty and Wasserman [7] and Bertin and Lecué [1] also do this. We compare our approach to these in some detail in Section 2.6. LABAVS can also be compared to other nonparametric techniques in use for low to moderate dimensions. These include generalised additive models, MARS and tree based methods (see [6]).

The earliest work on local polynomial regression dates back to that of Nadaraya [14] and Watson [21]. General references on the subject include [19, 16] and [9]. An adaptive approach to bandwidth selection may be found in [4], although this is not in the context of variable selection. Tibshirani [17] studies the LASSO, one of the most popular sparse solutions for the linear model; more recent related work on the linear model includes that of Candès and Tao [3] and Bickel et al. [2]. Zou [24] created the adaptive version of the LASSO and proved oracle performance for it. Lin and Zhang [8] and Yuan and Lin [22] have investigated sparse solutions to

smoothing spline models. The work of Tropp [18], Zhao and Yu [23], Meinshausen et al. [12], Meinshausen and Yu [11], and Wasserman and Roeder [20] is also relevant here.

The LABAVS algorithm also bears some similarity to the approach adopted by Hall et al. [5]. There the aim was to estimate the conditional density of a response using the predictors. Cross-validation was employed and the bandwidths in irrelevant dimensions diverged, thereby greatly downweighting those components. In the present paper the focus is more explicitly on variable selection, as well as attempting to capture local variable dependencies.

Further problems worthy of investigation include the use of local variable selection as a secondary step, following global selection; developing ways of dealing with heteroscedastic variability; and exploring the case where design points are deterministic.

## 2. Model and methodology

### 2.1. Model and definitions

Suppose that we have a continuous response  $Y_i$  and a  $d$ -dimensional random predictor vector  $X_i = (X_i^{(1)}, \dots, X_i^{(d)})$  which has support on some subspace  $\mathcal{C} \subset \mathbb{R}^d$ . Further, assume that the observation pairs  $(Y_i, X_i)$  are independent and identically distributed for  $i = 1, \dots, n$ , and that  $X_i$  has density function  $f$ . The response is related to the predictors through a function  $g$ ,

$$(2.1) \quad Y_i = g(X_i) + \epsilon_i,$$

with the error  $\epsilon_i$  having zero mean and fixed variance. Smoothness conditions for  $f$  and  $g$  will be discussed in the theory section.

Local polynomial regression makes use of a kernel and bandwidth to assign increased weight to neighbouring observations compared to those further away, which will often have zero weight. We take  $K(u) = \prod_{1 \leq j \leq d} K^*(u^{(j)})$  to be the  $d$ -dimensional rectangular kernel formed from a one dimensional kernel  $K^*$  such as the tricubic kernel,

$$K^*(u^{(j)}) = (35/32)(1 - x^2)^3 \mathbb{I}(|x| < 1).$$

Assume  $K^*$  is symmetric with support on  $[-1, 1]$ . For a  $d \times d$  bandwidth matrix  $H$  the kernel with bandwidth  $H$ , denoted  $K_H$ , is

$$(2.2) \quad K_H(u) = \frac{1}{|H|^{1/2}} K(H^{-1/2}u).$$

We assume that the bandwidth matrices are diagonal,  $H = \text{diag}(h_1^2, \dots, h_d^2)$ , with each  $h_j > 0$ , and write  $H(x)$  when  $H$  varies as a function of  $x$ . Asymmetric bandwidths can be defined as having both a lower and an upper (diagonal) bandwidth matrix,  $H^L$  and  $H^U$  respectively, for a given estimation point  $x$ , rather than a single bandwidth  $H$  for all  $x$ . The kernel weight of an observation  $X_i$  at estimation point  $x$  with asymmetrical bandwidth matrices  $H^L(x)$  and  $H^U(x)$ , is

$$(2.3) \quad K_{H^U(x), H^L(x)}(X_i - x) = \prod_{j: X_i^{(j)} < x^{(j)}} \frac{1}{h_j^L(x)} K^* \left( \frac{X_i^{(j)} - x^{(j)}}{h_j^L(x)} \right) \\ \times \prod_{j: X_i^{(j)} \geq x^{(j)}} \frac{1}{h_j^U(x)} K^* \left( \frac{X_i^{(j)} - x^{(j)}}{h_j^U(x)} \right).$$

This amounts to having (possibly) different window sizes above and below  $x$  in each direction. Although such unbalanced bandwidths would often lead to undesirable bias properties in local regression, here they will be used principally to extend bandwidths in dimensions considered redundant, so this issue is not a concern.

We also allow the possibility of infinite bandwidths  $h_j = \infty$ . In calculating the kernel in (2.2) when  $h_j$  is infinite, proceed as if the  $j$ th dimension did not exist (or equivalently, as if the  $j$ th factor in a rectangular kernel product is always equal to 1). If all bandwidths are infinite, consider the kernel weight to be 1 everywhere. Although the kernel and bandwidth conditions above have been defined fairly narrowly to promote simplicity in exposition, many of these assumptions are easily generalised.

Local polynomial regression estimates of the response at point  $x$ ,  $\hat{g}(x)$ , are found by fitting a polynomial  $q$  to the observed data, using the kernel and bandwidth to weight observations. This is usually done by minimising the weighted sum of squares,

$$(2.4) \quad \sum_{i=1}^n \{Y_i - q(X_i - x)\}^2 K_H(X_i - x).$$

Once the minimisation has been performed,  $q(0)$  becomes the point estimate for  $g(x)$ . The polynomial is of some fixed degree  $p$ , with larger values of  $p$  generally decreasing bias at the cost of increased variance. Of particular interest in the theoretical section will be the local linear fit, which minimises

$$(2.5) \quad \sum_{i=1}^n \left[ \left\{ Y_i - \gamma_0 - \sum_{j=1}^d (X_i^{(j)} - x^{(j)}) \gamma_j \right\}^2 K_H(X_i - x) \right],$$

over  $\gamma_0$  and  $\gamma = (\gamma_1, \dots, \gamma_d)$ .

## 2.2. The LABAVS algorithm

Below is the LABAVS algorithm that will perform local variable selection and vary the bandwidths accordingly. The choice of  $H$  in the first step can be local or global and should be selected as for a traditional polynomial regression, using cross-validation, a plug-in estimator or some other standard technique. Methods for assessing variable significance in Step 2, and the degree of shrinkage needed in Step 4, are discussed below.

