

Learning sparse gradients for variable selection and dimension reduction

Gui-Bo Ye · Xiaohui Xie

Received: 1 July 2010 / Accepted: 7 March 2012 / Published online: 22 March 2012
© The Author(s) 2012

Abstract Variable selection and dimension reduction are two commonly adopted approaches for high-dimensional data analysis, but have traditionally been treated separately. Here we propose an integrated approach, called sparse gradient learning (SGL), for variable selection and dimension reduction via learning the gradients of the prediction function directly from samples. By imposing a sparsity constraint on the gradients, variable selection is achieved by selecting variables corresponding to non-zero partial derivatives, and effective dimensions are extracted based on the eigenvectors of the derived sparse empirical gradient covariance matrix. An error analysis is given for the convergence of the estimated gradients to the true ones in both the Euclidean and the manifold setting. We also develop an efficient forward-backward splitting algorithm to solve the SGL problem, making the framework practically scalable for medium or large datasets. The utility of SGL for variable selection and feature extraction is explicitly given and illustrated on artificial data as well as real-world examples. The main advantages of our method include variable selection for both linear and nonlinear predictions, effective dimension reduction with sparse loadings, and an efficient algorithm for large p , small n problems.

Keywords Gradient learning · Variable selection · Effective dimension reduction · Forward-backward splitting

1 Introduction

Datasets with many variables have become increasingly common in biological and physical sciences. In biology, it is nowadays a common practice to measure the expression values of

Editor: P. Long.

G. Ye · X. Xie (✉)

School of Information and Computer Science, University of California, Irvine, CA 92697, USA
e-mail: xhx@ics.uci.edu

G. Ye

e-mail: yeg@uci.edu

tens of thousands of genes, genotypes of millions of SNPs, or epigenetic modifications at tens of millions of DNA sites in one single experiment. Variable selection and dimension reduction are increasingly viewed as a necessary step in dealing with these high-dimensional data.

Variable selection aims at selecting a subset of variables most relevant for predicting responses. Many algorithms have been proposed for variable selection (Guyon and Elisseeff 2003). They typically fall into two categories: Feature Ranking and Subset Selection. Feature Ranking scores each variable according to a metric, derived from various correlation or information theoretic criteria (Guyon and Elisseeff 2003; Weston et al. 2003; Dhillon et al. 2003), and eliminates variables below a threshold score. Because Feature Ranking methods select variables based on individual prediction power, they are ineffective in selecting a subset of variables that are marginally weak but in combination strong in prediction. Subset Selection aims to overcome this drawback by considering and evaluating the prediction power of a subset of variables as a group. One popular approach to subset selection is based on direct object optimization, which formalizes an objective function of variable selection and selects variables by solving an optimization problem. The objective function often consists of two terms: a *data fitting term* accounting for prediction accuracy, and a *regularization term* controlling the number of selected variables. LASSO proposed by Tibshirani (1996) and elastic net by Zou and Hastie (2005) are two examples of this type of approach. The two methods are widely used because of their implementation efficiency (Efron et al. 2004; Zou and Hastie 2005) and the ability of performing simultaneous variable selection and prediction; however, a linear prediction model is assumed by both methods. The component smoothing and selection operator method (COSSO) proposed in Lin and Zhang (2006) tries to overcome this shortcoming by using a functional LASSO penalty. However, COSSO is based on the framework of smoothing spline ANOVA which is difficult in handling high dimensional data.

Dimension reduction is another commonly adopted approach in dealing with high-dimensional data. Rooting in dimension reduction is the common belief that many real-world high-dimensional data are concentrated on a low-dimensional manifold embedded in the underlying Euclidean space. Therefore mapping the high-dimensional data into the underlying low-dimensional manifold should be able to improve prediction accuracy, to help visualize the data, and to construct better statistical models. A number of dimension reduction methods have been proposed, ranging from principle component analysis to manifold learning for non-linear settings (Belkin and Niyogi 2003; Zou et al. 2006; Mackey 2009; Roweis and Saul 2000; Tenenbaum et al. 2000; Donoho and Grimes 2003). However, most of these dimension reduction methods are unsupervised, and therefore are likely suboptimal with respect to predicting responses. In supervised settings, most recent work focuses on finding a subspace \mathcal{S} such that the projection of the high dimensional data \mathbf{x} onto \mathcal{S} captures the statistical dependency of the response y on \mathbf{x} . The space \mathcal{S} is called *effective dimension reduction* (EDR) space (Xia et al. 2002).

Several methods have been proposed to identify EDR space. The research goes back to sliced inverse regression (SIR) proposed by Li (1991), where the covariance matrix of the inverse regression is explored for dimension reduction. The main idea is that if the conditional distribution $\rho(y|\mathbf{x})$ concentrates on a subspace \mathcal{S} , then the inverse regression $E(\mathbf{x}|y)$ should lie in that same subspace. However, SIR imposes specific modeling assumptions on the conditional distribution $\rho(y|\mathbf{x})$ or the regression $E(y|\mathbf{x})$. These assumptions hold in particular if the distribution of \mathbf{x} is elliptic. In practice, however, we do not necessarily expect that \mathbf{x} will follow an elliptic distribution, nor is it easy to assess departures from ellipticity in a high-dimensional setting. A further limitation of SIR is that it

yields only a one-dimensional subspace for binary classifications. Other reverse regression based methods, including principal Hessian directions (pHd, Li 1992), sliced average variance estimation (SAVE, Cook and Yin 2001) and contour regression (Li et al. 2005), have been proposed, but they have similar limitations. To address these limitations, Xia et al. (2002) proposed a method called the (conditional) minimum average variance estimation (MAVE) to estimate the EDR directions. The assumption underlying MAVE is quite weak and only a semiparametric model is used. Under the semiparametric model, conditional covariance is estimated by linear smoothing and EDR directions are then estimated by minimizing the derived conditional covariance estimation. In addition, a simple outer product gradient (OPG) estimator is proposed as an initial estimator. Other related approaches include methods that estimate the derivative of the regression function (Hristache et al. 2001; Samarov 1993). Recently, Fukumizu et al. (2009) proposed a new methodology which derives EDR directly from a formulation of EDR in terms of the conditional independence of \mathbf{x} from the response y , given the projection of \mathbf{x} on the EDR space. The resulting estimator is shown to be consistent under weak conditions. However, all these EDR methods cannot be directly applied to the large p , small n case, where p is the dimension of the underlying Euclidean space in which the data lie, and n is the number of samples. To deal with the large p , small n case, Mukherjee and co-workers (2006) introduced a gradient learning method (which will be referred to as GL) for estimating EDR by introducing a Tikhonov regularization term on the gradient functions. The EDR directions were estimated using the eigenvectors of the empirical gradient covariance matrix.

Although both variable selection and dimension reduction offer valuable tools for statistical inference in high-dimensional space and have been prominently researched, few methods are available for combining them into a single framework where both variable selection and dimensional reduction can be done. One notable exception is the sparse principle component analysis (SPCA), which produces modified principle components with sparse loadings (Zou et al. 2006). However, SPCA is mainly used for unsupervised linear dimension reduction, our focus here is the variable selection and dimension reduction in supervised and potentially nonlinear settings. To motivate the reason why a combined approach might be interesting in a supervised setting, consider a microarray gene expression data measured in both normal and tumor samples. Out of 20,000 genes measured in microarray, only a small number of genes (e.g. oncogenes) are likely responsible for gene expression changes in tumor cells. Variable selection chooses more relevant genes and dimension reduction further extracts features based on the subset of selected genes. Taking a combined approach could potentially improve prediction accuracy by removing irrelevant noisy variables. Additionally, by focusing on a small number of most relevant genes and extracting features among them, it could also provide a more interpretable and manageable model regarding genes and biological pathways involved in the carcinogenesis.

In this article, we extend the gradient learning framework introduced by Mukherjee and co-workers (2006), and propose a sparse gradient learning approach (SGL) for integrated variable selection and dimension reduction in a supervised setting. The method adopts a direct object optimization approach to learn the gradient of the underlying prediction function with respect to variables, and imposes a regularization term to control the sparsity of the gradient. The gradient of the prediction function provides a natural interpretation of the geometric structure of the data (Guyon et al. 2002; Mukherjee and Zhou 2006; Mukherjee and Wu 2006; Mukherjee et al. 2010). If a variable is irrelevant to the prediction function, the partial derivative with respect to that variable is zero. Moreover, for non-zeros partial derivatives, the larger the norm of the partial derivative with respect to a variable is, the more important the corresponding variable is likely to be for prediction. Thus the norms

of partial derivatives give us a criterion for the importance of each variable and can be used for variable selection. Motivated by LASSO, we encourage the sparsity of the gradient by adding an ℓ^1 norm based regularization term to the objective vector function. Variable selection is automatically achieved by selecting variables with non-zero partial derivatives. The sparse empirical gradient covariance matrix (S-EGCM) constructed based on the learned sparse gradient reflects the variance of the data conditioned on the response variable. The eigenvectors of S-EGCM are then used to construct the EDR directions. A major innovation of our approach is that the variable selection and dimension reduction are achieved within a single framework. The features constructed by the eigenvectors of S-EGCM are sparse with non-zero entries corresponding only to selected variables.

The rest of this paper is organized as follows. In Sect. 2, we describe the sparse gradient learning algorithm for regression, where an automatic variable selection scheme is introduced. The derived sparse gradient is an approximation of the true gradient of regression function under certain conditions, which we give in Sect. 2.3 and their proofs are delayed in Sect. 3. We describe variable selection and feature construction using the learned sparse gradients in Sect. 2.4. As our proposed algorithm is an infinite dimensional minimization problem, it can not be solved directly. We provide an efficient implementation for solving it in Sect. 4. In Sect. 4.1, we give a representer theorem, which transfers the infinite dimensional sparse gradient learning problem to a finite dimensional one. In Sect. 4.3, we solve the transferred finite dimensional minimization problem by a forward-backward splitting algorithm. In Sect. 5, we generalize the sparse gradient learning algorithm to a classification setting. We illustrate the effectiveness of our gradient-based variable selection and feature extraction approach in Sect. 6 using both simulated and real-world examples.

2 Sparse gradient learning for regression

2.1 Basic definitions

Let y and \mathbf{x} be respectively \mathbb{R} -valued and \mathbb{R}^p -valued random variables. The problem of regression is to estimate the regression function $f_\rho(\mathbf{x}) = \mathbb{E}(y|\mathbf{x})$ from a set of observations $\mathcal{Z} := \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, where $\mathbf{x}_i := (x_i^1, \dots, x_i^p)^T \in \mathbb{R}^p$ is an input, and $y_i \in \mathbb{R}$ is the corresponding output.

We assume the data are drawn i.i.d. from a joint distribution $\rho(\mathbf{x}, y)$, and the response variable y depends only on a few directions in \mathbb{R}^p as follows

$$y = f_\rho(\mathbf{x}) + \epsilon = g(b_1^T \mathbf{x}, \dots, b_r^T \mathbf{x}) + \epsilon, \quad (1)$$

where ϵ is the noise, $B = (b_1, \dots, b_r)$ is a $p \times r$ orthogonal matrix with $r < p$, and $\mathbb{E}(\epsilon|\mathbf{x}) = 0$ almost surely. We call the r dimensional subspace spanned by $\{b_i\}_{i=1}^r$ the effective dimension reduction (EDR) space (Xia et al. 2002). For high-dimensional data, we further assume that B is a sparse matrix with many rows being zero vectors, i.e. the regression function only depends on a subset of variables in \mathbf{x} .

Suppose the regression function $f_\rho(\mathbf{x})$ is differentiable. The gradient of f_ρ with respect to variables is

$$\nabla f_\rho := \left(\frac{\partial f_\rho}{\partial x^1}, \dots, \frac{\partial f_\rho}{\partial x^p} \right)^T. \quad (2)$$

A quantity of particular interest is the gradient outer product matrix $G = (G_{ij})$, a $p \times p$ matrix with elements

$$G_{ij} := \left\langle \frac{\partial f_\rho}{\partial x^i}, \frac{\partial f_\rho}{\partial x^j} \right\rangle_{L^2_{\rho_X}}, \tag{3}$$

where ρ_X is the marginal distribution of \mathbf{x} . As pointed out by Li (1991) and Xia et al. (2002), under the assumption of the model in Eq. (1), the gradient outer product matrix G is at most of rank r , and the EDR spaces are spanned by the eigenvectors corresponding to non-zero eigenvalues of G . This observation has motivated the development of gradient-based methods for inferring the EDR directions (Xia et al. 2002; Mukherjee and Zhou 2006; Mukherjee and Wu 2006), and also forms the basis of our approach.

2.2 Regularization framework for sparse gradient learning

The optimization framework for sparse gradient learning includes a data fitting term and a regularization term. We first describe the data fitting term. Given a set of observations \mathcal{Z} , a commonly used data fitting term for regression is the mean square error $\frac{1}{n} \sum_{i=1}^n (y_i - f_\rho(\mathbf{x}_i))^2$. However, because our primary goal is to estimate the gradient of f_ρ , we adopts the framework of local linear regression (Ruppert and Ward 1994), which involves using a kernel to weight the data points so that only the data points within the neighborhood of an evaluation point effectively contribute on function estimation. More specifically, we use the first order Taylor expansion to approximate f_ρ by $f_\rho(\mathbf{x}) \approx f_\rho(\mathbf{x}_0) + \nabla f_\rho(\mathbf{x}_0) \cdot (\mathbf{x} - \mathbf{x}_0)$. When \mathbf{x}_j is close to \mathbf{x}_i , $f_\rho(\mathbf{x}_j) \approx y_i + \nabla f_\rho(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i)$. Define $\mathbf{f} := (f^1, \dots, f^p)$, where $f^j = \partial f_\rho / \partial x^j$ for $j = 1, \dots, p$. The mean square error used in our algorithm is

$$\mathcal{E}_{\mathcal{Z}}(\mathbf{f}) = \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + \mathbf{f}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i))^2 \tag{4}$$

considering Taylor expansion between all pairs of observations. Here $\omega_{i,j}^s$ is a weight kernel function that ensures the locality of the approximation, i.e. $\omega_{i,j}^s \rightarrow 0$ when $|\mathbf{x}_i - \mathbf{x}_j|$ is large. We can use, for example, the Gaussian with standard deviation s as a weight function. Let $\omega^s(\mathbf{x}) = \exp\{-\frac{|\mathbf{x}|^2}{2s^2}\}$. Then the weights are given by

$$\omega_{i,j}^s = \omega^s(\mathbf{x}_j - \mathbf{x}_i) = \exp\left\{-\frac{|\mathbf{x}_j - \mathbf{x}_i|^2}{2s^2}\right\}, \tag{5}$$

for all $i, j = 1, \dots, n$, with parameter s controlling the bandwidth of the weight function. In this paper, we view s as a parameter and is fixed in implementing our algorithm, although it is possible to tune s using a greedy algorithm as RODEO in Lafferty and Wasserman (2008).

At first glance, this data fitting term might not appear very meaningful for high-dimensional data as samples are typically distributed sparsely in a high dimensional space. However, the term can also be explained in the manifold setting (Mukherjee et al. 2010), in which case the approximation is well defined as long as the data lying in the low dimensional manifold are relatively dense. More specifically, assume X is a d -dimensional connected compact C^∞ submanifold of \mathbb{R}^p which is isometrically embedded. Then X is a metric space with a metric d_X and the inclusion map $\Phi : (X, d_X) \mapsto (\mathbb{R}^p, \|\cdot\|_2)$ is well defined and continuous (actually it is C^∞). Note that the empirical data $\{\mathbf{x}_i\}_{i=1}^n$ given in the Euclidean space \mathbb{R}^p are images of the points $\{\mathbf{q}_i\}_{i=1}^n \subset X$ under $\Phi : \mathbf{x}_i = \Phi(\mathbf{q}_i)$. Then the data fitting term (4) can be explained in the manifold setting. From the first order Taylor

expansion, if \mathbf{q}_i and \mathbf{q}_j are close enough, we can expect that $y_j \approx y_i + \langle \nabla_X f_\rho(\mathbf{q}_i), v_{ij} \rangle_{\mathbf{q}_i}$, where $v_{ij} \in T_{\mathbf{q}_i} X$ is the tangent vector such that $\mathbf{q}_j = \exp_{\mathbf{q}_i}(v_{ij})$. However, v_{ij} is not easy to compute, so we would like to represent the term $\langle \nabla_X f_\rho(\mathbf{q}_i), v_{ij} \rangle_{\mathbf{q}_i}$ in the Euclidean space \mathbb{R}^p . Suppose $\mathbf{x} = \Phi(\mathbf{q})$ and $\xi = \Phi(\exp_{\mathbf{q}}(v))$ for $\mathbf{q} \in X$ and $v \in T_{\mathbf{q}} X$. Since Φ is an isometric embedding, i.e. $d\Phi_{\mathbf{q}} : T_{\mathbf{q}} X \mapsto T_{\mathbf{x}} \mathbb{R}^p \cong \mathbb{R}^p$ is an isometry for every $\mathbf{q} \in X$, the following holds

$$\langle \nabla_X f(\mathbf{q}), v \rangle_{\mathbf{q}} = \langle d\Phi_{\mathbf{q}}(\nabla_X f(\mathbf{q})), d\Phi_{\mathbf{q}}(v) \rangle_{\mathbb{R}^p},$$

where $d\Phi_{\mathbf{q}}(v) \approx \phi(\exp_{\mathbf{q}}(v)) - \phi(\mathbf{q}) = \xi - \mathbf{x}$ for $v \approx 0$. Applying these relations to the observations $\mathcal{Z} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ and denoting $\mathbf{f} = d\Phi(\nabla_X f)$ yields

$$\mathcal{E}_{\mathcal{Z}}(\mathbf{f}) = \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + \mathbf{f}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i))^2. \tag{6}$$

This is exactly the same as the one in the Euclidean setting.

Now we turn to the regularization term on ∇f_ρ . As discussed above, we impose a sparsity constraint on the gradient vector \mathbf{f} . The motivation for the sparse constraint is based on the following two considerations: (1) Since most variables are assumed to be irrelevant for prediction, we expect the partial derivatives of f_ρ with respect to these variables to be zero; and (2) If variable x^j is important for prediction, we expect the function f_ρ should show significant variation along x^j , and as such the norm of $\frac{\partial f_\rho}{\partial x^j}$ should be large. Thus we will impose the sparsity constraint on the vector $(\|\frac{\partial f_\rho}{\partial x^1}\|, \dots, \|\frac{\partial f_\rho}{\partial x^p}\|)^T \in \mathbb{R}^p$, where $\|\cdot\|$ is a function norm, to regularize the number of non-zeros entries in the vector.

In this work, we specify the function norm $\|\cdot\|$ to be $\|\cdot\|_{\mathcal{K}}$, the norm in the reproducing kernel Hilbert space (RKHS) $\mathbb{H}_{\mathcal{K}}$ associated with a Mercer kernel $\mathcal{K}(\cdot, \cdot)$ (see Aronszajn 1950 and Sect. 4.1). The sparsity constraint on the gradient norm vector implies that the ℓ_0 norm of the vector $(\|f^1\|_{\mathcal{K}}, \dots, \|f^p\|_{\mathcal{K}})^T$ should be small. However, because the ℓ_0 norm is difficult to work with during optimization, we instead use the ℓ_1 norm of the vector (Donoho 1995, 2006; Daubechies et al. 2004; Micchelli and Pontil 2007) as our regularization term

$$\Omega(\mathbf{f}) := \lambda \sum_{j=1}^p \|f^j\|_{\mathcal{K}}, \tag{7}$$

where λ is a sparsity regularization parameter. This functional LASSO penalty has been used in Lin and Zhang (2006) as COSSO penalty. However, our component here is quite different from theirs, which makes our algorithm useful for high dimensional problems.

The norm $\|\cdot\|_{\mathcal{K}}$ is widely used in statistical inference and machine learning (see Vapnik 1998). It can ensure each approximated partial derivative $f^j \in \mathbb{H}_{\mathcal{K}}$, which in turn imposes some regularity on each partial derivative. It is possible to replace the hypothesis space $\mathbb{H}_{\mathcal{K}}^p$ for the vector \mathbf{f} in (7) by some other space of vector-valued functions (Micchelli and Pontil 2005) in order to learn the gradients.

Combining the data fidelity term (4) and the regularization term (7), we propose the following optimization framework, which will be referred as *sparse gradient learning*, to learn ∇f_ρ

$$\mathbf{f}_{\mathcal{Z}} := \arg \min_{\mathbf{f} \in \mathbb{H}_{\mathcal{K}}^p} \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + \mathbf{f}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i))^2 + \lambda \sum_{j=1}^p \|f^j\|_{\mathcal{K}}. \tag{8}$$

A key difference between our framework and the one in Mukherjee and Zhou (2006) is that our regularization is based on ℓ_1 norm, while the one in Mukherjee and Zhou (2006) is based on ridge regularization. The difference may appear minor, but makes a significant impact on the estimated ∇f_ρ . In particular, ∇f_ρ derived from Eq. (8) is sparse with many components potentially being zero functions, in contrast to the one derived from Mukherjee and Zhou (2006), which is comprised of all non-zero functions. The sparsity property is desirable for two primary reasons: (1) In most high-dimensional real-world data, the response variable is known to depend only on a subset of the variables. Imposing sparsity constraints can help eliminate noisy variables and thus improve the accuracy for inferring the EDR directions; (2) The resulting gradient vector provides a way to automatically select and rank relevant variables.

Remark 1 The OPG method introduced by Xia et al. (2002) to learn EDR directions can be viewed as a special case of the sparse gradient learning, corresponding to the case of setting $K(x, y) = \delta_{x,y}$ and $\lambda = 0$ in Eq. (8). Thus the sparse gradient learning can be viewed as an extension of learning gradient vectors only at observed points by OPG to a vector function of gradient over the entire space. Note that OPG cannot be directly applied to the data with $p > n$ since the problem is then underdetermined. Imposing a regularization term as in Eq. (8) removes such a limitation.

Remark 2 The sparse gradient learning reduces to a special case that is approximately LASSO (Tibshirani 1996) if we choose $K(x, y) = \delta_{x,y}$ and additionally require $\mathbf{f}(\mathbf{x}_i)$ to be invariant for different i (i.e. linearity assumption). Note that LASSO assumes the regression function is linear, which can be problematic for variable selection when the prediction function is nonlinear (Efron et al. 2004). The sparse gradient learning makes no linearity assumption, and can thus be viewed as an extension of LASSO for variable selection with nonlinear prediction functions.

Remark 3 A related framework is to learn the regression function directly, but impose constraints on the sparsity of the gradient as follows

$$\min_{f \in \mathbb{H}_{\mathcal{K}}} \frac{1}{n} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i)^2 + \lambda \sum_{i=1}^p \left\| \frac{\partial f}{\partial x^i} \right\|_{\mathcal{K}}. \tag{9}$$

This framework is however difficult to solve because the regularization term $\sum_{i=1}^p \left\| \frac{\partial f}{\partial x^i} \right\|_{\mathcal{K}}$ is both nonsmooth and inseparable, and the representer theorem introduced later to solve Eq. (8) cannot be applied here. Note that our primary goal is to select variables and identify the EDR directions. Thus we focus on learning gradient functions rather than the regression function itself.

Remark 4 We can also use the regularization term $(\sum_{j=1}^p \|f^j\|_{\mathcal{K}})^2$, which is more widely used in literature (Bach 2008). Our framework is equivalent to, with a different choice of λ ,

$$\mathbf{f}_{\mathcal{Z}} := \arg \min_{\mathbf{f} \in \mathbb{H}_{\mathcal{K}}^p} \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + \mathbf{f}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i))^2 + \lambda \left(\sum_{j=1}^p \|f^j\|_{\mathcal{K}} \right)^2. \tag{10}$$

2.3 Error analysis

Next we investigate the statistical performance of the sparse gradient learning with a Gaussian weight in Eq. (5). Throughout the paper on error analysis, we only consider the case where p is fixed and does not change with n . Assume that the data $\mathcal{Z} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ are i.i.d. drawn from a joint distribution ρ , which can be divided into a marginal distribution ρ_X and a conditional distribution $\rho(y|\mathbf{x})$. Denote f_ρ to be the regression function given by

$$f_\rho(x) = \int_Y y d\rho(y|\mathbf{x}).$$

We show that under certain conditions, $\mathbf{f}_Z \rightarrow \nabla f_\rho$ as $n \rightarrow \infty$ for suitable choices of the parameters λ and s that go to zero as $n \rightarrow \infty$. In order to derive the learning rate for the algorithm, some regularity conditions on both the marginal distribution and ∇f_ρ are required.

Denote ∂X be the boundary of X and $d(\mathbf{x}, \partial X)$ ($\mathbf{x} \in X$) be the shortest Euclidean distance from \mathbf{x} to ∂X , i.e., $d(\mathbf{x}, \partial X) = \inf_{\mathbf{y} \in \partial X} d(\mathbf{x}, \mathbf{y})$. Denote $\kappa = \sup_{\mathbf{x} \in X} \sqrt{\mathcal{K}(\mathbf{x}, \mathbf{x})}$.

Theorem 1 *Suppose the data $\mathcal{Z} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ are i.i.d. drawn from a joint distribution ρ and $|y_i| \leq M$ for all i for a positive constant M . Assume that for some constants $c_\rho > 0$ and $0 < \theta \leq 1$, the marginal distribution ρ_X satisfies*

$$\rho_X(\{\mathbf{x} \in X : d(\mathbf{x}, \partial X) < t\}) \leq c_\rho t \tag{11}$$

and the density $p(\mathbf{x})$ of ρ_X satisfies

$$\sup_{\mathbf{x} \in X} p(\mathbf{x}) \leq c_\rho \quad \text{and} \quad |p(\mathbf{x}) - p(\mathbf{u})| \leq c_\rho |\mathbf{x} - \mathbf{u}|^\theta, \quad \forall \mathbf{u}, \mathbf{x} \in X. \tag{12}$$

Let \mathbf{f}_Z be the estimated gradient function given by Eq. (8) and ∇f_ρ be the true gradient of the regression function f_ρ . Suppose that $\mathcal{K} \in C^2$ and $\nabla f_\rho \in \mathbb{H}_{\mathcal{K}}^p$. Choose $\lambda = \lambda(n) = n^{-\frac{\theta}{p+2+2\theta}}$ and $s = s(n) = n^{-\frac{1}{2(p+2+2\theta)}}$. Then there exists a constant $C > 0$ such that for any $0 < \eta \leq 1$ with confidence $1 - \eta$

$$\|\mathbf{f}_Z - \nabla f_\rho\|_{L_{\rho_X}^2} \leq C \log \frac{4}{\eta} \left(\frac{1}{n}\right)^{\frac{\theta}{4(p+2+2\theta)}}. \tag{13}$$

Condition (12) means the density of the marginal distribution is Hölder continuous with exponent θ . Condition (14) specifies the behavior of ρ_X near the boundary ∂X of X . Both are common assumptions for error analysis. When the boundary ∂X is piecewise smooth, Eq. (12) implies Eq. (14). Here we want to emphasize that our terminology sparse gradient for the derived \mathbf{f}_Z comes from this approximation property. Since we treat each component of the gradient separately in our estimation algorithm, \mathbf{f}_Z does not necessarily satisfy the gradient constraint $\frac{\partial^2 f}{\partial x^i \partial x^j} = \frac{\partial^2 f}{\partial x^j \partial x^i}$ for all i and j . However, we note that it is possible to add these constraints explicitly into the convex optimization framework that we will describe later.

