

# Conditional regression for single-index models

ALESSANDRO LANTERI<sup>1,a</sup>, MAURO MAGGIONI<sup>2,b</sup> and STEFANO VIGOGNA<sup>3,c</sup>

<sup>1</sup>DEMM, Università degli Studi di Milano, Milano, Italy. <sup>a</sup>[alessandro.lanteri@unimi.it](mailto:alessandro.lanteri@unimi.it)

<sup>2</sup>Department of Mathematics and Department of Applied Mathematics & Statistics, Johns Hopkins University, Baltimore, USA. <sup>b</sup>[mauromaggionijhu@icloud.com](mailto:mauromaggionijhu@icloud.com)

<sup>3</sup>MaLGa, DIBRIS, Università degli Studi di Genova, Genova, Italy. <sup>c</sup>[vigogna@dibris.unige.it](mailto:vigogna@dibris.unige.it)

The single-index model is a statistical model for intrinsic regression where responses are assumed to depend on a single yet unknown linear combination of the predictors, allowing to express the regression function as  $\mathbb{E}[Y|X] = f(\langle v, X \rangle)$  for some unknown *index* vector  $v$  and *link* function  $f$ . Conditional methods provide a simple and effective approach to estimate  $v$  by averaging moments of  $X$  conditioned on  $Y$ , but depend on parameters whose optimal choice is unknown and do not provide generalization bounds on  $f$ . In this paper we propose a new conditional method converging at  $\sqrt{n}$  rate under an explicit parameter characterization. Moreover, we prove that polynomial partitioning estimates achieve the 1-dimensional min-max rate for regression of Hölder functions when combined to any  $\sqrt{n}$ -convergent index estimator. Overall this yields an estimator for dimension reduction and regression of single-index models that attains statistical optimality in quasilinear time.

*Keywords:* Single-index model; dimension reduction; nonparametric regression; finite-sample bounds

## 1. Introduction

Consider the standard regression problem of estimating a function  $F : \mathbb{R}^d \rightarrow \mathbb{R}$  from  $n$  samples  $\{(X_i, Y_i)\}_{i=1}^n$ , where the  $X_i$ 's are independent realizations of a predictor variable  $X \in \mathbb{R}^d$ ,

$$Y_i = F(X_i) + \zeta_i, \quad i = 1, \dots, n, \quad (1)$$

and the  $\zeta_i$ 's are realizations, independent among themselves and of the  $X_i$ 's, of a random variable  $\zeta$  modeling noise. Under rather general assumptions on  $\zeta$  and the distribution  $\rho$  of  $X$ , if we only know that  $F$  is  $s$ -Hölder regular (and, say, compactly supported), it is well-known that the min-max nonparametric rate for estimating  $F$  in  $L^2(\rho)$  is  $n^{-s/(2s+d)}$  [22]. This is an instance of the *curse of dimensionality*: the rate slows down dramatically as the dimension  $d$  increases. Many regression models have been introduced throughout the decades to circumvent this phenomenon; see, for example, the classical reference [45]. When the covariates are intrinsically low-dimensional, concentrating on an unknown low-dimensional set, several estimators have been proved to converge at rates that are optimal with respect to the intrinsic dimension [3,34,35,41,42]. In other models, the domain may be high-dimensional, but the function itself is assumed to depend only on a small number of features. A classical case is the so-called *single-index model*, where  $F$  has the structure

$$F(x) = f(\langle v, x \rangle) \quad (2)$$

for some *index vector*  $v \in \mathbb{R}^d$  (that we may assume unitary without loss of generality) and *link function*  $f : \mathbb{R} \rightarrow \mathbb{R}$ . In this context one may consider different estimation problems, depending on whether  $f$  is known (e.g. in logistic regression) or both  $f$  and  $v$  are unknown. We are interested in the latter case. Clearly, if  $v$  was known we could learn  $f$  by solving a 1-dimensional regression problem, which may be done efficiently for large classes of functions  $f$ . So the question is: what is the price to pay for not knowing  $v$ ?

It was conjectured in [45] that the min-max rate for regression of single-index models is  $n^{-s/(2s+1)}$ , that is, the min-max rate for univariate functions: no statistical cost would have to be paid. This rate was achieved for pointwise convergence with kernel estimators in [25, Theorem 3.3] and [26, Section 2.5], observing that the index can be learned at the parametric rate  $n^{-1/2}$ . Based on these results or on similar heuristics, a wide part of literature focused on index estimation, setting aside the regression problem. From this perspective, the main point is that the estimation of the index  $v$  can be carried out at parametric rate in spite of the unknown nonparametric nonlinearity  $f$ . The existence of an estimator converging in  $L^2(\rho)$  at rate  $n^{-s/(2s+1)}$  was established in [22, Corollary 22.1]. A proof that  $n^{-s/(2s+1)}$  is also a lower min-max rate of convergence (in  $L^2(\rho)$ ), thus completing the proof of Stone's conjecture, can be found in [20, Theorem 2].

Granted that the estimation of the index does not entail additional *statistical* costs (in terms of regression rates), a different but no less important problem is determining the *computational* cost (expressed as the number of required elementary operations) to implement a statistically optimal estimator for the single-index model. The rate in [22, Corollary 22.1] is obtained by a least squares joint minimization over  $v$  and  $f$ , but no executable algorithm is provided. In [20] it was proposed to aggregate local polynomial estimators on a lattice of the unit sphere, yielding an adaptive, universal min-max estimator, although at the expense of a possibly exponential number of operations  $\Omega(n^{(d-1)/2})$ . While a heuristic faster algorithm is therein also proposed, its statistical effectiveness was not proved.

Several other methods for the estimation of  $v$  or  $f$  were developed over the years. A first category includes semiparametric methods based on maximum likelihood estimation [7,8,14–16,24,29]. M-estimators produce  $\sqrt{n}$ -consistent index estimates under general assumptions, but their implementation is cumbersome and computationally demanding, in that depends on sensitive bandwidth selections for kernel smoothing and relies on high-dimensional joint optimization. An attempt at avoiding the data sparsity problem was made by [12], which proposed a fixed-point iterative scheme only involving 1-dimensional nonparametric smoothers. Direct methods such as the average derivative estimation (ADE [25,44]) estimate the index vector exploiting its proportionality with the derivative of the regression function. Early implementations of this idea suffer from the curse of dimensionality due to kernel estimation of the gradient, while later iterative modifications [28] provide  $\sqrt{n}$ -consistency under mild assumptions, yet not eliminating the computational overhead. More recently, Isotron [30,31] and SILO [21] achieved linear complexity, but the proven regression rate, even if independent of  $d$ , is not min-max (albeit SILO focuses on the  $n \ll d$  regime, rather than the limit  $n \rightarrow \infty$  as here and most past work). In a different yet related direction, the even more recent [1] showed that convex neural networks can adapt to a large variety of statistical models, including single-index; however, they do not match the optimal learning rate (even for the single-index case), and at the same time do not have associated fast algorithms.

Meanwhile, a line of research addressed *sufficient dimension reduction* [37] in the more general *multi-index model* (or a slight extension thereof), where  $F$  depends on multiple  $k < d$  index directions spanning an unknown index subspace. Along this thread we can find nonparametric methods extending ADE to multiple indices, such as structural adaptation ([13,27]), the outer product of gradients (OPG [49]) and the minimum average variance estimation (MAVE [48,49]). Alternatively, *conditional methods* derive their estimates from statistics of the conditional distribution of the explanatory variable  $X$  given the response variable  $Y$ . Prominent examples are sliced inverse regression (SIR [18,40]), sliced average variance estimation (SAVE [9]), simple contour regression (SCR [39]) and its generalizations (e.g. GCR [39], DR [38]). Conditional methods are appealing for several reasons. Compared to nonparametric methods, their implementation is straightforward, consisting of noniterative computation of “sliced” empirical moments and having only one “slicing” parameter to tune. Moreover, they are computationally efficient and simple to analyze, enjoying  $\sqrt{n}$ -consistency and, in most cases, complexity linear in the sample size and quadratic in the ambient dimension. On the downside, this comes in general at the cost of stronger distributional assumptions, and no theoretically optimal choice of the slicing

parameter is known [11, p. 75]. Moreover, while conditional methods offer a provable, efficient solution for sufficient dimension reduction, they do not address the problem of estimating the link function on the estimated index space. We summarize the key properties for these and other aforementioned techniques in Table 1.

**Table 1.** Proven rate in  $L^2(\rho)$  (up to log factors) and computational cost of several methods for index estimation and/or regression in single-index models, together with salient assumptions on the model.

	Performance				Assumptions			
	Proven rate		Computational cost		$X$	$f$	$\zeta$	
	$\hat{v}$	$\hat{f}$	$\hat{v}$	$\hat{f}$				
<b>SIR</b>	[40]	$n^{-1/2}$	–	$d^2 n \log n$	–	linear $\mathbb{E}[X v^T X]$	N/A	N/A
<b>SAVE</b>	[9]	$n^{-1/2}$	–	$d^2 n \log n$	– const Cov $[X v^T X]$	linear $\mathbb{E}[X v^T X]$ ,	N/A	N/A
<b>SCR</b>	[39]	$n^{-1/2}$	–	$d^2 n^2 \log n$	– const Cov $[X v^T X]$	linear $\mathbb{E}[X v^T X]$ , monotone	stochastically density of $\zeta - \tilde{\zeta}$	decreasing
<b>DR</b>	[38]	$n^{-1/2}$	–	$d^2 n \log n$	– const Cov $[X v^T X]$	linear $\mathbb{E}[X v^T X]$ ,	N/A	N/A
<b>ADE</b>	[28]	$n^{-1/2}$	–	$d^2 n^2 \log n$	–	$C^0$ positive density	$C^2$	Gaussian
<b>rMAVE</b>	[48]	$n^{-1/2}$	N/A	$d^2 n^2$ per iteration $\mathbb{E} X ^6 < \infty$	–	$v^T X$ has $C^3$ density,	$C^3$	$\mathbb{E} Y ^3 < \infty$
<b>Aggregation</b>	[20]	–	$n^{-\frac{s}{2s+1}}$	$(n \log n)^d$ lower bounded density	–	compact supported	$C^s$	$\sigma(X)\mathcal{N}(0,1)$
<b>SIsotron</b>	[30]	N/A	$n^{-1/6}$	$(\frac{n}{d \log n})^{1/3} dn \log n$ Lipschitz	–	bounded	monotone,	bounded
<b>SILO</b>	[21]	$n^{-1/4}$	$n^{-1/8}$	$dn$	$n \log n$ Lipschitz	Gaussian	monotone,	bounded
<b>SVR</b>		$n^{-1/2}$	$n^{-\frac{s}{2s+1}}$	$d^2 n \log n$	$n \log n$ $\text{Var}[w^T X v^T X] \gtrsim 1$	linear $\mathbb{E}[X v^T X]$ , monotone, $C^s$	coarsely	sub-Gaussian

In this work we introduce a new estimator and a corresponding algorithm, called Smallest Vector Regression (SVR), that are statistically optimal and computationally efficient. Our dimension reduction technique falls in the category of conditional methods. Unlike existing studies for similar approaches, we are able to provide a characterization for the parameter selection, and bound both the index estimation and the regression errors. Since regression is performed using standard piecewise polynomial estimates on the projected samples after and independently of the index estimation step, our regression bounds hold conditioned to any index estimation method of sufficient accuracy. Our analysis yields convergence by proving finite-sample bounds in high probability. The resulting statements are stronger compared to the ones in the available literature on conditional methods, where typically only asymptotic convergence, at most, is established. As a side note, SVR has been empirically tested with success also in the multi-index model, but our analysis, and therefore our exposition, will be restricted to the single-index case (for a related analysis of the multi-index model, see [33]). In summary, the contributions of this work are:

1. We introduce a new conditional regression method that combines accuracy, robustness and low computational cost. This method is multiscale and sheds light on parameter choices that are important in theory and practice, and are mostly left unaddressed in other techniques.
2. We prove strong, finite-sample convergence bounds, both in probability and in expectation, for the index estimate of our conditional method.
3. We prove that the population statistics used to construct conditional polynomial partitioning estimation (up to degree 1) are Hölder continuous with respect to the index estimation. This allows to bridge the gap between a good estimator of the index subspace and the performance of regression on the estimated subspace.
4. We prove that all  $\sqrt{n}$ -convergent index estimation methods lead to the min-max 1-dimensional rate of convergence when combined with polynomial partitioning estimates (up to degree 1).
5. Using the above, we provide an algorithm for the estimation of the single-index model with theoretical guarantees of optimal convergence in quasilinear time.

The paper is organized as follows. In Section 2 we review several conditional regression methods, and introduce our new estimator; in Section 3 we analyze the converge of our method and establish min-max rates for regression conditioned on any sufficiently accurate index estimate; in Section 4 we conduct several numerical experiments, both validating the theory and exploring numerically aspects of various techniques that are not covered by theoretical results; in the supplementary material [36] we collect additional proofs and technical results.