### *LABAVS algorithm*

1. Find a starting  $d \times d$  bandwidth  $H = \text{diag}(h^2, \dots, h^2)$ .
2. For each point  $x$  of a representative grid in the data support, perform local variable selection to determine disjoint index sets  $\hat{\mathcal{A}}^+(x)$ ,  $\hat{\mathcal{A}}^-(x)$ , with  $\hat{\mathcal{A}}^+(x) \cup \hat{\mathcal{A}}^-(x) = \{1, \dots, d\}$ , for variables that are considered relevant and redundant respectively.
3. For any given  $x$ , derive new local bandwidth matrices  $H^L(x)$  and  $H^U(x)$  by extending the bandwidth in each dimension indexed in  $\hat{\mathcal{A}}^-(x)$ . The resulting space given nonzero weight by the kernel  $K_{H^L(x), H^U(x)}(u - x)$  is the rectangle of maximal area with all grid points  $x_0$  inside the rectangle satisfying  $\hat{\mathcal{A}}^+(x_0) \subset \hat{\mathcal{A}}^+(x)$ . Here  $\hat{\mathcal{A}}^+(x)$  is calculated explicitly as in Step 2, or taken as the set corresponding the closest grid point to  $x$ .

4. Shrink the bandwidth slightly for those variables in  $\hat{\mathcal{A}}^+(x)$  according to the amount that bandwidths have increased in the other variables. See Section 2.4 for details.
5. Compute the local polynomial estimator at  $x$ , excluding variables in  $\hat{\mathcal{A}}^-(x)$  and using adjusted asymmetrical bandwidths  $H^L(x)$  and  $H^U(x)$ . The expression to be minimised is

$$\sum_{i=1}^n \{Y_i - q(X_i - x)\}^2 K_{H^L(x), H^U(x)}(X_i - x),$$

where the minimisation runs over all polynomials  $q$  of appropriate degree. The value of  $q(0)$  in the minimisation is the final local linear estimator.

We refer to a rectangle in Step 3 of the algorithm since we are using a product kernel bandwidth, which has nonzero support on a rectangle. The key feature of the algorithm is that variable selection directly affects the bandwidth, increasing it in the direction of variables that have no influence on the point estimator. If a variable has no influence anywhere, it has the potential to be completely removed from the local regression, reducing the dimension of the problem. For variables that have no influence in certain areas, the algorithm achieves a partial dimension reduction. The increased bandwidths reduce the variance of the estimate and Step 4 swaps some of this reduction for a decrease in the bias to further improve the overall estimator.

As a concrete example of the approach, define the following one-dimensional ‘‘Huberised’’ linear function:

$$(2.6) \quad g(x) = x^2 \mathbb{I}(0 < x \leq 0.4) + (0.8x - 0.16) \mathbb{I}(x > 0.4),$$

and let  $g(X) = g\{([X^{(1)}]_+^2 + [X^{(2)}]_+^2)^{1/2}\}$  for 2-dimensional random variable  $X = (X^{(1)}, X^{(2)})$ . Assume that  $X$  is uniformly distributed on the space  $[-2, 2] \times [-2, 2]$ . Notice that when  $X^{(1)}, X^{(2)} < 0$  the response variable  $Y$  in (2.1) is independent of  $X^{(1)}$  and  $X^{(2)}$ ; when  $X^{(1)} < 0$  and  $X^{(2)} > 0$  the response depends on  $X^{(2)}$  only; when  $X^{(1)} > 0$  and  $X^{(2)} < 0$  the response depends on  $X^{(1)}$  only; when  $X^{(1)}, X^{(2)} > 0$  the response depends on both  $X^{(1)}$  and  $X^{(2)}$ . Thus in each of these quadrants a different subset of the predictors is significant. A local approach to variable significance can capture these different dependencies, while a global variable redundancy test would not eliminate any variables.

Now consider how the algorithm applies to this example, starting with a uniform initial bandwidth of  $h = 0.5$  in both dimensions. Assuming that variable significance is determined perfectly on a dense grid, Figure 1 illustrates the adjusted bandwidths for each of the quadrants. The dots are four sample estimation points, the surrounding unit squares indicate the initial bandwidths and the dashed lines indicate how the bandwidths are modified. In the bottom left quadrant both variables are considered redundant, and so the bandwidth expands to cover the entire quadrant. This is optimal behaviour, since the true function is constant over this region, implying that the best estimator will be produced by including the whole area. In the bottom right quadrant the first dimension is significant while the second is not. Thus the bandwidth for the second dimension is ‘‘stretched’’, while the first is shrunken somewhat. Again, this is desirable for improving the estimator. The stretching in the second dimension improves the estimator by reducing the variance as more points are considered. Then the shrunken first dimension swaps some of this reduction in variance for decreased bias. Finally, in the top right quadrant, there is no change in the bandwidth since both variables are considered to be significant.




FIG 1. Bandwidth adjustments under ideal circumstances in illustrative example.

**2.3. Variable selection step**

Below are three possible ways to effect variable selection at  $x_0$  in Step 2 of the algorithm, presented in the context of local linear regression. They all make use of a tuning parameter  $\lambda$  which controls how aggressive the model is in declaring variables as irrelevant. Cross validation can be used to select an appropriate level for  $\lambda$ . So that the tuning parameters are comparable at different points in the data domain, it is useful to consider a local standardisation of the data at  $x_0$ . Define  $\tilde{X}_{x_0} = (\tilde{X}_{x_0}^{(1)}, \dots, \tilde{X}_{x_0}^{(d)})$  and  $\tilde{Y}_{x_0}$  by

$$\tilde{X}_{x_0}^{(j)} = \frac{\sum_{i=1}^n X_i^{(j)} K_H(X_i - x_0)}{\sum_{i=1}^n K_H(X_i - x_0)}, \quad \tilde{Y}_{x_0} = \frac{\sum_{i=1}^n Y_i K_H(X_i - x_0)}{\sum_{i=1}^n K_H(X_i - x_0)},$$

and define  $\tilde{X}_i = (\tilde{X}_i^{(1)}, \dots, \tilde{X}_i^{(d)})$  and  $\tilde{Y}_i$  by

$$(2.7) \quad \tilde{X}_i^{(j)} = \frac{(X_i^{(j)} - \tilde{X}_{x_0}^{(j)}) \{K_H(X_i - x_0)\}^{1/2}}{\{\sum_{i=1}^n (X_i^{(j)} - \tilde{X}_{x_0}^{(j)})^2 K_H(X_i - x_0)\}^{1/2}},$$

$$\tilde{Y}_i = (Y_i - \tilde{Y}_{x_0}) \{K_H(X_i - x_0)\}^{1/2}.$$

Notice that  $\tilde{X}$  and  $\tilde{Y}$  incorporate the weight  $K_H(X_i - x_0)$  into the expression.