The convergence rate in Eq. (13) can be greatly improved if we assume that the data are lying in or near a low dimensional manifold (Ye and Zhou 2008; Mukherjee et al. 2010; Bickel and Li 2007). In this case, the learning rate in the exponent of $1/n$ depends only

on the dimension of the manifold, not the actual dimension of the Euclidean space. The improved convergence rate for local linear regression under manifold assumption appeared in Bickel and Li (2007). Here we would like to emphasize that our result is different from theirs in two points of view. First, our algorithm is different from the one discussed in Bickel and Li (2007). Second, we focus on the case where the distribution of the predictor variables is concentrated on a manifold and our criterion of performance is the integral of pointwise mean error with respect to the underlying distribution of the variables; by contrast, the discussion in Bickel and Li (2007) is more restrictive by applying only to predictors taking values in a low dimensional manifold and discussing estimation of the regression function at a point.

Denote d_X be the metric on X and dV be the Riemannian volume measure of X . Let ∂X be the boundary of X and $d_X(\mathbf{x}, \partial X)$ ($\mathbf{x} \in X$) be the shortest distance from \mathbf{x} to ∂X on the manifold X . Note that the inclusion map $\Phi : (X, d_X) \mapsto (\mathbb{R}^p, \|\cdot\|_2)$ is an isometric embedding and the empirical data $\{\mathbf{x}_i\}_{i=1}^n$ are given in the Euclidean space \mathbf{R}^p which are images of the points $\{\mathbf{q}_i\}_{i=1}^n \subset X$ under $\Phi : \mathbf{x}_i = \Phi(\mathbf{q}_i)$. Denote $(d\Phi)_{\mathbf{q}_i}^*$ the dual of $d\Phi_{\mathbf{q}_i}$ and $(d\Phi)^*$ maps a p -dimensional vector valued function \mathbf{f} to a vector field with $(d\Phi)^*\mathbf{f}(\mathbf{q}) = (d\Phi)_{\mathbf{q}}^*(\mathbf{f}(\mathbf{q}))$ (Do Carmo and Flaherty 1992).

Theorem 2 *Let X be a connected compact C^∞ submanifold of \mathbb{R}^p which is isometrically embedded and of dimension d . Suppose the data $\mathcal{Z} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ are i.i.d. drawn from a joint distribution ρ defined on $X \times Y$ and there exists a positive constant M such that $y_i \leq M$ for all i . Assume that for some constants $c_\rho > 0$ and $0 < \theta \leq 1$, the marginal distribution ρ_X satisfies*

$$\rho_X(\{\mathbf{x} \in X : d_X(\mathbf{x}, \partial X) < t\}) \leq c_\rho t \tag{14}$$

and the density $p(\mathbf{x}) = \frac{d\rho_X(\mathbf{x})}{dV}$ exists and satisfies

$$\sup_{\mathbf{x} \in X} p(\mathbf{x}) \leq c_\rho \quad \text{and} \quad |p(\mathbf{x}) - p(\mathbf{u})| \leq c_\rho d_X(\mathbf{x}, \mathbf{u})^\theta, \quad \forall \mathbf{u}, \mathbf{x} \in X. \tag{15}$$

Let \mathbf{f}_Z be the estimated gradient function given by Eq. (8) and $\nabla_X f_\rho$ be the true gradient of the regression function f_ρ . Suppose that $\mathcal{K} \in C^2(X \times X)$, $f_\rho \in C^2(X)$ and $d\Phi(\nabla_X f_\rho) \in \mathbb{H}_{\mathcal{K}}^p$. Choose $\lambda = \lambda(n) = n^{-\frac{\theta}{d+2+2\theta}}$ and $s = s(n) = n^{-\frac{1}{2(d+2+2\theta)}}$. Then there exists a constant $C > 0$ such that for any $0 < \eta \leq 1$ with confidence $1 - \eta$

$$\|(d\Phi)^*\mathbf{f}_Z - \nabla_X f_\rho\|_{L^2_{\rho_X}} \leq C \log \frac{4}{\eta} \left(\frac{1}{n}\right)^{\frac{\theta}{4(d+2+2\theta)}}. \tag{16}$$

Note that the convergence rate in Theorem 2 is exactly the same as the one in Theorem 1 except that we replaced the Euclidean dimension p by the intrinsic dimension d . The constraints $\nabla_X f_\rho \in \mathcal{H}_{\mathcal{K}}^p$ in Theorem 1 and $d\Phi(\nabla_X f_\rho) \in \mathcal{H}_{\mathcal{K}}^p$ are somewhat restrictive, and extension to mild conditions is possible (Mukherjee et al. 2010). Here we confine ourself to these conditions in order to avoid introducing more notations and conceptions. The proof of Theorems 1 and 2 are somewhat complicated and will be given in Sect. 3. The main idea behind the proof is to simultaneously control the sample error and the approximation error; see Sect. 3 for details.

Note that the main purpose of our method is for variable selection. If f_ρ depends only on a few coordinates of X , we can further improve the convergence rate to $n^{-\frac{\theta}{2(|\mathbf{J}|+2+3\theta)}}$, where $\mathbf{J} = \mathbf{J}(\nabla f_\rho) = \{j : \frac{\partial f_\rho}{\partial x^j} \neq 0\}$ and $|\mathbf{J}|$ is the number of elements in set \mathbf{J} , i.e., the number of

variables relevant to f_ρ (Bertin and Lecué 2008). Let $\mathbf{f} = (f^1, f^2, \dots, f^p)^T$ and $\mathbf{f}_\mathbf{J}$ be the concatenation of the loading function vectors indexed by \mathbf{J} , that is, $\mathbf{f}_\mathbf{J} = (f^j)_{j \in \mathbf{J}}$. Similarly, we define $\mathbf{x}_\mathbf{J} = (x^j)_{j \in \mathbf{J}}$ and $\mathbf{x}_{i,\mathbf{J}} = (x_i^j)_{j \in \mathbf{J}}$.

Next we show how the improved convergence rate can be derived using the framework based in (10), which is equivalent to (8) but easier to handle in proving the convergence results. To derive the results, we will adopt a two-step procedure: (1) we show that (10) selects, with high probability, only the relevant variables $x_j, j \in \mathbf{J}$; and (2) using the selected variables $\widehat{\mathbf{J}}$, we run the following algorithm

$$\mathbf{f}_{\mathcal{Z}, \widehat{\mathbf{J}}} := \arg \min_{\mathbf{f}_\mathbf{J} \in \mathbb{H}_{\mathcal{K}}^{\widehat{\mathbf{J}}}} \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + \mathbf{f}_\mathbf{J}(\mathbf{x}_{i,\widehat{\mathbf{J}}}) \cdot (\mathbf{x}_{j,\widehat{\mathbf{J}}} - \mathbf{x}_{i,\widehat{\mathbf{J}}}))^2 + \lambda \sum_{j \in \widehat{\mathbf{J}}} \|f^j\|_{\mathcal{K}}^2, \tag{17}$$

to obtain an improved convergence rate that depends only on $|\mathbf{J}|$. We will assume that the regression function and the distributions satisfy some regularity properties slightly different from the ones in Theorem 1:

Assumption 1 Assume $|y| \leq M$ almost surely. Suppose that for some $0 < \theta \leq 2/3$ and $C_\rho > 0$, the marginal distribution ρ_X satisfies

$$\rho_X \left(\left\{ \mathbf{x} \in X : \inf_{\mathbf{u} \in \mathbb{R}^p \setminus X} |\mathbf{u} - \mathbf{x}| \leq s \right\} \right) \leq C_\rho^2 s^{4\theta}, \quad \forall s > 0. \tag{18}$$

and the density $p(\mathbf{x})$ of $d\rho_X(\mathbf{x})$ exists and satisfies

$$\sup_{\mathbf{x} \in X} p(\mathbf{x}) \leq C_\rho, \quad |p(\mathbf{x}) - p(\mathbf{u})| \leq C_\rho |\mathbf{u} - \mathbf{x}|^\theta, \quad \forall \mathbf{u}, \mathbf{x} \in X. \tag{19}$$

Denote $V_\rho = \int_X (p(\mathbf{x}))^2 d\mathbf{x} > 0$ and $\mathbf{L}_\mathcal{K} \mathbf{f} = \int_X \mathcal{K}_\mathbf{x} \mathbf{f}(\mathbf{x}) \frac{p(\mathbf{x})}{V_\rho} d\rho_X(\mathbf{x})$. Suppose ∇f_ρ lies in the range of $\mathbf{L}'_\mathcal{K}$, $r > \frac{1}{2}$ (i.e. $\|\mathbf{L}'_\mathcal{K} \nabla f_\rho\|_{L^2(\rho_X)} < \infty$) and has a sparsity pattern $\mathbf{J} = \mathbf{J}(\nabla f_\rho) = \{j : \frac{\partial f_\rho}{\partial x^j} \neq 0\}$. For some $C'_\rho > 0$, f_ρ satisfies

$$|f_\rho(\mathbf{u}) - f_\rho(\mathbf{x}) - \nabla f_\rho(\mathbf{x}) \cdot (\mathbf{u} - \mathbf{x})| \leq C'_\rho |\mathbf{u} - \mathbf{x}|^2, \quad \forall \mathbf{u}, \mathbf{x} \in X.$$

Let $B_1 = \{\mathbf{f}_\mathbf{J} : \|\mathbf{f}_\mathbf{J}\|_{\mathcal{K}} \leq 1\}$ and $J_\mathcal{K}$ be the inclusion map from B_1 to $C(X)$. Let $0 < \eta < \frac{1}{2}$. We define the covering number $\mathcal{N}(J_\mathcal{K}(B_1), \eta)$ to be the minimal $\ell \in \mathbb{N}$ such that there exists ℓ disks in $J_\mathcal{K}(B_1)$ with radius η covering S . The following Theorem 3 tells us that, with probability tending to 1, Eq. (10) selects the true set of relevant variables. Theorem 4 shows the improved convergence rate which only depends on $|\mathbf{J}|$ if we use the two-step procedure to learn gradients. The proofs of Theorems 3 and 4 are postponed to Sect. 3.

Theorem 3 Suppose the data $\mathcal{Z} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ are i.i.d. drawn from a joint distribution ρ defined on $X \times Y$. Let $\mathbf{f}_\mathcal{Z}$ be defined as (10) and $\widehat{\mathbf{J}} = \{j : f_{\mathcal{Z}}^j \neq 0\}$. Choose $\lambda = \widetilde{C}_{M,\theta} s^{p+2+\theta}$, where $\widetilde{C}_{M,\theta}$ is a constant defined in (35). Then under Assumption 1, there exists a constant $0 < s_0 \leq 1$ such that, for all s, n satisfying $0 < s < s_0$ and $n^r s^{(2p+4+\theta)(1/2-r)+1-\theta} \geq \max\{C_{D,\theta}, \widetilde{C}_{D,\theta}\}$, where $C_{D,\theta}, \widetilde{C}_{D,\theta}$ are two constants defined in Propositions 7 and 9 separately, we have $\widehat{\mathbf{J}} = \mathbf{J}$ with probability greater than

$$1 - \widetilde{C}_1 \mathcal{N} \left(J_\mathcal{K}(B_1), \frac{s^{p+2+\theta}}{8(\kappa \text{Diam}(X))^2} \right) \exp\{-\widetilde{C}_2 n s^{p+4}\},$$

where \widetilde{C}_1 and \widetilde{C}_2 are two constants independent of n or s .

Remark 5 The estimation of covering number $\mathcal{N}(J_{\mathcal{K}}(B_1), \eta)$ is dependent of the smoothness of the Mercer Kernel \mathcal{K} (Cucker and Zhou 2007). If $\mathcal{K} \in C^\beta(X \times X)$ for some $\beta > 0$ and X has piecewise smooth boundary, then there is $C > 0$ independent of β such that $\mathcal{N}(J_{\mathcal{K}}(B_1), \eta) \leq C(\frac{1}{\eta})^{2p/\beta}$. In this case, if we choose $s > (\frac{1}{n})^{\frac{\beta}{(p+4)\beta+2p(p+2+\theta)}}$, then $1 - \mathcal{N}(J_{\mathcal{K}}(B_1), \frac{s^{p+2+\theta}}{8(\kappa \text{Diam}(X))^2}) \exp\{-\tilde{C}_2 ns^{p+4}\}$ will goes to 1. In particular, if we choose $s = (\frac{1}{n})^{\frac{\beta}{2(p+4)\beta+4p(p+2+\theta)}}$, then $1 - \tilde{C}_1 \mathcal{N}(J_{\mathcal{K}}(B_1), \frac{s^{p+2+\theta}}{8(\kappa \text{Diam}(X))^2}) \exp\{-\tilde{C}_2 ns^{p+4}\} \geq 1 - \tilde{C}'_1 \exp\{\tilde{C}'_2 \sqrt{n}\}$, where $\tilde{C}'_1, \tilde{C}'_2$ are two constants independent of n .

The following Theorem shows that the convergence rate depends on $|\mathbf{J}|$, instead of p , if we use the two-step procedure to learn the gradients.

Theorem 4 *Suppose the data $\mathcal{Z} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ are i.i.d. drawn from a joint distribution ρ defined on $X \times Y$. Let $\mathbf{f}_{\mathcal{Z}, \hat{\mathbf{J}}}$ be defined as in (17). Assume that there exists $\alpha > 0$ and $C_\alpha > 0$ such that $\ln \mathcal{N}(J_{\mathcal{K}}(B_1), \frac{\epsilon}{8\kappa^2(\text{Diam}(X))^2}) \leq C_\alpha (\frac{1}{\epsilon})^\alpha$. Let $\nabla_{\hat{\mathbf{J}}} f_\rho = (\frac{\partial f_\rho}{\partial x^j})_{j \in \hat{\mathbf{J}}}$. Under the same conditions as in Theorem 3, we have*

$$\text{Prob}\{\|\mathbf{f}_{\mathcal{Z}, \hat{\mathbf{J}}} - \nabla_{\hat{\mathbf{J}}} f_\rho\|_{L^2(\rho_X)} \geq \epsilon\} \leq \tilde{C}_3 \exp(-\tilde{C}_4 n^{\frac{\theta}{2(|\mathbf{J}|+3\theta)} \epsilon}), \quad \forall \epsilon > 0,$$

where \tilde{C}_3 and \tilde{C}_4 are two constants independent of n or s .

2.4 Variable selection and effective dimension reduction

Next we describe how to do variable selection and extract EDR directions based on the learned gradient $\mathbf{f}_{\mathcal{Z}} = (f_{\mathcal{Z}}^1, \dots, f_{\mathcal{Z}}^p)^T$.

As discussed above, because of the l_1 norm used in the regularization term, we expect many of the entries in the gradient vector $\mathbf{f}_{\mathcal{Z}}$ be zero functions. Thus, a natural way to select variables is to identify those entries with non-zeros functions. More specifically, we select variables based on the following criterion.

Definition 1 Variable selection via sparse gradient learning is to select variables in the set

$$\mathcal{S} := \{j : \|f_{\mathcal{Z}}^j\|_{\mathcal{K}} \neq 0, j = 1, \dots, p\} \tag{20}$$

where $\mathbf{f}_{\mathcal{Z}} = (f_{\mathcal{Z}}^1, \dots, f_{\mathcal{Z}}^p)^T$ is the estimated gradient vector.

To select the EDR directions, we focus on the empirical gradient covariance matrix defined below

$$\mathcal{E} := [(f_{\mathcal{Z}}^i, f_{\mathcal{Z}}^j)_{\mathcal{K}}]_{i,j=1}^p \tag{21}$$

The inner product $(f_{\mathcal{Z}}^i, f_{\mathcal{Z}}^j)_{\mathcal{K}}$ can be interpreted as the covariance of the gradient functions between coordinate i and j . The larger the inner product is, the more related the variables x^i and x^j are. Given a unit vector $\mathbf{u} \in \mathbb{R}^p$, the RKHS norm of the directional derivative $\|\mathbf{u} \cdot \mathbf{f}_{\mathcal{Z}}\|_{\mathcal{K}}$ can be viewed as a measure of the variation of the data \mathcal{Z} along the direction \mathbf{u} . Thus the direction \mathbf{u}_1 representing the largest variation in the data is the vector that maximizes $\|\mathbf{u} \cdot \mathbf{f}_{\mathcal{Z}}\|_{\mathcal{K}}^2$. Notice that

$$\|\mathbf{u} \cdot \mathbf{f}_{\mathcal{Z}}\|_{\mathcal{K}}^2 = \|\sum_i u_i f_{\mathcal{Z}}^i\|_{\mathcal{K}}^2 = \sum_{i,j} u_i u_j (f_{\mathcal{Z}}^i, f_{\mathcal{Z}}^j)_{\mathcal{K}} = \mathbf{u}^T \mathcal{E} \mathbf{u}.$$

So \mathbf{u}_1 is simply the eigenvector of \mathcal{E} corresponding to the largest eigenvalue. Similarly, to construct the second most important direction \mathbf{u}_2 , we maximize $\|\mathbf{u} \cdot \mathbf{f}_{\mathcal{Z}}\|_{\mathcal{K}}$ in the orthogonal complementary space of $\text{span}\{\mathbf{u}_1\}$. By Courant-Fischer Minimax Theorem (Golub and Van Loan 1989), \mathbf{u}_2 is the eigenvector corresponding to the second largest eigenvalue of \mathcal{E} . We repeat this procedure to construct other important directions. In summary, the effective dimension reduction directions are defined according to the following criterion.

Definition 2 The d EDR directions identified by the sparse gradient learning are the eigenvectors $\{\mathbf{u}_1, \dots, \mathbf{u}_d\}$ of \mathcal{E} corresponding to the d largest eigenvalues.

As we mentioned in Sect. 2.1, the EDR space is spanned by the eigenvectors of the gradient outer product matrix G defined in Eq. (3). However, because the distribution of the data is unknown, G cannot be calculated explicitly. The above definition provides a way to approximate the EDR directions based on the empirical gradient covariance matrix.

Because of the sparsity of the estimated gradient functions, matrix \mathcal{E} will appear to be block sparse. Consequently, the identified EDR directions will be sparse as well with non-zeros entries only at coordinates belonging to the set S . To emphasize the sparse property of both \mathcal{E} and the identified EDR directions, we will refer to \mathcal{E} as the sparse empirical gradient covariance matrix (S-EGCM), and the identified EDR directions as the sparse effective dimension reduction directions (S-EDRs).

3 Convergence analysis

In this section, we will give the proof of Theorems 1, 2, 3 and 4.

3.1 Convergence analysis in the Euclidean setting

Note that our energy functional in (8) involves a nonsmooth regularization term $\sum_i \|f^i\|_{\mathcal{K}}$. The method for the convergence analysis used in Mukherjee and Zhou (2006) can no longer be applied any more since it need explicit form of the solution which is only possible for the ℓ^2 regularization. However, we can still simultaneously control a sample or estimation error term and a regularization or approximation error term which is widely used in statistical learning theory (Vapnik 1998; Mukherjee and Wu 2006; Zhang 2004).

3.1.1 Comparative analysis

Recall the empirical error for a vector function $\mathbf{f} := (f^1, \dots, f^p)$,

$$\mathcal{E}_{\mathcal{Z}}(\mathbf{f}) = \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + \mathbf{f}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i))^2.$$

One can similarly define the expected error

$$\mathcal{E}(\mathbf{f}) = \int_{\mathcal{Z}} \int_{\mathcal{Z}} \omega^s(\mathbf{x} - \mathbf{u}) (y - v + \mathbf{f}(\mathbf{x}) \cdot (\mathbf{u} - \mathbf{x}))^2 d\rho(\mathbf{x}, y) d\rho(\mathbf{u}, v).$$

Denote

$$\sigma_s^2 = \int_{\mathcal{X}} \int_{\mathcal{Z}} \omega^s(\mathbf{x} - \mathbf{u}) (y - f_{\rho}(\mathbf{x}))^2 d\rho(\mathbf{x}, y) d\rho_{\mathcal{X}}(\mathbf{u}).$$

Then

$$\mathcal{E}(\mathbf{f}) = 2\sigma_s^2 + \int_X \int_X \omega(\mathbf{x} - \mathbf{u}) [f_\rho(\mathbf{x}) - f_\rho(\mathbf{u}) + \mathbf{f}(\mathbf{x}) \cdot (\mathbf{u} - \mathbf{x})]^2 d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}).$$

Note that our goal is to bound the $L^2_{\rho_X}$ differences of \mathbf{f} and ∇f_ρ . We have the following comparative theorem to bound the $L^2_{\rho_X}$ differences of \mathbf{f} and ∇f_ρ in terms of the excess error, $\mathcal{E}(\mathbf{f}) - 2\sigma_s^2$ using the following comparative theorem.

For $r > 0$, denote

$$\mathcal{F}_r = \left\{ \mathbf{f} \in \mathcal{H}_K^p : \sum_{i=1}^p \|f^i\|_K \leq r \right\}.$$

Theorem 5 Assume ρ_X satisfies the condition (11) and (12) and $\nabla f_\rho \in \mathcal{H}_K^p$. For $\mathbf{f} \in \mathcal{F}_r$ with some $r \geq 1$, there exists a constant $C_0 > 0$ such that

$$\|\mathbf{f} - \nabla f_\rho\|_{L^2_{\rho_X}} \leq C_0 \left(r^2 s^\theta + s^{2-\theta} + \frac{1}{s^{p+2+\theta}} (\mathcal{E}(\mathbf{f}) - 2\sigma_s^2) \right).$$

The proof of Theorem 5 is given in Appendix A.

3.1.2 Error decomposition

Now we turn to bound the quantity $\mathcal{E}(\mathbf{f}_Z) - 2\sigma_s^2$. Note that unlike the standard setting of regression and classification, $\mathcal{E}_Z(\mathbf{f})$ and $\mathcal{E}(\mathbf{f})$ are not respectively the empirical and expected mean of a random variable. This is due to the extra $d\rho(\mathbf{u}, v)$ in the expected error term. However, since

$$E_Z \mathcal{E}_Z(\mathbf{f}) = \frac{n-1}{n} \mathcal{E}(\mathbf{f}),$$

$\mathcal{E}_Z(\mathbf{f})$ and $\mathcal{E}(\mathbf{f})$ should be close to each other if the empirical error concentrates with n increasing. Thus, we can still decompose $\mathcal{E}(\mathbf{f}_Z) - 2\sigma_s^2$ into a sample error term and an approximation error term.

Note that $\Omega(\mathbf{f}) = \lambda \sum_i \|f^i\|_K$ with $\mathbf{f} = (f^1, \dots, f^p)$, so the minimizer of $\mathcal{E}(\mathbf{f}) + \Omega(\mathbf{f})$ in \mathbb{H}_K^p depends on λ . Let

$$\mathbf{f}_\lambda = \arg \min_{\mathbf{f} \in \mathbb{H}_K^p} \{ \mathcal{E}(\mathbf{f}) + \Omega(\mathbf{f}) \}. \tag{22}$$

By a standard decomposition procedure, we have the following result.

Proposition 1 Let

$$\varphi(Z) = (\mathcal{E}(\mathbf{f}_Z) - \mathcal{E}_Z(\mathbf{f}_Z)) + (\mathcal{E}_Z(\mathbf{f}_\lambda) - \mathcal{E}(\mathbf{f}_\lambda))$$

and

$$\mathcal{A}(\lambda) = \inf_{\mathbf{f} \in \mathbb{H}_K^p} \{ \mathcal{E}(\mathbf{f}) - 2\sigma_s^2 + \Omega(\mathbf{f}) \}.$$

Then, we have

$$\mathcal{E}(\mathbf{f}_Z) - 2\sigma_s^2 \leq \mathcal{E}(\mathbf{f}_Z) - 2\sigma_s^2 + \Omega(\mathbf{f}_Z) \leq \varphi(Z) + \mathcal{A}(\lambda)$$

The quantity $\varphi(Z)$ is called the sample error and $\mathcal{A}(\lambda)$ is the approximation error.

3.1.3 Sample error estimation

Note that the sample error $\varphi(\mathcal{Z})$ can be bounded by controlling

$$S(\mathcal{Z}, r) := \sup_{\mathbf{f} \in \mathcal{F}_r} |\mathcal{E}_{\mathcal{Z}}(\mathbf{f}) - \mathcal{E}(\mathbf{f})|.$$

In fact, if both $\mathbf{f}_{\mathcal{Z}}$ and \mathbf{f}_λ are in \mathcal{F}_r for some $r > 0$, then

$$\varphi(\mathcal{Z}) \leq 2S(\mathcal{Z}, r). \tag{23}$$

We use McDiarmid’s inequality in McDiarmid (1989) to bound $S(\mathcal{Z}, r)$. Denote $\kappa = \sup_{\mathbf{x} \in X} \sqrt{K(x, x)}$, $\text{Diam}(X) = \max_{\mathbf{x}, \mathbf{u} \in X} \|\mathbf{x} - \mathbf{u}\|$.