Notation			
symbol	definition	symbol	definition
$C, c$	positive absolute constants	$\ A\ $	spectral norm of a matrix $A$
$a \lesssim b$	$a \leq Cb$ for some positive absolute constant $C$	$\lambda_i(A)$	$i$ -th largest eigenvalue of a matrix $A$
$a \asymp b$	$a \lesssim b$ and $b \lesssim a$	$ I $	Lebesgue measure of an interval $I$
$\ u\ $	Euclidean norm of a vector $u$	$\#\mathcal{S}$	cardinality of a set $\mathcal{S}$
$B(x, r)$	Euclidean ball of center $x$ and radius $r$	$\mathbb{1}(E)$	indicator function of an event $E$
$\text{span}\{\mathcal{S}\}$	linear span of a set $\mathcal{S}$	$\mathbb{P}(E)$	probability of an event $E$
$\{\mathcal{S}\}^\perp$	orthogonal complement of a set $\mathcal{S}$	$\mathbb{E}[X]$	expectation of a random variable $X$
$P_u$	orthogonal projection onto $\text{span}\{u\}$	$\text{Var}[X]$	variance of a r.v. $X$
$\angle(u, v)$	angle between $\text{span} u$ and $\text{span} v$	$\text{Cov}[X]$	covariance matrix of a r.v. $X$
$\langle u, v \rangle$	inner product of vectors $u$ and $v$	$X   Y$	r.v. $X$ conditioned on r.v. $Y$

## 2. Conditional regression methods

We consider the regression problem as in (1), within the single-index model, with the definition and notation as in (2). When  $f$  is at least Lipschitz, (2) implies  $\nabla F(x) \in \text{span}\{v\}$  for a.e.  $x$ ; this is the reason why we may refer to  $v$  as the gradient direction. Given  $n$  independent copies  $(X_i, Y_i)$ ,  $i = 1, \dots, n$ , of the random pair  $(X, Y)$ , we will construct estimators  $\widehat{v}$  and  $\widehat{f}$ , and derive separate and compound non-asymptotic error bounds in probability and expectation. Our method is conditional in two ways:

- 1) the estimator  $\widehat{v}$  is a statistic of the conditional distributions of the  $X_i$ 's given the  $Y_i$ 's (restricted in suitable intervals);
- 2) the estimator  $\widehat{f}$  is conditioned on the estimate  $\widehat{v}$ .

Several conditional methods for step 1) have been previously introduced, see e.g. [9,38–40]. Our error bounds for step 2) are independent of the particular method used in 1), only requiring a minimal non-asymptotic convergence rate. For these reasons, one may as well consider other existing or new methods for 1), even non conditional, and check for each one the convergence rate needed to pair it with 2).

The common idea of all conditional methods is to compute statistics of the predictor  $X$  conditioned on the response  $Y$ . Conditioning on  $Y$ , one introduces anisotropy in the distribution of  $X$ , forcing it to reveal the index structure through its moments, be they means (SIR) or variances (SAVE, SCR, DR).

Before introducing SVR, we will review the two methods that have strongest connections with ours, namely SIR and SAVE. For consistency with SVR, we will present SIR and SAVE through a particular multiscale implementation. This will allow to progressively define the objects SVR is built upon, facilitating the comparison.

### 2.1. Sliced inverse regression

Sliced Inverse Regression [40] (SIR) estimates the index vector by a principal component analysis of a sliced empirical approximation of the inverse regression curve  $\mathbb{E}[X|Y]$ . Samples on this curve are obtained by slicing the range of the function and computing sample means of the corresponding approximate level sets. In our version of SIR, we take dyadic partitions  $\{S_{l,h}\}_{h=1}^{2^l}$ ,  $l \in \mathbb{Z}$ , of a subinterval of the range of  $Y$ , where each  $S_{l,h}$  is an interval of length  $\asymp 2^{-l}$ . Let  $n_{l,h} = \#\{Y_i \in S_{l,h}\}$ . After calculating the sample mean for each slice,

$$\widehat{\mu}_{l,h} = \frac{1}{n_{l,h}} \sum_i X_i \mathbb{1}\{Y_i \in S_{l,h}\}, \quad h = 1, \dots, 2^l,$$

SIR outputs  $\widehat{v}_l$  as the eigenvector of largest eigenvalue of the weighted covariance matrix

$$\widehat{M}_l = \sum_h \widehat{\mu}_{l,h} \widehat{\mu}_{l,h}^T \frac{n_{l,h}}{n}.$$

Note that the population limits of  $\widehat{\mu}_{l,h}$  and  $\widehat{M}_l$  are, respectively,

$$\mu_{l,h} = \mathbb{E}[X | Y \in S_{l,h}], \quad M_l = \sum_h \mu_{l,h} \mu_{l,h}^T \mathbb{P}\{Y \in S_{l,h}\}.$$

### 2.2. Sliced average variance estimation

Sliced Average Variance Estimation [9] (SAVE) generalizes SIR to second order moments. After slicing the range of  $Y$  and computing the centers  $\widehat{\mu}_{l,h}$ 's, it goes further and construct the sample covariance on each slice:

$$\widehat{\Sigma}_{l,h} = \frac{1}{n_{l,h}} \sum_i (X_i - \widehat{\mu}_{l,h})(X_i - \widehat{\mu}_{l,h})^T \mathbb{1}\{Y_i \in S_{l,h}\}.$$

Then, it averages the  $\widehat{\Sigma}_{l,h}$ 's and defines  $\widehat{v}_l$  as the eigenvector of largest eigenvalue of

$$\widehat{\Sigma}_l = \sum_h (I - \widehat{\Sigma}_{l,h})^2 \frac{n_{l,h}}{n}.$$

The matrices  $\widehat{\Sigma}_{l,h}$  and  $\widehat{\Sigma}_l$  are empirical estimates of

$$\Sigma_{l,h} = \text{Cov}[X | Y \in S_{l,h}], \quad \Sigma_l = \sum_h (I - \Sigma_{l,h})^2 \mathbb{P}\{Y \in S_{l,h}\}.$$

### 2.3. Smallest Vector Regression (SVR)

This is the new method we propose here. We perform a local principal component analysis on each approximate level set obtained by multiscale slices of  $Y$ . Because of the special structure (2), we expect each (approximate) level set to be narrow along  $v$  and spread out along the orthogonal directions. Thus, the smallest principal component should approximate  $v$ . Once we have an estimate for  $v$ , we can project down the  $d$ -dimensional samples and perform nonparametric regression of the 1-dimensional function  $f$ . The method consists of the following steps:

- 1.a) Construct a multiscale family of dyadic partitions of a subinterval  $S$  of the range of  $Y$ :

$$\{S_{l,h}\}_{h=1}^{2^l}, \quad l \in \mathbb{Z}.$$

For each  $l$ ,  $\{S_{l,h}\}_h$  is a partition of  $S$  with  $|S_{l,h}| = |S|2^{-l}$ .

- 1.b) Let  $\mathcal{H}_l$  be the set of  $h$ 's such that  $n_{l,h} \geq 2^{-l}n$ . For  $h \in \mathcal{H}_l$ , let  $\widehat{v}_{l,h}$  be the eigenvector corresponding to the *smallest* eigenvalue of  $\widehat{\Sigma}_{l,h}$ .
- 1.c) Compute the eigenvector  $\widehat{v}_l$  corresponding to the *largest* eigenvalue of

$$\widehat{V}_l = \frac{1}{\sum_{h \in \mathcal{H}_l} n_{l,h}} \sum_{h \in \mathcal{H}_l} \widehat{v}_{l,h} \widehat{v}_{l,h}^T n_{l,h}.$$

- 2) Regress  $f$  using a dyadic polynomial estimator  $\widehat{f}_{j|\widehat{v}_l}$  at scale  $j \geq 0$  on the samples  $(\langle \widehat{v}_l, X_i \rangle, Y_i)$ ,  $i = 1, \dots, n$  (more details in Section 2.4). Return  $\widehat{F}_{j|\widehat{v}_l}(x) = \widehat{f}_{j|\widehat{v}_l}(\langle \widehat{v}_l, x \rangle)$ .

While SVR shares step 1.a) with SIR, it differs from SIR in step 1.b), where it takes conditional (co)variance statistics in place of conditional means, and in step 1.c), where it averages smallest-variance directions rather than means. We may regard SAVE and SVR as two different modifications of SIR to higher order statistics, which allows in general for better and more robust estimates (see Section 4.1). The fundamental difference between SVR and SAVE is that SVR computes local estimates of the index vector which then aggregates in a global estimate, while SAVE first aggregates local information and then computes a single global estimate.

In our theoretical analysis we will assume  $Y$  sub-Gaussian, and choose  $S$  deterministically with radius multiple of a standard deviation proxy of  $|Y|$ , in order to ensure that  $S$  contains most (but possibly not all) samples (see Section 3.1). However, a practically good choice for  $S$  is the smallest closed interval containing all the observed data, namely  $[\min_i Y_i, \max_i Y_i]$ , in which case  $S$  would be a random interval. If  $Y$  is bounded in a known interval, one may also take  $S$  as that interval, so that  $S$  would be again deterministic.

### 2.4. Conditional partitioning estimators

In step 2) we use piecewise polynomial estimators in the spirit of [4,5]: these techniques are based on partitioning the domain (here, in a multiscale fashion), and constructing a local polynomial on each element of the partition by solving a least squares fitting problem. A global estimator is then obtained by summing the local polynomials over a certain partition (possibly using a partition of unity to obtain smoothness across the boundaries of the partition elements). The degree of the local polynomials needed to obtain optimal rates depends on the regularity of the function, and may be chosen adaptively if such regularity is unknown. A proper partition (or scale) is then chosen to minimize the expected mean squared error (MSE), by classical bias-variance trade-off.

In detail, given an estimated direction  $\widehat{v}$ , our step 2) consists of:

2.a) Construct a multiscale family of dyadic partitions of an interval  $I$ :

$$\{I_{j,k}\}_{k \in \mathcal{K}_j}, \quad j \in \mathbb{Z}.$$

For each  $j$ ,  $\{I_{j,k}\}_{k \in \mathcal{K}_j}$  is a partition of  $I$  with  $|I_{j,k}| = |I|2^{-j}$ .

2.b) For each  $I_{j,k}$ , compute the best fitting  $m$ -order polynomial

$$\widehat{f}_{j,k|\widehat{v}} = \arg \min_{\deg(p) \leq m} \sum_I |Y_i - p(\langle \widehat{v}, X_i \rangle)|^2 \mathbb{1}\{\langle \widehat{v}, X_i \rangle \in I_{j,k}\}.$$

2.c) Sum all (truncated) local estimates  $\widehat{f}_{j,k|\widehat{v}}$  over the partition  $\{I_{j,k}\}_{k \in \mathcal{K}_j}$ :

$$\widehat{f}_{j|\widehat{v}}(t) = \sum_{k \in \mathcal{K}_j} T_M[\widehat{f}_{j,k|\widehat{v}}(t)] \mathbb{1}\{t \in I_{j,k}\},$$

where  $T_M[t] = \text{sign}(t) \min\{|t|, M\}$  for some  $M \in (0, +\infty]$ .

The final estimator of  $F$  at scale  $j$  and conditioned on  $\widehat{v}$  is given by

$$\widehat{F}_{j|\widehat{v}}(x) = \widehat{f}_{j|\widehat{v}}(\langle \widehat{v}, x \rangle).$$

In SVR, step 2) is carried out on  $\widehat{v} = \widehat{v}_l$ , yielding for each  $l$  a multiscale family of local polynomials  $\widehat{f}_{j,k|l}$  and global estimators  $\widehat{f}_{j|l}$ ,  $\widehat{F}_{j|l}$ . However, we will prove results on the performance of 2) also when  $\widehat{v}$  is the output of any estimator with  $n^{-1/2}$  probabilistic convergence rate (Theorem 2).

Note that for SVR, but also for SIR and SAVE, the final estimator  $\widehat{F}_{j|l}$  depends on two scale parameters,  $l$  and  $j$ , which may be chosen independently. Our analysis yields optimal choices for these two scale parameters; the scale  $2^{-l}$  at which the direction  $v$  is estimated will not be finer than the noise level, while a possibly finer partition with  $j > l$  may be selected to improve the polynomial fit, allowing the estimator  $\widehat{F}_{j|l}$  to de-noise its predictions, provided that enough training samples are available (see Figure 1).

For theoretical purposes, we will assume  $X$  sub-Gaussian, and take  $I$  as a deterministic interval with radius proportional to a standard deviation proxy of  $\|X\|$  (see Section 3.2). Again, the practitioner may choose  $I$  as the random interval  $[\min_i \langle \widehat{v}, X_i \rangle, \max_i \langle \widehat{v}, X_i \rangle]$ , or as the projection of the support of  $X$ , if this is known. The truncation level  $M$  is ideally  $\|F\|_\infty (= \|f\|_\infty)$ , or a proxy thereof.


We report the complete sequence of steps run by SVR in Algorithm 1. The time complexity is shown in Table 2. Note that 2.c) has only an evaluation cost, i.e.  $\widehat{f}_{j|\widehat{v}}$  does not need to be constructed, but only evaluated.

### 3. Analysis of convergence

We next state and discuss several assumptions that are going to be useful for our theoretical analysis. We point out that, when stating an assumption, such an assumption it is not intended to be implicitly assumed once for all throughout the paper. Rather, in each result we will explicitly write which assumptions are being used, referring to the corresponding labels. We start by collecting a few basic requirements on the distributions of  $X$ ,  $Y$  and  $\zeta$ :

- (X)  $X$  has sub-Gaussian distribution with variance proxy  $R^2$ .
- (Y)  $Y$  has sub-Gaussian distribution with variance proxy  $R^2$ .





**Figure 1.** Local linear estimator (red) at different scales  $l$ , to regress the function  $f$  (green) from noisy samples (black). The horizontal axis is  $\langle v, x \rangle$ , while of course the estimator  $\hat{f}_{j|l}$  is a function of  $\langle \hat{v}, x \rangle$  and may appear multi-valued in  $\langle v, x \rangle$ . For small  $l$  (top row) the error in the estimation of the index  $v$  is large, leading to poor regression estimates regardless of the regression scale  $j$ . For larger  $l$  (bottom row) a good accuracy for the index vector  $v$  is achieved, and the estimator is able to approximate the function even below the noise level and the non-monotonicity scale (e.g. for  $j = 6$ ); overfitting occurs for  $j$  too large (e.g.  $j = 12$  in this case).

(Z)  $\zeta$  is sub-Gaussian with variance proxy  $\sigma^2$ .

(X), (Y) and (Z) are standard assumptions in regression analysis *tout court*. The following is instead typical of single-index models:

(LCM)  $\mathbb{E}[X | \langle v, X \rangle]$  is linear in  $\langle v, X \rangle$ .