1. **Hard thresholding:** Choose parameters to minimise the weighted least squares expression,

$$(2.8) \quad \sum_{i=1}^n \left\{ \tilde{Y}_i - \beta_0 - \sum_{j=1}^d \tilde{X}_i^{(j)} \beta_j \right\}^2,$$

and classify as redundant those variables for which  $|\hat{\beta}_j| < \lambda$ . This can be extended to higher degree polynomials, although performance tends to be more unstable.

2. **Backwards stepwise approach:** For each individual  $j$ , calculate the percentage increase in the sum of squares if the  $j$ th variable is excluded from the local fit. Explicitly, if  $\hat{q}$  is the optimal local fit using all variables and  $\hat{q}_j$  is the fit using all except the  $j$ th, we classify the  $j$ th variable as redundant if

$$(2.9) \quad \frac{\sum_{i=1}^n \left\{ Y_i - \hat{q}_j(\hat{X}_i) \right\}^2 K_H(\hat{X}_i) - \sum_{i=1}^n \left\{ Y_i - \hat{q}(\hat{X}_i) \right\}^2 K_H(\hat{X}_i)}{\sum_{i=1}^n \left\{ Y_i - \hat{q}(\hat{X}_i) \right\}^2 K_H(\hat{X}_i)} < \lambda,$$

where  $\hat{X}_i = X_i - x_0$ . This approach is so named as it is analogous to the first step of a backwards stepwise procedure.

3. **Local lasso:** Minimise the expression

$$(2.10) \quad \sum_{i=1}^n \left\{ \tilde{Y}_i - \gamma_0 - \sum_{j=1}^d \tilde{X}_i^{(j)} \gamma_j \right\}^2 + \lambda \sum_{j=1}^d |\gamma_j|.$$

Those variables for which  $\gamma_j$  is set to zero in this minimisation are then classified as redundant. While the normal lasso can have consistency problems [24], this local version does not since variables are asymptotically independent as  $h \rightarrow 0$ . The approach also scales naturally to higher order polynomials, provided all polynomial terms are locally standardised; a variable is considered redundant if all terms that include it have corresponding parameters set to zero by the lasso.

We have found that the first and second of the above approaches have produced the most compelling numerical results. The numerical work in Section 4 uses the first approach for linear polynomials, while the theoretical work in Section 3 establishes uniform consistency for both of the first two methods, guaranteeing oracle performance.

### 2.4. Variable shrinkage step

The variable shrinkage step depends on whether the initial bandwidth, and thus the shrunken bandwidth  $h'$ , is chosen locally or globally. Define

$$(2.11) \quad V[x, H] = \frac{\sum \{K_H(X_i - x)\}^2}{\{\sum K_H(X_i - x)\}^2},$$

where the bandwidth term in the function  $V$  is allowed to be asymmetrical, in which case we write the left hand side as  $V[x, \{H^L(x), H^U(x)\}]$ . Thus  $H$  has been replaced by the asymmetrical bandwidth  $\{H^L(x), H^U(x)\}$ , with  $H^L$  and  $H^U$  denoting the lower and upper bandwidths respectively. Then in the local case, letting  $d'(x)$  denote the cardinality of  $\hat{\mathcal{A}}(x)$ , define

$$(2.12) \quad M(x) = V[x, \{H^L(x), H^U(x)\}] / V[x, H].$$

The expression is asymptotically proportional to  $\{h'(x)\}^{-d'(x)}$  and estimates the degree of variance stabilisation resulting from the bandwidth adjustment. Using this, the correct amount of bandwidth needed in step 4 is  $h'(x) = h\{M(x)d'(x)/d\}^{1/4}$ . Since both sides of this expression depend on  $h'(x)$ , shrinkage can be approximated in the following way. Let

$$M^*(x) = V[x, \{\tilde{H}^L(x), \tilde{H}^L(x)\}] / V[x, H],$$

where  $\tilde{H}^L(x)$  and  $\tilde{H}^U(x)$  are the bandwidth matrices immediately after step 3. Then the shrunken bandwidths are  $h'(x) = h\{M^*(x)d'(x)/d\}^{1/(d'(x)+4)}$ .

In the global bandwidth case, we define

$$(2.13) \quad M[(H^L(X), H^U(X)), H] = \frac{E(V[X, (H^L(X), H^U(X))])}{E\{V[X, H]\}}.$$

This expression measures the average variance stabilisation across the domain. In this case, the shrinkage factor should satisfy

$$(2.14) \quad h' = h(M[(H^L(X), H^U(X)), H]E\{d'(X)\}/d)^{1/4}.$$

The theoretical properties in Section 3 deal with the global bandwidth scenario. The treatment for the local case is similar, except that care must be taken in regions of the domain where the function  $g$  behaves in a way that is exactly estimable by a local polynomial and thus has potentially no bias.