Lemma 1 For every $r > 0$,

$$\text{Prob}\{|S(\mathcal{Z}, r) - ES(\mathcal{Z}, r)| \geq \epsilon\} \leq 2 \exp\left(-\frac{n\epsilon^2}{32(M + \kappa \text{Diam}(X)r)^4}\right).$$

Proof Let (\mathbf{x}', y') be a sample i.i.d. drawn from the distribution $\rho(\mathbf{x}, y)$. Denote by \mathcal{Z}'_i the sample which coincides with \mathcal{Z} except that the i -th entry (\mathbf{x}_i, y_i) is replaced by (\mathbf{x}', y') . It is easy to verify that

$$\begin{aligned} S(\mathcal{Z}, r) - S(\mathcal{Z}'_i, r) &= \sup_{\mathbf{f} \in \mathcal{F}_r} |\mathcal{E}_{\mathcal{Z}}(\mathbf{f}) - \mathcal{E}(\mathbf{f})| - \sup_{\mathbf{f} \in \mathcal{F}_r} |\mathcal{E}_{\mathcal{Z}'_i}(\mathbf{f}) - \mathcal{E}(\mathbf{f})| \\ &\leq \sup_{\mathbf{f} \in \mathcal{F}_r} |\mathcal{E}_{\mathcal{Z}}(\mathbf{f}) - \mathcal{E}_{\mathcal{Z}'_i}(\mathbf{f})| \leq \frac{4(2n - 1)(M + \kappa \text{Diam}(X)r)^2}{n^2}. \end{aligned} \tag{24}$$

Interchange the roles of \mathcal{Z} and \mathcal{Z}'_i gives

$$|S(\mathcal{Z}, r) - S(\mathcal{Z}'_i, r)| \leq \frac{8(M + \kappa \text{Diam}(X)r)^2}{n}.$$

By McDiarmid’s inequality, we obtain the desired estimate. □

In order to bound $S(\mathcal{Z}, r)$ using Lemma 1, we need a bound of $ES(\mathcal{Z}, r)$.

Lemma 2 For every $r > 0$,

$$ES(\mathcal{Z}, r) \leq \frac{11(\kappa \text{Diam}(X)r + M)^2}{\sqrt{n}}.$$

The proof of Lemma 2 is given in Appendix B.

Now we can derive the following proposition by using inequality (23), Lemmas 1 and 2.

Proposition 2 Assume $r > 1$. There exists a constant $C_3 > 0$ such that with confidence at least $1 - \delta$,

$$\varphi(\mathcal{Z}) \leq C_3 \frac{(\kappa \text{Diam}(X)r + M)^2 \log \frac{2}{\delta}}{\sqrt{n}}.$$

Note that in order to use this Proposition, we still need a bound on $\Omega(\mathbf{f}_Z) = \lambda \sum_i \|f_Z^i\|_{\mathcal{K}}$. We first state a rough bound.

Lemma 3 For every $s > 0$ and $\lambda > 0$, $\Omega(\mathbf{f}_Z) \leq M^2$.

Proof The conclusion follows from the fact

$$\Omega(\mathbf{f}_Z) \leq \mathcal{E}_Z(\mathbf{f}_Z) + \Omega(\mathbf{f}_Z) \leq \mathcal{E}_Z(\mathbf{0}) \leq M^2. \quad \square$$

However, using this quantity the bound in Theorem 5 is at least of order $O(\frac{1}{\lambda^2 s^{2p+4-\theta}})$ which tends to ∞ as $s \rightarrow 0$ and $\lambda \rightarrow 0$. So a sharper bound is needed. It will be given in Sect. 3.1.5.

3.1.4 Approximation error estimation

We now bound the approximation error $\mathcal{A}(\lambda)$.

Proposition 3 If $\nabla f_\rho \in \mathbb{H}_{\mathcal{K}}^p$, then $\mathcal{A}(\lambda) \leq C_4(\lambda + s^{4+p})$ for some $C_4 > 0$.

Proof By the definition of $\mathcal{A}(\lambda)$ and the fact that $\nabla f_\rho \in \mathbb{H}_{\mathcal{K}}^p$,

$$\mathcal{A}(\lambda) \leq \mathcal{E}(\nabla f_\rho) - 2\sigma_s^2 + \Omega(\nabla f_\rho).$$

Since

$$\begin{aligned} \mathcal{E}(\nabla f_\rho) - 2\sigma_s^2 &= \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u})(f_\rho(\mathbf{x}) - f_\rho(\mathbf{u}) + \nabla f_\rho(\mathbf{x})(\mathbf{u} - \mathbf{x}))^2 d\rho_X(\mathbf{x})d\rho_X(\mathbf{u}) \\ &\leq (C_{\mathcal{K}})^2 c_\rho \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u})|\mathbf{u} - \mathbf{x}|^4 d\mathbf{u}d\rho_X(\mathbf{x}) \\ &\leq (C_{\mathcal{K}})^2 c_\rho M_4 s^{4+p}. \end{aligned}$$

Taking $C_4 = \max\{(C_{\mathcal{K}})^2 c_\rho M_4, \sum_{i=1}^p \|(\nabla f_\rho)^i\|_{\mathcal{K}}\}$, we get the desired result. □

3.1.5 Convergence rate

Following directly from Propositions 1, 2 and 3, we get

Theorem 6 If $\nabla f_\rho \in \mathbb{H}_{\mathcal{K}}^p$, \mathbf{f}_Z and \mathbf{f}_λ are in \mathcal{F}_r for some $r \geq 1$, then with confidence $1 - \delta$

$$\mathcal{E}(\mathbf{f}_Z) - 2\sigma_s^2 \leq C_2 \left(\frac{(M + \kappa \text{Diam}(X)r)^2 \log \frac{2}{\delta}}{\sqrt{n}} + s^{4+p} + \lambda \right),$$

where C_2 is a constant independent of r, s or λ .

In order to apply Theorem 5, we need a sharp bound on $\Omega(\mathbf{f}_Z) := \lambda \sum_i \|f_Z^i\|_{\mathcal{K}}$.

Lemma 4 Under the assumptions of Theorem 1,

$$\Omega(\mathbf{f}_\lambda) \leq C_4(\lambda + s^{4+p}),$$

where C_4 is a constant independent of λ or s .

Proof Since $\mathcal{E}(\mathbf{f}_\lambda) - 2\sigma_s^2$ is non-negative for all \mathbf{f} , we have

$$\Omega(\mathbf{f}_\lambda) \leq \mathcal{E}(\mathbf{f}_\lambda) - 2\sigma_s^2 + \lambda\Omega(\mathbf{f}_\lambda) = \mathcal{A}(\lambda).$$

This in conjunction with Proposition 3 implies the conclusion. □

Lemma 5 *Under the assumptions of Theorem 1, with confidence at least $1 - \delta$*

$$\Omega(\mathbf{f}_Z) \leq C_5 \left(\lambda + s^{4+p} + \left(1 + \frac{\kappa \text{Diam}(X)M}{\lambda} \right)^2 \frac{M^2 \log \frac{2}{\delta}}{\sqrt{n}} \right)$$

for some $C_5 > 0$ independent of s or λ .

Proof By the fact $\mathcal{E}(\mathbf{f}_Z) - 2\sigma_s^2 > 0$ and Proposition 1, we have $\Omega(\mathbf{f}_Z) \leq \frac{1}{\lambda}(\varphi(Z) + \mathcal{A}(\lambda))$. Since both \mathbf{f}_Z and \mathbf{f}_λ are in $\mathcal{F}_{\frac{M^2}{\lambda}}$, using Proposition 2, we have with probability at least $1 - \delta$,

$$\varphi(Z) \leq C_3 \left(1 + \frac{\kappa \text{Diam}(X)M}{\lambda} \right)^2 \frac{M^2 \log \frac{2}{\delta}}{\sqrt{n}}.$$

Together with Proposition 3, we obtain the desired estimate with $C_5 = \max\{C_3, C_4\}$. □

Now we will use Theorems 5 and 6 to prove Theorem 1.

Proof of Theorem 1 By Theorems 5 and 6, we have with at least probability $1 - \frac{\delta}{2}$,

$$\|\mathbf{f}_Z - \nabla f_\rho\|_{L^2_{\rho_X}}^2 \leq C_0 \left\{ r^2 s^\theta + s^{2-\theta} + \frac{C_2}{s^{p+2+\theta}} \left(\frac{(M + \kappa \text{Diam}(X)r)^2 \log \frac{4}{\delta}}{\sqrt{n}} + s^{4+p} + \lambda \right) \right\},$$

if both \mathbf{f}_Z and \mathbf{f}_λ are in \mathcal{F}_r for some $r > 1$. By Lemmas 4 and 5, we can state that both \mathbf{f}_Z and \mathbf{f}_λ are in \mathcal{F}_r with probability $1 - \frac{\delta}{2}$ if

$$r = \max \left\{ 1 + \frac{s^{4+p}}{\lambda}, \left(1 + \frac{\kappa \text{Diam}(X)M}{\lambda} \right)^2 \frac{M^2 \log \frac{4}{\delta}}{\lambda\sqrt{n}} \right\}.$$

Choose $s = \left(\frac{1}{n}\right)^{\frac{1}{2(p+2+2\theta)}}$, $\lambda = \left(\frac{1}{n}\right)^{\frac{\theta}{p+2+2\theta}}$, we obtain with confidence $1 - \delta$,

$$\|\mathbf{f}_Z - \nabla f_\rho\|_{L^2_{\rho_X}} \leq C \left(\frac{1}{n}\right)^{\frac{\theta}{4(p+2+2\theta)}}. \quad \square$$

3.2 Convergence analysis in the manifold setting

The convergence analysis in the Manifold setting can be derived in a similar way as the one in the Euclidean setting. The idea behind the proof for the convergence of the gradient consists of simultaneously controlling a sample or estimation error term and a regularization or approximation error term.

As done in the convergence analysis in the Euclidean setting, we first use the excess error, $\mathcal{E}(\mathbf{f}) - 2\sigma_s^2$, to bound the $L^2_{\rho_X}$ differences of $\nabla_X f_\rho$ and $(d\Phi)^*(\mathbf{f})$.

Recall

$$\mathcal{F}_r = \left\{ \mathbf{f} \in \mathcal{H}^p_{\mathcal{K}} : \sum_{i=1}^p \|f^i\|_{\mathcal{K}} \leq r \right\}, \quad r > 0.$$

Theorem 7 Assume ρ_X satisfies the condition (14) and (15) and $\nabla_X f_\rho \in C^2(X)$. For $\mathbf{f} \in \mathcal{F}_r$ with some $r \geq 1$, there exist a constant $C_0 > 0$ such that

$$\|(d\Phi)^*(\mathbf{f}) - \nabla_X f_\rho\|_{L^2_{\rho_X}}^2 \leq C_0 \left(r^2 s^\theta + \frac{1}{s^{d+2+\theta}} (\mathcal{E}(\mathbf{f}) - 2\sigma_s^2) \right).$$

Proof It can be directly derived from Lemma B.1 in Mukherjee et al. (2010) by using the inequality $\sum_{i=1}^n |v_i|^2 \leq (\sum_{i=1}^n |v_i|)^2$. □

3.2.1 Excess error estimation

In this subsection, we will bound $\mathcal{E}(\mathbf{f}_Z) - 2\sigma_s^2$. First, we decompose the excess error into sample error and approximation error.

Proposition 4 Let \mathbf{f}_λ be defined as (22),

$$\varphi(Z) = (\mathcal{E}(\mathbf{f}_Z) - \mathcal{E}_Z(\mathbf{f}_Z)) + (\mathcal{E}_Z(\mathbf{f}_\lambda) - \mathcal{E}(\mathbf{f}_\lambda))$$

and

$$\mathcal{A}(\lambda) = \inf_{\mathbf{f} \in \mathbb{H}_{\mathcal{K}}^p} \{ \mathcal{E}(\mathbf{f}) - 2\sigma_s^2 + \Omega(\mathbf{f}) \}.$$

Then, we have

$$\mathcal{E}(\mathbf{f}_Z) - 2\sigma_s^2 + \Omega(\mathbf{f}_Z) \leq \varphi(Z) + \mathcal{A}(\lambda).$$

Since the proof of Proposition 2 doesn't need any structure information of X , it is still true in the manifold setting. Thus we have the same sample error bound as the one in the Euclidean setting. What left is to give an estimate for the approximation error $\mathcal{A}(\lambda)$ in the manifold setting.

Proposition 5 Let X be a connected compact C^∞ submanifold of \mathbb{R}^p which is isometrically embedded and of dimension d . If $f_\rho \in C^2(X)$ and $d\Phi(\nabla_X f_\rho) \in \mathcal{H}_{\mathcal{K}}^p$, then

$$\mathcal{A}(\lambda) \leq C_6(\lambda + s^{4+d})$$

for some $C_6 > 0$.

Proof By the definition of $\mathcal{A}(\lambda)$ and the fact that $d\Phi(\nabla_X f_\rho) \in \mathcal{H}_{\mathcal{K}}^p$,

$$\mathcal{A}(\lambda) \leq \mathcal{E}(d\Phi(\nabla_X f_\rho)) - 2\sigma_s^2 + \Omega(d\Phi(\nabla_X f_\rho)).$$

Note that $f_\rho \in C^2(X)$ and $d\Phi(\nabla_X f_\rho) \in \mathcal{H}_{\mathcal{K}}^p$. By Lemma B.2 in Mukherjee et al. (2010), we have

$$\mathcal{E}(d\Phi(\nabla_X f_\rho)) - 2\sigma_s^2 \leq C_7 s^{4+d},$$

where C_7 is a constant independent of s . Taking $C_6 = \max\{C_7, \sum_{i=1}^p \|(d\Phi(\nabla_X f_\rho))^i\|_{\mathcal{K}}\}$, we get the desired result. □

Combining Propositions 4, 2 and 5, we get the estimate for the excess error.

Theorem 8 *If $d\Phi(\nabla f_\rho) \in \mathbb{H}_{\mathcal{K}}^p$, $\mathbf{f}_{\mathcal{Z}}$ and \mathbf{f}_λ are in \mathcal{F}_r for some $r \geq 1$, then with confidence $1 - \delta$,*

$$\mathcal{E}(\mathbf{f}_{\mathcal{Z}}) - 2\sigma_s^2 \leq C_8 \left(\frac{(M + \kappa \text{Diam}(X)r)^2 \log \frac{2}{\delta}}{\sqrt{n}} + s^{d+4} + \lambda \right),$$

where C_8 is a constant independent of s, λ, δ or r .

3.2.2 Convergence rate

In order to use Theorems 7 and 8, we need sharp estimations for $\sum_{i=1}^p \|(d\Phi(\nabla_X f_\rho))^i\|_{\mathcal{K}}$ and $\sum_{i=1}^p \|f_\lambda^i\|_{\mathcal{K}}$. This can be done using the same argument as the one in the Euclidean setting, we omit the proof here.

Lemma 6 *Under the assumptions of Theorem 2, with confidence at least $1 - \delta$,*

$$\Omega(\mathbf{f}_{\mathcal{Z}}) \leq C_9 \left(\lambda + s^{4+d} + \left(1 + \frac{\kappa \text{Diam}(X)M}{\lambda} \right)^2 \frac{M^2 \log \frac{2}{\delta}}{\sqrt{n}} \right)$$

and

$$\Omega(\mathbf{f}_\lambda) \leq C_9(\lambda + s^{4+d}),$$

where C_9 is a constant independent of λ or s .

Now we prove Theorem 2.

Proof of Theorem 2 By the same argument as the one in proving Theorem 1, we can derive the convergence rate using Theorems 7, 8 and Lemma 6. □

3.3 Proof of Theorems 3 and 4

In order to prove Theorem 3, we need to characterize the solution of (10).

Proposition 6 *A vector function $\mathbf{f} \in \mathbb{H}_{\mathcal{Z}}^p$ with sparsity pattern $\mathbf{J} = \mathbf{J}(\mathbf{f}) = \{j : f^j \neq 0\}$ is optimal for problem (10) if and only if*

$$\frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + \mathbf{f}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i)) (\mathbf{x}_j^k - \mathbf{x}_i^k) \mathcal{K}_{\mathbf{x}_i} + \lambda \left(\sum_{j=1}^p \|f^j\|_{\mathcal{K}} \right) \frac{f^k}{\|f^k\|_{\mathcal{K}}} = 0, \tag{25}$$

$k \in \mathbf{J},$

$$\left\| \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + \mathbf{f}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i)) (\mathbf{x}_j^k - \mathbf{x}_i^k) \mathcal{K}_{\mathbf{x}_i} \right\|_{\mathcal{K}} \leq \lambda \sum_{j=1}^p \|f^j\|_{\mathcal{K}}, \tag{26}$$

$k \in \mathbf{J}^c.$

The proof of the Proposition can be derived as the same way as Proposition 10 in Bach (2008), we omit the details here.

It is easy to see that the problem (10) has a unique solution. If we can construct a solution $\tilde{\mathbf{f}}_{\mathcal{Z}}$ satisfies the two conditions in Proposition 6 with high probability, then Theorem 3 holds.

Let $\mathbf{J} = \mathbf{J}(\nabla f_{\rho}) = \{j : \frac{\partial f_{\rho}}{\partial x^j} \neq 0\}$ and $\tilde{\mathbf{f}}_{\mathcal{Z},\mathbf{J}}$ be any minimizer of

$$\min_{\mathbf{f}_{\mathbf{J}} \in \mathbb{H}_{\mathcal{K}}^{|\mathbf{J}|}} \frac{1}{n(n-1)} \sum_{i=1}^n (y_i - y_j + \mathbf{f}_{\mathbf{J}}(\mathbf{x}_i) \cdot (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}}))^2 + \lambda \left(\sum_{j \in \mathbf{J}} \|f^j\|_{\mathcal{K}} \right)^2. \tag{27}$$

We extend it by zeros on \mathbf{J}^c and denote it by $\tilde{\mathbf{f}}_{\mathcal{Z}}$. Then $\tilde{\mathbf{f}}_{\mathcal{Z}}$ satisfies (25) by the first optimality condition for $\tilde{\mathbf{f}}_{\mathcal{Z},\mathbf{J}}$. So we only need to prove that $\tilde{\mathbf{f}}_{\mathcal{Z}}$ satisfies (26) with probability tending to 1. For this, we construct the following events.

$$\begin{aligned} \Omega_0 &= \left\{ \|\tilde{\mathbf{f}}_{\mathcal{Z},\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \leq \frac{1}{2} \min_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}} \right\}, \\ \Omega_1 &= \left\{ \|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D}_n)^{-1} \mathbf{D}_n \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \leq \frac{1}{6} \|\nabla f_{\rho}\|_{\mathcal{K}} \right\}, \\ \Omega_2 &= \left\{ \frac{1}{\lambda} \|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \tilde{Y}_{\mathbf{J}}\|_{\mathcal{K}} \leq \frac{1}{6} \|\nabla f_{\rho}\|_{\mathcal{K}} \right\}, \\ \Omega_3 &= \left\{ \frac{1}{\lambda} \|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D}_n)^{-1} (Y_{\mathbf{J}} - \mathbf{S}_{\mathcal{Z},\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho})\|_{\mathcal{K}} \leq \frac{1}{6} \|\nabla f_{\rho}\|_{\mathcal{K}} \right\}, \\ \Omega_4 &= \left\{ \|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}\|_{\mathcal{K}} < (\kappa^2 C_{\rho} M_{|\mathbf{J}|,1} M_{|\mathbf{J}^c|,1} + 1) s^{p+2+\theta} \right\}, \end{aligned}$$

where

$$\mathbf{S}_{\mathcal{Z},\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) = \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s(\mathbf{f}_{\mathbf{J}}(\mathbf{x}_i) \cdot (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}})) (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}}) \mathcal{K}_{\mathbf{x}_i}, \quad \forall \mathbf{f}_{\mathbf{J}} \in \mathbb{H}_{\mathcal{K}}^{|\mathbf{J}|}, \tag{28}$$

$$\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) = \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s(\mathbf{f}_{\mathbf{J}}(\mathbf{x}_i) \cdot (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}})) (\mathbf{x}_{j,\mathbf{J}^c} - \mathbf{x}_{i,\mathbf{J}^c}) \mathcal{K}_{\mathbf{x}_i}, \quad \forall \mathbf{f}_{\mathbf{J}} \in \mathbb{H}_{\mathcal{K}}^{|\mathbf{J}|}, \tag{29}$$

$$Y_{\mathbf{J}} = \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s(y_i - y_j) (\mathbf{x}_{i,\mathbf{J}} - \mathbf{x}_{j,\mathbf{J}}) \mathcal{K}_{\mathbf{x}_i}, \tag{30}$$

$$\tilde{Y}_{\mathbf{J}} = \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s(y_i - y_j) (\mathbf{x}_{i,\mathbf{J}^c} - \mathbf{x}_{j,\mathbf{J}^c}) \mathcal{K}_{\mathbf{x}_i}, \tag{31}$$

and $\mathbf{D}_n = \sum_{i=1}^p \|\tilde{\mathbf{f}}_{\mathcal{Z}}\|_{\mathcal{K}} \text{diag}(1/\|\tilde{\mathbf{f}}_{\mathcal{Z}}\|_{\mathcal{K}})_{j \in \mathbf{J}}$. The main idea of the proof of Theorem 3 is how to bound the probability of those five sets since the probability of (26) holds can be lower-bounded by the probability of events $\Omega_0, \Omega_1, \Omega_2, \Omega_3$ and Ω_4 .

Note that, on event Ω_0 ,

$$\|\partial f_{\rho} / \partial x^j\|_{\mathcal{K}} - \frac{1}{2} \min_{j \in \mathbf{J}} \|\partial f_{\rho} / \partial x^j\|_{\mathcal{K}} \leq \|\tilde{\mathbf{f}}_{\mathcal{Z}}\|_{\mathcal{K}} \leq \frac{1}{2} \min_{j \in \mathbf{J}} \|\partial f_{\rho} / \partial x^j\|_{\mathcal{K}} + \|\partial f_{\rho} / \partial x^j\|_{\mathcal{K}}, \quad \forall j \in \mathbf{J}$$

holds. Then on the event Ω_0 , \mathbf{D}_n is well-defined and satisfies $\frac{1}{3} D_{\min} \mathbf{I} \leq \mathbf{D}_n \leq 3 D_{\max} \mathbf{I}$, where

$$D_{\min} = \min_{j' \in \mathbf{J}} \left(\sum_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}} \right) / \|\partial f_{\rho} / \partial x^{j'}\|_{\mathcal{K}}$$

and

$$D_{\max} = \max_{j' \in \mathbf{J}} \left(\sum_{j \in \mathbf{J}} \left\| \frac{\partial f_\rho}{\partial x^j} \right\|_{\mathcal{K}} \right) / \|\partial f_\rho / \partial x^{j'}\|_{\mathcal{K}}.$$

In addition,

$$\frac{1}{2} \|\nabla f_\rho\|_{\mathcal{K}} \leq \sum_{j=1}^p \|\tilde{\mathbf{f}}_{\mathcal{Z}}^j\|_{\mathcal{K}} \leq \frac{3\sqrt{|\mathbf{J}|}}{2} \|\nabla f_\rho\|_{\mathcal{K}}.$$

Using those notations, we only need to prove that $\tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}}$ satisfies

$$\mathbf{S}_{\mathcal{Z}, \mathbf{J}} \tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}} - \mathbf{Y}_{\mathbf{J}} + \lambda \mathbf{D}_n \tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}} = 0, \tag{32}$$

$$\|\tilde{\mathbf{S}}_{\mathcal{Z}, \mathbf{J}} \tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}} - \tilde{\mathbf{Y}}_{\mathbf{J}}\|_{\mathcal{K}} \leq \lambda \sum_{j=1}^p \|\mathbf{f}_{\mathcal{Z}}^j\|_{\mathcal{K}}. \tag{33}$$

according to Proposition 6 and the uniqueness of the solution. By the first optimality condition for $\tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}}$, the condition (32) holds. We can now put that back into $\tilde{\mathbf{S}}_{\mathcal{Z}, \mathbf{J}} \tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}} - \tilde{\mathbf{Y}}_{\mathbf{J}}$ and get

$$\begin{aligned} \tilde{\mathbf{S}}_{\mathcal{Z}, \mathbf{J}} \tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}} - \tilde{\mathbf{Y}}_{\mathbf{J}} &= \tilde{\mathbf{S}}_{\mathcal{Z}, \mathbf{J}} (\mathbf{S}_{\mathcal{Z}, \mathbf{J}} + \lambda \mathbf{D}_n)^{-1} \mathbf{Y}_{\mathbf{J}} - \tilde{\mathbf{Y}}_{\mathbf{J}} \\ &= -\lambda \tilde{\mathbf{S}}_{\mathcal{Z}, \mathbf{J}} (\mathbf{S}_{\mathcal{Z}, \mathbf{J}} + \lambda \mathbf{D}_n)^{-1} \mathbf{D}_n \nabla_{\mathbf{J}} f_\rho + \tilde{\mathbf{S}}_{\mathcal{Z}, \mathbf{J}} \nabla_{\mathbf{J}} f_\rho - \tilde{\mathbf{Y}}_{\mathbf{J}} \\ &\quad + \tilde{\mathbf{S}}_{\mathcal{Z}, \mathbf{J}} (\mathbf{S}_{\mathcal{Z}, \mathbf{J}} + \lambda \mathbf{D}_n)^{-1} (\mathbf{Y}_{\mathbf{J}} - \mathbf{S}_{\mathcal{Z}, \mathbf{J}} \nabla_{\mathbf{J}} f_\rho). \end{aligned}$$

Together with the fact that on the event Ω_0 , $\frac{1}{2} \|\nabla f_\rho\|_{\mathcal{K}} \leq \sum_{j=1}^p \|\tilde{\mathbf{f}}_{\mathcal{Z}}^j\|_{\mathcal{K}} \leq \frac{3\sqrt{|\mathbf{J}|}}{2} \|\nabla f_\rho\|_{\mathcal{K}}$,

$$\Omega_1 \cap \Omega_2 \cap \Omega_3 | \mathcal{Z} \in \Omega_0 \subseteq \left\{ \|\tilde{\mathbf{S}}_{\mathcal{Z}, \mathbf{J}} \tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}} - \tilde{\mathbf{Y}}_{\mathbf{J}}\|_{\mathcal{K}} \leq \lambda \sum_{j=1}^p \|\mathbf{f}_{\mathcal{Z}}^j\|_{\mathcal{K}} | \mathcal{Z} \in \Omega_0 \right\}$$

holds. Therefore,

$$\begin{aligned} &\text{Prob} \left\{ \|\tilde{\mathbf{S}}_{\mathcal{Z}, \mathbf{J}} \tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}} - \tilde{\mathbf{Y}}_{\mathbf{J}}\|_{\mathcal{K}} > \lambda \sum_{j=1}^p \|\mathbf{f}_{\mathcal{Z}}^j\|_{\mathcal{K}} \right\} \\ &= \text{Prob} \left\{ \|\tilde{\mathbf{S}}_{\mathcal{Z}, \mathbf{J}} \tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}} - \tilde{\mathbf{Y}}_{\mathbf{J}}\|_{\mathcal{K}} > \lambda \sum_{j=1}^p \|\mathbf{f}_{\mathcal{Z}}^j\|_{\mathcal{K}} | \mathcal{Z} \in \Omega_0 \right\} \text{Prob}\{\Omega_0\} \\ &\leq \text{Prob}\{\Omega_1^c | \Omega_0 \cap \Omega_4\} \text{Prob}(\Omega_0 \cap \Omega_4) + \text{Prob}\{\Omega_2^c\} \text{Prob}\{\Omega_0\} \\ &\quad + \text{Prob}\{\Omega_3^c | \Omega_0 \cap \Omega_4\} \text{Prob}(\Omega_0 \cap \Omega_4). \end{aligned} \tag{34}$$

The following Propositions provide the estimates for probability of those events.