(LCM) is commonly referred to as the *linear conditional mean* assumption [40, Condition 3.1]. For standardized  $X$ , it is equivalent to requiring  $\mathbb{E}[X | \langle v, X \rangle] = P_v X$  [37, Lemma 1.1]. Every elliptical distribution satisfies (LCM) for every  $v$  [6, Corollary 5], and conversely [19]. While it does introduce some symmetry, it is less restrictive than it may seem. It has been shown to hold approximately in high dimension, where most low-dimensional projections are nearly normal [17,23]. (LCM) is introduced to ensure that  $\hat{v}$  is an unbiased estimate of  $v$  [10, Theorem 1]. As an alternative to (LCM), one can rely on the stronger assumption

(SMD) for every  $w \in \text{span}\{v\}^\perp$ ,  $\langle w, X \rangle | \langle v, X \rangle$  has symmetric distribution (i.e. has same distribution as  $-\langle w, X \rangle | \langle v, X \rangle$ ).

The restriction (SMD) to *symmetrical marginal distributions* will be in fact our working assumption. Assuming (SMD) rather than (LCM) is for us a purely technical choice, dictated by our goal of disentangling the dependence of the SVR estimator  $\hat{v}_l$  on the slicing parameter  $l$ . To be more specific, our motivation stems from the application of the Bernstein inequality to concentrate the sample statistics



---

**Algorithm 1:** SVR

---

**Input :** samples  $\{(X_i, Y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \mathbb{R}$ , intervals  $S, I$ , polynomial degree  $m \in \mathbb{N}$ , truncation level  $M \in (0, \infty]$ .

**Output:**  $\widehat{v}_l$  estimate of  $v$ ,  $\widehat{f}_{j|l}$  estimate of  $f$ .

---

- 1.a) construct  $\{S_{l,h}\}_{l,h}$ , dyadic decomposition of  $S$ ;
  - 1.b) for all  $h \in \mathcal{H}_l = \{h : n_{l,h} \geq 2^{-l}n\}$ , where  $n_{l,h} = \#\{Y_i \in S_{l,h}\}$ , compute
 
$$\widehat{\mu}_{l,h} = \frac{1}{n_{l,h}} \sum_i X_i \mathbb{1}\{Y_i \in S_{l,h}\},$$

$$\widehat{\Sigma}_{l,h} = \frac{1}{n_{l,h}} \sum_i (X_i - \widehat{\mu}_{l,h})(X_i - \widehat{\mu}_{l,h})^T \mathbb{1}\{Y_i \in S_{l,h}\},$$

$$\widehat{v}_{l,h}, \text{ the eigenvector of } \widehat{\Sigma}_{l,h} \text{ corresponding to the smallest eigenvalue;}$$
  - 1.c) compute  $\widehat{v}_l$ , the eigenvector of  $\widehat{V}_l = \frac{1}{\sum_{h \in \mathcal{H}_l} n_{l,h}} \sum_{h \in \mathcal{H}_l} \widehat{v}_{l,h} \widehat{v}_{l,h}^T n_{l,h}$  corresponding to the largest eigenvalue;
  - 2.a) construct  $\{I_{j,k}\}_{j,k}$ , dyadic decomposition of  $I$ ;
  - 2.b) compute  $\widehat{f}_{j,k|l} = \arg \min_{\deg(p) \leq m} \sum_i |Y_i - p(\langle \widehat{v}_l, X_i \rangle)|^2 \mathbb{1}\{\langle \widehat{v}_l, X_i \rangle \in I_{j,k}\}$ ;
  - 2.c) define  $\widehat{f}_{j|l}(t) = \sum_k T_M[\widehat{f}_{j,k|l}(t)] \mathbb{1}\{t \in I_{j,k}\}$ .
- 

**Table 2.** Computational cost breakdown for SVR.

	<b>task</b>	<b>computational cost</b>
1.a)	dyadic decomposition of the range	$O(n \log n)$
1.b)	PCA on level sets	$O(d^2 n \log n)$
1.c)	PCA of local directions	$O(d^2 n \log n)$
2.a)	dyadic decomposition of the domain	$O(n \log n)$
2.b)	$m$ -order polynomial regression	$O(m^2 n \log n)$
	total	$O((d^2 + m^2)n \log n)$

to their population limits. Since  $X$  and  $Y$  will be assumed to be sub-Gaussian, and therefore possibly unbounded, we may not use the standard Bernstein inequality for bounded variables [46, Lemma 2.2.9], and we should rather use the sub-exponential version [46, Lemma 2.2.11]. As it turns out, the slicing parameter would be encoded in the sub-Gaussian norm of  $X$  conditioned on a slice, which is however an implicit quantity, difficult to characterize in general. On the other hand, Bernstein for bounded variables, although less general, provides a sharper bound, with probability scaling with the variance rather than the sub-Gaussian norm. Consequently, it allows to encode the slicing parameter in an explicit variance term. But in order to apply Bernstein for bounded variables, and thus obtain a more interpretable bound, we need to condition the statistics of interest in suitable balls of constant radius (see Section 3.1). Such conditioning would in general break (LCM), but not (SMD), hence the need to assume (SMD).

Although stronger than (LCM), the condition (SMD) is still satisfied by all elliptical distributions, as well as other types of symmetric distributions along the index direction. While ellipticity serves often as a simplified model for the analysis of conditional methods, it is not necessarily crucial when it comes to applications. In particular, the performance of SVR seems to be robust to even considerable deviations from such assumptions, as we illustrate experimentally in Section 4.1.

In addition to (LCM) or (SMD), second order methods usually require the so-called *constant conditional variance* assumption [9, p. 2117]:

(CCV)  $\text{Cov}[X | \langle v, X \rangle]$  is nonrandom.

For standardized  $X$ , and assuming **(LCM)**, **(CCV)** is equivalent to  $\text{Cov}[X \mid \langle v, X \rangle] = P_{v^\perp}$  [37, Corollary 5.1]. **(CCV)** is true for the normal distribution [37, Proposition 5.1], and again approximately true in high dimension [17,23]. Some care is required when assuming both **(LCM)** and **(CCV)**: imposing **(LCM)** for every  $v$  is equivalent to assuming spherical symmetry [19], and the only spherical distribution satisfying **(CCV)** is the normal distribution [32, Theorem 7]. For this reason we will introduce a relaxation of **(CCV)** in the next subsection.

We present separately bounds on the estimation of  $v$  in the next Section 3.1, and on the regression of  $f$  in Section 3.2. Our main result, Theorem 2, will give in particular a near-optimal high probability bound on the SVR estimator for  $F$ , in the form

$$\left(\mathbb{E}_X[|\widehat{F}(X) - F(X)|^2]\right)^{1/2} \leq K(d\sqrt{\log d} + d^s)(\log n)^{s/2} \left(\frac{\log n}{n}\right)^{\frac{s}{2s+1}}$$

where  $F$  is  $s$ -Hölder continuous with  $s \in [1/2, 2]$ , and  $K$  is a constant independent of  $n$  and  $d$ .

### 3.1. Index estimation bounds

The goal of this section is to obtain finite-sample bounds on the SVR index estimate  $\widehat{v}_l$ . In order to obtain an explicit characterization with respect to the parameter  $l$ , we introduce a variant of  $\widehat{v}_l$  based on bounded statistics. As discussed earlier in Section 3, similar but less interpretable bounds for SVR can be proved for unbounded statistics.

Consider the event

$$\mathcal{B} = \{\|X\| \leq C_X \sqrt{d}R, |Y| \leq C_Y R\}$$

for some  $C_X, C_Y \geq 1$  independent of  $n$  and that will be considered as fixed constant from now on. We define the following bounded versions of  $\mu_{l,h}, \Sigma_{l,h}$ :

$$\mu_{l,h}^b = \mathbb{E}[X \mid Y \in S_{l,h}, \mathcal{B}], \quad \Sigma_{l,h}^b = \text{Cov}[X \mid Y \in S_{l,h}, \mathcal{B}].$$

Given the event

$$\mathcal{B}_i = \{\|X_i\| \leq C_X \sqrt{d}R, |Y_i| \leq C_Y R\},$$

we define

$$n^b = \sum_i \mathbb{1}(\mathcal{B}_i), \quad n_{l,h}^b = \sum_i \mathbb{1}(\{Y_i \in S_{l,h}\} \cap \mathcal{B}_i).$$

Note that  $n^b$  is a random number, but assuming **(X)** and **(Y)** it is larger than a constant fraction of  $n$ , with high probability (see [36, Lemma 4]). The sample counterparts of  $\mu_{l,h}^b, \Sigma_{l,h}^b$  are

$$\widehat{\mu}_{l,h}^b = \frac{1}{n_{l,h}^b} \sum_i X_i \mathbb{1}(\{Y_i \in S_{l,h}\} \cap \mathcal{B}_i), \quad \widehat{\Sigma}_{l,h}^b = \frac{1}{n_{l,h}^b} \sum_i (X_i - \widehat{\mu}_{l,h}^b)(X_i - \widehat{\mu}_{l,h}^b)^T \mathbb{1}(\{Y_i \in S_{l,h}\} \cap \mathcal{B}_i).$$

We denote by  $v_{l,h}^b$  and  $\widehat{v}_{l,h}^b$  the eigenvectors of smallest eigenvalue of  $\Sigma_{l,h}^b$  and  $\widehat{\Sigma}_{l,h}^b$ , respectively. Finally, let  $\mathcal{H}_l^b = \{h : n_{l,h}^b \geq 2^{-l}n^b\}$ . We define  $\widehat{v}_l^b$  as the eigenvector of largest eigenvalue of

$$\widehat{V}_l^b = \frac{1}{\sum_{h \in \mathcal{H}_l^b} n_{l,h}^b} \sum_{h \in \mathcal{H}_l^b} \widehat{v}_{l,h}^b (\widehat{v}_{l,h}^b)^T n_{l,h}^b.$$

In accordance with  $\mathcal{B}$ , we pick the interval  $S$  in SVR as

$$S = [-C_Y R, C_Y R].$$

Furthermore, we will be assuming lower bounded conditional variance on the distribution conditioned on  $\mathcal{B}$ :

(LCV) There is  $\alpha \geq 1$  such that  $\text{Var}[\langle w, X \rangle \mid \langle v, X \rangle, \mathcal{B}] \geq R^2/\alpha$  almost surely for all  $w \in \text{span}\{v\}^\perp \cap \mathbb{S}^{d-1}$ .

This assumption can be seen as a relaxation of the standard (CCV) we mentioned in Section 3. It should be contrasted with (X): while (X) imposes a uniform upper bound on the variance of all marginals of  $X$ , (LCV) requires that the variance of all (conditional) marginals orthogonal to  $v$  is uniformly lower bounded by a quantity of the same order. This clearly helps SVR in not mistaking a generic direction of small variance with the actual direction of the index. A practical way towards meeting (X) and (LCV) (but also (SMD)) is standardizing the data, which is the main reason why conditional methods commonly include this preprocessing step. Since we directly assume (X) and (LCV) (as well as (SMD)), we did not include standardization as part of SVR, and therefore of our theoretical analysis.

Besides the distributional assumptions discussed so far, we introduce the following functional property:

( $\Omega$ ) There are  $\omega \geq 0$  and  $L > 0$  such that, for every interval  $T$  with  $|T| \geq \omega$ ,  $|\min f^{-1}(T), \max f^{-1}(T)| \leq L|T|$ .

Assumption ( $\Omega$ ) may be regarded as a large scale sub-Lipschitz property. Note that, if  $f$  is bi-Lipschitz, then ( $\Omega$ ) is satisfied with  $\omega = 0$ . However, ( $\Omega$ ) for  $\omega > 0$  does not imply that  $f$  is monotone; it relaxes monotonicity to monotonicity “at scales larger than  $\omega$ ”.

We may now state the main result for the SVR estimator of  $v$ :

**Theorem 1 (SVR).** *Suppose (X), (Y), (Z), ( $\Omega$ ), (SMD) and (LCV) hold true. Let  $l$  be such that  $2^{-l} \lesssim (L\sqrt{\alpha})^{-1}$  and  $|S_{l,h}| \gtrsim \max\{\sigma, \omega\}$  for all  $h \in \mathcal{H}_l^b$ . Then, for  $n$  large enough so that  $\frac{n}{\sqrt{\tau \log n}} \gtrsim \alpha^2 L^2 (t + l + \log d) d 2^l$ , we have*

$$(a) \quad \|P_{\hat{v}_l^b} - P_v\| \lesssim \alpha L \sqrt{t + l + \log d} \sqrt{\frac{d 2^{-l}}{n/\sqrt{\tau \log n}}} \text{ with probability higher than } 1 - e^{-t} - n^{-\tau}.$$

Moreover, if  $\frac{n}{\log n \sqrt{\log n}} \gtrsim \alpha^2 L^2 d 2^l$ , then

$$(b) \quad \mathbb{E}[\|P_{\hat{v}_l^b} - P_v\|^2] \lesssim \alpha^2 L^2 (l + \log d) \frac{d 2^{-l}}{n/\sqrt{\log n}}.$$

If, furthermore,  $|\zeta| \leq \sigma$  a.s., then (a) and (b) hold with  $n/\sqrt{\log n}$  replaced by  $n$  and without  $n^{-\tau}$  in (a).

Theorem 1 not only proves convergence for SVR, but also shows that finer scales give more accurate estimates, provided the number of local samples  $n_{l,h}^b$  is not too small and we stay above the critical scales  $\sigma$  and  $\omega$ , representing the noise and the non-monotonicity levels, respectively. Without assumption (LCM), both SIR and SVR provide biased estimates of the index vector; it is not known if such bias is removable. Nevertheless, Theorem 1 suggests that the estimation error of SVR could be driven to 0 by increasing  $l$ , only limited by the constraint of keeping the scale larger than  $\max\{\sigma, \omega\}$ . On the other hand, for distributions not satisfying the assumptions above, the inverse regression curve can deviate considerably from the direction  $v$ , regardless of the size of the noise (see Figure 3). In SVR, assuming

for a moment monotonicity ( $\omega = 0$ ) and zero noise ( $\sigma = 0$ ), choosing the scale parameter  $l$  according to the lower bound on  $n$  yields a  $O(n^{-2})$  convergence rate for the MSE, disregarding logarithmic factors.