### 2.5. Further remarks

1. The choice of distance between grid points in Step 2 is somewhat arbitrary, but should be taken as less than  $h$  so that all data points are considered in calculations. In the asymptotic theory we let this length decrease faster than the rate of the bandwidth, and in numerical experimentation the choice impacts only slightly on the results.
2. Step 5 of the algorithm forces the estimate at point  $x$  to exclude variables indexed in  $\hat{A}^-(x)$ . An alternative is to still use all variables in the final fit. This may be advantageous in situations with significant noise, where variable admission and omission is more likely to have errors. Despite including these extra variables, the adjusted bandwidths still ensure that estimation accuracy is increased.
3. Finding the maximal rectangle for each representative point, as suggested in step 3 of the algorithm, can be a fairly intensive computational task. In our numerical work we simplified this by expanding the rectangle equally until the boundary met a “bad” grid point (i.e. a point  $x'$  such that  $\hat{A}^+(x') \not\subseteq \hat{A}^+(x)$ ). The corresponding direction was then held constant while the others continued to increase uniformly. We continued until each dimension stopped expanding or grew to be infinite. This approach does not invalidate the asymptotic results in Section 3, but there may be some deterioration in numerical performance associated with this simplification.
4. If a variable is redundant everywhere, results in Section 3 demonstrate that the algorithm is consistent; the probability that the variable is classified as redundant everywhere tends to 1 as  $n$  grows. However, the exact probability is not easy to calculate and for fixed  $n$  we may want greater control over the ability to exclude a variable completely. In such circumstances a global variable selection approach may be appropriate.
5. As noted at the start of Section 2.2, the initial bandwidth in Step 1 does not necessarily have to be fixed over the domain. For instance, a nearest neighbour bandwidth, where  $h$  at  $x$  is roughly proportional to  $f(x)^{-1}$ , could be used. Employing this approach offers many practical advantages and the theoretical basis is similar to that for the constant bandwidth. The numerical work makes use of nearest neighbour bandwidths throughout. In addition, we



could use an initial bandwidth that was allowed to vary for each variable,  $H = \text{diag}(h_1^2, \dots, h_p^2)$ . So long as, asymptotically, each  $h_j$  was equal to  $C_j h$  for some controlling bandwidth  $h$  and constant  $C_j$ , the theory would hold, although details are not pursued here.

## 2.6. Comparison to other local variable selection approaches

As mentioned in the introduction, two recent papers take a similar approach to this problem. Firstly Lafferty and Wasserman [7] introduce the rodeo procedure. This attempts to assign adaptive bandwidths based on the derivative with respect to the bandwidth for each dimension,  $\partial \hat{g}(x) / \partial h_j$ . This has the attractive feature of bypassing the actual local shape and instead focussing on whether an estimate is improved by shrinking the bandwidths. It is also a greedy approach, starting with large bandwidths in each direction and shrinking only those that cause a change in the estimator at a point. The second paper is by Bertin and Lecué [1], who implement a two step procedure to reduce the dimensionality of a local estimate. The first step fits a local linear estimate with an  $L_1$  or lasso type penalty, which identifies the relevant variables. This is followed by a second local linear fit using this reduced dimensionality. The lasso penalty they use is precisely the same as the third approach suggested in Section 2.3.

We comment on the similarities and differences of these two approaches compared to the current presentation, which are summarised in Table 1. Firstly the theoretical framework of the two other papers focus exclusively on the performance at a single point (although global oracle behaviour may be provable), while the LABAVS approach ensures uniformly oracle performance on the whole domain. The framework for the other two also assumes that variables are either active on the whole domain or redundant everywhere, while we have already discussed the usefulness of an approach that can adapt to variables that are redundant on various parts of the data. We believe this is particularly important, since local tests of variable significance will give the same results everywhere. Related to this, our method does not require an assumption of nonzero gradients (whether with respect to the bandwidth or variables) to obtain adequate theoretical performance, in contrast to the other methods. On the other hand, it is quite challenging to ensure uniform performance while allowing  $d$  to be increasing, so our presentation assumes  $d$  is fixed; the others do not. It is also worth noting that the greedy approach of Lafferty and Wasserman potentially gives it an advantage in higher dimensional situations.

While all approaches work in a similar framework, the above discussion demonstrates that there are significant differences. Our methodology may be viewed as a generalisation of the work of Bertin and Lecué, save for imposing fixed dimensionality. It can also be viewed as a competitor to the rodeo, and some numerical examples comparing the two are provided.

TABLE 1  
*Summary of locally adaptive bandwidth approaches*

	LABAVS	Rodeo	[1]
Oracle performance on entire domain	✓	✗	✗
Allows for locally redundant variables	✓	✗	✗
Relevant variables allowed to have zero gradient	✓	✗	✗
Theory allows dimension $d$ to increase with $n$	✗	✓	✓
Greedy algorithm applicable for higher dimensions	✗	✓	✗

With regards to computation time, for estimation at a single point, the rodeo is substantially faster, since calculating variable significance on a large grid of points is not required. If however we need to make predictions at a reasonable number of points, then Labavs is likely to be more efficient, since the grid calculations need only be done once, while the rodeo requires a new set of bandwidth calculations for each point.

### 3. Theoretical properties

As mentioned in the introduction, a useful means of establishing the power of a model that includes variable selection is to compare it with an oracle model, where the redundant variables are removed before the modelling is undertaken. In the linear (and the parametric) context, we interpret the oracle property as satisfying two conditions as  $n \rightarrow \infty$ :

1. the probability that the correct variables are selected converges to 1, and
2. the nonzero parameters are estimated at the same asymptotic rate as they would be if the correct variables were known in advance.

We wish to extend this notion of an oracle property to the nonparametric setting, where some predictors may be redundant. Here there are no parameters to estimate, so attention should instead be given to the error associated with estimating  $g$ . Below we define weak and strong forms of these oracle properties:

**Definition 1.** The *weak* oracle property in nonparametric regression is:

1. the probability that the correct variables are selected converges to 1, and
2. at each point  $x$  the error of the estimator  $\hat{g}(x)$  decreases at the same asymptotic rate as it would if the correct variables were known in advance.

**Definition 2.** The *strong* oracle property in nonparametric regression is:

1. the probability that the correct variables are selected converges to 1, and
2. at each point  $x$  the error of the estimator  $\hat{g}(x)$  has the same first-order asymptotic properties as it would if the correct variables were known in advance.

Observe that the weak oracle property achieves the correct rate of estimation while the strong version achieves both the correct rate and the same asymptotic distribution. The first definition is most analogous to its parametric counterpart, while the second is more ambitious in scope.