Proposition 7 *Let $\mathbf{J} = \mathbf{J}(\nabla f_\rho)$ and $\tilde{\mathbf{f}}_{\mathcal{Z}, \mathbf{J}}$ be any minimizer of (27). Assume Assumption 1 holds. Let $\lambda = \tilde{C}_{M, \theta} s^{p+2+\theta}$, where*

$$\begin{aligned} \tilde{C}_{M, \theta} &= \max \left\{ 2\kappa C'_\rho C_\rho M_{|\mathbf{J}|, 3} M_{|\mathbf{J}^c|, 2} + 2, 12\kappa C'_\rho C_\rho M_{|\mathbf{J}|, 2} M_{|\mathbf{J}^c|, 3} / \|\nabla f_\rho\|_{\mathcal{K}}, \right. \\ &\quad \left. \sqrt{\frac{18(\kappa C'_\rho C_\rho M_{|\mathbf{J}|, 3} M_{|\mathbf{J}^c|, 2} + 1)(\kappa^2 C_\rho M_{|\mathbf{J}|, 1} M_{|\mathbf{J}^c|, 1} + 1)}{D_{\min} \|\nabla f_\rho\|_{\mathcal{K}}}} \right\}. \end{aligned} \tag{35}$$

Then there exists a constant $0 < s_0 \leq 1$ such that, for all s, n satisfying $0 < s < s_0$ and $n^r s^{(2p+4+\theta)(1/2-r)+1-\theta} \geq C_{D,\theta} := 4D_{\max}(D_{\min})^{r-\frac{1}{2}}(\tilde{C}_{M,\theta})^{r-\frac{1}{2}}C_{\rho,r}$, we have

$$\text{Prob}(\Omega_0) \geq 1 - 4 \exp\{-C_{\mathbf{J}}ns^{p+4}\}, \tag{36}$$

where $C_{\mathbf{J}}$ is a constant independent of n or s .

Proposition 8 Let $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1}^n$ be i.i.d. draws from a probability distribution ρ on Z . Under Assumption 1,

$$\text{Prob}\{\Omega_4\} \geq 1 - \mathcal{N}\left(J_{\mathcal{K}}(B_1), \frac{s^{p+2+\theta}}{8(\kappa \text{Diam}(X))^2}\right) \exp\{-C_{\tilde{s}}ns^{p+2+\theta}\}, \tag{37}$$

where $C_{\tilde{s}}$ is a constant independent of n or s .

Proposition 9 Let $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1}^n$ be i.i.d. draws from a probability distribution ρ on Z . Choose $\lambda = \tilde{C}_{M,\theta}s^{p+2+\theta}$. Under Assumption 1, if

$$n^r s^{(2p+4+\theta)(\frac{1}{2}-r)+p+2+\theta} \geq \tilde{C}_{D,\theta} := 36D_{\max}(\kappa^2 C_{\rho} M_{|\mathbf{J}|,1} M_{|\mathbf{J}^c|,1} + 1)C_{\rho,r} \left(\frac{1}{3}D_{\min}\tilde{C}_{M,\theta}\right)^{r-\frac{3}{2}} / \|\nabla f_{\rho}\|_{\infty}$$

then there exists a constant C_{Ω_1} such that

$$\text{Prob}(\Omega_1^c | \mathcal{Z} \in \Omega_0 \cap \Omega_4) \leq 2 \exp\{-C_{\Omega_1}ns^{p+2+\theta}\}.$$

Proposition 10 Let $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1}^n$ be i.i.d. draws from a probability distribution ρ on Z . Choose $\lambda = \tilde{C}_{M,\theta}s^{p+2+\theta}$ with $\tilde{C}_{M,\theta}$ defined in (35). Then under Assumption 1, there exists a constant $C_{\Omega_2} > 0$ such that

$$\text{Prob}(\Omega_2) \geq 1 - 2 \exp\{-C_{\Omega_2}ns^{p+2+\theta}\}.$$

Proposition 11 Let $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1}^n$ be i.i.d. draws from a probability distribution ρ on Z . Choose $\lambda = \tilde{C}_{M,\theta}s^{p+2+\theta}$ with $\tilde{C}_{M,\theta}$ defined in (35). Then under Assumption 1, there exists a constant $C_{\Omega_3} > 0$ such that

$$\text{Prob}(\Omega_3^c | \mathcal{Z} \in \Omega_0 \cap \Omega_4) \leq 4 \exp\{-C_{\Omega_3}ns^{p+2+\theta}\}.$$

Proof of Theorem 3 The result of Theorem 3 follows directly from inequality (34), Propositions 7, 8, 9, 10 and 11. □

Proof of Theorem 4 For any $\epsilon > 0$, we have

$$\begin{aligned} &\text{Prob}\{\|\mathbf{f}_{\mathcal{Z},\hat{\mathbf{J}}} - \nabla_{\hat{\mathbf{J}}}f_{\rho}\|_{L^2(\rho_X)} \geq \epsilon\} \\ &= \text{Prob}\{\|\mathbf{f}_{\mathcal{Z},\hat{\mathbf{J}}} - \nabla_{\hat{\mathbf{J}}}f_{\rho}\|_{L^2(\rho_X)} \geq \epsilon | \hat{\mathbf{J}} = \mathbf{J}\} \text{Prob}(\hat{\mathbf{J}} = \mathbf{J}) \\ &\quad + \text{Prob}\{\|\mathbf{f}_{\mathcal{Z},\hat{\mathbf{J}}} - \nabla_{\hat{\mathbf{J}}}f_{\rho}\|_{L^2(\rho_X)} \geq \epsilon | \hat{\mathbf{J}} \neq \mathbf{J}\} \text{Prob}(\hat{\mathbf{J}} \neq \mathbf{J}) \\ &\leq \text{Prob}\{\|\mathbf{f}_{\mathcal{Z},\mathbf{J}} - \nabla_{\mathbf{J}}f_{\rho}\|_{L^2(\rho_X)} \geq \epsilon\} + \text{Prob}(\hat{\mathbf{J}} \neq \mathbf{J}). \end{aligned}$$

Using Proposition 9 in Mukherjee and Zhou (2006), we have

$$\text{Prob}\{\|\mathbf{f}_{\mathcal{Z},\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{L^2(\rho_X)} \geq \epsilon\} \leq \exp(-C_{\rho,K} n^{\frac{\theta}{2(|\mathbf{J}|+2+3\theta)}} \epsilon),$$

where $C_{\rho,K}$ is a constant independent of n or s . Theorem 3 together with the assumption $\ln \mathcal{N}(J_{\mathcal{K}}(B_1), \frac{\epsilon}{8(\kappa \text{Diam}(X))^2}) \leq C_{\alpha} (\frac{1}{\epsilon})^{\alpha}$ implies

$$\text{Prob}(\widehat{\mathbf{J}} \neq \mathbf{J}) \leq C_1 \exp\left\{\left(\frac{8(\kappa \text{Diam}(X))^2}{s^{p+2+\theta}}\right)^{\alpha} - C_2 n s^{p+4}\right\}.$$

Choosing $s = (\frac{1}{n})^{\frac{1}{2((p+2+\theta)\alpha+p+4)}}$, the desired result follows. □

4 Algorithm for solving sparse gradient learning

In this section, we describe how to solve the optimization problem in Eq. (8). Our overall strategy is to first transfer the convex functional from the infinite dimensional to a finite dimensional space by using the reproducing property of RHKS, and then develop a forward-backward splitting algorithm to solve the reduced finite dimensional problem.

4.1 From infinite dimensional to finite dimensional optimization

Let $\mathcal{K} : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ be continuous, symmetric and positive semidefinite, i.e., for any finite set of distinct points $\{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathbb{R}^p$, the matrix $[\mathcal{K}(\mathbf{x}_i, \mathbf{x}_j)]_{i,j=1}^n$ is positive semidefinite (Aronszajn 1950). Such a function is called a Mercer kernel. The RKHS $\mathbb{H}_{\mathcal{K}}$ associated with the Mercer kernel \mathcal{K} is defined to be the completion of the linear span of the set of functions $\{\mathcal{K}_{\mathbf{x}} := \mathcal{K}(\mathbf{x}, \cdot) : \mathbf{x} \in \mathbb{R}^n\}$ with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{K}}$ satisfying $\langle \mathcal{K}_{\mathbf{x}}, \mathcal{K}_{\mathbf{u}} \rangle_{\mathcal{K}} = \mathcal{K}(\mathbf{x}, \mathbf{u})$. The reproducing property of $\mathbb{H}_{\mathcal{K}}$ states that

$$\langle \mathcal{K}_{\mathbf{x}}, h \rangle_{\mathcal{K}} = h(\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{R}^p, h \in \mathbb{H}_{\mathcal{K}}. \tag{38}$$

By the reproducing property (38), we have the following representer theorem, which states that the solution of (8) exists and lies in the finite dimensional space spanned by $\{\mathcal{K}_{\mathbf{x}_i}\}_{i=1}^n$. Hence the sparse gradient learning in Eq. (8) can be converted into a finite dimensional optimization problem. The proof of the theorem is standard and follows the same line as done in Schölkopf and Smola (2002), Mukherjee and Zhou (2006).

Theorem 9 *Given a data set \mathcal{Z} , the solution of Eq. (8) exists and takes the following form*

$$f_{\mathcal{Z}}^j(\mathbf{x}) = \sum_{i=1}^n c_{i,\mathcal{Z}}^j \mathcal{K}(\mathbf{x}, \mathbf{x}_i), \tag{39}$$

where $c_{i,\mathcal{Z}}^j \in \mathbb{R}$ for $j = 1, \dots, p$ and $i = 1, \dots, n$.

Proof The existence follows from the convexity of functionals $\mathcal{E}_{\mathcal{Z}}(\mathbf{f})$ and $\Omega(\mathbf{f})$. Suppose $\mathbf{f}_{\mathcal{Z}}$ is a minimizer. We can write functions $\mathbf{f}_{\mathcal{Z}} \in \mathcal{H}_{\mathcal{K}}^p$ as

$$\mathbf{f}_{\mathcal{Z}} = \mathbf{f}_{\parallel} + \mathbf{f}_{\perp},$$

where each element of \mathbf{f}_{\parallel} is in the span of $\{K_{\mathbf{x}_1}, \dots, K_{\mathbf{x}_n}\}$ and \mathbf{f}_{\perp} are functions in the orthogonal complement. The reproducing property yields $\mathbf{f}(\mathbf{x}_i) = \mathbf{f}_{\parallel}(\mathbf{x}_i)$ for all \mathbf{x}_i . So the functions \mathbf{f}_{\perp} do not have an effect on $\mathcal{E}_{\mathcal{Z}}(\mathbf{f})$. But $\|\mathbf{f}_{\mathcal{Z}}\|_K = \|\mathbf{f}_{\parallel} + \mathbf{f}_{\perp}\|_K > \|\mathbf{f}_{\parallel}\|_K$ unless $\mathbf{f}_{\perp} = 0$. This implies that $\mathbf{f}_{\mathcal{Z}} = \mathbf{f}_{\parallel}$, which leads to the representation of $\mathbf{f}_{\mathcal{Z}}$ in Eq. (39). \square

Using Theorem 9, we can transfer the infinite dimensional minimization problem (8) to an finite dimensional one. Define the matrix $C_{\mathcal{Z}} := [c_{i,\mathcal{Z}}^j]_{j=1,i=1}^{p,n} \in \mathbb{R}^{p \times n}$. Therefore, the optimization problem in (8) has only $p \times n$ degrees of freedom, and is actually an optimization problem in terms of a coefficient matrix $C := [c_i^j]_{j=1,i=1}^{p,n} \in \mathbb{R}^{p \times n}$. Write C into column vectors as $C := (\mathbf{c}_1, \dots, \mathbf{c}_n)$ with $\mathbf{c}_i \in \mathbb{R}^p$ for $i = 1, \dots, n$, and into row vectors as $C := (\mathbf{c}^1, \dots, \mathbf{c}^p)^T$ with $\mathbf{c}^j \in \mathbb{R}^n$ for $j = 1, \dots, p$. Let the kernel matrix be $K := [\mathcal{K}(\mathbf{x}_i, \mathbf{x}_j)]_{i=1,j=1}^{n,n} \in \mathbb{R}^{n \times n}$. After expanding each component f^j of \mathbf{f} in (8) as $f^j(\mathbf{x}) = \sum_{i=1}^n c_i^j \mathcal{K}(\mathbf{x}, \mathbf{x}_i)$, the objective function in Eq. (8) becomes a function of C as

$$\begin{aligned} \Phi(C) &= \mathcal{E}_{\mathcal{Z}}(\mathbf{f}) + \Omega(\mathbf{f}) \\ &= \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s \left(y_i - y_j + \sum_{k=1}^p \sum_{\ell=1}^n c_{\ell}^k \mathcal{K}(\mathbf{x}_i, \mathbf{x}_{\ell})(x_j^k - x_i^k) \right)^2 + \lambda \sum_{j=1}^p \sqrt{\sum_{i,k=1}^n c_i^j \mathcal{K}(\mathbf{x}_i, \mathbf{x}_k) c_k^j} \\ &= \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s \left(y_i - y_j + \sum_{\ell=1}^n \mathcal{K}(\mathbf{x}_{\ell}, \mathbf{x}_i)(\mathbf{x}_j - \mathbf{x}_i)^T \mathbf{c}_{\ell} \right)^2 + \lambda \sum_{j=1}^p \sqrt{(\mathbf{c}^j)^T K \mathbf{c}^j} \\ &= \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + (\mathbf{x}_j - \mathbf{x}_i)^T C \mathbf{k}_i)^2 + \lambda \sum_{j=1}^p \sqrt{(\mathbf{c}^j)^T K \mathbf{c}^j}, \end{aligned} \tag{40}$$

where $\mathbf{k}_i \in \mathbb{R}^n$ is the i -th column of K , i.e., $K = (\mathbf{k}_1, \dots, \mathbf{k}_n)$. Then, by Theorem 9,

$$C_{\mathcal{Z}} = \arg \min_{C \in \mathbb{R}^{p \times n}} \Phi(C). \tag{41}$$

4.2 Change of optimization variables

The objective function $\Phi(C)$ in the reduced finite dimensional problem convex is a non-smooth function. As such, most of the standard convex optimization techniques, such as gradient descent, Newton’s method, etc, cannot be directly applied. We will instead develop a forward-backward splitting algorithm to solve the problem. For this purpose, we first convert the problem into a simpler form by changing the optimization variables.

Note that K is symmetric and positive semidefinite, so its square root $K^{1/2}$ is also symmetric and positive semidefinite, and can be easily calculated. Denote the i -th column of $K^{1/2}$ by $\mathbf{k}_i^{1/2}$, i.e., $K^{1/2} = (\mathbf{k}_1^{1/2}, \dots, \mathbf{k}_n^{1/2})$. Let $\tilde{C} = CK^{1/2}$ and write $\tilde{C} = (\tilde{\mathbf{c}}_1, \dots, \tilde{\mathbf{c}}_n) = (\tilde{\mathbf{c}}^1, \dots, \tilde{\mathbf{c}}^p)^T$, where $\tilde{\mathbf{c}}_i$ and $\tilde{\mathbf{c}}^j$ are the i -th column vector and j -th row vector respectively. Then $\Phi(C)$ in Eq. (40) can be rewritten as a function of \tilde{C}

$$\Psi(\tilde{C}) = \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{C} \mathbf{k}_i^{1/2})^2 + \lambda \sum_{j=1}^p \|\tilde{\mathbf{c}}^j\|_2, \tag{42}$$

where $\|\cdot\|_2$ is the Euclidean norm of \mathbb{R}^p . Thus finding a solution $C_{\mathcal{Z}}$ of (41) is equivalent to identifying

$$\tilde{C}_{\mathcal{Z}} = \arg \min_{\tilde{C} \in \mathbb{R}^{p \times n}} \Psi(\tilde{C}), \tag{43}$$

followed by setting $C_Z = \tilde{C}_Z K^{-1/2}$, where $K^{-1/2}$ is the (pseudo) inverse of $K^{1/2}$ when K is (not) invertible.

Note that the problem we are focusing on is of large p small n , so the computation of $K^{\frac{1}{2}}$ is trivial as it is an $n \times n$ matrix. However, if we meet with the case that n is large, we can still solve (41) by adopting other algorithms such as the one used in Micchelli et al. (2010).

Given matrix \tilde{C}_Z , the variables selected by the sparse gradient learning as defined in Eq. (20) is simply

$$S = \{j : \|\tilde{\mathbf{c}}^j\|_2 \neq 0, j = 1, \dots, n\}. \tag{44}$$

And similarly, the S-EDR directions can also be directly derived from \tilde{C}_Z by noting that the sparse gradient covariance matrix is equal to

$$E = C_Z^T K C_Z = \tilde{C}_Z^T \tilde{C}_Z. \tag{45}$$

4.3 Forward-backward splitting algorithm

Next we propose a forward-backward splitting to solve Eq. (43). The forward-backward splitting is commonly used to solve the ℓ_1 related optimization problems in machine learning (Langford et al. 2009) and image processing (Daubechies et al. 2004; Cai et al. 2008). Our algorithm is derived from the general formulation described in Combettes and Wajs (2005).

We first split the objective function Ψ into a smooth term and a non-smooth term. Let $\Psi = \Psi_1 + \Psi_2$, where

$$\Psi_1(\tilde{C}) = \lambda \sum_{i=1}^p \|\tilde{\mathbf{c}}^i\|_2 \quad \text{and} \quad \Psi_2(\tilde{C}) = \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{C} \mathbf{k}_i^{1/2})^2.$$

The forward-backward splitting algorithm works by iteratively updating \tilde{C} . Given a current estimate $\tilde{C}^{(k)}$, the next one is updated according to

$$\tilde{C}^{(k+1)} = \text{prox}_{\delta\Psi_1}(\tilde{C}^{(k)} - \delta\nabla\Psi_2(\tilde{C}^{(k)})), \tag{46}$$

where $\delta > 0$ is the step size, and $\text{prox}_{\delta\Psi_1}$ is a proximity operator defined by

$$\text{prox}_{\delta\Psi_1}(D) = \arg \min_{\tilde{C} \in \mathbb{R}^{p \times n}} \frac{1}{2} \|D - \tilde{C}\|_F^2 + \delta\Psi_1(\tilde{C}), \tag{47}$$

where $\|\cdot\|_F$ is the Frobenius norm of $\mathbb{R}^{p \times n}$.

To implement the algorithm (46), we need to know both $\nabla\Psi_2$ and $\text{prox}_{\delta\Psi_1}(\cdot)$. The term $\nabla\Psi_2$ is relatively easy to obtain,

$$\nabla\Psi_2(\tilde{C}) = \frac{2}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{C} \mathbf{k}_i^{1/2}) (\mathbf{x}_j - \mathbf{x}_i) (\mathbf{k}_i^{1/2})^T. \tag{48}$$

The proximity operator $\text{prox}_{\delta\Psi_1}$ is given in the following lemma.

Lemma 7 Let $T_{\lambda\delta}(D) = \text{prox}_{\delta\Psi_1}(D)$, where $D = (\mathbf{d}^1, \dots, \mathbf{d}^p)^T$ with \mathbf{d}^j being the j -th row vector of D . Then

$$T_{\lambda\delta}(D) = (t_{\lambda\delta}(\mathbf{d}^1), \dots, t_{\lambda\delta}(\mathbf{d}^p))^T, \tag{49}$$

where

$$t_{\lambda\delta}(\mathbf{d}^j) = \begin{cases} \mathbf{0}, & \text{if } \|\mathbf{d}^j\|_2 \leq \lambda\delta, \\ \frac{\|\mathbf{d}^j\|_2 - \lambda\delta}{\|\mathbf{d}^j\|_2} \mathbf{d}^j, & \text{if } \|\mathbf{d}^j\|_2 > \lambda\delta. \end{cases} \tag{50}$$

Proof From (47), one can easily see that the row vectors $\tilde{\mathbf{c}}^j$, $j = 1, \dots, n$, of $\tilde{\mathbf{C}}$ are independent of each others. Therefore, we have

$$t_{\lambda\delta}(\mathbf{d}^j) = \arg \min_{\mathbf{c} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{d}^j - \mathbf{c}\|_2^2 + \lambda\delta \|\mathbf{c}\|_2. \tag{51}$$

The energy function in the above minimization problem is strongly convex, hence has a unique minimizer. Therefore, by the subdifferential calculus (c.f. Hiriart-Urruty and Lemaréchal 1993), $t_{\lambda\delta}(\mathbf{d}^j)$ is the unique solution of the following equation with unknown \mathbf{c}

$$\mathbf{0} \in \mathbf{c} - \mathbf{d}^j + \lambda\delta \partial(\|\mathbf{c}\|_2), \tag{52}$$

where

$$\partial(\|\mathbf{c}\|_2) = \{ \mathbf{p} : \mathbf{p} \in \mathbb{R}^n; \|\mathbf{u}\|_2 - \|\mathbf{c}\|_2 - (\mathbf{u} - \mathbf{c})^T \mathbf{p} \geq 0, \forall \mathbf{u} \in \mathbb{R}^n \}$$

is the subdifferential of the function $\|\mathbf{c}\|_2$. If $\|\mathbf{c}\|_2 > 0$, the function $\|\mathbf{c}\|_2$ is differentiable, and its subdifferential contains only its gradient, i.e., $\partial(\|\mathbf{c}\|_2) = \{ \frac{\mathbf{c}}{\|\mathbf{c}\|_2} \}$. If $\|\mathbf{c}\|_2 = 0$, then $\partial(\|\mathbf{c}\|_2) = \{ \mathbf{p} : \mathbf{p} \in \mathbb{R}^n; \|\mathbf{u}\|_2 - \mathbf{u}^T \mathbf{p} \geq 0, \forall \mathbf{u} \in \mathbb{R}^n \}$. One can check that $\partial(\|\mathbf{c}\|_2) = \{ \mathbf{p} : \mathbf{p} \in \mathbb{R}^n; \|\mathbf{p}\|_2 \leq 1 \}$ for this case. Indeed, for any vector $\mathbf{p} \in \mathbb{R}^n$ with $\|\mathbf{p}\|_2 \leq 1$, $\|\mathbf{u}\|_2 - \mathbf{u}^T \mathbf{p} \geq 0$ by the Cauchy-Schwartz inequality. On the other hand, if there is an element \mathbf{p} of $\partial(\|\mathbf{c}\|_2)$ such that $\|\mathbf{p}\|_2 > 1$, then, by setting $\mathbf{u} = \mathbf{p}$, we get $\|\mathbf{p}\|_2 - \mathbf{p}^T \mathbf{p} = \|\mathbf{p}\|_2(1 - \|\mathbf{p}\|_2) < 0$, which contradicts the definition of $\partial(\|\mathbf{c}\|_2)$. In summary,

$$\partial(\|\mathbf{c}\|_2) = \begin{cases} \{ \frac{\mathbf{c}}{\|\mathbf{c}\|_2} \}, & \text{if } \|\mathbf{c}\|_2 > 0, \\ \{ \mathbf{p} : \mathbf{p} \in \mathbb{R}^n; \|\mathbf{p}\|_2 \leq 1 \}, & \text{if } \|\mathbf{c}\|_2 = 0. \end{cases} \tag{53}$$

With (53), we see that $t_{\lambda\delta}(\mathbf{d}^j)$ in (50) is a solution of (52) hence (49) is verified. □

Now, we obtain the following forward-backward splitting algorithm to find the optimal $\tilde{\mathbf{C}}$ in Eq. (41). After choosing a random initialization, we update $\tilde{\mathbf{C}}$ iteratively until convergence according to

$$\begin{cases} D^{(k+1)} = \tilde{\mathbf{C}}^{(k)} - \frac{2\delta}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{\mathbf{C}}^{(k)} \mathbf{k}_i^{1/2}) (\mathbf{x}_j - \mathbf{x}_i) (\mathbf{k}_i^{1/2})^T, \\ \tilde{\mathbf{C}}^{(k+1)} = T_{\lambda\delta}(D^{(k+1)}). \end{cases} \tag{54}$$

The iteration alternates between two steps: (1) an empirical error minimization step, which minimizes the empirical error $\mathcal{E}_{\mathcal{Z}}(\mathbf{f})$ along gradient descent directions; and (2) a variable selection step, implemented by the proximity operator $T_{\lambda\delta}$ defined in (49). If the norm of the j -th row of $D^{(k)}$, or correspondingly the norm $\|f^j\|_{\mathcal{K}}$ of the j -th partial derivative, is smaller than a threshold $\lambda\delta$, the j -th row of $D^{(k)}$ will be set to 0, i.e., the j -th variable is not selected. Otherwise, the j -th row of $D^{(k)}$ will be kept unchanged except to reduce its norm by the threshold $\lambda\delta$.

Since $\Psi_2(\tilde{\mathbf{C}})$ is a quadratic function of the entries of $\tilde{\mathbf{C}}$, the operator norm of its Hessian $\|\nabla^2 \Psi_2\|$ is a constant. Furthermore, since the function Ψ_2 is coercive, i.e., $\|\tilde{\mathbf{C}}\|_F \rightarrow \infty$

implies that $\Psi(\tilde{C}) \rightarrow \infty$, there exists at least one solution of (43). By applying the convergence theory for the forward-backward splitting algorithm in Combettes and Wajs (2005), we obtain the following theorem.

Theorem 10 *If $0 < \delta < \frac{2}{\|\nabla^2 \psi_2\|}$, then the iteration (54) is guaranteed to converge to a solution of Eq. (43) for any initialization $\tilde{C}^{(0)}$.*

The regularization parameter λ controls the sparsity of the optimal solution. When $\lambda = 0$, no sparsity constraint is imposed, and all variables will be selected. On the other extreme, when λ is sufficiently large, the optimal solution will be $\tilde{C} = 0$, and correspondingly none of the variables will be selected. The following theorem provides an upper bound of λ above which no variables will be selected. In practice, we choose λ to be a number between 0 and the upper bound usually through cross-validation.