To prove Theorem 1, we first establish bounds on the local statistics involved in the computation of the estimator of  $v$ :

**Proposition 1.** *Suppose (Z) and (Ω) hold true. Let  $T \subset S$  be a bounded interval with  $|T| \geq \omega$ . Then:*

(a) *For every  $i = 1, \dots, n$  and every  $\tau \geq 1$ ,*

$$\mathbb{P}\{|\langle v, X_i \rangle - \mathbb{E}[\langle v, X \rangle \mid Y \in T, \mathcal{B}]| \gtrsim L(|T| + \sqrt{\tau \log n} \sigma) \mid Y_i \in T, \mathcal{B}_i\} \leq 2n^{-\tau}.$$

*If  $|\zeta| \leq \sigma$  a.s., then*

$$\mathbb{P}\{|\langle v, X_i \rangle - \mathbb{E}[\langle v, X \rangle \mid Y \in T, \mathcal{B}]| \lesssim L(|T| + \sigma) \mid Y_i \in T, \mathcal{B}_i\} = 1.$$

(b)  $\text{Var}[\langle v, X \rangle \mid Y \in T, \mathcal{B}] \lesssim L^2(|T|^2 + \sigma^2).$

**Proof.** Let  $Z_t = (-\sqrt{2(t+1)}\sigma, -\sqrt{2t}\sigma) \cup [\sqrt{2t}\sigma, \sqrt{2(t+1)}\sigma)$  for  $t \in \mathbb{N}$ . To prove (a) we first note that, thanks to (Z), we have  $\zeta_i \in \bigcup_{t \leq \tau \log n} Z_t$  with probability higher than  $1 - 2n^{-\tau}$ , for every  $i$ . Conditioned on this event and on  $Y_i \in T$ ,  $\langle v, X_i \rangle \in f^{-1}(T + \bigcup_{t \leq \tau \log n} Z_t)$ . On the other hand,  $\mathbb{E}[\langle v, X \rangle \mid Y \in T, \mathcal{B}, \zeta \in Z_t] \in [\min f^{-1}(T + Z_t), \max f^{-1}(T + Z_t)]$ . It follows from assumption (Ω) that

$$|\langle v, X_i \rangle - \mathbb{E}[\langle v, X \rangle \mid Y \in T, \mathcal{B}, \zeta \in Z_t]| \lesssim L(|T| + \sqrt{\max\{t, \tau \log n\}}\sigma).$$

Thus, by the law of total expectation,

$$\begin{aligned} |\langle v, X_i \rangle - \mathbb{E}[\langle v, X \rangle \mid Y \in T, \mathcal{B}]| &\lesssim \sum_{t=0}^{\infty} |\langle v, X_i \rangle - \mathbb{E}[\langle v, X \rangle \mid Y \in T, \mathcal{B}, \zeta \in Z_t]| \mathbb{P}\{\zeta \in Z_t\} \\ &\lesssim L\left(|T| + \sqrt{\tau \log n} \sigma + \sigma \sum_{t > \tau \log n} \sqrt{t} e^{-t}\right) \\ &\lesssim L(|T| + \sqrt{\tau \log n} \sigma). \end{aligned}$$

The case where  $|\zeta| \leq \sigma$  almost surely is similar and simpler. For (b), we write

$$\begin{aligned} \text{Var}[\langle v, X \rangle \mid Y \in T, \mathcal{B}] &= \mathbb{E}[(\langle v, X \rangle - \mathbb{E}[\langle v, X \rangle \mid Y \in T, \mathcal{B}])^2 \mid Y \in T, \mathcal{B}] \\ &= \sum_{t=0}^{\infty} \mathbb{E}[(\langle v, X \rangle - \mathbb{E}[\langle v, X \rangle \mid Y \in T, \mathcal{B}])^2 \mid Y \in T, \mathcal{B}, \zeta \in Z_t] \mathbb{P}\{\zeta \in Z_t\}. \end{aligned}$$

Conditioned on  $\zeta \in Z_t$ , assumption (Ω) gives

$$\begin{aligned} |\langle v, X \rangle - \mathbb{E}[\langle v, X \rangle \mid Y \in T, \mathcal{B}]| &\lesssim \sum_{s=0}^{\infty} |\langle v, X \rangle - \mathbb{E}[\langle v, X \rangle \mid Y \in T, \mathcal{B}, \zeta \in Z_s]| \mathbb{P}\{\zeta \in Z_s\} \\ &\lesssim L\left(|T| + \sqrt{t}\sigma + \sigma \sum_{s=0}^{\infty} \sqrt{s} e^{-s}\right) \\ &\lesssim L(|T| + \sqrt{t}\sigma), \end{aligned}$$

whence

$$\text{Var}[\langle v, X \rangle \mid Y \in T, \mathcal{B}] \lesssim L^2 \left( |T|^2 + \sigma^2 \sum_{t=0}^{\infty} t e^{-t} \right) \lesssim L^2 (|T|^2 + \sigma^2). \quad \square$$

**Proposition 2.** Suppose (X), (Y), (Z), (Ω), (SMD) and (LCV) hold true. Then, for every  $l$  such that  $2^{-l} \lesssim (L\sqrt{\alpha})^{-1}$  and  $|S_{l,h}| \geq \max\{\sigma, \omega\}$ ,  $h \in \mathcal{H}_l^b$ ,  $v$  is the eigenvector of smallest eigenvalue of  $\Sigma_{l,h}^b$ , and

$$\lambda_{d-1}(\Sigma_{l,h}^b) - \lambda_d(\Sigma_{l,h}^b) \gtrsim R^2/\alpha.$$

**Proof.** We have

$$\Sigma_{l,h}^b = \text{Cov}[X \mid Y \in S_{l,h}, \mathcal{B}] = \mathbb{E}[XX^T \mid Y \in S_{l,h}, \mathcal{B}] - \mathbb{E}[X \mid Y \in S_{l,h}, \mathcal{B}]\mathbb{E}[X^T \mid Y \in S_{l,h}, \mathcal{B}].$$

Since  $X$  is independent of  $\zeta$ , (2) implies that  $X$  is independent of  $Y$  given  $\langle v, X \rangle$ , hence

$$\begin{aligned} \mathbb{E}[XX^T \mid Y \in S_{l,h}, \mathcal{B}] &= \mathbb{E}[\mathbb{E}[XX^T \mid \langle v, X \rangle, Y \in S_{l,h}, \mathcal{B}] \mid Y \in S_{l,h}, \mathcal{B}] \\ &= \mathbb{E}[\mathbb{E}[XX^T \mid \langle v, X \rangle, \mathcal{B}] \mid Y \in S_{l,h}, \mathcal{B}]. \end{aligned}$$

For the same reason, we have

$$\mathbb{E}[X \mid Y \in S_{l,h}, \mathcal{B}] = \mathbb{E}[\mathbb{E}[X \mid \langle v, X \rangle, Y \in S_{l,h}, \mathcal{B}] \mid Y \in S_{l,h}, \mathcal{B}] = \mathbb{E}[\mathbb{E}[X \mid \langle v, X \rangle, \mathcal{B}] \mid Y \in S_{l,h}, \mathcal{B}].$$

Since  $X$  satisfies (SMD), so does  $X \mid \mathcal{B}$ . Thus, in particular,

$$\mathbb{E}[X \mid \langle v, X \rangle, \mathcal{B}] = v \mathbb{E}[\langle v, X \rangle \mid \langle v, X \rangle, \mathcal{B}].$$

Now, applying [10, Theorem 1.a] to  $X \mid \mathcal{B}$ , we get that  $v$  is an eigenvector of  $\Sigma_{l,h}^b$ . Let  $w$  be a unitary vector orthogonal to  $v$ . Then

$$w^T \mathbb{E}[X \mid Y \in S_{l,h}, \mathcal{B}] = \mathbb{E}[\langle w, v \rangle \mathbb{E}[\langle v, X \rangle \mid \langle v, X \rangle, \mathcal{B}] \mid Y \in S_{l,h}, \mathcal{B}] = 0.$$

Moreover, assumption (LCV) gives

$$w^T \mathbb{E}[XX^T \mid Y \in S_{l,h}, \mathcal{B}]w = \mathbb{E}[\text{Var}[\langle w, X \rangle \mid \langle v, X \rangle, \mathcal{B}] \mid Y \in S_{l,h}, \mathcal{B}] \geq R^2/\alpha.$$

Therefore,

$$\min_{\substack{w \in \text{span}\{v\}^\perp \\ \|w\|=1}} w^T \text{Cov}[X \mid Y \in S_{l,h}, \mathcal{B}]w \geq R^2/\alpha.$$

To upper bound  $v^T \Sigma_{l,h}^b v$ , note that  $|S_{l,h}| = |S|2^{-l} \lesssim R2^{-l}$ . Thus, assumption (Ω) implies by Proposition 1(b) that

$$v^T \Sigma_{l,h}^b v \lesssim L^2 R^2 2^{-2l}.$$

We finally put together lower and upper bound. Taking  $2^l \gtrsim L\sqrt{\alpha}$  implies that  $\lambda_d(\Sigma_{l,h}^b)$  is the eigenvalue associated to  $v$  and yields the desired inequality.  $\square$

We now establish convergence in probability for the local estimators  $\hat{v}_{l,h}^b$ .

**Proposition 3 (local SVR).** *Suppose  $(\mathbf{X})$ ,  $(\mathbf{Y})$ ,  $(\mathbf{Z})$ ,  $(\Omega)$ ,  $(\text{SMD})$  and  $(\text{LCV})$  hold true. Then, for every  $l$  such that  $2^{-l} \lesssim (L\sqrt{\alpha})^{-1}$  and  $|S_{l,h}| \gtrsim \max\{\sigma, \omega\}$  for all  $h \in \mathcal{H}_l^b$ , for every  $\varepsilon > 0$  and  $\tau \geq 1$ ,*

$$\mathbb{P}\{\|P_{\widehat{v}_{l,h}^b} - P_v\| > \varepsilon \mid n_{l,h}^b\} \lesssim d \left[ \exp\left(-\frac{cn_{l,h}^b \varepsilon^2}{\alpha^2 L^2 d \sqrt{\tau} \log n(2^{-2l} + 2^{-l} \varepsilon)}\right) + \exp\left(-\frac{cn_{l,h}^b}{\alpha^2 d}\right) \right] + n^{-\tau}.$$

If  $|\zeta| \leq \sigma$  a.s., then

$$\mathbb{P}\{\|\widehat{v}_{l,h} - v\| > \varepsilon \mid n_{l,h}^b\} \lesssim d \left[ \exp\left(-\frac{cn_{l,h}^b \varepsilon^2}{\alpha^2 L^2 d(2^{-2l} + 2^{-l} \varepsilon)}\right) + \exp\left(-\frac{cn_{l,h}^b}{\alpha^2 d}\right) \right].$$

**Proof.** Since  $\|P_{\widehat{v}_{l,h}^b} - P_v\| = \sin \angle(\widehat{v}_{l,h}^b, v) \leq 1$  (see [43, Ch. 1, Sec. 5.3, Theorem 5.5]), we can assume  $\varepsilon^2 \leq \varepsilon \leq 1$  whenever needed. The Davis–Kahan Theorem [2, Theorem VII.3.1], together with [43, Ch. 1, Sec. 5.3, Theorem 5.5] and Proposition 2, gives

$$\|P_{\widehat{v}_{l,h}^b} - P_v\| \leq \frac{\|v^T(\widehat{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)\|}{|\lambda_{d-1}(\widehat{\Sigma}_{l,h}^b) - \lambda_d(\Sigma_{l,h}^b)|}.$$

By Proposition 2 and Weyl’s inequality [47, Theorem 4.5.3] we get

$$\begin{aligned} |\lambda_{d-1}(\widehat{\Sigma}_{l,h}^b) - \lambda_d(\Sigma_{l,h}^b)| &\geq \lambda_d(\Sigma_{l,h}^b) - \lambda_{d-1}(\Sigma_{l,h}^b) - |\lambda_{d-1}(\widehat{\Sigma}_{l,h}^b) - \lambda_{d-1}(\Sigma_{l,h}^b)| \\ &\gtrsim R^2/\alpha - \|\widehat{\Sigma}_{l,h}^b - \Sigma_{l,h}^b\|. \end{aligned}$$

We bound  $\|\widehat{\Sigma}_{l,h}^b - \Sigma_{l,h}^b\|$  using the Bernstein inequality [47, Theorem 5.4.1]. First, we introduce the intermediate term

$$\widetilde{\Sigma}_{l,h}^b = \frac{1}{n_{l,h}^b} \sum_i (X_i - \mu_{l,h}^b)(X_i - \mu_{l,h}^b)^T \mathbb{1}(\{Y_i \in S_{l,h}\} \cap \mathcal{B}_i),$$

and split  $\widehat{\Sigma}_{l,h}^b - \Sigma_{l,h}^b$  into

$$\widehat{\Sigma}_{l,h}^b - \Sigma_{l,h}^b = \widetilde{\Sigma}_{l,h}^b - \Sigma_{l,h}^b - (\widehat{\mu}_{l,h}^b - \mu_{l,h}^b)(\widehat{\mu}_{l,h}^b - \mu_{l,h}^b)^T.$$

Conditioned on  $\mathcal{B}_i$  we have  $\|X_i - \mu_{l,h}^b\|^2 \lesssim R^2 d$ , hence

$$\mathbb{P}\{\|\widetilde{\Sigma}_{l,h}^b - \Sigma_{l,h}^b\| \gtrsim R^2/\alpha \mid n_{l,h}^b\} \lesssim d \exp\left(-c \frac{n_{l,h}^b}{\alpha^2 d}\right).$$

Also, by similar calculations,  $\|\widehat{\mu}_{l,h}^b - \mu_{l,h}^b\|^2 \lesssim R^2/\alpha$  with same probability. We now apply the Bernstein inequality to concentrate  $v^T(\widehat{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)$ . We have

$$\begin{aligned} \|v^T(\widehat{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)\| &\leq \|v^T(\widetilde{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)\| + |v^T(\widehat{\mu}_{l,h}^b - \mu_{l,h}^b)| \|\widehat{\mu}_{l,h}^b - \mu_{l,h}^b\| \\ &\lesssim \|v^T(\widetilde{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)\| + \sqrt{d}R |v^T(\widehat{\mu}_{l,h}^b - \mu_{l,h}^b)|. \end{aligned}$$