Here we establish the strong version of the nonparametric oracle property for the LABAVS algorithm, with technical details found in the Appendix [13]. We shall restrict attention to the case of fixed dimension. The case of increasing dimension can be treated similarly, provided dimension does not diverge too rapidly. Our treatment focuses on local linear polynomials, partly for convenience but also recognising that the linear factors dominate higher order terms in the asymptotic local fit. Thus our initial fit is found by minimising the expression (2.5). We impose further conditions on the kernel  $K$ :

$$(3.1) \quad \int K(z) dz = 1, \int z^{(j)} K(z) dz = 0 \text{ for each } j, \int z^{(j)} z^{(k)} K(z) dz = 0 \text{ when } j \neq k \text{ and } \int (z^{(j)})^2 K(z) dz = \mu_2(K) > 0, \text{ with } \mu_2(K) \text{ independent of } j.$$

The useful quantity  $R(K)$ , depending on the choice of kernel, is defined as

$$R(K) = \int K(z)^2 dz = \left\{ \int K^*(z^{(j)})^2 dz^{(j)} \right\}^d,$$

where  $K^*$  is the univariate kernel introduced on page 218. Let  $a_n \asymp b_n$  denote the property that  $a_n = O(b_n)$  and  $b_n = O(a_n)$ . We also require the following conditions (3.2), needed to ensure uniform consistency of our estimators.

1. For some  $\eta > 0$ , the dilation of a given compact set  $\mathcal{C}$  by balls of radius  $\eta$  is contained in the set  $\{x : f(x) > 0\}$ . Further,  $f$  and its first order partial derivatives are bounded and uniformly continuous on  $\mathcal{C}$ .
2. The kernel function  $K$  is bounded with compact support and satisfies  $|p(u)K(u) - p(v)K(v)| \leq C_1\|u - v\|$  for some  $C_1 > 0$  and all points  $u, v$  in  $\mathcal{C}$ . Here  $p(u)$  denotes a single polynomial term of the form  $\prod (u^{(j)})^{a_j}$  with the nonnegative integers  $a_j$  satisfying  $\sum a_j \leq 4$ . The bound  $C_1$  should hold for all such choices of  $p$ .
3. The function  $g$  has bounded and uniformly continuous partial derivatives up to and including order  $p$ , with  $p \geq 2$ . If  $(D^k g)(x)$  denotes the partial derivative

$$(3.2) \quad \frac{\partial^{|k|} g(x)}{\partial (x^{(1)})^{k_1} \dots \partial (x^{(d)})^{k_d}},$$

with  $|k| = \sum k_j$ , then we assume that these derivatives exist in the  $\eta$ -dilation of  $\mathcal{C}$ , when  $|k| = 2$ , and satisfy, for some constant  $C_2$  and all  $u$  and  $v$  in the dilation,

$$|h(u) - h(v)| \leq C_2\|u - v\|.$$

4.  $E(|Y|^c) < \infty$  for some  $c > 2$ .
5. The conditional density  $f_{X|Y}(x|y)$  of  $X$ , conditional on  $Y$ , exists and is bounded.
- 6.

$$\text{For some } 0 < \delta < 1, \frac{n^{1-2/\sigma} h^d}{\log n \{\log n (\log \log n)^{1+\delta}\}^{2/\sigma}} \rightarrow \infty.$$

7. The Hessian of  $g$ ,  $\mathcal{H}_g$ , is nonzero on a set of nonzero measure in  $\mathcal{C}$ .

The conditions in (3.2), except perhaps the first, are fairly natural and not overly constrictive. For example, the sixth will occur naturally for any reasonable choice of  $h$ , while the second follows easily if  $K$  has a bounded derivative. The last condition is purely for convenience in the asymptotics; if  $\mathcal{H}_g$  were zero almost everywhere then  $g$  would be linear and there would be no bias in the estimate, improving accuracy. The first condition addresses the issue of density approaching zero; in these areas the uniform convergence results will not apply, so we restrict attention to compact subset  $\mathcal{C}$ , chosen so that the density did not fall below a specified minimum. Performance inside this region would then conform to the optimal accuracies presented, while estimation outside this region would be poorer. This distinction is unavoidable, since estimation in the tails is usually problematic and it would be unusual to guarantee uniformly good performance there.

Step 1 of the LABAVS Algorithm allows the initial bandwidth to be chosen globally or locally. Here we shall focus on the global case, where an initial bandwidth  $H = \text{diag}(h^2, \dots, h^2)$  is used. Further, we assume that this  $H$  is chosen to minimise

the mean integrated squared error (MISE):

$$E \left[ \int \{\hat{g}(x) - g(x)\}^2 f(x) dx \right],$$

where the outer expectation runs over the estimator  $\hat{g}$ . It is possible to show under our assumptions that

$$(3.3) \quad h = \left[ \frac{d\sigma^2 R(K) A_C}{n\mu_2(K)^2 A_{\mathcal{H}_g}} \right]^{1/(d+4)},$$

where  $A_C$  and  $A_{\mathcal{H}_g}$  are constants, defined in the Appendix, depending only on  $\mathcal{C}$  and the function  $g$  respectively. Notice in particular that  $h \asymp n^{-1/(d+4)}$ . Details are given in Lemma A.1 in the Appendix [13].

A key result in establishing good performance, in Theorem 3.1 below, is uniform consistency of the local polynomial parameter estimates. It is a simplified version of a result of Masry [10], and no proof is included.

**Theorem 3.1.** *Suppose the conditions in (3.2) hold and we use parameter estimates from a degree  $p$  polynomial regression to estimate the partial derivatives of  $g$ . Then for each  $k$  with  $0 \leq |k| \leq p$  we have*

$$\sup_{x \in \mathcal{C}} |(\widehat{D^k g})(x) - (D^k g)(x)| = O \left[ \left( \frac{\log n}{nh^{d+2|k|}} \right)^{1/2} \right] + O(h^{p-|k|+1}) \text{ almost surely.}$$

Since the partial derivative estimate at  $x$  is proportional to the corresponding local polynomial coefficient, Theorem 1 ensures that the local polynomial coefficients are consistently estimated uniformly for suitable  $h$ . The scaling applied in (2.7) does not impact on this, as the proof of Theorem 3.2 demonstrates.