Theorem 11 *Consider the sparse gradient learning in Eq. (43). Let*

$$\lambda_{\max} = \max_{1 \leq k \leq p} \frac{2}{n^2} \left\| \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j) (x_i^k - x_j^k) \mathbf{k}_i^{1/2} \right\|_2 \tag{55}$$

Then the optimal solution is $\tilde{C} = 0$ for all $\lambda \geq \lambda_{\max}$, that is, none of the variables will be selected.

Proof Obviously, if $\lambda = \infty$, the minimizer of Eq. (42) is a $p \times n$ zero matrix.

When $\lambda < \infty$, the minimizer of Eq. (42) could also be a $p \times n$ zero matrix as long as λ is large enough. Actually, from iteration (54), if we choose $C^{(0)} = 0$, then

$$D^{(1)} = -\frac{2\delta}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j) (\mathbf{x}_j - \mathbf{x}_i) (\mathbf{k}_i^{\frac{1}{2}})^T$$

and $\tilde{C}^{(1)} = T_{\lambda\delta}(D^{(1)})$.

Let

$$\lambda_{\max} = \max_{1 \leq k \leq p} \frac{2}{n^2} \left\| \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j) (\mathbf{x}_j^k - \mathbf{x}_i^k) (\mathbf{k}_i^{\frac{1}{2}})^T \right\|_2$$

Then for any $\lambda \geq \lambda_{\max}$, we have $\tilde{C}^{(1)} = \mathbf{0}_{p \times n}$ by the definition of $T_{\lambda\delta}$. By induction, $\tilde{C}^{(k)} = \mathbf{0}_{p \times n}$ and the algorithm converge to $\tilde{C}^{(\infty)} = \mathbf{0}_{p \times n}$ which is a minimizer of Eq. (42) when $0 < \delta < \frac{2}{\|\nabla^2 \psi_2\|}$. We get the desired result. \square

Remark 6 In the proof of Theorem 11, we choose $C^{(0)} = \mathbf{0}_{p \times n}$ as the initial value of iteration (54) for simplicity. Actually, our argument is true for any initial value as long as $0 < \delta < \frac{2}{\|\nabla^2 \psi_2\|}$ since the algorithm converges to the minimizer of Eq. (42) when $0 < \delta < \frac{2}{\|\nabla^2 \psi_2\|}$. Note that the convergence is independent of the choice of the initial value.

It is not the first time to combine an iterative algorithm with a thresholding step to derive solutions with sparsity (see, e.g., Daubechies et al. 2004). However, different from the previous work, the sparsity we focus here is a block sparsity, that is, the row vectors of C (corresponding to partial derivatives f^j) are zero or nonzero vector-wise. As such, the

thresholding step in (49) is performed row-vector-wise, not entry-wise as in the usual soft-thresholding operator (Donoho 1995).

4.4 Matrix size reduction

The iteration in Eq. (54) involves a weighted summation of n^2 number of $p \times n$ matrices as defined by $(\mathbf{x}_j - \mathbf{x}_i)(\mathbf{k}_i^{1/2})^T$. When the dimension of the data is large, these matrices are big, and could greatly influence the efficiency of the algorithm. However, if the number of samples is small, that is, when $n \ll p$, we can improve the efficiency of the algorithm by introducing a transformation to reduce the size of these matrices.

The main motivation is to note that the matrix

$$\mathbf{M}_x := (\mathbf{x}_1 - \mathbf{x}_n, \mathbf{x}_2 - \mathbf{x}_n, \dots, \mathbf{x}_{n-1} - \mathbf{x}_n, \mathbf{x}_n - \mathbf{x}_n) \in \mathbb{R}^{p \times n}$$

is of low rank when n is small. Suppose the rank of \mathbf{M}_x is t , which is no higher than $\min(n - 1, p)$.

We use singular value decomposition to matrix \mathbf{M}_x with economy size. That is, $\mathbf{M}_x = U \Sigma V^T$, where U is a $p \times n$ unitary matrix, V is $n \times n$ unitary matrix, and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_t, 0, \dots, 0) \in \mathbb{R}^{n \times n}$. Let $\beta = \Sigma V^T$, then

$$\mathbf{M}_x = U\beta. \tag{56}$$

Denote $\beta = (\beta_1, \dots, \beta_n)$. Then $\mathbf{x}_j - \mathbf{x}_i = U(\beta_j - \beta_i)$. Using these notations, Eq. (54) is equivalent to

$$\begin{cases} D^{(k+1)} = \tilde{C}^{(k)} - \frac{2\delta U}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{C}^{(k)} \mathbf{k}_i^{1/2}) (\beta_j - \beta_i) (\mathbf{k}_i^{1/2})^T, \\ \tilde{C}^{(k+1)} = T_{\lambda,\delta}(D^{(k+1)}). \end{cases} \tag{57}$$

Note that now the second term in the right hand side of the first equation in (57) involves the summation of n^2 number of $n \times n$ matrix rather than $p \times n$ matrices. More specifically, $\omega_{i,j}^s (y_i - y_j + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{C}^{(k)} \mathbf{k}_i^{1/2})$ is a scalar in both (39) and (42). So the first equation in (39) involves the summation of n^2 matrix $(\mathbf{x}_j - \mathbf{x}_i)(\mathbf{k}_i^{1/2})^T$ which is $p \times n$, while the ones in (42) are $(\beta_j - \beta_i)(\mathbf{k}_i^{1/2})^T$ which is $n \times n$. Furthermore, we calculate the first iteration of Eq. (57) using two steps: (1) we calculate $y_i - y_j + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{C}^{(k)} \mathbf{k}_i^{1/2}$ and store it in an $n \times n$ matrix r ; (2) we calculate the first iteration of Eq. (57) using the value $r(i, j)$. These two strategies greatly improve the efficiency of the algorithm when $p \gg n$. More specifically, we reduce the update for $D^{(k)}$ in Eq. (54) of complexity $O(n^3 p)$ into a problem of complexity $O(n^2 p + n^4)$. A detailed implementation of the algorithm is shown in Algorithm 1.

Remark 7 Each update in Eq. (54) involves the summation of n^2 terms, which could be inefficient for datasets with large number of samples. A strategy to reduce the number of computations is to use a truncated weight function, e.g.,

$$\omega_{ij}^s = \begin{cases} \exp(-\frac{2\|\mathbf{x}_i - \mathbf{x}_j\|^2}{s^2}), & \mathbf{x}_j \in \mathcal{N}_i^k, \\ 0, & \text{otherwise,} \end{cases} \tag{58}$$

where $\mathcal{N}_i^k = \{\mathbf{x}_j : \mathbf{x}_j \text{ is in the } k \text{ nearest neighborhood of } \mathbf{x}_i\}$. This can reduce the number of summations from n^2 to kn .

Algorithm 1 Forward-backward splitting algorithm to solve sparse gradient learning for regression

Input: data $\{\mathbf{x}_i, y_i\}_{i=1}^n$, kernel $\mathcal{K}(\mathbf{x}, \mathbf{y})$, weight function $\omega^s(\mathbf{x}, \mathbf{y})$, parameters δ, λ and matrix $\tilde{C}^{(0)}$.

Output: the selected variables \mathcal{S} and S-EDRs.

1. Compute $K, K^{1/2}$. Do the singular value decomposition with economy size for the matrix $\mathbf{M}_x = (\mathbf{x}_1 - \mathbf{x}_n, \dots, \mathbf{x}_n - \mathbf{x}_n)$ and get $\mathbf{M}_x = U \Sigma V^T$. Denote $\beta = (\beta_1, \dots, \beta_n) = \Sigma V^T$. Compute $G_{ij} = \omega_{i,j}^s(\beta_j - \beta_i)(\mathbf{k}_i^{1/2})^T, i = 1, \dots, n, j = 1, \dots, n$ and let $k = 0$.
2. **While** the convergence condition is not true **do**
 - (a) Compute the residual $r^{(k)} = (r_{ij}^{(k)}) \in \mathbb{R}^{n \times n}$, where $r_{ij}^{(k)} = y_i - y_j + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{C}^{(k)} \mathbf{k}_i^{1/2}$.
 - (b) Compute $g^{(k)} = \frac{2}{n^2} \sum_{i,j=1}^n r_{ij}^{(k)} G_{ij}$.
 - (c) Set $D^{(k)} = \tilde{C}^{(k)} - \delta U g^{(k)}$. For the row vectors $(\mathbf{d}^i)^{(k)}, i = 1, \dots, p$, of $D^{(k)}$, perform the variable selection procedure according to (50) to get row vectors $(\tilde{\mathbf{c}}^i)^{(k+1)}$ of $\tilde{C}^{(k+1)}$.
 - i. If $\|(\mathbf{d}^i)^{(k)}\|_2 \leq \lambda\delta$, the variable is not selected, and we set $(\tilde{\mathbf{c}}^i)^{(k+1)} = 0$.
 - ii. If $\|(\mathbf{d}^i)^{(k)}\|_2 > \lambda\delta$, the variable is selected, and we set

$$(\tilde{\mathbf{c}}^i)^{(k+1)} = \frac{\|(\mathbf{d}^i)^{(k)}\|_2 - \lambda\delta}{\|(\mathbf{d}^i)^{(k)}\|_2} (\mathbf{d}^i)^{(k)}.$$
 - (d) Update $\tilde{C}^{(k+1)} = ((\tilde{\mathbf{c}}^1)^{(k+1)}, \dots, (\tilde{\mathbf{c}}^n)^{(k+1)})^T$, and set $k = k + 1$.
3. Variable selection: $\mathcal{S} = \{i : (\tilde{\mathbf{c}}^i)^{(k+1)} \neq 0\}$.
4. Feature extraction: let S-EGCM $\mathcal{E} = \tilde{C}^{(k+1)} \cdot (\tilde{C}^{(k+1)})^T$ and compute its eigenvectors via singular value decomposition of $\tilde{C}^{(k+1)}$, we get the desired S-EDRs.

5 Sparse gradient learning for classification

In this section, we extend the sparse gradient learning algorithm from regression to classification problems. We will also briefly introduce an implementation.

5.1 Defining objective function

Let \mathbf{x} and $y \in \{-1, 1\}$ be respectively \mathbb{R}^p -valued and binary random variables. The problem of classification is to estimate a classification function $f_C(\mathbf{x})$ from a set of observations $\mathcal{Z} := \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, where $\mathbf{x}_i := (x_i^1, \dots, x_i^p)^T \in \mathbb{R}^p$ is an input, and $y_i \in \{-1, 1\}$ is the corresponding output. A real valued function $f_\rho^\phi : X \mapsto \mathbb{R}$ can be used to generate a classifier $f_C(\mathbf{x}) = \text{sgn}(f_\rho^\phi(\mathbf{x}))$, where

$$\text{sgn}(f_\rho^\phi(\mathbf{x})) = \begin{cases} 1, & \text{if } f_\rho^\phi(\mathbf{x}) > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Similar to regression, we also define an objective function, including a data fitting term and a regularization term, to learn the gradient of f_ρ^ϕ . For classical binary classification, we commonly use a convex loss function $\phi(t) = \log(1 + e^{-t})$ to learn f_ρ^ϕ and define the data fitting term to be $\frac{1}{n} \sum_{i=1}^n \phi(y_i f_\rho^\phi(\mathbf{x}_i))$. The usage of loss function $\phi(t)$ is mainly motivated

by the fact that the optimal $f_\rho^\phi(\mathbf{x}) = \log[P(y = 1|\mathbf{x})/P(y = -1|\mathbf{x})]$, representing the log odds ratio between the two posterior probabilities. Note that the gradient of f_ρ^ϕ exists under very mild conditions.

As in the case of regression, we use the first order Taylor expansion to approximate the classification function f_ρ^ϕ by $f_\rho^\phi(\mathbf{x}) \approx f_\rho^\phi(\mathbf{x}_0) + \nabla f_\rho^\phi(\mathbf{x}_0) \cdot (\mathbf{x} - \mathbf{x}_0)$. When \mathbf{x}_j is close to \mathbf{x}_i , $f_\rho^\phi(\mathbf{x}_j) \approx f^0(\mathbf{x}_i) + \mathbf{f}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i)$, where $\mathbf{f} := (f^1, \dots, f^p)$ with $f^j = \partial f_\rho^\phi / \partial x^j$ for $j = 1, \dots, p$, and f^0 is a new function introduced to approximate $f_\rho^\phi(\mathbf{x}_i)$. The introduction of f^0 is unavoidable since y_i is valued -1 or 1 and not a good approximation of f_ρ^ϕ at all. After considering Taylor expansion between all pairs of samples, we define the following empirical error term for classification

$$\mathcal{E}_Z^\phi(f^0, \mathbf{f}) := \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s \phi(y_j(f^0(\mathbf{x}_i) + \mathbf{f}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i))), \tag{59}$$

where $\omega_{i,j}^s$ is the weight function as in (5).

For the regularization term, we introduce

$$\Omega(f^0, \mathbf{f}) = \lambda_1 \|f^0\|_{\mathcal{K}}^2 + \lambda_2 \sum_{i=1}^p \|f^i\|_{\mathcal{K}}. \tag{60}$$

Comparing with the regularization term for regression, we have included an extra term $\lambda_1 \|f^0\|_{\mathcal{K}}^2$ to control the smoothness of the f^0 function. We use two regularization parameters λ_1 and λ_2 for the trade-off between $\|f^0\|_{\mathcal{K}}^2$ and $\sum_{i=1}^p \|f^i\|_{\mathcal{K}}$.

Combining the data fidelity term and regularization term, we formulate the sparse gradient learning for classification as follows

$$(f_Z^\phi, \mathbf{f}_Z^\phi) = \arg \min_{(f^0, \mathbf{f}) \in \mathbb{H}_{\mathcal{K}}^{p+1}} \mathcal{E}_Z^\phi(f^0, \mathbf{f}) + \Omega(f^0, \mathbf{f}). \tag{61}$$

5.2 Forward-backward splitting for classification

Using representer theorem, the minimizer of the infinite dimensional optimization problem in Eq. (61) has the following finite dimensional representation

$$f_Z^\phi = \sum_{i=1}^n \alpha_{i,Z} \mathcal{K}(\mathbf{x}, \mathbf{x}_i), \quad (\mathbf{f}_Z^\phi)^j = \sum_{i=1}^n c_{i,Z}^j \mathcal{K}(\mathbf{x}, \mathbf{x}_i)$$

where $\alpha_{i,Z}, c_{i,Z}^j \in \mathbb{R}$ for $i = 1, \dots, n$ and $j = 1, \dots, p$.

Then using the same technique as in the regression setting, the objective functional in minimization problem (61) can be reformulated as a finite dimensional convex function of vector $\alpha = (\alpha_1, \dots, \alpha_n)^T$ and matrix $\tilde{C} = (\tilde{c}_i^j)_{i=1,j=1}^{n,p}$. That is,

$$\Psi(\alpha, \tilde{C}) = \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s \phi(y_j(\alpha^T \mathbf{k}_i + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{C} \mathbf{k}_i^{\frac{1}{2}})) + \lambda_1 \alpha^T K \alpha + \lambda_2 \sum_{j=1}^p \|\tilde{c}^j\|_2.$$

Then the corresponding finite dimensional convex

$$(\tilde{\alpha}_Z^\phi, \tilde{C}_Z^\phi) = \arg \min_{\alpha \in \mathbb{R}^n, \tilde{C} \in \mathbb{R}^{p \times n}} \Psi(\tilde{C}) \tag{62}$$

can be solved by the forward-backward splitting algorithm.

We split $\Psi(\alpha, \tilde{C}) = \Psi_1 + \Psi_2$ with

$$\Psi_1 = \lambda_2 \sum_{j=1}^p \|\tilde{\mathbf{c}}^j\|_2$$

and

$$\Psi_2 = \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s \phi(y_j(\alpha^T \mathbf{k}_i + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{\mathbf{C}} \mathbf{k}_i^{1/2})) + \lambda_1 \alpha^T K \alpha.$$

Then the forward-backward splitting algorithm for solving (62) becomes

$$\begin{cases} \alpha^{(k+1)} = \alpha^{(k)} - \delta \left(\frac{1}{n^2} \sum_{i,j=1}^n \frac{-\omega_{ij} y_j \mathbf{k}_i}{1 + \exp(y_j((\alpha^{(k)})^T \mathbf{k}_i + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{\mathbf{C}}^{(k)} \mathbf{k}_i^{1/2}))} + 2\lambda_1 K \alpha^{(k)} \right), \\ D^{(k+1)} = \tilde{\mathbf{C}}^{(k)} - \frac{\delta U}{n^2} \sum_{i,j=1}^n \frac{-\omega_{i,j}^s y_j (\beta_j - \beta_i) (\mathbf{k}_i^{1/2})^T}{1 + \exp(y_j((\alpha^{(k)})^T \mathbf{k}_i + (\mathbf{x}_j - \mathbf{x}_i)^T \tilde{\mathbf{C}}^{(k)} \mathbf{k}_i^{1/2}))}, \\ \tilde{\mathbf{C}}^{(k+1)} = T_{\lambda_2 \delta}(D^{(k+1)}), \end{cases} \tag{63}$$

where U, β satisfy equation (56) with U being a $p \times n$ unitary matrix.

With the derived $\tilde{\mathbf{C}}_{\mathcal{Z}}^{\phi}$, we can do variable selection and dimension reduction as done for the regression setting. We omit the details here.

6 Examples

Next we illustrate the effectiveness of variable selection and dimension reduction by sparse gradient learning algorithm (SGL) on both artificial datasets and a gene expression dataset. As our method is a kernel-based method, known to be effective for nonlinear problems, we focus our experiments on nonlinear settings for the artificial datasets, although the method can be equally well applied to linear problems.

Before we report the detailed results, we would like to mention that our forward-backward splitting algorithm is very efficient for solving the sparse gradient learning problem. For the simulation studies, it takes only a few minutes to obtain the results to be described next. For the gene expression data involving 7129 variables, it takes less than two minutes to learn the optimal gradient functions on an Intel Core 2 Duo desktop PC (E7500, 2.93 GHz).

6.1 Simulated data for regression

In this example, we illustrate the utility of sparse gradient learning for variable selection by comparing it to the popular variable selection method LASSO. We pointed out in Sect. 2 that LASSO, assuming the prediction function is linear, can be viewed as a special case of sparse gradient learning. Because sparse gradient learning makes no assumption on the linearity of the prediction function, we expect it to be better equipped than LASSO for selecting variables with nonlinear responses.

We simulate 100 observations from the model

$$y = (2x^1 - 1)^2 + x^2 + x^3 + x^4 + x^5 + \epsilon,$$

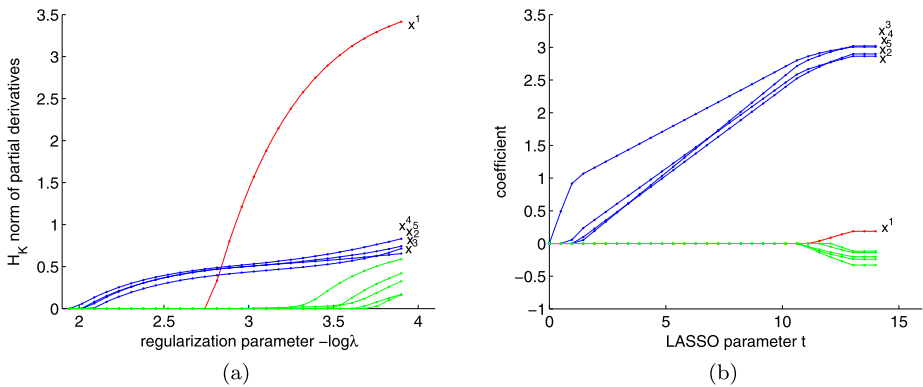


Fig. 1 (Color online) Regularization path for SGL and LASSO. Red line represents the variable x^1 , blue lines represent the variables x^2, x^3, x^4, x^5 and green lines represent noisy variables $x^6, x^7, x^8, x^9, x^{10}$. (a) H_K norm of each partial derivatives derived by SGL with respect to regularization parameter, where regularization parameter is scaled to be $-\log \lambda$ with base 10. (b) LASSO shrinkage of coefficients with respect to LASSO parameter t

Table 1 Frequencies of variables x^1, x^2, \dots, x^{10} selected by SGL and LASSO in 100 repeats

Variable	x^1	x^2	x^3	x^4	x^5	x^6	x^7	x^8	x^9	x^{10}
SGL	78	100	100	100	100	7	4	6	5	2
LASSO	16	100	100	100	100	25	14	13	13	19

where $x^i, i = 1, \dots, 5$ are i.i.d. drawn from uniform distribution on $[0, 1]$ and ϵ is drawn from standard normal distribution with variance 0.05. Let $x^i, i = 6, \dots, 10$ be additional five noisy variables, which are also i.i.d. drawn from uniform distribution on $[0, 1]$. We assume the observation dataset is given in the form of $\mathcal{Z} := \{\mathbf{x}_i, y_i\}_{i=1}^{100}$, where $\mathbf{x}_i = (x_i^1, x_i^2, \dots, x_i^{10})$ and $y_i = (2x_i^1 - 1)^2 + x_i^2 + x_i^3 + x_i^4 + x_i^5 + \epsilon$. It is easy to see that only the first 5 variables contribute the value of y .

This is a well-known example as pointed out by B. Turlach (2004) to show the deficiency of LASSO. As the ten variables are uncorrelated, LASSO will select variables based on their correlation with the response variable y . However, because $(2x^1 - 1)^2$ is a symmetric function with respect to symmetric axis $x^1 = \frac{1}{2}$ and the variable x^1 is drawn from a uniform distribution on $[0, 1]$, the correlation between x^1 and y is 0. Consequently, x^1 will not be selected by LASSO. Because SGL selects variables based on the norm of the gradient functions, it has no such a limitation.

To run the SGL algorithm in this example, we use the truncated Gaussian in Eq. (58) with 10 neighbors as our weight function. The bandwidth parameter s is chosen to be half of the median of the pairwise distances of the sampling points. As the gradients of the regression function with respect to different variables are all linear, we choose $\mathcal{K}(\mathbf{x}, \mathbf{y}) = 1 + \mathbf{xy}$.

Figure 1 shows the variables selected by SGL and LASSO for the same dataset when the regularization parameter varies. Both methods are able to successfully select the four linear variables (i.e. x^2, \dots, x^4). However, LASSO failed to select x^1 and treated x^1 as if it were one of five noisy term x^6, \dots, x^{10} (Fig. 1(b)). In contrast, SGL is clearly able to differentiate x^1 from the group of five noisy variables (Fig. 1(a)).

To summarize how often each variable will be selected, we repeat the simulation 100 times. For each simulation, we choose a regularization parameter so that each algorithm

returns exactly five variables. Table 1 shows the frequencies of variables x^1, x^2, \dots, x^{10} selected by SGL and LASSO in 100 repeats. Both methods are able to select the four linear variables, x^2, x^3, x^4, x^5 , correctly. But, LASSO fails to select x^1 and treats it as the same as the noisy variables $x^6, x^7, x^8, x^9, x^{10}$. This is in contrast to SGL, which is able to correctly select x^1 in 78% of the times, much greater than the frequencies (median 5%) of selecting the noisy variables. This example illustrates the advantage of SGL for variable selection in nonlinear settings.

6.2 Simulated data for classification

Next we apply SGL to an artificial dataset that has been commonly used to test the efficiency of dimension reduction methods in the literature. We consider a binary classification problem in which the sample data are lying in a 200 dimensional space with only the first 2 dimensions being relevant for classification and the remaining variables being noises. More specifically, we generate 40 samples with half from +1 class and the other half from -1 class. For the samples from +1 class, the first 2-dimensions of the sample data correspond to points drawn uniformly from a 2-dimensional spherical surface with radius 3. The remaining 198 dimensions are noisy variables with each variable being i.i.d. drawn from Gaussian distribution $N(0, \sigma)$. That is,

$$x^j \sim N(0, \sigma), \quad \text{for } j = 3, 4, \dots, 200. \quad (64)$$

For the samples from -1 class, the first 2-dimensions of the sample data correspond to points drawn uniformly from a 2-dimensional spherical surface with radius 3×2.5 and the remaining 198 dimensions are noisy variables with each variable x^j i.i.d. drawn from $N(0, \sigma)$ as (64). Obviously, this data set can be easily separated by a sphere surface if we project the data to the Euclidean space spanned by the first two dimensions.

In what follows, we illustrate the effectiveness of SGL on this data set for both variable selection and dimension reduction. In implementing SGL, both the weight function and the kernel are all chosen to be $\exp(-\frac{\|x-u\|^2}{2s^2})$ with s being half of the median of pairwise distance of the sampling points.

We generated several datasets with different noise levels by varying σ from 0.1 to 3. SGL correctly selected x^1 and x^2 as the important variables for all cases we tested. Furthermore, SGL also generated two S-EDRs that captured the underlying data structure for all these cases (Fig. 2). It is important to emphasize that the two S-EDRs generated by SGL are the only two features the algorithm can possibly obtain, since the derived S-EGCM are supported on a 2×2 matrix. As a result, both of the derived S-EDRs are linear combinations of the first two variables. By contrast, using the gradient learning method (GL) reported in Mukherjee et al. (2010), the first two returned dimension reduction directions (called ESFs) are shown to be able to capture the correct underlying structure only when $\sigma < 0.7$. In addition, the derived ESFs are linear combinations of all 200 original variables instead of only two variables as in S-EDRs. Figure 2(b), (e) shows the training data and the test data projected on the derived two S-EDRs for a dataset with large noise ($\sigma = 3$). Comparing to the data projected on the first two dimensions (Fig. 2(a), (d)), the derived S-EDRs preserves the structure of the original data. In contrast, the gradient learning algorithm without sparsity constraint performed much poorer (Fig. 2(c), (f)).