We start with  $\|v^T(\tilde{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)\|$ . By Proposition 1(a), conditioned on  $Y_i \in S_{l,h}$  and  $\mathcal{B}_i$  we have, with probability no lower than  $1 - 2n^{-\tau}$ ,

$$|v^T(X_i - \mu_{l,h}^b)| \|X_i - \mu_{l,h}^b\| \lesssim LR^2 \sqrt{d\tau \log n} 2^{-l},$$

or  $|v^T(X_i - \mu_{l,h}^b)| \|X_i - \mu_{l,h}^b\| \lesssim LR^2 \sqrt{d} 2^{-l}$  when  $|\zeta| \leq \sigma$ . Next, we estimate the variance. We have

$$\begin{aligned} \|v^T(\tilde{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)\|^2 &= v^T(\tilde{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)^2 v \\ &= v^T(\tilde{\Sigma}_{l,h}^b)^2 v - v^T \tilde{\Sigma}_{l,h}^b \Sigma_{l,h}^b v - v^T \Sigma_{l,h}^b \tilde{\Sigma}_{l,h}^b v + v^T(\Sigma_{l,h}^b)^2 v, \end{aligned}$$

hence, taking the expectation,

$$\mathbb{E}[\|v^T(\tilde{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)\|^2 | n_{l,h}^b] = \mathbb{E}[v^T(\tilde{\Sigma}_{l,h}^b)^2 v | n_{l,h}^b] - v^T(\Sigma_{l,h}^b)^2 v,$$

where

$$\begin{aligned} \mathbb{E}[v^T(\tilde{\Sigma}_{l,h}^b)^2 v | n_{l,h}^b] &= \frac{1}{(n_{l,h}^b)^2} v^T \mathbb{E} \left[ \left( \sum_i (X_i - \mu_{l,h}^b)(X_i - \mu_{l,h}^b)^T \right)^2 \middle| n_{l,h}^b \right] v \\ &\leq \frac{1}{n_{l,h}^b} v^T \mathbb{E}[(X - \mu_{l,h}^b) \|X - \mu_{l,h}^b\|^2 (X - \mu_{l,h}^b)^T | Y \in S_{l,h}, \mathcal{B}] v + v^T(\Sigma_{l,h}^b)^2 v \\ &\leq \frac{1}{n_{l,h}^b} dR^2 \mathbb{E}[(v^T(X - \mu_{l,h}^b))^2 | Y \in S_{l,h}, \mathcal{B}] + v^T(\Sigma_{l,h}^b)^2 v \\ &= \frac{1}{n_{l,h}^b} dR^2 \text{Var}[v^T X | Y \in S_{l,h}, \mathcal{B}] + v^T(\Sigma_{l,h}^b)^2 v. \end{aligned}$$

Thus, Proposition 1(b) gives

$$\mathbb{E}[\|v^T(\tilde{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)\|^2 | n_{l,h}^b] \leq \frac{1}{n_{l,h}^b} L^2 dR^4 2^{-2l}.$$

Therefore, by the Bernstein inequality we obtain

$$\mathbb{P}\{\|v^T(\tilde{\Sigma}_{l,h}^b - \Sigma_{l,h}^b)\| > \alpha^{-1} R^2 \varepsilon | n_{l,h}^b\} \lesssim d \exp \left( -c \frac{n_{l,h}^b \varepsilon^2}{\alpha^2 L^2 d \sqrt{\tau \log n} (2^{-2l} + 2^{-l} \varepsilon)} \right),$$

without  $\sqrt{\tau \log n}$  if  $|\zeta| \leq \sigma$ . We are now left to bound  $v^T(\tilde{\mu}_{l,h}^b - \mu_{l,h}^b)$ . By Proposition 1(a) we have that, conditioned on  $Y_i \in S_{l,h}$  and  $\mathcal{B}_i$ , with probability higher than  $1 - 2n^{-\tau}$ ,

$$|v^T X_i| \lesssim LR \sqrt{\tau \log n} 2^{-l},$$

or  $|v^T X_i| \lesssim LR 2^{-l}$  if  $|\zeta| \leq \sigma$ . Moreover, by Proposition 1(b),

$$\text{Var}[v^T X | Y \in S_{l,h}, \mathcal{B}] \lesssim L^2 R^2 2^{-2l}.$$

Thus, the Bernstein inequality gives

$$\mathbb{P}\{\|v^T(\tilde{\mu}_{l,h}^b - \mu_{l,h}^b)\| > \alpha^{-1} d^{-1/2} R \varepsilon | n_{l,h}^b\} \lesssim d \exp \left( -c \frac{n_{l,h}^b \varepsilon^2}{\alpha^2 L^2 d \sqrt{\tau \log n} (2^{-2l} + 2^{-l} \varepsilon)} \right),$$



where again the factor  $\sqrt{\tau \log n}$  can be dropped in the case of bounded noise. □

We are finally in a position to prove Theorem 1.

**Proof of Theorem 1.** Since  $\lambda_1(vv^T) = 1$  and  $\lambda_2(vv^T) = 0$ , the Davis–Kahan Theorem [50, Corollary 1], along with [43, Ch. 1, Sec. 5.3, Theorem 5.5]), yields

$$\begin{aligned} \|P_{\widehat{v}_l^b} - P_v\| &\lesssim (\lambda_1(vv^T) - \lambda_2(vv^T))^{-1} \left\| \frac{1}{\sum_{h \in \mathcal{H}_{l,h}^b} n_{l,h}^b} \sum_{h \in \mathcal{H}_{l,h}^b} \widehat{v}_{l,h}^b (\widehat{v}_{l,h}^b)^T n_{l,h}^b - vv^T \right\| \\ &= \left\| \frac{1}{\sum_{h \in \mathcal{H}_{l,h}^b} n_{l,h}^b} \sum_{h \in \mathcal{H}_{l,h}^b} (P_{\widehat{v}_{l,h}^b} - P_v) n_{l,h}^b \right\| \\ &\leq \frac{1}{\sum_{h \in \mathcal{H}_{l,h}^b} n_{l,h}^b} \sum_{h \in \mathcal{H}_{l,h}^b} \|P_{\widehat{v}_{l,h}^b} - P_v\| n_{l,h}^b. \end{aligned}$$

For each  $h \in \mathcal{H}_l$  we have  $n_{l,h}^b \geq 2^{-l} n^b$ , where  $n^b \gtrsim n$  with probability higher than  $1 - 2e^{-cn}$  (by [36, Lemma 4]), hence  $n_{l,h}^b \gtrsim 2^{-l} n$ . We thus apply Proposition 3. Taking the union bound over  $h \in \mathcal{H}_l$  gives  $\|P_{\widehat{v}_{l,h}^b} - P_v\| \leq \varepsilon$  for all  $h \in \mathcal{H}_l$  with probability higher than

$$1 - C2^l d \left[ \exp\left(-\frac{c2^{-l} n \varepsilon^2}{\alpha^2 L^2 d \sqrt{\tau} \log n (2^{-2l} + 2^{-l} \varepsilon)}\right) + \exp\left(-\frac{c2^{-l} n}{\alpha^2 d}\right) \right] - n^{-\tau}.$$

We now constrain  $\varepsilon \leq 2^{-l}$ , so that the two exponentials are bounded by

$$\begin{aligned} \exp\left(-\frac{c2^{-l} n \varepsilon^2}{\alpha^2 L^2 d \sqrt{\tau} \log n (2^{-2l} + 2^{-l} \varepsilon)}\right) + \exp\left(-\frac{c2^{-l} n}{\alpha^2 d}\right) &\leq \exp\left(-\frac{c2^{-l} n \varepsilon^2}{\alpha^2 L^2 d \sqrt{\tau} \log n 2^{-2l}}\right) + \exp\left(-\frac{c2^{-l} n \varepsilon^2}{\alpha^2 d 2^{-2l}}\right) \\ &\lesssim \exp\left(-\frac{cn \varepsilon^2}{\alpha^2 L^2 d \sqrt{\tau} \log n 2^{-l}}\right). \end{aligned} \tag{*}$$

Then, for  $t > 0$ , we take  $\varepsilon = c^{-1/2} \alpha L \sqrt{t + l + \log d} 2^{-l/2} \sqrt{\frac{d}{n/\sqrt{\tau} \log n}}$ , with the constraint  $\varepsilon \leq 2^{-l}$  translating to  $\frac{n}{\sqrt{\tau} \log n} \gtrsim \alpha^2 L^2 (t + l + \log d) d 2^l$ , which leads to (a). For (b), we first condition on  $\mathcal{Z} = \{|\zeta_i| \leq \sqrt{2\tau \log n} \sigma \text{ for all } i\}$  and calculate

$$\begin{aligned} \mathbb{E}[\|P_{\widehat{v}_l^b} - P_v\|^2] - n^{-\tau} &\lesssim \int_0^1 \varepsilon \mathbb{P}\{\|P_{\widehat{v}_l^b} - P_v\| > \varepsilon \mid \mathcal{Z}\} d\varepsilon \\ &= \int_0^{2^{-l}} \varepsilon \mathbb{P}\{\|P_{\widehat{v}_l^b} - P_v\| > \varepsilon \mid \mathcal{Z}\} d\varepsilon + \int_{2^{-l}}^1 \varepsilon \mathbb{P}\{\|P_{\widehat{v}_l^b} - P_v\| > \varepsilon \mid \mathcal{Z}\} d\varepsilon \\ &\leq \int_0^{2^{-l}} \min\left\{1, 2^l d \exp\left(-\frac{c(n/\sqrt{\tau} \log n) \varepsilon^2}{\alpha^2 L^2 d 2^{-l}}\right)\right\} \varepsilon d\varepsilon \\ &\quad + \int_{2^{-l}}^1 2^l d \exp\left(-\frac{c(n/\sqrt{\tau} \log n)}{\alpha^2 L^2 d 2^l}\right) \varepsilon d\varepsilon \end{aligned}$$

$$\lesssim \alpha^2 L^2 d \log(2^l d) \frac{2^{-l}}{n/\sqrt{\tau} \log n} + 2^l d \exp\left(-\frac{c(n/\sqrt{\tau} \log n)}{\alpha^2 L^2 d 2^l}\right), \tag{**}$$

where for the integral  $\int_0^{2^{-l}}$  we have used (\*), for the integral  $\int_{2^{-l}}^1$  we have bounded

$$\begin{aligned} \exp\left(-\frac{c2^{-l}n\epsilon^2}{\alpha^2 L^2 d \sqrt{\tau} \log n(2^{-2l} + 2^{-l} \epsilon)}\right) &\leq \exp\left(-\frac{c2^{-l}n\epsilon^2}{2\alpha^2 L^2 d \sqrt{\tau} \log n 2^{-l} \epsilon}\right) \\ &= \exp\left(-\frac{cn\epsilon}{2\alpha^2 L^2 d \sqrt{\tau} \log n}\right) \\ &\leq \exp\left(-\frac{cn2^{-l}}{2\alpha^2 L^2 d \sqrt{\tau} \log n}\right), \end{aligned}$$

and the last inequality follows from [36, Lemma 5]. For  $\tau = 2$  and  $n$  large enough so that  $\frac{n}{\sqrt{\log n}} \gtrsim \alpha^2 L^2 d 2^{-l} \log n$ , the second term in (\*\*) is bounded by the first one, and we obtain (b). Analogous computations for the case where  $|\zeta| \leq \sigma$  lead to the final claim.  $\square$

### 3.2. Conditional regression bounds

In this section we study how partitioning polynomial regression of the link function in a single-index model is affected by an estimate  $\hat{v}$  of the index vector, where regression estimators are viewed as conditioned on  $\hat{v}$ . For simplicity, we assume that  $\hat{v}$  was trained on a separate sample of  $(X, Y)$  of cardinality  $n$ . As a consequence,  $\hat{v}$  is independent of the samples  $\{(X_i, Y_i)\}_{i=1}^n$  used to regress the link function. If we are given a single training set, independence can be easily enforced by randomly splitting the data in half. In this case, all rates are understood to scale with  $n/2$ , where the factor  $1/2$  is absorbed by the other absolute constants.

We will focus on one standard class of priors for regression functions, namely the class  $C^s$  of Hölder continuous functions. We recall that a function  $g : \mathbb{R}^d \rightarrow \mathbb{R}$  is  $C^s$  Hölder continuous ( $g \in C^s$ ) if, for  $s = m + \alpha$ ,  $m \geq 0$  an integer and  $\alpha \in (0, 1]$ ,  $g$  has bounded continuous derivatives up to order  $m$  and

$$|g|_{C^s} = \max_{|\lambda|=m} \sup_{x \neq z} \frac{\partial^\lambda g(x) - \partial^\lambda g(z)}{\|x - z\|^\alpha} < \infty.$$

The Hölder norm is defined by

$$\|g\|_{C^s} = \sum_{|\lambda| \leq m} \|\partial^\lambda g\|_\infty + |g|_{C^s}.$$

In expectation, piecewise polynomial estimators of degree  $m$  are min-max optimal on the class  $C^s$  for all  $s = m + \alpha$  with  $m \geq 0$  and  $\alpha \in (0, 1]$  [22, Corollary 11.2]. Our analysis will be limited to polynomials of degree  $m \in \{0, 1\}$ , and thus to functions of smoothness  $s \leq 2$ . On the other hand, piecewise polynomials of order greater than zero are in general not optimal in high probability [4, Section 3]. For this reason, in the case of a  $C^s$  Hölder regression function with  $s \in (1, 2]$  we will assume the following regularity condition:

- (R) Let  $\rho$  be the distribution of  $X$ , and let  $\rho_v$  denote the push-forward measure of  $\rho$  along the map  $x \in \mathbb{R}^d \mapsto \langle x, v \rangle \in \mathbb{R}$ . For every interval  $I \subset \text{supp } \rho_v$ ,  $\text{Var}[\langle v, X \rangle \mid \langle v, X \rangle \in I] \gtrsim |I|^2$ .