Let  $\mathcal{C}^-$  denote the points  $x \in \mathcal{C}$  satisfying  $\partial g(x)/\partial x^{(j)} = 0$  and  $\partial^2 g(x)/\partial x^{(j)2} \neq 0$ . That is,  $\mathcal{C}^-$  denotes the points where the true set of relevant variables changes. Notice that in the illustrative example in Section 2.2 we had  $\mathcal{C}^- = \{x | x^{(1)} = 0, x^{(2)} = 0\}$ . The smoothness assumed of  $g$  implies that  $\mathcal{C}^-$  has Lebesgue measure 0. Let  $\delta > 0$  and let  $\mathcal{O}_\delta$  be the smallest open set in  $\mathcal{C}^-$  such that

$$(3.4) \quad \inf_{x \in \mathcal{C} \setminus \mathcal{O}_\delta, j \in \mathcal{A}^+(x)} |\partial g(x)/\partial x^{(j)}| = \delta.$$

Intuitively this means that on the set  $\mathcal{C} \setminus \mathcal{O}_\delta$  the relevant variables have the absolute value of their corresponding parameters  $|\partial g(x)/\partial x^{(j)}|$  bounded below by  $\delta > 0$ , while irrelevant variables have  $\partial g(x)/\partial x^{(j)} = 0$ . Thus we have a ‘‘gap’’ between the true and irrelevant variables in this region that we may exploit. The volume of  $\mathcal{O}_\delta$  may be made arbitrarily small by choosing  $\delta$  small. Call the set  $\hat{\mathcal{A}}^+(x)$  in the algorithm correct if the variables in it are the same as the set of variables  $j$  with  $\partial g(x)/\partial x^{(j)} \neq 0$ . Denote the latter correct set by  $\mathcal{A}^+(x)$ .

**Theorem 3.2.** *Suppose  $\delta$  is given,  $h$  is chosen to minimise squared error as in (3.3),  $\hat{\mathcal{A}}^+(x)$  is formed using the first approach in Section 2.3, and  $\lambda$  has a growth rate between arbitrary constant multiples of  $h^2(n \log n)^{1/2}$  and  $hn^{1/2}$ . If  $f$  has bounded and uniformly continuous derivatives of degree 2, then for each  $x$  the probability that  $\hat{\mathcal{A}}^+(x)$  is correct on  $\mathcal{C} \setminus \mathcal{O}_\delta$  tends to 1 as  $n \rightarrow \infty$ . Furthermore, variables that are genuinely redundant everywhere will be correctly classified as such with probability tending to 1.*

The property (3.4) ensures that the coefficients in the local linear fit are consistently estimated with error of order  $O\{h(\log n)^{1/2}\}$ . The adjustment in (2.7) means that the actual coefficients estimated are of order  $hn^{1/2}$  times this, so the range of  $\lambda$  given is correct for separating true and redundant variables. The definition of  $\mathcal{O}_\delta$  ensures that the classification is correct on  $\mathcal{C} \setminus \mathcal{O}_\delta$ , while variables that are redundant everywhere will be recognised as such.

The next result ensures consistency for the second approach in Section 2.3. We make one further assumption, concerning the error  $\epsilon_i$ . Observe that this holds trivially if  $\epsilon_i$  is bounded. Assume that:

$$(3.5) \quad \text{there exists } C_3 \text{ such that } E(|\epsilon_i|^\alpha) \leq C_3^\alpha \text{ for } \alpha = 1, 2, 3, 4, \dots$$

**Theorem 3.3.** *Suppose  $\delta$  is given,  $h$  is chosen to minimise squared error as in (3.3), and  $\hat{A}^+(x)$  is formed using the second approach in Section 2.3. Provided that  $\lambda = o(h^2)$  and  $h^4 \log n = o(\lambda)$ , the probability that  $\hat{A}^+(x)$  is correct on  $\mathcal{C} \setminus \mathcal{O}_\delta$  tends to 1 as  $n \rightarrow \infty$ . Furthermore, variables that are genuinely redundant everywhere will be correctly classified as such with probability tending to 1.*

The previous two results ensure that we have consistent variable selection for the first two approaches in Section 2.3. Finally we can state and prove the strong oracle property for  $\mathcal{C} \setminus \mathcal{O}_\delta$ . Although the result does not cover the whole space  $\mathcal{C}$ , recall that we may make the area  $\mathcal{O}_\delta$  arbitrarily small by decreasing  $\delta$ . Furthermore, the proof implies that if we restricted attention to removing only those variables that are redundant everywhere, we would actually have the oracle property on the whole of  $\mathcal{C}$ ; however we sacrifice this performance on  $\mathcal{O}_\delta$  to improve the fit elsewhere by adjusting for locally redundant variables. In the following theorem the matrix  $\ddot{H}$  is the diagonal bandwidth matrix with bandwidth  $\infty$  for globally redundant variables and  $\ddot{h}$  for the other variables, where

$$\ddot{h} = h \left( M(\ddot{H}, H) \ddot{d}/d \right)^{1/4}.$$

Here  $\ddot{d}$  denotes the number of variables that are not globally redundant.

**Theorem 3.4.** *The estimates produced by the algorithm, where variable selection is performed using the first or second approach in Section 2.3, satisfy the strong definition of the nonparametric oracle property on  $\mathcal{C}$ . Further, when there are locally redundant variables, the expected squared error of the point estimate is decreased relative to oracle performance by a factor of  $M[(H^L(X), H^U(X)), \ddot{H}] < 1$ . That is,*

$$E\{\ddot{g}(x) - g(x)\}^2 = M[(H^L(X), H^U(X)), \ddot{H}] E\{\hat{g}(x) - g(x)\}^2$$

where  $\ddot{g}$  denotes the estimator arising from the LABAVS algorithm and  $\hat{g}$  is the oracle estimator.

In closing we should mention that, in a range of statistical problems, there is generally a trade-off between oracle performance, or even simply statistical efficiency, and robustness. In particular, methodologies that enjoy optimal or near optimal properties in relatively narrowly defined settings tend not to quite as effective as certain other, non-optimal approaches in terms of their robustness against heavy-tailed noise, or heteroscedasticity, or related departures from conventional models. There is no reason to believe that the same tendency would not be observable in the context of variable selection methods that have oracle properties.