To explain why SGL performed better than GL without sparsity constraint, we plotted the norms of the derived empirical gradients from both methods in Fig. 3. Note that although the

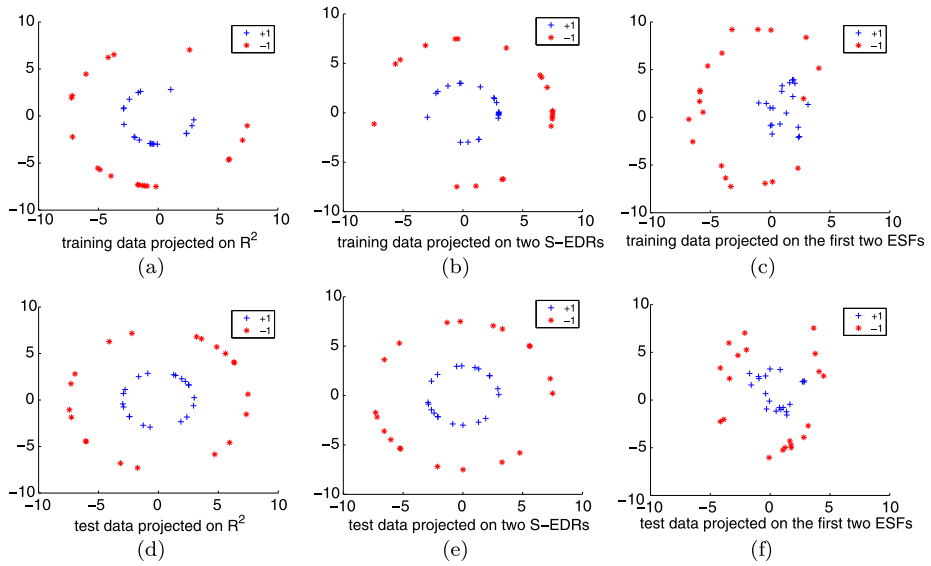


Fig. 2 (Color online) Nonlinear classification simulation with $\sigma = 3$. (a) Training data projected on the first two dimensions, (b) Training data projected on two S-EDRs derived by SGL. (c) Training data projected on first two ESFs derived by GL. (d) Test data projected on the first two dimensions. (e) Test data projected on two S-EDRs derived by SGL. (f) Test data projected on first two ESFs derived by GL

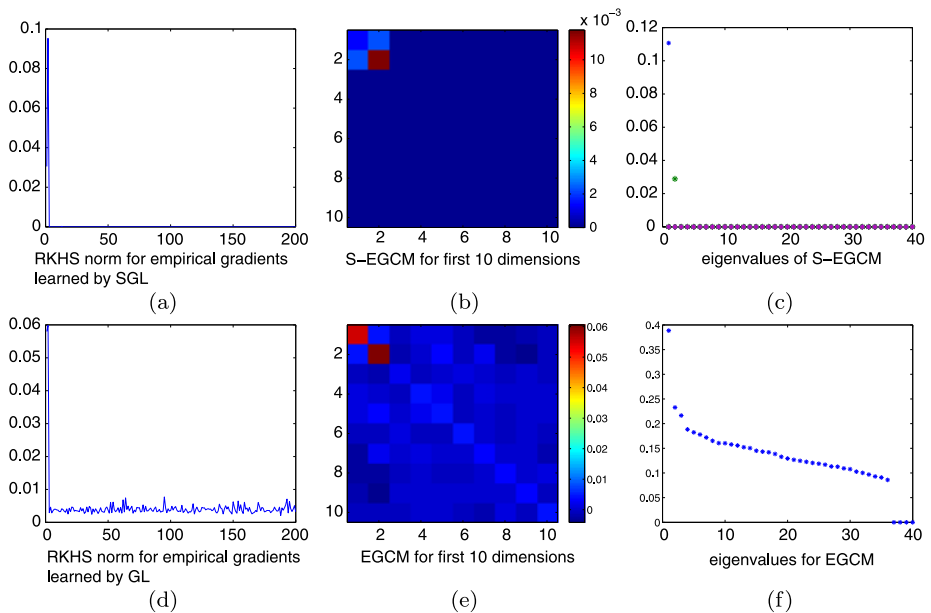


Fig. 3 (Color online) Nonlinear classification simulation with $\sigma = 3$ (continued). (a) RKHS norm of empirical gradient derived by SGL. (b) S-EGCM for first 10 dimension. (c) Eigenvalues of S-EGCM. (d) RKHS norm of empirical gradient derived by GL, (e) EGCM for first 10 dimension. (f) Eigenvalues of EGCM

Table 2 Summary of the Leukemia classification results

Method	SGL (variable selection)	SGL (S-EDRs)	GL (ESFs)	Linear SVM	LASSO
Number of variables or features	106	1	6	7129 (all)	33
Leave one out error (LOO)	0/38	0/38	0/38	3/38	1/38
Test errors	0/34	0/34	2/34	2/34	1/34

norms of partial derivatives of unimportant variables derived from the method without sparsity constraint are small, they are not exactly zero. As a result, all variables contributed and, consequently, introduced noise to the empirical gradient covariance matrix (Fig. 3(e), (f)).

We also tested LASSO for this artificial data set, and not surprisingly it failed to identify the right variables in all cases we tested. We omit the details here.

6.3 Leukemia classification

Next we apply SGL to do variable selection and dimension reduction on gene expression data. A gene expression data typically consists of the expression values of tens of thousands of mRNAs from a small number of samples as measured by microarrays. Because of the large number of genes involved, the variable selection step becomes especially important both for the purpose of generating better prediction models, and also for elucidating biological mechanisms underlying the data.

The gene expression data we will use is a widely studied dataset, consisting of the measurements of 7129 genes from 72 acute leukemia samples (Golub et al. 1999). The samples are labeled with two leukemia types according to the precursor of the tumor cells—one is called acute lymphoblastic leukemia (ALL), and the other one is called acute myelogenous leukemia (AML). The two tumor types are difficult to distinguish morphologically, and the gene expression data is used to build a classifier to classify these two types.

Among 72 samples, 38 are training data and 34 are test data. We coded the type of leukaemia as a binary response variable y , with 1 and -1 representing ALL and AML respectively. The variables in the training samples $\{\mathbf{x}_i\}_{i=1}^{38}$ are normalized to be zero mean and unit length for each gene. The test data are similarly normalized, but only using the empirical mean and variance of the training data.

We applied three methods (SGL, GL and LASSO) to the dataset to select variables and extract the dimension reduction directions. To compare the performance of the three methods, we used linear SVM to build a classifier based on the variables or features returned by each method, and evaluated the classification performance using both leave-one-out (LOO) error on the training data and the testing error. To implement SGL, the bandwidth parameter s is chosen to be half of the median of the pairwise distances of the sampling points, and $\mathcal{K}(\mathbf{x}, \mathbf{y}) = \mathbf{xy}$. The regularization parameters for the three methods are all chosen according to their prediction power measured by leave-one-out error.

Table 2 shows the results of the three methods. We implemented two SVM classifiers for SGL using either only the variables or the features returned by SGL. Both classifiers are able to achieve perfect classification for both leave-one-out and testing samples. The performance of SGL is better than both GL and LASSO, although only slightly. All three methods performed significantly better than the SVM classifier built directly from the raw data in terms of LOO error and test error.

In addition to the differences in prediction performance, we note a few other observations. First, SGL selects more genes than LASSO, which likely reflects the failure of LASSO to choose genes with nonlinear relationships with the response variable, as we illustrated in our first example. Second, The S-EDRs derived by SGL are linear combinations of 106 selected variables rather than all original variables as in the case of ESFs derived by GL. This is a desirable property since an important goal of the gene expression analysis is to identify regulatory pathways underlying the data, e.g. those distinguishing the two types of tumors. By associating only a small number of genes, S-EDRs provide better and more manageable candidate pathways for further experimental testing.

7 Discussion

Variable selection and dimension reduction are two common strategies for high-dimensional data analysis. Although many methods have been proposed before for variable selection or dimension reduction, few methods are currently available for simultaneous variable selection and dimension reduction. In this work, we described a sparse gradient learning algorithm that integrates automatic variable selection and dimension reduction into the same optimization framework. The algorithm can be viewed as a generalization of LASSO from linear to non-linear variable selection, and a generalization of the OPG method for learning EDR directions from a non-regularized to regularized estimation. We showed that the integrated framework offers several advantages over the previous methods by using both simulated and real-world examples.

The SGL method can be refined by using an adaptive weight function rather than a fixed one as in our current implementation. The weight function $\omega_{i,j}^s$ is used to measure the distance between two sample points. If the data are lying in a lower dimensional space, the distance would be more accurately captured by using only variables related to the lower dimensional space rather than all variables. One way to implement this is to calculate the distance using only selected variables. Note that the forward-backward splitting algorithm eliminates variables at each step of the iteration. We can thus use an adaptive weight function that calculates the distances based only on selected variables returned after each iteration. More specifically, let $\mathcal{S}^{(k)} = \{i : \|(\tilde{\mathbf{c}}^i)^{(k)}\|_2 \neq 0\}$ represent the variables selected after iteration k . An adaptive approach is to use $\sum_{i \in \mathcal{S}^{(k)}} (x_i^i - x_j^i)^2$ to measure the distance $\|\mathbf{x}_i - \mathbf{x}_j\|^2$ after iteration k .

An interesting area for future research is to extend SGL for semi-supervised learning. In many applications, it is often much easier to obtain unlabeled data with a larger sample size $u \gg n$. Most natural (human or animal) learning seems to occur in semi-supervised settings (Belkin et al. 2006). It is possible to extend SGL for the semi-supervised learning along several directions. One way is to use the unlabeled data $\mathcal{X} = \{\mathbf{x}_i\}_{i=n+1}^{n+u}$ to control the approximate norm of \mathbf{f} in some Sobolev spaces and introduce a semi-supervised learning algorithm as

$$\mathbf{f}_{\mathcal{Z}, \mathcal{X}, \lambda, \mu} = \arg \min_{\mathbf{f} \in \mathcal{H}_K^p} \left\{ \frac{1}{n^2} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + \mathbf{f}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i))^2 + \frac{\mu}{(n+u)^2} \sum_{i,j=1}^{n+u} W_{i,j} \|\mathbf{f}(\mathbf{x}_i) - \mathbf{f}(\mathbf{x}_j)\|_{\tilde{\ell}^2(\mathbb{R}^p)}^2 + \lambda \|\mathbf{f}\|_K \right\},$$

where $\|\mathbf{f}\|_K = \sum_{i=1}^p \|f^i\|_K$, $W_{i,j}$ are edge weights in the data adjacency graph, μ is another regularization parameter and often satisfies $\lambda = o(\mu)$. In order to make the algorithm efficient, we can use truncated weight in implementation as done in Sect. 6.1.

The regularization term $\sum_{i,j=1}^{n+u} W_{i,j} \|\mathbf{f}(\mathbf{x}_i) - \mathbf{f}(\mathbf{x}_j)\|_{\ell^2(\mathbb{R}^p)}^2$ is mainly motivated by the recent work of Belkin and Niyogi (2006). In that paper, they have introduced a regularization term $\sum_{i,j=1}^{n+u} W_{i,j} (f(\mathbf{x}_i) - f(\mathbf{x}_j))^2$ for semi-supervised regression and classification problems. The term $\sum_{i,j=1}^{n+u} W_{i,j} (f(\mathbf{x}_i) - f(\mathbf{x}_j))^2$ is well-known to be related to graph Laplacian operator. It is used to approximate $\int_{\mathbf{x} \in \mathcal{M}} \|\nabla_{\mathcal{M}} f\|^2 d\rho_X(\mathbf{x})$, where \mathcal{M} is a compact submanifold which is the support of marginal distribution $\rho_X(\mathbf{x})$, and $\nabla_{\mathcal{M}}$ is the gradient of f defined on \mathcal{M} (Do Carmo and Flaherty 1992). Intuitively, $\int_{\mathbf{x} \in \mathcal{M}} \|\nabla_{\mathcal{M}} f\|^2 d\rho_X(\mathbf{x})$ is a smoothness penalty corresponding to the probability distribution. The idea behind $\int_{\mathbf{x} \in \mathcal{M}} \|\nabla_{\mathcal{M}} f\|^2 d\rho_X(\mathbf{x})$ is that it reflects the intrinsic structure of $\rho_X(\mathbf{x})$. Our regularization term $\sum_{i,j=1}^{n+u} W_{i,j} \|\mathbf{f}(\mathbf{x}_i) - \mathbf{f}(\mathbf{x}_j)\|_{\ell^2(\mathbb{R}^p)}^2$ is a corresponding vector form of $\sum_{i,j=1}^{n+u} W_{i,j} (f(\mathbf{x}_i) - f(\mathbf{x}_j))^2$ in Belkin et al. (2006). The regularization framework of the SGL for semi-supervised learning can thus be viewed as a generalization of this previous work.

Appendix A: Proof of Theorem 5

To prove Theorem 5, we need several lemmas which require the notations of the following quantities. Denote

$$Q(\mathbf{f}) = \int_X \int_X \omega(\mathbf{x} - \mathbf{u}) ((\mathbf{f}(\mathbf{x}) - \nabla f_\rho(\mathbf{x}))(\mathbf{u} - \mathbf{x}))^2 d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}),$$

the border set

$$X_s = \{\mathbf{x} \in X : d(\mathbf{x}, \partial X) > s \text{ and } p(\mathbf{x}) \geq (1 + c_\rho)s^\theta\}$$

and the moments for $0 \leq q < \infty$

$$M_q = \int_{\mathbb{R}^p} e^{-\frac{|\mathbf{t}|^2}{2}} |\mathbf{t}|^q d\mathbf{t}, \quad \tilde{M}_q = \int_{|\mathbf{t}| \leq 1} e^{-\frac{|\mathbf{t}|^2}{2}} |\mathbf{t}|^q d\mathbf{t}.$$

Note that X_s is nonempty when s is small enough.

Lemma 8 Under assumptions of Theorem 5,

$$\frac{\tilde{M}_2 s^{p+2+\theta}}{p} \int_{X_s} |\mathbf{f}(\mathbf{x}) - \nabla f_\rho(\mathbf{x})|^2 d\rho_X(\mathbf{x}) \leq Q(\mathbf{f})$$

Proof For $\mathbf{x} \in X_s$, we have $d(\mathbf{x}, \partial X) > s$ and $p(\mathbf{x}) \geq (1 + c_\rho)s^\theta$. Thus $\{\mathbf{u} \in X : \|\mathbf{u} - \mathbf{x}\| \leq s\} \subset X$ and for $\mathbf{u} \in \{\mathbf{u} \in X : \|\mathbf{u} - \mathbf{x}\| \leq s\}$, $p(\mathbf{u}) = p(\mathbf{x}) - (p(\mathbf{x}) - p(\mathbf{u})) \geq (1 + c_\rho)s^\theta - c_\rho \|\mathbf{u} - \mathbf{x}\|^\theta \geq s^\theta$. Therefore,

$$\begin{aligned} Q(\mathbf{f}) &\geq \int_{X_s} \int_{\|\mathbf{x}-\mathbf{u}\| \leq s} \omega^s(\mathbf{x} - \mathbf{u}) ((\mathbf{f}(\mathbf{x}) - \nabla f_\rho(\mathbf{x}))(\mathbf{x} - \mathbf{u}))^2 p(\mathbf{u}) d\mathbf{u} d\rho_X(\mathbf{x}) \\ &\geq s^\theta \int_{X_s} \int_{\|\mathbf{x}-\mathbf{u}\| \leq s} \omega^s(\mathbf{x} - \mathbf{u}) ((\mathbf{f}(\mathbf{x}) - \nabla f_\rho(\mathbf{x}))(\mathbf{x} - \mathbf{u}))^2 d\mathbf{u} d\rho_X(\mathbf{x}). \end{aligned}$$

Denote the i -th entry of a vector \mathbf{x} by x^i . Then $((\mathbf{f}(\mathbf{x}) - \nabla f_\rho(\mathbf{x}))(\mathbf{x} - \mathbf{u}))^2$ equals to

$$\sum_{i=1}^p \sum_{j=1}^p \left(f^i(\mathbf{x}) - \frac{\partial f_\rho}{\partial x^i}(\mathbf{x}) \right) \left(f^j(\mathbf{x}) - \frac{\partial f_\rho}{\partial x^j}(\mathbf{x}) \right) (x^i - u^i)(x^j - u^j).$$

For the case $i \neq j$, we have

$$\int_{\|\mathbf{u}-\mathbf{x}\| \leq s} \omega^s(\mathbf{x} - \mathbf{u})(x^i - u^i)(x^j - u^j) d\mathbf{u} = s^{p+2} \int_{|t| \leq 1} e^{-\frac{\|t\|^2}{2}} t^i t^j dt = 0.$$

Therefore,

$$\begin{aligned} Q(\mathbf{f}) &\geq s^{p+2+\theta} \sum_{i=1}^p \int_{X_s} \left(f^i(\mathbf{x}) - \frac{\partial f_\rho}{\partial x^i}(\mathbf{x}) \right)^2 d\rho_X(\mathbf{x}) \int_{|t| \leq 1} e^{-\frac{\|t\|^2}{2}} (t^i)^2 dt \\ &= \frac{\tilde{M}_2 s^{p+2+\theta}}{p} \int_{X_s} \|\mathbf{f}(\mathbf{x}) - \nabla f_\rho(\mathbf{x})\|^2 d\rho_X(\mathbf{x}), \end{aligned}$$

which yields the desired estimate. □

Lemma 9 *Under the assumption of Theorem 5, we have*

$$Q(\mathbf{f}) \leq C_1 (s^{4+p} + \mathcal{E}(\mathbf{f}) - 2\sigma_s^2),$$

where C_1 is a constant independent of s or \mathbf{f} .

Proof Denote $a_1 = (\mathbf{f}(\mathbf{x}) - \nabla f_\rho(\mathbf{x}))(\mathbf{u} - \mathbf{x})$ and $a_2 = f_\rho(\mathbf{x}) - f_\rho(\mathbf{u}) + \nabla f_\rho(\mathbf{x})(\mathbf{u} - \mathbf{x})$. We have

$$Q(\mathbf{f}) = \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u})(a_1)^2 d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u})$$

and

$$\mathcal{E}(\mathbf{f}) = \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u})(a_1 + a_2)^2 d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}) + 2\sigma_s^2.$$

Note that $(a_1 + a_2)^2 \geq (a_1)^2 - 2|a_1||a_2|$. Thus

$$\mathcal{E}(\mathbf{f}) - 2\sigma_s^2 \geq Q(\mathbf{f}) - 2 \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u})|a_1||a_2| d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}).$$

By the fact $\nabla f_\rho \in \mathbb{H}_{\mathcal{K}}^p$ and Lemma 19 in Mukherjee and Wu (2006), there exists a constant $C_{\mathcal{K}} > 0$ depending on \mathcal{K} and f_ρ such that

$$|a_2| \leq C_{\mathcal{K}} \|\mathbf{x} - \mathbf{u}\|^2.$$

Together with the assumption $p(\mathbf{x}) \leq c_\rho$, we have

$$\begin{aligned} &\int_X \int_X \omega^s(\mathbf{x} - \mathbf{u})|a_1||a_2| d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}) \\ &\leq \sqrt{Q(\mathbf{f})} \left(\int_X \int_X \omega^s(\mathbf{x} - \mathbf{u})|a_2|^2 d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}) \right)^{\frac{1}{2}} \end{aligned}$$

$$\begin{aligned} &\leq C_{\mathcal{K}}\sqrt{Q(\mathbf{f})}\left(c_{\rho}\int_X\int_{\mathbb{R}^p}\omega^s(\mathbf{x}-\mathbf{u})|\mathbf{x}-\mathbf{u}|^4d\mathbf{x}d\rho_X(\mathbf{u})\right)^{\frac{1}{2}} \\ &\leq C_{\mathcal{K}}\sqrt{c_{\rho}M_4s^{2+p/2}}\sqrt{Q(\mathbf{f})}. \end{aligned} \tag{65}$$

Combining the above arguments, we obtain

$$Q(\mathbf{f}) - C_{\mathcal{K}}\sqrt{c_{\rho}M_4s^{2+p/2}}\sqrt{Q(\mathbf{f})} \leq \mathcal{E}(\mathbf{f}) - 2\sigma_s^2.$$

This implies the conclusion with $C_1 = 2 \max\{C_{\mathcal{K}}^2c_{\rho}M_4, 1\}$. □

Proof of Theorem 5 Write

$$\|\mathbf{f} - \nabla f_{\rho}\|_{L^2_{\rho_X}}^2 = \int_{X \setminus X_s} |\mathbf{f}(\mathbf{x}) - \nabla f_{\rho}(\mathbf{x})|^2 d\rho_X(\mathbf{x}) + \int_{X_s} |\mathbf{f}(\mathbf{x}) - \nabla f_{\rho}(\mathbf{x})|^2 d\rho_X(\mathbf{x}). \tag{66}$$

We have

$$\rho_X(X \setminus X_s) \leq c_{\rho}s + (1 + c_{\rho})c_{\rho}|X|s^{\theta} \leq (c_{\rho} + (1 + c_{\rho})c_{\rho}|X|)s^{\theta},$$

where $|X|$ is the Lebesgue measure of X . So the first term on the right of (66) is bounded by

$$\kappa^2(r + \|\nabla f_{\rho}\|_{\mathcal{K}})^2(c_{\rho} + (1 + c_{\rho})c_{\rho}|X|)s^{\theta}.$$

By Lemmas 8 and 9, the second term on the right hand of (66) is bounded by

$$\frac{pC_1}{\widetilde{M}_2} \frac{1}{s^{p+2+\theta}}(s^{4+p} + \mathcal{E}(\mathbf{f}) - 2\sigma_s^2).$$

Combining these two estimates finishes the proof of the claim with

$$C_0 = \kappa^2(1 + \|\nabla f_{\rho}\|_{\mathcal{K}})^2(c_{\rho} + (1 + c_{\rho})c_{\rho}|X|) + \frac{pC_1}{\widetilde{M}_2}.$$

This is the end of the proof. □

Appendix B: Proof of Lemma 2

Proof of Lemma 2 Denote $\xi(\mathbf{x}, y, \mathbf{u}, v) = \omega^s(\mathbf{x} - \mathbf{u})(y - v + \mathbf{f}(\mathbf{x})(\mathbf{u} - \mathbf{x}))$. Then $\mathcal{E}(\mathbf{f}) = E_{(\mathbf{u},v)}E_{(\mathbf{x},y)}\xi(\mathbf{x}, y, \mathbf{u}, v)$ and $\mathcal{E}_{\mathcal{Z}}(\mathbf{f}) = \frac{1}{n^2} \sum_{i,j=1}^n \xi(\mathbf{x}_i, y_i, \mathbf{x}_j, y_j)$. One can easily check that

$$\begin{aligned} S(\mathcal{Z}, r) &\leq \sup_{\mathbf{f} \in \mathcal{F}_r} \left| \mathcal{E}(\mathbf{f}) - \frac{1}{n} \sum_{j=1}^n E_{(\mathbf{x},y)}\xi(\mathbf{x}, y, \mathbf{x}_j, y_j) \right| \\ &\quad + \sup_{\mathbf{f} \in \mathcal{F}_r} \left| \frac{1}{n} \sum_{j=1}^n E_{(\mathbf{x},y)}\xi(\mathbf{x}, y, \mathbf{x}_j, y_j) - \mathcal{E}_{\mathcal{Z}}(\mathbf{f}) \right| \\ &\leq \sup_{\mathbf{f} \in \mathcal{F}_r} E_{(\mathbf{x},y)} \left| E_{(\mathbf{u},v)}\xi(\mathbf{x}, y, \mathbf{u}, v) - \frac{1}{n} \sum_{j=1}^n \xi(\mathbf{x}, y, \mathbf{x}_j, y_j) \right| \end{aligned}$$

$$\begin{aligned}
 & + \frac{1}{n} \sum_{j=1}^n \sup_{\mathbf{f} \in \mathcal{F}_r} \sup_{(\mathbf{u}, v) \in \mathcal{Z}} \left| E_{(\mathbf{x}, y)} \xi(\mathbf{x}, y, \mathbf{u}, v) - \frac{1}{n-1} \sum_{i=1, i \neq j}^n \xi(\mathbf{x}_i, y_i, \mathbf{u}, v) \right| \\
 & + \frac{1}{n^2(n-1)} \sum_{j=1}^n \sum_{i \neq j, i=1}^n \xi(\mathbf{x}_i, y_i, \mathbf{x}_j, y_j) \\
 & := S_1 + S_2 + S_3.
 \end{aligned}$$

Let $\epsilon_i, i = 1, \dots, n$ be independent Rademacher variables. For S_1 , by using the properties of Rademacher complexities (van der Vaart and Wellner 1996), we have

$$\begin{aligned}
 ES_1(\mathcal{Z}) & = E_{(\mathbf{x}, y)} \sup_{\mathbf{f} \in \mathcal{F}_r} \left| E_{(\mathbf{u}, v)} \xi(\mathbf{x}, y, \mathbf{u}, v) - \frac{1}{n} \sum_{j=1}^n \xi(\mathbf{x}, y, \mathbf{x}_j, y_j) \right| \\
 & \leq 2 \sup_{(\mathbf{x}, y) \in \mathcal{Z}} E \sup_{\mathbf{f} \in \mathcal{F}_r} \left| \frac{1}{n} \sum_{j=1}^n \epsilon_j \omega^s(\mathbf{x} - \mathbf{x}_j) (y_j - y + \mathbf{f}(\mathbf{x}_j)(\mathbf{x} - \mathbf{x}_j))^2 \right| \\
 & \leq 4(M + \kappa \text{Diam}(X)r) \left(\sup_{(\mathbf{x}, y) \in \mathcal{Z}} E \sup_{\mathbf{f} \in \mathcal{F}_r} \frac{1}{n} \sum_{j=1}^n \epsilon_j (y_j - y + \mathbf{f}(\mathbf{x}_j)(\mathbf{x} - \mathbf{x}_j)) + \frac{M}{\sqrt{n}} \right) \\
 & \leq \frac{5(\kappa \text{Diam}(X)r + M)^2}{\sqrt{n}}.
 \end{aligned}$$

Similarly, we can verify $ES_2(\mathcal{Z}) \leq \frac{5(\kappa \text{Diam}(X)r + M)^2}{\sqrt{n}}$. Obviously, $S_3 \leq \frac{(M + \kappa \text{Diam}(X)r)^2}{n}$. Combining the estimates for S_1, S_2 and S_3 , we can get the desired estimate

$$ES(\mathcal{Z}, r) \leq \frac{10(\kappa \text{Diam}(X)r + M)^2}{\sqrt{n}} + \frac{(M + \kappa \text{Diam}(X)r)^2}{n} \leq \frac{11(M + \kappa \text{Diam}(X)r)^2}{\sqrt{n}}.$$

□

Appendix C: Proof of Propositions in Sect. 3.3

Firstly, we would like to denote some operators and constants that would be used frequently in this section.

$$\mathbf{f}_{\rho, s, \mathbf{J}} = \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u}) (f_\rho(\mathbf{u}) - f_\rho(\mathbf{x})) (\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{\mathbf{J}}) \mathcal{K}_x d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}), \tag{67}$$

$$\tilde{\mathbf{f}}_{\rho, s, \mathbf{J}} = \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u}) (f_\rho(\mathbf{u}) - f_\rho(\mathbf{x})) (\mathbf{u}_{\mathbf{J}^c} - \mathbf{x}_{\mathbf{J}^c}) \mathcal{K}_x d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}), \tag{68}$$

$$\mathbf{L}_{\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) = \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u}) (\mathbf{f}_{\mathbf{J}}(\mathbf{x}) \cdot (\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{\mathbf{J}})) (\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{\mathbf{J}}) \mathcal{K}_x d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}), \quad \forall \mathbf{f}_{\mathbf{J}} \in \mathbb{H}_{\mathcal{K}}^{|\mathbf{J}|}, \tag{69}$$

$$\tilde{\mathbf{L}}_{\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) = \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u}) (\mathbf{f}_{\mathbf{J}}(\mathbf{x}) \cdot (\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{\mathbf{J}})) (\mathbf{u}_{\mathbf{J}^c} - \mathbf{x}_{\mathbf{J}^c}) \mathcal{K}_x d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}), \quad \forall \mathbf{f}_{\mathbf{J}} \in \mathbb{H}_{\mathcal{K}}^{|\mathbf{J}|}. \tag{70}$$

For any $d, q \in \mathbb{N}_+$, denote

$$M_{d, q} = \int_{\mathbb{R}^d} e^{-\frac{\|\mathbf{t}\|^2}{2}} \|\mathbf{t}\|^q d\mathbf{t}.$$

Denote $\text{Diam}(X) = \sup_{\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2 \in X} |\tilde{\mathbf{x}}_1 - \tilde{\mathbf{x}}_2|$.