To control the distributional mismatch of the projection  $\langle \widehat{\nu}, X \rangle$ , we will also make one of the following two assumptions:

(P1)  $X$  has spherical distribution.

(P2) Let  $\rho$  be the distribution of  $X$ , and let  $\rho_\nu$  denote the push-forward measure of  $\rho$  along the map  $x \in \mathbb{R}^d \mapsto \langle x, \nu \rangle \in \mathbb{R}$ . The distribution  $\rho$  has upper bounded density with bounded support, and for every interval  $I \subset \text{supp } \rho_\nu$ ,  $\rho_\nu(I) \gtrsim |I|$ .

As discussed early in Section 3, spherical distributions (P1) provide a customary model for conditional regression estimators and cover, but are not limited to, the Gaussian distribution. On the other hand, the class (P2) includes a variety of regular densities on compact normal domains, where no special symmetry is required.

We now state our main theorem:

**Theorem 2.** *Assume (X) and (Z). Furthermore, assume either (P1) or (P2). Let  $\widehat{\nu}$  be an estimator for  $\nu$  such that, for every  $\varepsilon > 0$ ,*

$$(\widehat{\mathbf{V}}) \mathbb{P}\{\|P_{\widehat{\nu}} - P_\nu\| > \varepsilon\} \leq A \exp(-n\varepsilon^2/B)$$

for some  $A, B \geq 1$  possibly dependent on  $d$  and specific parameters. Suppose  $f \in C^s$  with  $s \in [1/2, 2]$ , and assume (R) in the case  $s > 1$ . Let  $\widehat{F}_{j|\widehat{\nu}}$  be a piecewise constant ( $s \leq 1$ ) or linear ( $s > 1$ ) estimator of  $F$  at scale  $j$  conditioned on  $\widehat{\nu}$  and truncated at  $M = \|F\|_\infty$ , as defined in Section 2.4 for  $\langle \widehat{\nu}, x \rangle \in I = [-r, r]$ , and 0 outside. Then, setting  $2^{-j} \asymp \sqrt{B}(\log n/n)^{1/(2s+1)}$  and  $r \asymp \sqrt{2d \log n} R$  we have:

(a) For every  $\nu > 0$  there is  $c_\nu(d, B, R, \|f\|_{C^s}, s) \geq 1$  such that

$$\mathbb{P}\left\{(\mathbb{E}_X[|\widehat{F}_{j|\widehat{\nu}}(X) - F(X)|^2])^{\frac{1}{2}} > c_\nu(\log n)^{s/2} \left(\frac{\log n}{n}\right)^{\frac{s}{2s+1}}\right\} \lesssim An^{-\nu}.$$

(b)  $\mathbb{E}[|\widehat{F}_{j|\widehat{\nu}}(X) - F(X)|^2] \leq K(\log n)^s \left(\frac{\log n}{n}\right)^{\frac{2s}{2s+1}}$   
 for some  $K = K(d, A, B, R, \|f\|_{C^s}, s)$ .

The dependence of all constants upon  $d$ ,  $A$  and  $B$  is polynomial.

Theorem 2 shows that partitioning polynomials estimators (of degree 0 and 1) achieve the 1-dimensional min-max convergence rate (up to logarithmic factors) (for  $s$ -Hölder functions with  $s \in [1/2, 2]$ ) when conditioned on any  $\sqrt{n}$ -convergent estimate of  $\nu$ , and thus, in particular, on the estimate  $\widehat{\nu}$  obtained with SVR (under its assumptions). The same conclusion follows for other prominent index estimators, including conditional methods such as SIR, SAVE, SCR and DR. Although formally our theorem requires non-asymptotic  $\sqrt{n}$ -convergence to the index, and only asymptotic  $\sqrt{n}$ -consistency has been established for the aforementioned methods, finite-sample bounds can be derived as well with similar arguments to the ones employed in Section 3.1.

The reader will note that the parameter choices made in Theorem 2 depend on quantities which might be unknown, notably the bound  $\|F\|_\infty$  in the truncation level  $M$  and the smoothness  $s$  in the polynomial degree  $m$  and scale  $2^{-j}$ . While such results are customary in statistical learning theory, as they still provides the existence of an optimal choice for the parameters, and thus of a suitable range for cross-validation, we briefly discuss here this limitation. The choice  $M = \|F\|_\infty$  is made mostly for simplicity. In many applications, at least a proxy  $M \geq \|F\|_\infty$  is known. For tight proxies  $\|F\|_\infty \leq M \lesssim \|F\|_\infty$ , the bounds in Theorem 2 hold true up to an additional constant factor. Otherwise, the constants  $c_\nu$  and  $K$

would also depend on  $M$ . For what concerns the knowledge of  $s$ , one could use again an upper bound. An algorithmic solution would instead be trying polynomials of increasing degree and cross-validate the scale. Smoothness-independent theoretical bounds are possible choosing the partition adaptively by wavelet-thresholding techniques [4,5,42].

To prove Theorem 2 we first show that, with high probability, a conditional polynomial estimator (of degree 0 or 1) differs from an oracle estimator (possessing knowledge of  $v$ ) by the angle between  $\widehat{v}$  and  $v$ :

**Proposition 4.** *Assume (X) and (Z). Furthermore, assume either (P1) or (P2). Suppose  $f \in C^\alpha$  with  $\alpha \in [1/2, 1]$ . Let  $\widehat{v}$  be an estimate of  $v$  with  $\langle \widehat{v}, v \rangle \geq 0$ . For  $u \in \{v, \widehat{v}\}$ , let  $\widehat{F}_{j|u}$  be a piecewise constant ( $\alpha \leq 1$ ) or linear ( $\alpha = 1$ ) estimator of  $F$  at scale  $j$  conditioned on  $u$  and truncated at  $M = \|F\|_\infty$ , as defined in Section 2.4 for  $\langle u, x \rangle \in I = [-r, r]$ , and 0 outside. Assume (R) in case of a linear estimator. Then, for every  $\varepsilon > 0$ ,  $r \geq 1$  and conditioned on  $2^{-j} \geq \|\widehat{v} - v\|/t$  for some  $t \geq 1$ , we have*

$$(\mathbb{E}_X[|\widehat{f}_{j|\widehat{v}}(\langle \widehat{v}, X \rangle) - \widehat{f}_{j|v}(\langle \widehat{v}, X \rangle)|^2 \mathbb{1}\{X \in B(0, r)\}])^{\frac{1}{2}} \lesssim t|f|_{C^\alpha} r^{\frac{1}{2-\alpha}} \|\widehat{v} - v\|^{\frac{1}{2-\alpha}} + \varepsilon$$

with probability higher than  $1 - C\#\mathcal{K}_j \exp\left(-\frac{c n\varepsilon^2}{\#\mathcal{K}_j t^2 \|f\|_{C^\alpha}^2 r^{2\alpha}}\right) - 2n \exp(-r^2/2dR^2)$ .

The proof of Proposition 4 can be found in the supplementary material [36]. The key tool to obtain the dependence on  $\|\widehat{v} - v\|$  in the upper bound is the Wasserstein distance, which enables to control the difference between statistics computed on the conditional distribution given  $\widehat{v}$  rather than  $v$ . We now proceed to prove Theorem 2.

**Proof of Theorem 2.** First, we observe that, since  $F(x) = f(\langle v, x \rangle)$ , defining  $\check{f}(t) = f(-t)$  we have  $F(x) = \check{f}(\langle -v, x \rangle)$  as well. Moreover,  $f \in C^s$  if and only if  $\check{f} \in C^s$ , with  $|f|_{C^t} = |\check{f}|_{C^t}$  for all  $t \leq s$ . Thus, since we can always replace  $v$  with  $-v$  and  $f$  with  $\check{f}$ , in the following we may assume without loss of generality that  $\langle \widehat{v}, v \rangle \geq 0$ , hence  $\|\widehat{v} - v\| \leq \sqrt{2}\|P_{\widehat{v}} - P_v\|$ .

Let  $m = \lceil s \rceil - 1 \in \{0, 1\}$  and  $\alpha = s \wedge 1 \in [\frac{1}{2}, 1]$ . We start by isolating the error outside a ball  $B(0, r)$ :

$$\mathbb{E}_X[|F(X) - \widehat{F}_{j|\widehat{v}}(X)|^2] \lesssim \mathbb{E}_X[|F(X) - \widehat{F}_{j|\widehat{v}}(X)|^2 \mathbb{1}\{\|X\| \leq r\}] + |f|_{C^0}^2 \mathbb{P}\{\|X\| > r\},$$

whence, using [36, Lemma 3], we get the tail bound

$$\mathbb{P}\{\|X\| > r\} \lesssim \exp(-r^2/2dR^2). \tag{T}$$

For  $x \in B(0, r)$  we decompose

$$|F(x) - \widehat{F}_{j|\widehat{v}}(x)| = |f(\langle v, x \rangle) - \widehat{f}_{j|\widehat{v}}(\langle \widehat{v}, x \rangle)| \leq \underbrace{|f(\langle v, x \rangle) - f(\langle \widehat{v}, x \rangle)|}_{(\theta_1)} + |f(\langle \widehat{v}, x \rangle) - \widehat{f}_{j|\widehat{v}}(\langle \widehat{v}, x \rangle)|,$$

and bound (θ1) by the angle  $\|\widehat{v} - v\|$ :

$$\mathbb{E}_X[|f(\langle v, X \rangle) - f(\langle \widehat{v}, X \rangle)|^2 \mathbb{1}\{\|X\| \leq r\}] \leq |f|_{C^\alpha}^2 r^{2\alpha} \|\widehat{v} - v\|^{2\alpha}.$$

Hence, from assumption  $(\widehat{V})$  and [36, Lemma 5] we get

$$\mathbb{P}\{\mathbb{E}_X[|f(\langle v, X \rangle) - f(\langle \widehat{v}, X \rangle)|^2 \mathbb{1}\{\|X\| \leq r\}] > \varepsilon^2\} \leq A \exp\left(-c \frac{n\varepsilon^{2/\alpha}}{Br^2 |f|_{C^\alpha}^{2/\alpha}}\right) \tag{\Theta 1a}$$

$$\mathbb{E}[|f(\langle v, X \rangle) - f(\langle \widehat{v}, X \rangle)|^2 \mathbb{1}\{\|X\| \leq r\}] \leq (\log(A)B)^\alpha |f|_{C^\alpha}^2 r^{2\alpha} n^{-\alpha}. \tag{\Theta 1b}$$

We further decompose

$$\begin{aligned}
 & |f(\langle \widehat{v}, x \rangle) - \widehat{f}_{j|\widehat{v}}(\langle \widehat{v}, x \rangle)| \\
 & \leq \underbrace{|f(\langle \widehat{v}, x \rangle) - f_{j|v}(\langle \widehat{v}, x \rangle)|}_{(b)} + \underbrace{|f_{j|v}(\langle \widehat{v}, x \rangle) - \widehat{f}_{j|v}(\langle \widehat{v}, x \rangle)|}_{(v)} + \underbrace{|\widehat{f}_{j|v}(\langle \widehat{v}, x \rangle) - \widehat{f}_{j|\widehat{v}}(\langle \widehat{v}, x \rangle)|}_{(\theta)},
 \end{aligned}$$

where  $f_{j|v}$  is the  $m$ -order population estimator of  $f$  at scale  $j$  conditioned on  $v$ , namely

$$f_{j,k|v} = \arg \min_{\deg(p) \leq m} \mathbb{E}[|Y - p(\langle v, X \rangle)|^2 \mathbb{1}\{\langle v, X \rangle \in I_{j,k}\}], \quad f_{j|v}(t) = \sum_{k \in \mathcal{K}_j} T_M[f_{j,k|v}(t)] \mathbb{1}\{t \in I_{j,k}\}.$$

Integrating (b) we get a bias term, that we control exploiting the Hölder continuity of  $f$  (see [42, Appendix A, Example 1]<sup>1</sup>):

$$\mathbb{E}_X[|f(\langle \widehat{v}, X \rangle) - f_{j|v}(\langle \widehat{v}, X \rangle)|^2 \mathbb{1}\{\|X\| \leq r\}] \lesssim |f|_{C^s}^2 r^{2s} 2^{-2js}. \tag{B}$$

The variable (v) leads to a variance term, which can be concentrated with known calculations (see [42, Proposition 2 and Lemma 5]<sup>1</sup>):

$$\mathbb{P}\{\mathbb{E}_X[|f_{j|v}(\langle \widehat{v}, X \rangle) - \widehat{f}_{j|v}(\langle \widehat{v}, X \rangle)|^2 \mathbb{1}\{\|X\| \leq r\}] > \varepsilon^2\} \lesssim \#\mathcal{K}_j \exp\left(-c \frac{n\varepsilon^2}{\#\mathcal{K}_j |f|_{C^0}^2}\right) \tag{Va}$$

$$\mathbb{E}[|f_{j|v}(\langle \widehat{v}, X \rangle) - \widehat{f}_{j|v}(\langle \widehat{v}, X \rangle)|^2 \mathbb{1}\{\|X\| \leq r\}] \lesssim |f|_{C^0}^2 \frac{\log(\#\mathcal{K}_j) \#\mathcal{K}_j}{n}. \tag{Vb}$$

For (θ) we condition on the event  $\|\widehat{v} - v\| \leq t2^{-j}$ , taking into account that, thanks to assumption  $(\widehat{V})$ , the complement has probability

$$\mathbb{P}\{\|\widehat{v} - v\| > t2^{-j}\} \leq A \exp(-nt^2 2^{-2j} / B). \tag{Θ2}$$