#### 4. Numerical properties

The examples presented in this section compare the performance of two versions of the LABAVS algorithm with ordinary least squares, a traditional local linear fit, generalised additive models, tree-based gradient boosting and MARS. Table 2 describes the approaches used. The implementations of the latter four methods were from the R packages *locfit*, *gam*, *gbm* and *polspline* respectively. Tuning parameters such as bandwidths for local methods,  $\lambda$  in LABAVS, number of trees in boosting, and MARS model complexity, were chosen to give best performance for each method. The LABAVS models used the first variable selection approach of Section 2.3. All the local methods used nearest neighbour bandwidths, with the initial bandwidth chosen each time so as to minimise cross-validated squared error. The OLS linear model was included as a standard benchmark, but obviously will fail to adequately detect nonlinear features of a dataset.

**Example 1.** The example introduced in Section 2.2 was simulated with  $n = 500$ . The error for  $Y_i$  was normal with standard deviation 0.3. We first compare LABAVS to the rodeo and the methodology of Bertin and Lécué (2008) at the four representative points in Figure 1. Table 3 shows the mean squared error of the prediction compared to the true value over 100 simulations. In all cases parameters were chosen to minimise this average error. At all points the LABAVS approach performed strongest. The method of Bertin and Lécué [1] performed poorly in situations where at least one variable is redundant; this is to be expected, since it excludes the variable completely and so will incorporate regions where it is actually important, causing significant bias. The rodeo also did not perform as well; we found it tended to overestimate the optimal bandwidths in redundant directions.

We then compared LABAVS with the other model approaches which are designed to make multiple predictions, rather than a specific point. For each simulation all the models were fitted and the average squared error was estimated using a separate test set of 500 observations. The simulation was run 100 times and the average error and its associated standard deviation for each model are recorded in Table 4.

Inspection of the results shows that the LABAVS models performed best, able to allow for the different dependencies on the variables. In particular the algorithm improved on the performance of the local linear model on which it is based. The local linear regression, the boosted model and MARS also performed reasonably, while the GAM struggled with the nonadditive nature of the problem, and a strict linear model is clearly unsuitable here.

To show how effectively variable selection is for LABAVS, Figure 2 graphically represents the sets  $\hat{\mathcal{A}}^+$  at each grid point for one of the simulations, with the darkest representing  $\{\}$ , the next darkest  $\{1\}$ , the next darkest  $\{2\}$  and finally the lightest  $\{1, 2\}$ . Here the variable selection has performed well; there is some encroachment

TABLE 2  
*Approaches included in computational comparisons*

Name	Description
LABAVS-A	LABAVS with linear fit, all vars in final fit
LABAVS-B	LABAVS with linear fit, relevant vars only in final fit
LOC1	Local linear regression
OLS	Ordinary least squares linear regression
GBM	Boosting with trees, depth equal to three
GAM	Generalised additive models with splines
MARS	Multivariate adaptive regression splines

TABLE 3  
*Mean squared prediction error on sample points in Example 1*

Test Point	LABAVS-A	LABAVS-B	rodeo	Bertin and Lécué
(1, 1)	0.0022	0.0022	0.0065	0.0023
(1, -1)	0.0011	0.0013	0.0015	0.0018
(-1, 1)	0.0009	0.0011	0.0015	0.0013
(-1, -1)	0.0006	0.0007	0.0008	0.0013




FIG 2. *Plot of detected variable significance across subspace in Example 1.*

of irrelevant variables into the wrong quadrants but the selection pattern is broadly correct. The encroachment is more prevalent near the boundaries since the bandwidths are slightly larger there, to cover the same number of neighbouring points.

**Example 2.** We next show that LABAVS can effectively remove redundant variables completely. Retain the setup of Example 1, except that we add  $d^* = d - 2$  variables similarly distributed (uniform on  $[-2, 2]$ ), which have no influence on the response. Also, keep the parameters relating to the LABAVS fit the same as in the previous example, except that the cutoff for hard threshold variable selection,  $\lambda$ , is permitted to vary. Table 5 shows the proportion of times from 500 simulations that LABAVS effected complete removal of the redundant dimensions, for various  $\lambda$  and  $d^*$ . Note that the cutoff level of 0.55 is that used in the previous example, and the two genuine variables were never completely removed in any of the simulations. The results suggest that to properly exclude redundant variables, a higher threshold is

TABLE 4  
*Mean squared error sum of test dataset in Example 1*

Approach	Error	Std Dev
LABAVS-A	2.18	(0.71)
LABAVS-B	1.87	(0.65)
LOC1	2.31	(0.73)
OLS	42.85	(2.64)
GBM	2.47	(0.67)
GAM	5.93	(0.57)
MARS	2.35	(0.90)

TABLE 5  
*Proportion of simulations where redundant variables completely removed by LABAVS*

$\lambda$	Number of redundant dimensions			
	1	2	3	4
0.55	0.394	0.086	0.034	0.038
0.65	0.800	0.542	0.456	0.506
0.75	0.952	0.892	0.874	0.864
0.85	0.996	0.984	0.994	0.974
0.95	0.998	1.000	1.000	0.992




FIG 3. *Plot of detected variable significance across subspace in Example 2, under various choices for  $\lambda$ .*

needed than would otherwise be the case. This causes the final model to be slightly underfitted when compared to the oracle model, but this effect is not too severe; Figure 3 shows how the variable significance plots change for a particular simulation with different values of the cutoff. It is clear that the patterns are still broadly correct, and the results still represent a significant improvement to traditional linear regression.

**Example 3.** The first real data example used is the ozone dataset from [6], p. 175. It is the same as the air dataset in S-PLUS, up to a cube root transformation in the response. The dataset contains meteorological measurements for New York collected from May to September 1973. There are 111 observations in the dataset, a fairly moderate size. Our aim here is to predict the ozone concentration using two of the other variables, temperature and wind, scaled to unit variance when fitting the models. The smoothed perspective plot of the data in Figure 4 shows strong dependence on each of the two variables in some parts of the domain, but some sections appear flat in one or both directions in other parts. For example, the area surrounding a temperature of 70 and wind speed of 15 appears to be flat, implying that for reasonably low winds and high temperatures the ozone concentration is fairly stable. This suggests that LABAVS, by expanding the bandwidths here, could be potentially useful in reducing error. We performed a similar comparative analysis to that in Example 1, except that error rates were calculated using leave-one-out cross validation, where an estimate for each individual observations was made after using all other observations to build the model. The resulting mean squared errors and corresponding standard deviations are presented in Table 6.