Lemma 10 *Let $\mathbf{J} = \mathbf{J}(\nabla f_\rho)$ and $\nabla_{\mathbf{J}} f_\rho = (\frac{\partial f_\rho}{\partial \mathbf{x}^j})_{j \in \mathbf{J}}$. Under Assumption 1, we have*

$$\|\mathbf{f}_{\rho,s,\mathbf{J}} - \mathbf{L}_{\mathbf{J}}(\nabla_{\mathbf{J}} f_\rho)\|_{\mathcal{K}} \leq s^{p+3} \kappa C'_\rho C_\rho M_{|\mathbf{J}|,3} M_{|\mathbf{J}^c|,2}.$$

Proof According to the definition of $\mathbf{f}_{\rho,s,\mathbf{J}}$ in (67) and $\mathbf{L}_{\mathbf{J}}$ in (69), we have

$$\begin{aligned} & \mathbf{f}_{\rho,s,\mathbf{J}} - \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_\rho \\ &= \int_X \int_X w(\mathbf{x} - \mathbf{u}) (f_\rho(\mathbf{u}) - f_\rho(\mathbf{x}) - \nabla_{\mathbf{J}} f_\rho(\mathbf{x}) \cdot (\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{\mathbf{J}})) (\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{\mathbf{J}}) \mathcal{K}_{\mathbf{x}} d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}). \end{aligned}$$

Assumption 1 implies

$$|f_\rho(\mathbf{u}) - f_\rho(\mathbf{x}) - \nabla_{\mathbf{J}} f_\rho(\mathbf{x}) \cdot (\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{\mathbf{J}})| \leq C'_\rho |\mathbf{u} - \mathbf{x}|^2, \quad \forall \mathbf{u}, \mathbf{x} \in X.$$

Therefore,

$$\begin{aligned} \|\mathbf{f}_{\rho,s,\mathbf{J}} - \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_\rho\| &\leq C'_\rho \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u}) |\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{\mathbf{J}}| |\mathbf{u} - \mathbf{x}|^2 \|\mathcal{K}_{\mathbf{x}}\|_{\mathcal{K}} d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}) \\ &\leq \kappa C'_\rho C_\rho \int_X \int_X e^{-\frac{|\mathbf{x}-\mathbf{u}|^2}{2s^2}} |\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{\mathbf{J}}| |\mathbf{u} - \mathbf{x}|^2 d\mathbf{x} d\rho_X(\mathbf{u}) \\ &\leq s^{p+3} \kappa C'_\rho C_\rho M_{|\mathbf{J}|,3} M_{|\mathbf{J}^c|,2}. \end{aligned} \quad \square$$

The following Proposition which will be used frequently in this paper can be found in Mukherjee and Zhou (2006).

Proposition 12 *Let $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1}^n$ be i.i.d. draws from a probability distribution ρ on Z , $(\mathbb{H}, \|\cdot\|)$ be a Hilbert space, and $F : Z^n \mapsto \mathbb{H}$ be measurable. If there is $\tilde{M} \geq 0$ such that $\|F(\mathcal{Z}) - E_{\mathbf{z}_i}(F(\mathcal{Z}))\| \leq \tilde{M}$ for each $1 \leq i \leq n$ and almost every $\mathcal{Z} \in Z^n$, then for every $\epsilon > 0$*

$$\text{Prob}_{\mathcal{Z} \in Z^n} \left\{ \|F(\mathcal{Z}) - E_{\mathcal{Z}}(F(\mathcal{Z}))\| \geq \epsilon \right\} \leq 2 \exp \left\{ -\frac{\epsilon^2}{2(\tilde{M}\epsilon + \sigma^2)} \right\},$$

where $\sigma^2 := \sum_{i=1}^n \sup_{\mathcal{Z} \setminus \{\mathbf{z}_i\} \in Z^{n-1}} E_{\mathbf{z}_i} \{ \|F(\mathcal{Z}) - E_{\mathbf{z}_i}(F(\mathcal{Z}))\|^2 \}$.

Lemma 11 *Let $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1}^n$ be i.i.d. draws from a probability distribution ρ on Z . Assume $\mathbf{f}_{\mathbf{J}} \in \mathbb{H}_{\mathcal{K}}^{|\mathbf{J}|}$ and independent of \mathcal{Z} . For any $\epsilon > 0$, we have*

$$\text{Prob} \left\{ \|\mathbf{S}_{\mathcal{Z},\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) - \mathbf{L}_{\mathbf{J}}(\mathbf{f}_{\mathbf{J}})\|_{\mathcal{K}} \geq \epsilon \right\} \leq 2 \exp \left\{ -\frac{\epsilon^2}{2(M_s \epsilon + \sigma_s^2)} \right\},$$

where $M_s = \frac{4\kappa^2(\text{Diam}(X))^2 \|\mathbf{f}_{\mathbf{J}}\|_\infty}{n}$, $\sigma_s = 16\kappa^4 C_\rho M_{|\mathbf{J}|,4} M_{|\mathbf{J}^c|,0} \|\mathbf{f}_{\mathbf{J}}\|_\infty^2 \frac{s^{p+4}}{n}$.

Proof Let

$$F(\mathcal{Z}) = \mathbf{S}_{\mathcal{Z},\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) = \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s(\mathbf{f}_{\mathbf{J}}(\mathbf{x}_i)) \cdot (\mathbf{x}_{i,\mathbf{J}} - \mathbf{x}_{j,\mathbf{J}}) (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}}) \mathcal{K}_{\mathbf{x}_i}.$$

By independence, the expected value of $F(\mathcal{Z})$ equals

$$\begin{aligned} & \frac{1}{n(n-1)} \sum_{i=1}^n E_{z_i} \sum_{i \neq j} E_{z_j} \left\{ \omega_{i,j}^s(\mathbf{f}_J(\mathbf{x}_i) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_{i,J})) (\mathbf{x}_{j,J} - \mathbf{x}_{i,J}) \mathcal{K}_{\mathbf{x}_i} \right\} \\ &= \frac{1}{n} \sum_{i=1}^n E_{z_i} \int_X \omega^s(\mathbf{x}_i - \mathbf{u}) (\mathbf{f}_J(\mathbf{x}_i) \cdot (\mathbf{u}_J - \mathbf{x}_{i,J})) (\mathbf{u}_J - \mathbf{x}_{i,J}) \mathcal{K}_{\mathbf{x}_i} d\rho_X(\mathbf{u}). \end{aligned}$$

It follows that $E_{\mathcal{Z}} F(\mathcal{Z}) = \mathbf{L}_J(\mathbf{f}_J)$. Now we would apply Proposition 12 to the function $F(\mathcal{Z})$ to get our error bound on $F(\mathcal{Z}) - E_{\mathcal{Z}} F(\mathcal{Z})$. Let $i \in \{1, \dots, n\}$, we know that

$$\begin{aligned} F(\mathcal{Z}) - E_{z_i} F(\mathcal{Z}) &= \frac{1}{n(n-1)} \sum_{j \neq i} \omega_{i,j}^s(\mathbf{f}_J(\mathbf{x}_i) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_{i,J})) (\mathbf{x}_{j,J} - \mathbf{x}_{i,J}) \mathcal{K}_{\mathbf{x}_i} \\ &+ \frac{1}{n(n-1)} \sum_{j \neq i} \omega_{i,j}^s(\mathbf{f}_J(\mathbf{x}_j) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_{i,J})) (\mathbf{x}_{j,J} - \mathbf{x}_{i,J}) \mathcal{K}_{\mathbf{x}_j} \\ &- \frac{1}{n(n-1)} \sum_{j \neq i} \int_X e^{-\frac{|\mathbf{x} - \mathbf{x}_j|^2}{2s^2}} (\mathbf{f}_J(\mathbf{x}) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_J)) (\mathbf{x}_{j,J} - \mathbf{x}_J) \mathcal{K}_{\mathbf{x}} d\rho_X(\mathbf{x}) \\ &- \frac{1}{n(n-1)} \sum_{j \neq i} \int_X e^{-\frac{|\mathbf{x} - \mathbf{x}_j|^2}{2s^2}} (\mathbf{f}_J(\mathbf{x}_j) \cdot (\mathbf{x}_J - \mathbf{x}_{j,J})) (\mathbf{x}_J - \mathbf{x}_{j,J}) \mathcal{K}_{\mathbf{x}_j} d\rho_X(\mathbf{x}). \end{aligned}$$

Note that $\text{Diam}(X) = \sup_{\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2 \in X} |\tilde{\mathbf{x}}_1 - \tilde{\mathbf{x}}_2|$. Therefore, $|\mathbf{x}_J - \mathbf{x}_{j,J}| \leq |\mathbf{x} - \mathbf{x}_j| \leq \text{Diam}(X)$. For any $\mathbf{x} \in X$, we see that

$$\|F(\mathcal{Z}) - E_{z_i}(F(\mathcal{Z}))\|_{\mathcal{K}} \leq M_s = \frac{4(\kappa \text{Diam}(X))^2}{n} \|\mathbf{f}_J\|_{\infty}.$$

Furthermore,

$$\begin{aligned} (E_{z_i} \|F(\mathcal{Z}) - E_{z_i} F(\mathcal{Z})\|_{\mathcal{K}}^2)^{1/2} &\leq \frac{4\kappa}{n(n-1)} \sum_{j \neq i} \left(\int_X \left(e^{-\frac{|\mathbf{x} - \mathbf{x}_j|^2}{2s^2}} \kappa \|\mathbf{f}_J\|_{\infty} |\mathbf{x}_{j,J} - \mathbf{x}_J|^2 \right)^2 d\rho_X(\mathbf{x}) \right)^{\frac{1}{2}} \\ &\leq 4\kappa^2 \sqrt{C_{\rho} M_{|J|,4} M_{|J^c|,0}} \|\mathbf{f}_J\|_{\infty} \frac{s^{\frac{p}{2}+2}}{n}. \end{aligned}$$

Therefore,

$$\sigma^2 = \sum_{i=1}^n \sup_{\mathcal{Z} \setminus \{z_i\} \in \mathcal{Z}^{n-1}} E_{z_i} \|F(\mathcal{Z}) - E_{z_i}(F(\mathcal{Z}))\|_{\mathcal{K}}^2 \leq \sigma_s^2 := 16\kappa^4 C_{\rho} M_{|J|,4} M_{|J^c|,0} \|\mathbf{f}_J\|_{\infty}^2 \frac{s^{p+4}}{n}.$$

The lemma follows directly from Proposition 12. □

Lemma 12 *Let $\mathcal{Z} = \{z_i\}_{i=1}^n$ be i.i.d. draws from a probability distribution ρ on Z . Assume $|y| \leq M$ almost surely. Let Y_J and $\mathbf{f}_{\rho,s,J}$ be defined as in (30) and (67). For any $\epsilon > 0$, we have*

$$\text{Prob}\{\|\mathbf{f}_{\rho,s,J} - Y_J\|_{\mathcal{K}} \geq \epsilon\} \leq 2 \exp\left\{-\frac{\epsilon^2}{2(M_Y \epsilon + \sigma_Y^2)}\right\},$$

where $M_Y = \frac{8\kappa M \text{Diam}(X)}{n}$, $\sigma_Y = 64\kappa^2 M^2 C_{\rho} M_{|J|,2} M_{|J^c|,2} \frac{s^{p+2}}{n}$.

Proof Let

$$F(\mathcal{Z}) = Y_{\mathbf{J}} = \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s(y_i - y_j)(\mathbf{x}_{i,\mathbf{J}} - \mathbf{x}_{j,\mathbf{J}}) \mathcal{K}_{\mathbf{x}_i}.$$

By independence, the expected value of $F(\mathcal{Z})$ equals

$$\begin{aligned} & \frac{1}{n(n-1)} \sum_{i=1}^n E_{z_i} \sum_{j \neq i} E_{z_j} \omega_{i,j}^s(y_i - y_j)(\mathbf{x}_{i,\mathbf{J}} - \mathbf{x}_{j,\mathbf{J}}) \mathcal{K}_{\mathbf{x}_i} \\ &= \frac{1}{n} \sum_{i=1}^n E_{z_i} \int_X \omega(\mathbf{x}_i - \mathbf{u})(y_i - v)(\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}}) \mathcal{K}_{\mathbf{x}_i} d\rho_X(\mathbf{u}, v). \end{aligned}$$

It follows that $E_{\mathcal{Z}} F(\mathcal{Z}) = \mathbf{f}_{\rho,s,\mathbf{J}}$.

Now we would apply Proposition 12 to the function $F(\mathcal{Z})$ to get our error bound on $F(\mathcal{Z}) - E_{\mathcal{Z}} F(\mathcal{Z})$. Let $i \in \{1, 2, \dots, n\}$, we know that

$$\begin{aligned} F(\mathcal{Z}) - E_{z_i} F(\mathcal{Z}) &= \frac{1}{n(n-1)} \sum_{j \neq i} \omega_{i,j}(y_i - y_j)(\mathbf{x}_{i,\mathbf{J}} - \mathbf{x}_{j,\mathbf{J}}) (\mathcal{K}_{\mathbf{x}_i} + \mathcal{K}_{\mathbf{x}_j}) \\ &\quad - \frac{1}{n(n-1)} \sum_{j \neq i} \int_{\mathcal{Z}} \omega(\mathbf{x} - \mathbf{x}_j)(y - y_j)(\mathbf{x}_{\mathbf{J}} - \mathbf{x}_{j,\mathbf{J}}) (\mathcal{K}_{\mathbf{x}} + \mathcal{K}_{\mathbf{x}_j}) d\rho(\mathbf{x}, y). \end{aligned}$$

Using $\|\mathbf{x} - \mathbf{x}_j\| \leq \text{Diam}(X)$ for any $\mathbf{x} \in X$, we see that

$$\|F(\mathcal{Z}) - E_{z_i} F(\mathcal{Z})\|_{\mathcal{K}} \leq M_{\mathbf{Y}} = \frac{8\kappa M \text{Diam}(X)}{n}.$$

Furthermore, $(E_{z_i} \|F(\mathcal{Z}) - E_{z_i}(F(\mathcal{Z}))\|_{\mathcal{K}}^2)^{\frac{1}{2}}$ is bounded by

$$\begin{aligned} & \frac{4}{n(n-1)} \sum_{j \neq i} \left\{ \int_X (\omega(\mathbf{x} - \mathbf{x}_j))^2 \|\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{\mathbf{J}}\|^2 (2\kappa M)^2 d\rho_X(\mathbf{x}) \right\}^{1/2} \\ &= \frac{8\kappa M}{n(n-1)} \sum_{j \neq i} \left\{ \int_X e^{-\frac{\|\mathbf{x} - \mathbf{x}_j\|^2}{s^2}} \|\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{\mathbf{J}}\|^2 C_{\rho} d\mathbf{x} \right\}^{1/2} \\ &\leq \frac{8\kappa M \sqrt{s^{p+2}}}{n} \sqrt{C_{\rho} M_{|\mathbf{J}|,2} M_{|\mathbf{J}^c|,0}}. \end{aligned}$$

Therefore,

$$\sigma^2 = \sum_{i=1}^n \sup_{\mathcal{Z} \setminus \{z_i\} \in \mathcal{Z}^{n-1}} E_{z_i} \{ \|F(\mathcal{Z}) - E_{z_i}(F(\mathcal{Z}))\|^2 \} \leq \sigma_{\mathbf{Y}}^2 := 64\kappa^2 M^2 C_{\rho} M_{|\mathbf{J}|,2} M_{|\mathbf{J}^c|,0} \frac{s^{p+2}}{n}.$$

The lemma follows directly from Proposition 12. □

The following lemma can be easily derived from Theorem 19 in Mukherjee and Zhou (2006).

Lemma 13 *Suppose Assumption 1 holds. Let $0 < s \leq \min\{1, \lambda^{1/(p+2+\theta)}\}$. If $\|\mathbf{L}_{\mathcal{K}}^{-r} \nabla f_{\rho}\|_{\rho} < \infty$ for some $r \geq \frac{1}{2}$, then we have for any $\lambda > 0$,*

$$\|\lambda(\mathbf{L}_{\mathbf{J}} + \lambda \mathbf{I})^{-1} \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \leq C_{\rho,r} (\lambda s^{p+2})^{r-1/2} n^{-r},$$

where $C_{\rho,r} = 2(V_{\rho}(2\pi)^{n/2})^{-r} \|\mathbf{L}_{\mathcal{K}}^{-r} \nabla f_{\rho}\|_{\rho}$.

Proof of Proposition 7 Note that from Cauchy-Schwarz inequality, we have

$$\begin{aligned} \left(\sum_{j \in \mathbf{J}} \|f_j\|_{\mathcal{K}}\right)^2 &= \left(\sum_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}}^{\frac{1}{2}} \frac{\|f^j\|_{\mathcal{K}}}{\|\partial f_{\rho} / \partial x^j\|_{\mathcal{K}}^{\frac{1}{2}}}\right)^2 \\ &\leq \sum_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}} \sum_{j \in \mathbf{J}} \frac{\|f^j\|_{\mathcal{K}}^2}{\|\partial f_{\rho} / \partial x^j\|_{\mathcal{K}}}. \end{aligned}$$

We consider the unique minimizer $\tilde{\mathbf{f}}_{\mathbf{Z},\mathbf{J}}$ of the following cost function, built by replacing the regularization by its upper bound,

$$F(\mathbf{f}_{\mathbf{J}}) = \frac{1}{n(n-1)} \sum_{i=1}^n (y_i - y_j + \mathbf{f}_{\mathbf{J}}(\mathbf{x}_i) \cdot (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}}))^2 + \lambda \sum_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}} \sum_{j \in \mathbf{J}} \frac{\|f^j\|_{\mathcal{K}}^2}{\|\partial f_{\rho} / \partial x^j\|_{\mathcal{K}}}. \tag{71}$$

Firstly, we would prove that there exists a constant $0 < s_0 \leq 1$ such that for all $0 < s \leq s_0$, we have

$$\text{Prob} \left\{ \|\tilde{\mathbf{f}}_{\mathbf{Z},\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq \frac{1}{2} \min_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}} \right\} \leq \text{Prob} \{ \|\tilde{\mathbf{f}}_{\mathbf{Z},\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq s^{1-\theta} \}. \tag{72}$$

That is, we need to show that for all $0 < s \leq s_0$, $\|\tilde{\mathbf{f}}_{\mathbf{Z},\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} < s^{1-\theta}$ implies

$$\|\tilde{\mathbf{f}}_{\mathbf{Z},\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} < \frac{1}{2} \min_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}}.$$

Consider the cost function defining $\tilde{\mathbf{f}}_{\mathbf{Z},\mathbf{J}}$.

$$\tilde{F}(\mathbf{f}_{\mathbf{J}}) = \frac{1}{n(n-1)} \sum_{i=1}^n (y_i - y_j + \mathbf{f}_{\mathbf{J}}(\mathbf{x}_i) \cdot (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}}))^2 + \lambda \left(\sum_{j \in \mathbf{J}} \|f^j\|_{\mathcal{K}}\right)^2.$$

Denote C_F to be the uniform lower bound on the second derivative of $(\sum_{j \in \mathbf{J}} \|f^j\|_{\mathcal{K}})^2$. Then for all $\mathbf{f}_{\mathbf{J}} \in \mathbb{H}_{\mathcal{K}}^{|\mathbf{J}|}$, we have

$$\tilde{F}(\mathbf{f}_{\mathbf{J}}) \geq \tilde{F}(\tilde{\mathbf{f}}_{\mathbf{Z},\mathbf{J}}) + \langle \mathbf{f}_{\mathbf{J}} - \tilde{\mathbf{f}}_{\mathbf{Z},\mathbf{J}}, \nabla_{\mathbf{f}_{\mathbf{J}}} \tilde{F}(\tilde{\mathbf{f}}_{\mathbf{Z},\mathbf{J}}) \rangle + C_F \lambda \|\mathbf{f}_{\mathbf{J}} - \tilde{\mathbf{f}}_{\mathbf{Z},\mathbf{J}}\|_{\mathcal{K}}^2. \tag{73}$$

Note that

$$\begin{aligned} \tilde{F}(\mathbf{f}_{\mathbf{J}}) - F(\mathbf{f}_{\mathbf{J}}) &= \lambda \left(\sum_{j \in \mathbf{J}} \|f^j\|_{\mathcal{K}}\right)^2 - \lambda \sum_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}} \sum_{j \in \mathbf{J}} \frac{\|f^j\|_{\mathcal{K}}^2}{\|\partial f_{\rho} / \partial x^j\|_{\mathcal{K}}}, \\ \nabla_{f^i} \tilde{F}(\mathbf{f}_{\mathbf{J}}) - \nabla_{f^i} F(\mathbf{f}_{\mathbf{J}}) &= 2\lambda \left(\sum_{j \in \mathbf{J}} \|f^j\|_{\mathcal{K}}\right) \frac{f^i}{\|f^i\|_{\mathcal{K}}} - 2\lambda \left(\sum_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}}\right) \frac{f^i}{\|\partial f_{\rho} / \partial x^j\|_{\mathcal{K}}}. \end{aligned}$$

There exists a constant $C'_F > 0$ such that

$$\begin{aligned} \|\nabla_{f_i} \tilde{F}(\tilde{\mathbf{f}}_{Z,\mathbf{J}})\|_{\mathcal{K}} &= \|\nabla_{f_i} \tilde{F}(\tilde{\mathbf{f}}_{Z,\mathbf{J}}) - \nabla_{f_i} F(\tilde{\mathbf{f}}_{Z,\mathbf{J}}) - (\nabla_{f_i} \tilde{F}(\nabla_{\mathbf{J}} f_{\rho}) - \nabla_{f_i} F(\nabla_{\mathbf{J}} f_{\rho}))\|_{\mathcal{K}} \\ &\leq C'_F \lambda \|\tilde{\mathbf{f}}_{Z,\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}}. \end{aligned}$$

Together with (73) and the fact $\|\tilde{\mathbf{f}}_{Z,\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} < s^{1-\theta}$, for any $\mathbf{f}_{\mathbf{J}} \in \{\mathbf{f}_{\mathbf{J}} : \|\mathbf{f}_{\mathbf{J}} - \tilde{\mathbf{f}}_{Z,\mathbf{J}}\|_{\mathcal{K}} \leq \sqrt{s^{1-\theta}}\}$,

$$\tilde{F}(\mathbf{f}_{\mathbf{J}}) \geq \tilde{F}(\tilde{\mathbf{f}}_{Z,\mathbf{J}}) - C'_F \lambda (s^{1-\theta})^{\frac{3}{2}} + C_F \lambda s^{1-\theta}.$$

Choose $s_0 = \min\{(\frac{C_F}{C'_F})^{\frac{2}{1-\theta}}, (\frac{1}{4} \min_{j \in \mathbf{J}} \|\frac{\partial f_{\rho}}{\partial x^j}\|_{\mathcal{K}})^{\frac{2}{1-\theta}}, 1\}$, then for all $s < s_0$, we would have $\tilde{F}(\mathbf{f}_{\mathbf{J}}) \geq \tilde{F}(\tilde{\mathbf{f}}_{Z,\mathbf{J}})$. Hence, we must have all minima inside the ball $\mathbf{f}_{\mathbf{J}} \in \{\mathbf{f}_{\mathbf{J}} : \|\mathbf{f}_{\mathbf{J}} - \tilde{\mathbf{f}}_{Z,\mathbf{J}}\|_{\mathcal{K}} \leq \sqrt{s^{1-\theta}}\}$ which implies that

$$\|\tilde{\mathbf{f}}_{Z,\mathbf{J}} - \tilde{\mathbf{f}}_{Z,\mathbf{J}}\|_{\mathcal{K}} \leq \sqrt{s^{1-\theta}}.$$

Together with the fact that $0 < s < \min\{s_0, 1\}$ and $\|\tilde{\mathbf{f}}_{Z,\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} < s^{1-\theta}$, we have

$$\begin{aligned} \|\tilde{\mathbf{f}}_{Z,\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} &\leq \|\tilde{\mathbf{f}}_{Z,\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} + \|\tilde{\mathbf{f}}_{Z,\mathbf{J}} - \tilde{\mathbf{f}}_{Z,\mathbf{J}}\|_{\mathcal{K}} < s^{1-\theta} + \sqrt{s^{1-\theta}} \\ &\leq 2\sqrt{s^{1-\theta}} \leq \frac{1}{2} \min_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}}. \end{aligned}$$

Therefore, the inequality (72) holds.