Thus, Proposition 4 along with assumption  $(\widehat{V})$  and [36, Lemma 5] gives

$$\begin{aligned}
 & \mathbb{P}\{\mathbb{E}_X[|\widehat{f}_{j|v}(\langle \widehat{v}, X \rangle) - \widehat{f}_{j|\widehat{v}}(\langle \widehat{v}, X \rangle)|^2 \mathbb{1}\{\|X\| \leq r\}] > \varepsilon^2 \mid \|\widehat{v} - v\| \leq t2^{-j}\} \\
 & \lesssim \#\mathcal{K}_j \exp\left(-c \frac{n\varepsilon^2}{\#\mathcal{K}_j t^2 |f|_{C^\alpha}^2 r^{2\alpha}}\right) + A \exp\left(-\frac{n\varepsilon^{2(2-\alpha)}}{B t^{2(2-\alpha)} |f|_{C^\alpha}^{2(2-\alpha)} r^2}\right), \tag{Θ3a}
 \end{aligned}$$

$$\begin{aligned}
 & \mathbb{E}[|\widehat{f}_{j|v}(\langle \widehat{v}, X \rangle) - \widehat{f}_{j|\widehat{v}}(\langle \widehat{v}, X \rangle)|^2 \mathbb{1}\{\|X\| \leq r\} \mid \|\widehat{v} - v\| \leq 2^{-j}] \\
 & \lesssim |f|_{C^\alpha}^2 r^{2\alpha} \frac{\log(\#\mathcal{K}_j) \#\mathcal{K}_j}{n} + (\log(A)B)^{\frac{1}{2-\alpha}} |f|_{C^\alpha}^2 r^{\frac{2}{2-\alpha}} n^{-\frac{1}{2-\alpha}}. \tag{Θ3b}
 \end{aligned}$$

In order to balance the tail (T), the bias (B), the variance (Vb) and the angle terms (Θ1b), (Θ2) and (Θ3b), we choose

$$r = \sqrt{2d \log n} R, \quad 2^{-j} \asymp \sqrt{B} (\log n / n)^{1/(2s+1)}, \quad t = 1,$$

<sup>1</sup> [42] deals more generally with piecewise polynomial regression on an unknown manifold  $\mathcal{M}$  of dimension  $d$ . In our setting,  $\mathcal{M}$  is a known Euclidean domain of dimension  $d = 1$ , in fact an interval. In particular, the dyadic decomposition 2.a) satisfies assumptions (A1)÷(A4) in [42] with  $\theta_1 = 1$  and  $\theta_2 = 2r$ . Assumption (A5)(i) is satisfied thanks to (R), while (A5)(ii) is trivial.

and plug in  $\#\mathcal{K}_j = 2^j$ , which leads to

$$\mathbb{E}[|\widehat{F}_j|_{\widehat{\mathcal{F}}}(X) - F(X)|^2] \lesssim |f|_{C^0}^2 \left(\frac{1}{n}\right) \tag{T}$$

$$+ |f|_{C^\alpha}^2 R^{2\alpha} (\log(A)Bd)^\alpha \left(\frac{\log n}{n}\right)^\alpha \tag{\Theta1b}$$

$$+ |f|_{C^s}^2 R^{2s} (Bd)^s (\log n)^s \left(\frac{\log n}{n}\right)^{\frac{2s}{2s+1}} \tag{B}$$

$$+ |f|_{C^0}^2 \frac{1}{2s+1} \left(\frac{\log n}{n}\right)^{\frac{2s}{2s+1}} \tag{Vb}$$

$$+ A \exp\left(-n \left(\frac{\log n}{n}\right)^{\frac{2}{2s+1}}\right) \tag{\Theta2}$$

$$+ |f|_{C^\alpha}^2 \frac{1}{2s+1} R^{2\alpha} d^\alpha (\log n)^\alpha \left(\frac{\log n}{n}\right)^{\frac{2s}{2s+1}} \tag{\Theta3b}$$

$$+ |f|_{C^\alpha}^2 R^{\frac{2}{2-\alpha}} (\log(A)B)^{\frac{1}{2-\alpha}} d^{\frac{1}{2-\alpha}} \left(\frac{\log n}{n}\right)^{\frac{1}{2-\alpha}} .$$

In **(\Theta1b)**,  $(\log n/n)^\alpha \leq (\log n/n)^{2s/2s+1}$  for  $\alpha = s \in [1/2, 1]$ , and for  $\alpha = 1$  and  $s > 1$ . In **(\Theta2)**,  $\exp(-n(\log n/n)^{2/(2s+1)}) \leq 1/n$  for  $s \geq 1/2$ . In **(\Theta3b)**,  $(\log n/n)^{1/(2-\alpha)} \leq (\log n/n)^{2s/(2s+1)}$  for  $\alpha = s \in (0, 1]$ , and for  $\alpha = 1$  and  $s > 1$ . Collecting the constants we obtain **(b)**. The bound **(a)** follows similarly from **(T)**, **(B)**, **(Va)**, **(\Theta1a)**, **(\Theta2)** and **(\Theta3a)** setting

$$r = \sqrt{2d \log n} R, \quad 2^{-j} \asymp \sqrt{B} (\log n/n)^{1/(2s+1)}, \quad \varepsilon = c_\nu \log^{\frac{s}{2}} n (\log n/n)^{s/(2s+1)}, \quad t = \sqrt{c_\nu},$$

and taking  $c_\nu$  large enough. □


The additional logarithmic factors in **(a)** and **(b)** are exclusively due to the unboundedness of the distribution and can be avoided in the bounded case. While here we are restricting to constant and linear estimators, hence to the smoothness range  $s \leq 2$ , the results may be extended to higher order polynomials, and thus to smoother regression functions, with similar proofs. However, from our decomposition, and specifically from **(\Theta1b)**, it seems not possible to maintain min-max optimality in the range  $s < 1/2$ . We remark that this smoothness constraint does not depend on the regression technique, that is, the term **(\Theta1b)** would still arise when using different regression methods than partitioning polynomials.

### 4. Numerical experiments

In this section we conduct numerical experiments to demonstrate that the theoretical results above have practical relevance and to investigate how relaxations of the assumptions affect the estimators. In order to highlight specific aspects of different algorithms we use three different functions to conduct our experiments. The first two are

$$F_1(x) = \exp(\langle v, x \rangle / 3), \quad F_2(x) = F_1(x) + \sin(20 \langle v, x \rangle) / 15 .$$

Both functions are smooth.  $F_1$  is monotone and thus we may choose  $\omega = 0$ , while  $F_2$  is non-monotone, thus condition  $(\Omega)$  is satisfied only for  $\omega > 0$ . This allows us to explore the behavior of  $\widehat{v}$  under monotonicity or lack thereof, and how the estimators are effected by the choice of the scales  $l$  and  $j$ . To investigate the convergence rate of the regression estimator  $\widehat{F}$ , we use a monotone function  $F_3$  which is piecewise quadratic on a random partition and continuous. The domain of  $x$ , and its dimension  $d$ , will be specified in each experiment.  $F_1, F_2, F_3$  are shown in Figure 2.




**Figure 2.** Different functions used in the experiments, with horizontal axis representing  $\langle v, x \rangle$ .

**Table 3.** Performance of the different algorithms in different settings, with  $\text{err} = \log_{10}(\|\widehat{v} - v\|^2)$ , corresponding standard error, and average computational time in seconds/100.

		S1			S2			
	$\sigma$	err	se	time	err	se	time	
$F_1$	0%	SIR	-3.04	-2.83	0.60	-0.68	-1.83	0.60
		SAVE	-5.97	-1.06	1.30	-7.41	-7.14	1.40
		SVR	-6.42	-6.06	1.00	-7.60	-6.65	0.70
$F_2$		SIR	-3.11	-2.99	0.50	-0.67	-1.83	0.50
		SAVE	-4.39	-4.10	1.30	-4.13	-4.01	1.30
		SVR	-4.41	-4.08	0.90	-4.16	-3.92	0.70
$F_1$	1%	SIR	-2.97	-2.69	0.50	-0.68	-1.81	0.50
		SAVE	-4.57	-4.24	1.40	-4.16	-4.03	1.30
		SVR	-4.58	-4.28	1.00	-4.12	-3.75	0.80
$F_2$		SIR	-2.94	-2.72	0.50	-0.68	-1.77	0.50
		SAVE	-4.06	-3.57	1.40	-3.42	-3.45	1.40
		SVR	-4.08	-3.87	0.90	-3.45	-3.46	0.80
$F_1$	2%	SIR	-2.92	-2.67	0.50	-0.68	-1.38	0.60
		SAVE	-3.92	-3.58	1.30	-3.08	-3.01	1.40
		SVR	-3.91	-3.61	0.90	-3.21	-3.06	0.80
$F_2$		SIR	-2.87	-2.64	0.60	-0.68	-1.55	0.60
		SAVE	-3.64	-1.98	1.40	-2.90	-2.85	1.40
		SVR	-3.66	-3.45	0.90	-3.02	-2.80	0.80






**Figure 3.** Left column displays the ingredients for the estimates: the empirical inverse regression curve (green) used by SIR and the local gradients (blue), with length proportional to the number of samples in the corresponding level set, used by SVR. Estimates of  $v$  using SVR (blue) and SIR (green) are displayed on the right column. The methods are applied on setting S1 (top row) and S2 (bottom row). The black line indicates  $v$ , while data points are colored according to the value of the corresponding response variable, generated with  $F_2$  and  $\sigma = 0$ , using a red-to-yellow color scale.

### 4.1. Estimating the index vector

Here we compare the performances of SIR, SAVE and SVR in estimating the index vector  $v$ . We consider two settings S1, S2, corresponding to two different non-elliptical, and thus non-spherical, distributions for  $X$ :  $\rho_{X,1}, \rho_{X,2}$ ;  $\rho_{X,1}$  is a standard normal  $\mathcal{N}(0,1)$  in one coordinate, and a skewed normal with shape parameter  $\alpha = 5$  in the other coordinate;  $\rho_{X,2}$  is uniform on the triangle with vertices  $(0,0),(1,1),(0,1)$ ; all distributions are normalized to have zero mean and standard deviation equal to one. Note that in both settings conditions (LCM) and (SMD) are not satisfied. For each setting we draw  $n = 1000$  i.i.d. samples and generate the response variable  $Y_i = F(X_i) + \zeta_i$  using functions  $F_1$  and  $F_2$ , where  $\zeta_i \sim \mathcal{N}(0, \sigma^2)$ . We use different levels of noise setting  $\sigma$  equal to the 0%, 1% and 2% of  $|f(-4) - f(4)|$ . We chose  $v = (1/\sqrt{5}, 2/\sqrt{5})$  for setting S1, and  $v = (1, 0)$  for setting S2. The results in Table 3 show the detailed performance of SIR, SAVE and SVR for all settings, functions, and noise levels. First, we note that the cases of  $F_1$  with 1% noise and  $F_2$  with zero noise produce similar results. This is consistent with the intuition manifested from Theorem 1 that noise and non-monotonicity levels play a similar role in the accuracy of the estimators. In all settings SIR performs worse than the two other methods, SVR and SAVE, which on the other hand have similar performance, although SVR produces most of the times slightly better estimates. The poor performance of SIR in these settings requires a better explanation. In Figure 3 we show graphically how the empirical inverse regression curve may drift away from  $v$ , resulting in a poor SIR estimate. On the other hand, the local gradients used by SVR provide good local estimates. This example shows how methods with higher order statistics are in general more robust to assumptions relaxations.



**Figure 4.** Behavior of the SVR estimate  $\widehat{v}_l$  with respect to scale and sample size, for regression of  $F_2$  (see text). Left: error versus scale  $l$ . Right: error versus sample size  $n$ .


To investigate more extensively the performance of SVR in estimating  $v$ , we perform another experiment: we draw  $X$  from a 10-dimensional standard normal distribution, and to generate the response variable we use function  $F_2$  plus an additive Gaussian noise with standard deviation  $\sigma = 0.01|f_2(-4) - f_2(4)|$ . We repeat the experiment for different values of the sample size  $n$ . Results are shown in Figure 4. The left inset shows that the error in  $\widehat{v}$  stabilizes at scales comparable to the noise level  $\sigma$ , which suggests that the assumption  $|S_{l,h}| \gtrsim \sigma$  is needed. The right plot shows that the rate of the error of  $\widehat{v}$ , for scales  $l$  coarser than the noise level, is approximately  $-\frac{1}{2}$ , which is again consistent with Theorem 1.

### 4.2. Estimating the regression function

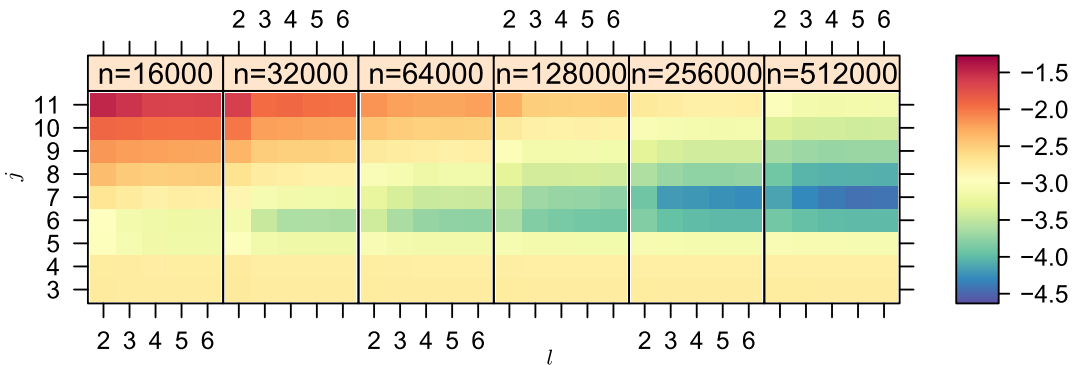
In this section we perform some experiments to support our theoretical results regarding the regression estimator obtained with SVR. The first experiment we perform consists on drawing  $X_i, i = 1, \dots, n$ , from a  $d$ -dimensional standard normal distribution and obtain  $Y_i = F_3(X_i) + \zeta_i$  where  $\zeta_i \sim \mathcal{N}(0, \sigma^2)$ . Here we use function  $F_3$  because we want to limit the function smoothness in order to obtain concentration rates comparable with the min-max rate with  $s = 1$ . We vary the dimension  $d = 5, 10, 50, 100$ , the size of the noise  $\sigma$ , equal to the 5% and 10% of  $|f_3(-4) - f_3(4)|$ . To investigate the convergence rates of the estimator we repeat each experiment for different sample sizes  $n$ . In Figure 5 we show the empirical MSE, averaged over 10 repetitions, as a function of the sample size, in logarithmic scale, for both our estimator and the k-Nearest-Neighbor (kNN) regression. We see that the MSE of the SVR estimator decays with a rate slightly better than the optimal value  $-2/3$ , independently from the dimension  $d$  and the noise level  $\sigma$ : this is all consistent with Theorem 2. As expected, kNN-regression has a convergence rate which severely deteriorates with the dimension (curse of dimensionality). We can also notice that the MSE drops far below the noise level, which confirms the de-noising feature of the SVR estimator.