The results suggest that the data is best modelled using local linear methods, and that LABAVS offers a noticeable improvement over a traditional local fit, due to its ability to improve the estimate in the presence of redundant variables. The perspective plot in left panel of Figure 4 suggests a highly non-additive model, which may explain why GAM performs poorly. There is also a large amount of



TABLE 6  
*Cross-validated mean squared error sum for the ozone dataset*

Approach	Error	Std Dev
LABAVS-A	277	(53)
LABAVS-B	284	(55)
LOC1	290	(55)
OLS	491	(110)
GBM	403	(118)
GAM	391	(98)
MARS	457	(115)

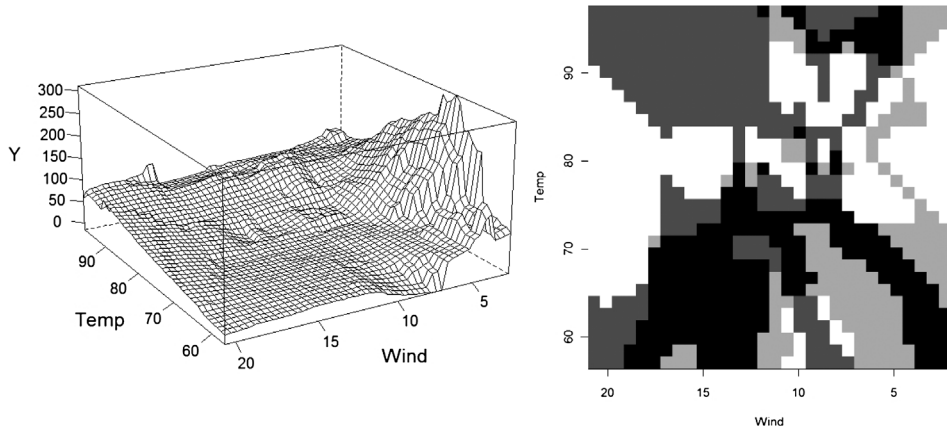


FIG 4. *Ozone dataset smoothed perspective plot and variable selection plot.*

local curvature, which hinders the OLS, GBM and MARS fits. The right panel of Figure 4 shows the variable selection results for the linear version of LABAVS across the data support, using the same shading as in Figure 1. We see that variable dependence is fairly complex, with all combinations of variables being significant in different regions. In particular, notice that the procedure has labelled both variables redundant in the region around (70, 15), confirming our initial suspicions. This plot is also highly suggestive, revealing further interesting features. For instance, there is also little dependence on wind when temperatures are relatively high. Such observations are noteworthy and potentially useful.

## References

- [1] BERTIN, K. and LECUÉ, G. (2008). Selection of variables and dimension reduction in high-dimensional non-parametric regression. *Electron. J. Statist.* **2** 1224–1241. MR2461900
- [2] BICKEL, P. J., RITOV, Y. and TSYBAKOV, A. B. (2009). Simultaneous analysis of lasso and Dantzig selector. *Ann. Statist.* **37** 1705–1732. MR2533469
- [3] CANDÈS, E. and TAO, T. (2007). The Dantzig selector: Statistical estimation when  $p$  is much larger than  $n$ . *Ann. Statist.* **35** 2313–2351. MR2382644
- [4] FAN, J. and GIJBELS, I. (1995). Data-driven bandwidth selection in local polynomial fitting: Variable bandwidth and spatial adaptation. *J. Roy. Statist. Soc. Ser. B* **57** 371–394. MR1323345
- [5] HALL, P., RACINE, J. and LI, QI (2004). Cross-validation and the estimation

- of conditional probability densities. *J. Amer. Statist. Assoc.* **99** 1015–1026. MR2109491
- [6] HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2001). *The Elements of Statistical Learning*. Springer, New York. MR1851606
- [7] LAFFERTY, J. and WASSERMAN, L. (2008). Rodeo: sparse, greedy nonparametric regression. *J. Amer. Statist. Assoc.* **36** 28–63. MR2387963
- [8] LIN, Y. and ZHANG, H. H. (2006). Component selection and smoothing in smoothing spline analysis of variance models. *Ann. Statist.* **34** 2272–2297.
- [9] LOADER, C. (1999). *Local Regression and Likelihood*. Springer, New York. MR1704236
- [10] MASRY, E. (1996). Multivariate local polynomial regression for time series: Uniform strong consistency and rates. *J. Time Series Anal.* **17** 571–600.
- [11] MEINSHAUSEN, N. and YU, B. (2009). Lasso-type recovery of sparse representations for high-dimensional data. *Ann. Statist.* **37** 246–270. MR2488351
- [12] MEINSHAUSEN, N., ROCHA, G. and YU, B. (2007). A tale of three cousins: Lasso,  $L_2$  Boosting and Dantzig. Discussion of Candès and Tao (2007). *Ann. Statist.* **35** 2373–2384. MR2382649
- [13] MILLER, H. and HALL, P. (2010). Local polynomial regression and variable selection (long version). Manuscript.
- [14] NADARAYA, E. A. (1964). On estimating regression. *Theory Probab. Appl.* **9** 141–142.
- [15] RUPPERT, D. and WAND, M. P. (1994). Multivariate locally weighted least squares regression. *Ann. Statist.* **22** 1346–1370. MR1311979
- [16] SIMONOFF, J. S. (1996). *Smoothing Methods in Statistics*. Springer, New York. MR1391963
- [17] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. MR1379242
- [18] TROPP, J. A. (2004). Greed is good: Algorithmic results for sparse approximation. *IEEE Trans. Inform. Theory* **50** 2231–2242. MR2097044
- [19] WAND, M. P. and JONES, M. C. (1995). *Kernel Smoothing*. Chapman and Hall, London. MR1319818
- [20] WASSERMAN, L. and ROEDER, K. (2009). High-dimensional variable selection. *Ann. Statist.* **37** 2178–2201. MR2543689
- [21] WATSON, G. S. (1964). Smooth regression analysis. *Sankhyā Ser. A* **26** 359–372.
- [22] YUAN, M. and LIN, YI (2006). Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **68** 49–67. MR2212574
- [23] ZHAO, P. and YU, B. (2007). Stagewise lasso. *J. Mach. Learn. Res.* **8** 2701–2726. MR2383572
- [24] ZOU, H. (2006). The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.* **101** 1418–1429. MR2279469