To explore the behavior of the empirical MSE as a function of the scales  $l$  and  $j$  we conduct another experiment: we draw  $X$  from a 10-dimensional standard normal distribution, and obtain the response variable  $Y = F_2(X) + \zeta$ , with  $\zeta$  Gaussian noise with standard deviation  $\sigma = 0.01|f_2(-4) - f_2(4)|$ . Figure

6 shows the behavior of the  $\log_{10}(\text{MSE})$ , obtained with SVR, for different values of  $l$ ,  $j$  and  $n$ . To obtain robust estimates in regions with high Monte Carlo variability, in regimes where our results do not hold, the errors are averaged over 50 repetition of each setting with a 10% trimming. By observing each row, we notice that the MSE reaches its minimum for low values of  $l$  and stays constant for larger  $l$ . By looking at the plot column-wise, we observe the bias variance trade-off, with coarse scales giving rough estimates, and fine scales resulting in overfitting. As expected, as the sample size grows, the optimal scale  $j$  increases.



**Figure 5.** Comparison of convergence rates for the regression estimator with SVR and KNN-regression in different settings.



**Figure 6.** Empirical MSE versus sample size  $n$  and scales  $l$  and  $j$ .

## Acknowledgements

S.V. thanks Timo Klock for the discussion and the useful exchange of views about this and related problems.

## Funding

This research was partially supported by AFOSR FA9550-17-1-0280, NSF-DMS-1821211, NSF-ATD-1737984.

## Supplementary Material

**Supplement to “Conditional regression for single-index models”** (DOI: [10.3150/22-BEJ1482SUPP](https://doi.org/10.3150/22-BEJ1482SUPP); .pdf). This supplement contains the proof of Proposition 4 and some auxiliary results.

## References

- [1] Bach, F. (2017). Breaking the curse of dimensionality with convex neural networks. *J. Mach. Learn. Res.* **18** Paper No. 19, 53. [MR3634886](#)
- [2] Bhatia, R. (1997). *Matrix Analysis. Graduate Texts in Mathematics* **169**. New York: Springer. [MR1477662](#) <https://doi.org/10.1007/978-1-4612-0653-8>
- [3] Bickel, P.J. and Li, B. (2007). Local polynomial regression on unknown manifolds. In *Complex Datasets and Inverse Problems. Institute of Mathematical Statistics Lecture Notes—Monograph Series* **54** 177–186. Beachwood, OH: IMS. [MR2459188](#) <https://doi.org/10.1214/074921707000000148>
- [4] Binev, P., Cohen, A., Dahmen, W. and DeVore, R. (2007). Universal algorithms for learning theory. II. Piecewise polynomial functions. *Constr. Approx.* **26** 127–152. [MR2327596](#) <https://doi.org/10.1007/s00365-006-0658-z>
- [5] Binev, P., Cohen, A., Dahmen, W., DeVore, R. and Temlyakov, V. (2005). Universal algorithms for learning theory. I. Piecewise constant functions. *J. Mach. Learn. Res.* **6** 1297–1321. [MR2249856](#)
- [6] Cambanis, S., Huang, S. and Simons, G. (1981). On the theory of elliptically contoured distributions. *J. Multivariate Anal.* **11** 368–385. [MR0629795](#) [https://doi.org/10.1016/0047-259X\(81\)90082-8](https://doi.org/10.1016/0047-259X(81)90082-8)
- [7] Carroll, R.J., Fan, J., Gijbels, I. and Wand, M.P. (1997). Generalized partially linear single-index models. *J. Amer. Statist. Assoc.* **92** 477–489. [MR1467842](#) <https://doi.org/10.2307/2965697>
- [8] Carroll, R.J., Ruppert, D. and Welsh, A.H. (1998). Local estimating equations. *J. Amer. Statist. Assoc.* **93** 214–227. [MR1614624](#) <https://doi.org/10.2307/2669618>
- [9] Cook, R.D. (1998). Save: a method for dimension reduction and graphics in regression. *Communications in Statistics - Theory and Methods* **29** 2109–2121.
- [10] Cook, R.D. and Lee, H. (1999). Dimension reduction in binary response regression. *J. Amer. Statist. Assoc.* **94** 1187–1200. [MR1731482](#) <https://doi.org/10.2307/2669934>
- [11] Coudret, R., Liquet, B. and Saracco, J. (2014). Comparison of sliced inverse regression approaches for underdetermined cases. *J. SFdS* **155** 72–96. [MR3211755](#)
- [12] Cui, X., Härdle, W.K. and Zhu, L. (2011). The EFM approach for single-index models. *Ann. Statist.* **39** 1658–1688. [MR2850216](#) <https://doi.org/10.1214/10-AOS871>
- [13] Dalalyan, A.S., Juditsky, A. and Spokoiny, V. (2008). A new algorithm for estimating the effective dimension-reduction subspace. *J. Mach. Learn. Res.* **9** 1648–1678. [MR2438819](#)
- [14] Delecroix, M., Härdle, W. and Hristache, M. (1997). Efficient estimation in single-index regression. Interdisciplinary Research Project 373: Quantification and Simulation of Economic Processes. Humboldt University of Berlin. SFB 373 Discussion Paper 37.

- [15] Delecroix, M. and Hristache, M. (1999).  $M$ -estimateurs semi-paramétriques dans les modèles à direction révélatrice unique. *Bull. Belg. Math. Soc. Simon Stevin* **6** 161–185. [MR1705152](#)
- [16] Delecroix, M., Hristache, M. and Patilea, V. (2006). On semiparametric  $M$ -estimation in single-index regression. *J. Statist. Plann. Inference* **136** 730–769. [MR2181975](#) <https://doi.org/10.1016/j.jspi.2004.09.006>
- [17] Diaconis, P. and Freedman, D. (1984). Asymptotics of graphical projection pursuit. *Ann. Statist.* **12** 793–815. [MR0751274](#) <https://doi.org/10.1214/aos/1176346703>
- [18] Duan, N. and Li, K.-C. (1991). Slicing regression: A link-free regression method. *Ann. Statist.* **19** 505–530. [MR1105834](#) <https://doi.org/10.1214/aos/1176348109>
- [19] Eaton, M.L. (1986). A characterization of spherical distributions. *J. Multivariate Anal.* **20** 272–276. [MR0866075](#) [https://doi.org/10.1016/0047-259X\(86\)90083-7](https://doi.org/10.1016/0047-259X(86)90083-7)
- [20] Gaïffas, S. and Lecué, G. (2007). Optimal rates and adaptation in the single-index model using aggregation. *Electron. J. Stat.* **1** 538–573. [MR2369025](#) <https://doi.org/10.1214/07-EJS077>
- [21] Ganti, R., Rao, N., Willett, R.M. and Nowak, R. (2015). Learning single index models in high dimensions. [arXiv:1506.08910](#).
- [22] Györfi, L., Kohler, M., Krzyżak, A. and Walk, H. (2002). *A Distribution-Free Theory of Nonparametric Regression. Springer Series in Statistics*. New York: Springer. [MR1920390](#) <https://doi.org/10.1007/b97848>
- [23] Hall, P. and Li, K.-C. (1993). On almost linearity of low-dimensional projections from high-dimensional data. *Ann. Statist.* **21** 867–889. [MR1232523](#) <https://doi.org/10.1214/aos/1176349155>
- [24] Härdle, W., Hall, P. and Ichimura, H. (1993). Optimal smoothing in single-index models. *Ann. Statist.* **21** 157–178. [MR1212171](#) <https://doi.org/10.1214/aos/1176349020>
- [25] Härdle, W. and Stoker, T.M. (1989). Investigating smooth multiple regression by the method of average derivatives. *J. Amer. Statist. Assoc.* **84** 986–995. [MR1134488](#)
- [26] Horowitz, J.L. (1998). *Semiparametric Methods in Econometrics. Lecture Notes in Statistics* **131**. New York: Springer. [MR1624936](#) <https://doi.org/10.1007/978-1-4612-0621-7>
- [27] Hristache, M., Juditsky, A., Polzehl, J. and Spokoiny, V. (2001). Structure adaptive approach for dimension reduction. *Ann. Statist.* **29** 1537–1566. [MR1891738](#) <https://doi.org/10.1214/aos/1015345954>
- [28] Hristache, M., Juditsky, A. and Spokoiny, V. (2001). Direct estimation of the index coefficient in a single-index model. *Ann. Statist.* **29** 595–623. [MR1865333](#) <https://doi.org/10.1214/aos/1009210681>
- [29] Ichimura, H. (1993). Semiparametric least squares (SLS) and weighted SLS estimation of single-index models. *J. Econometrics* **58** 71–120. [MR1230981](#) [https://doi.org/10.1016/0304-4076\(93\)90114-K](https://doi.org/10.1016/0304-4076(93)90114-K)
- [30] Kakade, S.M., Kanade, V., Shamir, O. and Kalai, A.T. (2011). Efficient learning of generalized linear and single index models with isotonic regression. *Adv. Neural Inf. Process. Syst.* **24** 927–935.
- [31] Kalai, A.T. and Sastry, R. (2009). The isotron algorithm: High-dimensional isotonic regression. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT)*.
- [32] Kelker, D. (1970). Distribution theory of spherical distributions and a location-scale parameter generalization. *Sankhyā Ser. A* **32** 419–438. [MR0287628](#)
- [33] Klock, T., Lanteri, A. and Vigogna, S. (2021). Estimating multi-index models with response-conditional least squares. *Electronic Journal of Statistics* **15** Paper No. 1, 589–629.
- [34] Kpotufe, S. (2011).  $k$ -nn regression adapts to local intrinsic dimension. *Adv. Neural Inf. Process. Syst.* **24** 729–737.
- [35] Kpotufe, S. and Garg, V. (2013). Adaptivity to local smoothness and dimension in kernel regression. *Adv. Neural Inf. Process. Syst.* **26** 3075–3083.
- [36] Lanteri, A., Maggioni, M. Vigogna, S. (2022). Supplement to “Conditional regression for single-index models.” <https://doi.org/10.3150/22-BEJ1482SUPP>
- [37] Li, B. (2018). *Sufficient Dimension Reduction: Methods and Applications with R. Chapman & Hall/CRC Monographs on Statistics and Applied Probability*. Boca Raton: CRC Press.
- [38] Li, B. and Wang, S. (2007). On directional regression for dimension reduction. *J. Amer. Statist. Assoc.* **102** 997–1008. [MR2354409](#) <https://doi.org/10.1198/016214507000000536>
- [39] Li, B., Zha, H. and Chiaromonte, F. (2005). Contour regression: A general approach to dimension reduction. *Ann. Statist.* **33** 1580–1616. [MR2166556](#) <https://doi.org/10.1214/009053605000000192>
- [40] Li, K.-C. (1991). Sliced inverse regression for dimension reduction. *J. Amer. Statist. Assoc.* **86** 316–342. [MR1137117](#)

- [41] Liao, W., Maggioni, M. and Vigogna, S. (2016). Learning adaptive multiscale approximations to data and functions near low-dimensional sets. In *2016 IEEE Information Theory Workshop (ITW)* 226–230. IEEE.
- [42] Liao, W., Maggioni, M. and Vigogna, S. (2022). Multiscale regression on unknown manifolds. *Math. Eng.* **4** Paper No. 4, 1–25. [MR4311326](#) <https://doi.org/10.3934/mine.2022028>
- [43] Stewart, G.W. and Sun, J.G. (1990). *Matrix Perturbation Theory*. *Computer Science and Scientific Computing*. Boston, MA: Academic Press. [MR1061154](#)
- [44] Stoker, T.M. (1986). Consistent estimation of scaled coefficients. *Econometrica* **54** 1461–1481. [MR0868152](#) <https://doi.org/10.2307/1914309>
- [45] Stone, C.J. (1982). Optimal global rates of convergence for nonparametric regression. *Ann. Statist.* **10** 1040–1053. [MR0673642](#)
- [46] van der Vaart, A.W. and Wellner, J.A. (1996). *Weak Convergence and Empirical Processes: With applications to statistics*. *Springer Series in Statistics*. New York: Springer. [MR1385671](#) <https://doi.org/10.1007/978-1-4757-2545-2>
- [47] Vershynin, R. (2018). *High-Dimensional Probability: An introduction with applications in data science*. *Cambridge Series in Statistical and Probabilistic Mathematics* **47**. Cambridge: Cambridge Univ. Press. [MR3837109](#) <https://doi.org/10.1017/9781108231596>
- [48] Xia, Y. (2006). Asymptotic distributions for two estimators of the single-index model. *Econometric Theory* **22** 1112–1137. [MR2328530](#) <https://doi.org/10.1017/S0266466606060531>
- [49] Xia, Y., Tong, H., Li, W.K. and Zhu, L.-X. (2002). An adaptive estimation of dimension reduction space. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **64** 363–410. [MR1924297](#) <https://doi.org/10.1111/1467-9868.03411>
- [50] Yu, Y., Wang, T. and Samworth, R.J. (2015). A useful variant of the Davis-Kahan theorem for statisticians. *Biometrika* **102** 315–323. [MR3371006](#) <https://doi.org/10.1093/biomet/asv008>

*Received December 2020 and revised February 2022*