Now we would prove that there exists constants $C_{\mathbf{J}}$ such that

$$\text{Prob}\{\|\tilde{\mathbf{f}}_{Z,\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq s^{1-\theta}\} \leq 2 \exp\{-C_{\mathbf{J}} n s^{p+2}\}. \tag{74}$$

Since (71) is a regularized least-square problem, we have

$$\tilde{\mathbf{f}}_{Z,\mathbf{J}} = (\mathbf{S}_{Z,\mathbf{J}} + \lambda \mathbf{D})^{-1} Y_{\mathbf{J}}, \tag{75}$$

where $\mathbf{D} = (\sum_{j \in \mathbf{J}} \|\frac{\partial f_{\rho}}{\partial x^j}\|_{\mathcal{K}}) \text{diag}(1/\|\partial f_{\rho}/\partial x^j\|_{\mathcal{K}})_{j \in \mathbf{J}}$. Let

$$D_{\min} = \min_{j' \in \mathbf{J}} \left(\sum_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}} \right) / \|\partial f_{\rho}/\partial x^{j'}\|_{\mathcal{K}}$$

and

$$D_{\max} = \max_{j' \in \mathbf{J}} \left(\sum_{j \in \mathbf{J}} \left\| \frac{\partial f_{\rho}}{\partial x^j} \right\|_{\mathcal{K}} \right) / \|\partial f_{\rho}/\partial x^{j'}\|_{\mathcal{K}}.$$

Then \mathbf{D} is upperbounded and lowerbounded, as an auto-adjoint operator on $\mathbb{H}_{\mathcal{K}}$, by strictly positive constants times the identity operator, that is, $D_{\max} \mathbf{I} \geq \mathbf{D} \geq D_{\min} \mathbf{I}$. Note that $\tilde{\mathbf{f}}_{Z,\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho} = (\mathbf{S}_{Z,\mathbf{J}} + \lambda \mathbf{D})^{-1} (Y_{\mathbf{J}} - \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho}) + (\mathbf{S}_{Z,\mathbf{J}} + \lambda \mathbf{D})^{-1} \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \nabla_{\mathbf{J}} f_{\rho}$. Hence

$$\begin{aligned} &\text{Prob}\{\|\tilde{\mathbf{f}}_{Z,\mathbf{J}} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq s^{1-\theta}\} \\ &\leq \text{Prob}\left\{\|Y_{\mathbf{J}} - \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq \frac{1}{2} s^{1-\theta} \lambda\right\} \\ &\quad + \text{Prob}\left\{\|(\mathbf{S}_{Z,\mathbf{J}} + \lambda \mathbf{D})^{-1} \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq \frac{1}{2} s^{1-\theta}\right\}. \end{aligned} \tag{76}$$

Choosing $\lambda = \tilde{C}_{M,\theta} s^{p+2+\theta}$ and using Lemma 10, we have

$$\begin{aligned} & \text{Prob} \left\{ \|Y_{\mathbf{J}} - \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq \frac{1}{2} s^{1-\theta} \lambda \right\} \\ & \leq \text{Prob} \left\{ \|Y_{\mathbf{J}} - \mathbf{f}_{\rho,s,\mathbf{J}}\|_{\mathcal{K}} \geq \frac{1}{2} s^{1-\theta} \lambda - \|\mathbf{f}_{\rho,s,\mathbf{J}} - \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \right\} \\ & \leq \text{Prob} \left\{ \|Y_{\mathbf{J}} - \mathbf{f}_{\rho,s,\mathbf{J}}\|_{\mathcal{K}} \geq s^{p+3} \right\} \\ & \leq \exp \left\{ -\frac{s^{2p+6}}{2(M_Y s^{p+3} + \sigma_Y)} \right\} \leq \exp \{-C_Y n s^{p+4}\}, \end{aligned} \tag{77}$$

where $C_Y = \frac{1}{16\kappa M \text{Diam } X + 32\kappa^2 M^2 C_{\rho} M_{|\mathbf{J}|,2} M_{|\mathbf{J}^c|,0}}$. On the other hand,

$$\begin{aligned} & (\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D})^{-1} \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \nabla_{\mathbf{J}} f_{\rho} \\ & = \{(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D})^{-1} - (\mathbf{L}_{\mathbf{J}} + \lambda \mathbf{D})^{-1}\} \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} + (\mathbf{L}_{\mathbf{J}} + \lambda \mathbf{D})^{-1} \lambda \mathbf{D} \nabla_{\mathbf{J}} f_{\rho}. \end{aligned}$$

Since $(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D})^{-1} - (\mathbf{L}_{\mathbf{J}} + \lambda \mathbf{D})^{-1} = (\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D})^{-1} (\mathbf{L}_{\mathbf{J}} - \mathbf{S}_{\mathcal{Z},\mathbf{J}}) (\mathbf{L}_{\mathbf{J}} + \lambda \mathbf{D})^{-1}$, we have

$$\| \{(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D})^{-1} - (\mathbf{L}_{\mathbf{J}} + \lambda \mathbf{D})^{-1}\} \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} \|_{\mathcal{K}} \leq \frac{1}{\lambda} \| \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \mathbf{S}_{\mathcal{Z},\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} \|_{\mathcal{K}}.$$

Together with Lemma 13 and the fact $n^r s^{(2p+4+\theta)(1/2-r)+1-\theta} \geq 4D_{\max} (D_{\min})^{r-\frac{3}{2}} (\tilde{C}_{M,\theta})^{r-\frac{1}{2}} C_{\rho,r}$, we have

$$\begin{aligned} & \text{Prob} \left\{ \|(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D})^{-1} \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq \frac{1}{2} s^{1-\theta} \right\} \\ & \leq \text{Prob} \left\{ \|\mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \mathbf{S}_{\mathcal{Z},\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq \frac{1}{2} \lambda s^{1-\theta} - D_{\max} C_{\rho,r} (D_{\min})^{r-\frac{3}{2}} \lambda (\lambda s^{p+2})^{r-1/2} n^{-r} \right\} \\ & \leq \text{Prob} \left\{ \|\mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \mathbf{S}_{\mathcal{Z},\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq \frac{1}{4} \lambda s^{1-\theta} \right\}. \end{aligned}$$

Then using Lemma 11, we can easily get

$$\text{Prob} \left\{ \|(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D})^{-1} \mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq \frac{1}{2} s^{1-\theta} \right\} \leq 2 \exp \{-\tilde{C}_s n s^{p+4}\}, \tag{78}$$

where

$$\tilde{C}_s = \frac{(\tilde{C}_{M,\theta})^2}{32((\kappa \text{Diam}(X))^2 \|\nabla f_{\rho}\|_{\infty} \tilde{C}_{M,\theta} + 16\kappa^4 C_{\rho} M_{|\mathbf{J}|,4} M_{|\mathbf{J}^c|,0} \|\nabla f_{\rho}\|_{\infty}^2)}.$$

Combining (76), (77) and (78), we get (74) by letting $C_{\mathbf{J}} = \min\{C_Y, \tilde{C}_s\}$. Together with (72), we get the desired result. \square

Lemma 14 Let $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1}^n$ be i.i.d. draws from a probability distribution ρ on Z . Assume $\mathbf{f}_{\mathbf{J}} \in \mathbb{H}_{\mathcal{K}}^{|\mathbf{J}|}$ and independent of \mathcal{Z} . For any $\epsilon > 0$, we have

$$\text{Prob} \left\{ \|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) - \tilde{\mathbf{L}}_{\mathbf{J}}(\mathbf{f}_{\mathbf{J}})\|_{\mathcal{K}} \geq \epsilon \right\} \leq \exp \left\{ -\frac{\epsilon^2}{2(\tilde{M}_s \epsilon + \tilde{\sigma}_s^2)} \right\},$$

where $\tilde{M}_s = \frac{4(\kappa \text{Diam}(X))^2 \|\mathbf{f}_{\mathbf{J}}\|_{\infty}}{n}$ and $\tilde{\sigma}_s^2 = 16\kappa^4 C_{\rho} M_{|\mathbf{J}|,2} M_{|\mathbf{J}^c|,2} \|\mathbf{f}_{\mathbf{J}}\|_{\infty}^2 \frac{s^{p+4}}{n}$.

Proof Let

$$F(\mathcal{Z}, \mathbf{f}_J) = \tilde{S}_{\mathcal{Z}, J}(\mathbf{f}_J) = \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s(\mathbf{f}_J(\mathbf{x}_i) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_{i,J}))(\mathbf{x}_{j,J^c} - \mathbf{x}_{i,J^c})\mathcal{K}_{\mathbf{x}_i}.$$

By independence, the expected value of $F(\mathcal{Z}, \mathbf{f}_J)$ equals

$$\begin{aligned} & \frac{1}{n(n-1)} \sum_{i=1}^n E_{z_i} \sum_{i \neq j} E_{z_j} \{ \omega_{i,j}^s(\mathbf{f}_J(\mathbf{x}_i) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_{i,J}))(\mathbf{x}_{j,J^c} - \mathbf{x}_{i,J^c})\mathcal{K}_{\mathbf{x}_i} \} \\ &= \frac{1}{n} \sum_{i=1}^n E_{z_i} \int_X \omega^s(\mathbf{x}_i - \mathbf{u})(\mathbf{f}_J(\mathbf{x}_i) \cdot (\mathbf{u}_J - \mathbf{x}_{i,J}))(\mathbf{u}_{J^c} - \mathbf{x}_{i,J^c})\mathcal{K}_{\mathbf{x}_i} d\rho_X(\mathbf{u}). \end{aligned}$$

It follows that $E_{\mathcal{Z}} F(\mathcal{Z}, \mathbf{f}_J) = \tilde{L}_J(\mathbf{f}_J)$. Now we would apply Proposition 12 to the function $F(\mathcal{Z}, \mathbf{f}_J)$ to get our error bound on $F(\mathcal{Z}, \mathbf{f}_J) - E_{\mathcal{Z}} F(\mathcal{Z}, \mathbf{f}_J)$. Let $i \in \{1, \dots, n\}$, we know that

$$\begin{aligned} & F(\mathcal{Z}, \mathbf{f}_J) - E_{z_i} F(\mathcal{Z}, \mathbf{f}_J) \\ &= \frac{1}{n(n-1)} \sum_{j \neq i} \omega_{i,j}^s(\mathbf{f}_J(\mathbf{x}_i) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_{i,J}))(\mathbf{x}_{j,J^c} - \mathbf{x}_{i,J^c})\mathcal{K}_{\mathbf{x}_i} \\ & \quad + \frac{1}{n(n-1)} \sum_{j \neq i} \omega_{i,j}^s(\mathbf{f}_J(\mathbf{x}_j) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_{i,J}))(\mathbf{x}_{j,J^c} - \mathbf{x}_{i,J^c})\mathcal{K}_{\mathbf{x}_j} \\ & \quad - \frac{1}{n(n-1)} \sum_{j \neq i} \int_X e^{-\frac{|\mathbf{x}-\mathbf{x}_j|^2}{2s^2}} (\mathbf{f}_J(\mathbf{x}) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_J))(\mathbf{x}_{j,J^c} - \mathbf{x}_{j,J^c})\mathcal{K}_{\mathbf{x}} d\rho_X(\mathbf{x}) \\ & \quad - \frac{1}{n(n-1)} \sum_{j \neq i} \int_X e^{-\frac{|\mathbf{x}-\mathbf{x}_j|^2}{2s^2}} (\mathbf{f}_J(\mathbf{x}_j) \cdot (\mathbf{x}_J - \mathbf{x}_{j,J}))(\mathbf{x}_{j,J^c} - \mathbf{x}_{j,J^c})\mathcal{K}_{\mathbf{x}_j} d\rho_X(\mathbf{x}). \end{aligned}$$

Note that $\text{Diam}(X) = \sup_{\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2 \in X} |\tilde{\mathbf{x}}_1 - \tilde{\mathbf{x}}_2|$. Therefore, $\max\{|\mathbf{x}_J - \mathbf{x}_{j,J}|, |\mathbf{x}_{J^c} - \mathbf{x}_{j,J^c}|\} \leq |\mathbf{x} - \mathbf{x}_j| \leq \text{Diam}(X)$. For any $\mathbf{x} \in X$, we see that

$$\|F(\mathcal{Z}, \mathbf{f}_J) - E_{z_i}(F(\mathcal{Z}, \mathbf{f}_J))\|_{\mathcal{K}} \leq \tilde{M}_s = \frac{4(\kappa \text{Diam}(X))^2 \|\mathbf{f}_J\|_{\infty}}{n}.$$

Furthermore,

$$\begin{aligned} & (E_{z_i} \|F(\mathcal{Z}, \mathbf{f}_J) - E_{z_i} F(\mathcal{Z}, \mathbf{f}_J)\|_{\mathcal{K}}^2)^{1/2} \\ & \leq \frac{4\kappa}{n^2} \sum_{j \neq i} \left(\int_X (e^{-\frac{|\mathbf{x}-\mathbf{x}_j|}{2s^2}} \kappa \|\mathbf{f}_J\|_{\infty} |\mathbf{x}_{j,J} - \mathbf{x}_J| |\mathbf{x}_{j,J^c} - \mathbf{x}_{j,J^c}|)^2 d\rho_X(\mathbf{x}) \right)^{\frac{1}{2}} \\ & \leq 4\kappa^2 \sqrt{C_{\rho} M_{|J|,2} M_{|J^c|,2}} \|\mathbf{f}_J\|_{\infty} \frac{s^{\frac{\rho}{2}+2}}{n}. \end{aligned}$$

Therefore,

$$\sigma^2 = \sum_{i=1}^n \sup_{\mathcal{Z} \setminus \{z_i\} \in \mathcal{Z}^{n-1}} E_{z_i} \|F(\mathcal{Z}, \mathbf{f}_J) - E_{z_i} (F(\mathcal{Z}, \mathbf{f}_J))\|_{\mathcal{K}}^2 \leq \tilde{\sigma}_s^2 := 16\kappa^4 C_{\rho} M_{|J|,2} M_{|J^c|,2} \|\mathbf{f}_J\|_{\infty}^2 \frac{s^{\rho+4}}{n}.$$

According to Proposition 12, for any $\epsilon > 0$, we have

$$\text{Prob}\{\|\tilde{\mathbf{S}}_{Z,J}(\mathbf{f}_J) - \tilde{\mathbf{L}}_J(\mathbf{f}_J)\|_{\mathcal{K}} \geq \epsilon\} \leq \exp\left\{-\frac{\epsilon^2}{2(\tilde{M}_s\epsilon + \tilde{\sigma}_s^2)}\right\},$$

where $\tilde{M}_s = \frac{4(\kappa \text{Diam}(X))^2 \|\mathbf{f}_J\|_{\infty}}{n}$ and $\tilde{\sigma}_s^2 = 16\kappa^4 C_{\rho} M_{|J|,2} M_{|J^c|,2} \|\mathbf{f}_J\|_{\infty}^2 \frac{s^{p+4}}{n}$. □

Note that $B_1 = \{\mathbf{f}_J : \|\mathbf{f}_J\|_{\mathcal{K}} \leq 1\}$ and $J_{\mathcal{K}}$ be the inclusion map from B_1 to $C(X)$. Let $0 < \eta < \frac{1}{2}$. The covering number $\mathcal{N}(J_{\mathcal{K}}(B_1), \eta)$ is the minimal $\ell \in \mathbb{N}$ such that there exists ℓ disks in $J_{\mathcal{K}}(B_1)$ with radius η covering S .

Proposition 13 *Let $Z = \{z_i\}_{i=1}^n$ be i.i.d. draws from a probability distribution ρ on Z . For any $\epsilon > 0$, we have*

$$\text{Prob}\{\|\tilde{\mathbf{S}}_{Z,J} - \tilde{\mathbf{L}}_J\|_{\mathcal{K}} < \epsilon\} \geq 1 - \mathcal{N}\left(J_{\mathcal{K}}(B_1), \frac{\epsilon}{8(\kappa \text{Diam}(X))^2}\right) \exp\left\{-\frac{\epsilon^2}{8(\tilde{M}_s\epsilon + 2\tilde{\sigma}_s^2)}\right\},$$

where $\tilde{M}_s = \frac{4(\kappa \text{Diam}(X))^2}{n}$ and $\tilde{\sigma}_s^2 = 16\kappa^4 C_{\rho} M_{|J|,2} M_{|J^c|,2} \frac{s^{p+4}}{n}$. In particular,

$$\text{Prob}\{\|\tilde{\mathbf{S}}_{Z,J} - \tilde{\mathbf{L}}_J\|_{\mathcal{K}} < s^{p+2+\theta}\} \geq 1 - \mathcal{N}\left(J_{\mathcal{K}}(B_1), \frac{s^{p+2+\theta}}{8(\kappa \text{Diam}(X))^2}\right) \exp\{-C_{\tilde{s}} n s^{p+2+\theta}\}, \tag{79}$$

where $C_{\tilde{s}} = 1/(16(\kappa \text{Diam}(X))^2 + 256\kappa^4 C_{\rho} M_{|J|,2} M_{|J^c|,2})$.

Proof (1) Let $N = \mathcal{N}(J_{\mathcal{K}}(B_1), \frac{\epsilon}{8(\kappa \text{Diam}(X))^2})$. Note that B_1 is dense in $J_{\mathcal{K}}(B_1)$, there exists $\mathbf{f}_{J,j} \in \mathbb{H}_{\mathcal{K}}^{[J]}$, $j = 1, \dots, N$ such that disks D_j center at $\mathbf{f}_{J,j}$ with radius $\frac{\epsilon}{4(\kappa \text{Diam}(X))^2}$ cover $J_{\mathcal{K}}(B_1)$. It is easy to see that, for any $\mathbf{f}_J \in D_j \cap \mathbb{H}_{\mathcal{K}}^{[J]}$, we have

$$\|\tilde{\mathbf{S}}_{Z,J}(\mathbf{f}_J) - \tilde{\mathbf{L}}_J(\mathbf{f}_J) - \tilde{\mathbf{S}}_{Z,J}(\mathbf{f}_{J,j}) + \tilde{\mathbf{L}}_J(\mathbf{f}_{J,j})\|_{\mathcal{K}} \leq 2(\kappa \text{Diam}(X))^2 \|\mathbf{f}_J - \mathbf{f}_{J,j}\|_{\infty} \leq \frac{\epsilon}{2}.$$

Since this holds for all $\mathbf{f}_J \in D_j \cap \mathbb{H}_{\mathcal{K}}^{[J]}$, we get

$$\sup_{\mathbf{f}_J \in D_j \cap \mathbb{H}_{\mathcal{K}}^{[J]}} \|\tilde{\mathbf{S}}_{Z,J}(\mathbf{f}_J) - \tilde{\mathbf{L}}_J(\mathbf{f}_J)\|_{\mathcal{K}} \geq \epsilon \implies \|\tilde{\mathbf{S}}_{Z,J}(\mathbf{f}_{J,j}) - \tilde{\mathbf{L}}_J(\mathbf{f}_{J,j})\|_{\mathcal{K}} \geq \frac{\epsilon}{2}.$$

We conclude that, for $j = 1, \dots, N$,

$$\text{Prob}\left\{\sup_{\mathbf{f}_J \in D_j \cap \mathbb{H}_{\mathcal{K}}^{[J]}} \|\tilde{\mathbf{S}}_{Z,J}(\mathbf{f}_J) - \tilde{\mathbf{L}}_J(\mathbf{f}_J)\|_{\mathcal{K}} \geq \epsilon\right\} \leq \text{Prob}\left\{\|\tilde{\mathbf{S}}_{Z,J}(\mathbf{f}_{J,j}) - \tilde{\mathbf{L}}_J(\mathbf{f}_{J,j})\|_{\mathcal{K}} \geq \frac{\epsilon}{2}\right\}.$$

Note that $B_1 = \{\mathbf{f}_J : \|\mathbf{f}_J\|_{\mathcal{K}} \leq 1\} = (D_1 \cap \mathbb{H}_{\mathcal{K}}^{[J]}) \cup \dots \cup (D_N \cap \mathbb{H}_{\mathcal{K}}^{[J]})$, we have

$$\begin{aligned} & \text{Prob}\left\{\sup_{\|\mathbf{f}_J\|_{\mathcal{K}} \leq 1} \|\tilde{\mathbf{S}}_{Z,J}(\mathbf{f}_J) - \tilde{\mathbf{L}}_J(\mathbf{f}_J)\|_{\mathcal{K}} \geq \epsilon\right\} \\ & \leq \sum_{j=1}^N \text{Prob}\left\{\sup_{\mathbf{f}_J \in D_j \cap \mathbb{H}_{\mathcal{K}}^{[J]}} \|\tilde{\mathbf{S}}_{Z,J}(\mathbf{f}_J) - \tilde{\mathbf{L}}_J(\mathbf{f}_J)\|_{\mathcal{K}} \geq \epsilon\right\}. \end{aligned}$$

Together with Lemma 14, we have

$$\text{Prob}\{\|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}} - \tilde{\mathbf{L}}_{\mathbf{J}}\|_{\mathcal{K}} < \epsilon\} \geq 1 - 2\mathcal{N}\left(J_{\mathcal{K}}(B_1), \frac{\epsilon}{8(\kappa \text{Diam}(X))^2}\right) \exp\left\{-\frac{\epsilon^2}{8(\tilde{M}_s \epsilon + 2\tilde{\sigma}_s^2)}\right\},$$

where $\tilde{M}_s = \frac{2(\kappa \text{Diam}(X))^2}{n}$ and $\tilde{\sigma}_s^2 = 16\kappa^4 C_\rho M_{|\mathbf{J}|,2} M_{|\mathbf{J}^c|,2} \frac{s^{p+4}}{n}$. Letting $\epsilon = s^{p+2+\theta}$, (79) follows directly. \square

Proof of Proposition 8 It is easy to see that $\|\tilde{\mathbf{L}}_{\mathbf{J}}(\mathbf{f}_{\mathbf{J}})\|_{\mathcal{K}} \leq \int_X \int_X \kappa^2 \|\mathbf{f}_{\mathbf{J}}\|_{\mathcal{K}} \omega(\mathbf{x} - \mathbf{u}) |\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{\mathbf{J}}| |\mathbf{u}_{\mathbf{J}^c} - \mathbf{x}_{\mathbf{J}^c}| p(\mathbf{x}) d\mathbf{x} d\rho(\mathbf{u})$. Using the fact that $\omega(\mathbf{x}) |\mathbf{x}_{\mathbf{J}}| |\mathbf{x}_{\mathbf{J}^c}| = \omega(-\mathbf{x}) |\mathbf{x}_{\mathbf{J}}| |\mathbf{x}_{\mathbf{J}^c}|$ and $|p(\mathbf{x}) - p(\mathbf{u})| \leq C_\rho |\mathbf{x} - \mathbf{u}|^\theta$, we have

$$\|\tilde{\mathbf{L}}_{\mathbf{J}}\|_{\mathcal{K}} \leq \kappa^2 C_\rho M_{|\mathbf{J}|,1} M_{|\mathbf{J}^c|,1} s^{p+2+\theta}.$$

Note that $\|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}} - \tilde{\mathbf{L}}_{\mathbf{J}}\|_{\mathcal{K}} \geq \|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}\|_{\mathcal{K}} - \|\tilde{\mathbf{L}}_{\mathbf{J}}\|_{\mathcal{K}}$. The inequality $\|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}\|_{\mathcal{K}} \geq (\kappa^2 C_\rho M_{|\mathbf{J}|,1} M_{|\mathbf{J}^c|,1} + 1) s^{p+2+\theta}$ implies $\|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}} - \tilde{\mathbf{L}}_{\mathbf{J}}\|_{\mathcal{K}} \geq s^{p+2+\theta}$. Together with (79), we have

$$\begin{aligned} \text{Prob}(\Omega_4^c) &\leq \text{Prob}\{\|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}} - \tilde{\mathbf{L}}_{\mathbf{J}}\|_{\mathcal{K}} \geq s^{p+2+\theta}\} \\ &\leq \mathcal{N}\left(J_{\mathcal{K}}(B_1), \frac{s^{p+2+\theta}}{8(\kappa \text{Diam}(X))^2}\right) \exp\{-C_{\tilde{s}} n s^{p+2+\theta}\}. \end{aligned}$$

The desired result follows by using $\text{Prob}(\Omega_4^c) = 1 - \text{Prob}(\Omega_4)$. \square

Proof of Proposition 9 On the event Ω_0 , we have $\frac{1}{3} D_{\min} \mathbf{I} \leq \mathbf{D}_n \leq 3 D_{\max} \mathbf{I}$. Therefore,

$$\begin{aligned} &\|(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D}_n)^{-1} \mathbf{D}_n \nabla_{\mathbf{J}} f_\rho\|_{\mathcal{K}} \\ &\leq 3 D_{\max} \left\| \left(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \frac{1}{3} D_{\min} \lambda \mathbf{I}\right)^{-1} \nabla_{\mathbf{J}} f_\rho \right\|_{\mathcal{K}} \\ &\leq 3 D_{\max} \left\| \left(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \frac{1}{3} D_{\min} \lambda \mathbf{I}\right)^{-1} \nabla_{\mathbf{J}} f_\rho - \left(\mathbf{L}_{\mathbf{J}} + \frac{1}{3} D_{\min} \lambda \mathbf{I}\right)^{-1} \nabla_{\mathbf{J}} f_\rho \right\|_{\mathcal{K}} \\ &\quad + 3 D_{\max} \left\| \left(\mathbf{L}_{\mathbf{J}} + \frac{1}{3} D_{\min} \lambda \mathbf{I}\right)^{-1} \nabla_{\mathbf{J}} f_\rho \right\|_{\mathcal{K}} \\ &\leq \frac{27 D_{\max}}{D_{\min}^2 \lambda^2} \|(\mathbf{L}_{\mathbf{J}} - \mathbf{S}_{\mathcal{Z},\mathbf{J}}) \nabla_{\mathbf{J}} f_\rho\|_{\mathcal{K}} + \frac{9 D_{\max}}{D_{\min} \lambda} \left\| \left(\mathbf{L}_{\mathbf{J}} + \frac{1}{3} D_{\min} \lambda \mathbf{I}\right)^{-1} \frac{1}{3} D_{\min} \lambda \nabla_{\mathbf{J}} f_\rho \right\|_{\mathcal{K}}. \end{aligned}$$

Together with Lemma 13, on the event $\Omega_0 \cap \Omega_4$,

$$\begin{aligned} &\|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}} (\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D}_n)^{-1} \mathbf{D}_n \nabla_{\mathbf{J}} f_\rho\|_{\mathcal{K}} \\ &\leq \|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}\|_{\mathcal{K}} \|(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D}_n)^{-1} \mathbf{D}_n \nabla_{\mathbf{J}} f_\rho\|_{\mathcal{K}} \\ &\leq \frac{27 D_{\max} (\kappa^2 C_\rho M_{|\mathbf{J}|,1} M_{|\mathbf{J}^c|,1} + 1) s^{p+2+\theta}}{D_{\min}^2 \lambda^2} \|(\mathbf{L}_{\mathbf{J}} - \mathbf{S}_{\mathcal{Z},\mathbf{J}}) \nabla_{\mathbf{J}} f_\rho\|_{\mathcal{K}} \\ &\quad + \frac{9 D_{\max}}{D_{\min} \lambda} (\kappa^2 C_\rho M_{|\mathbf{J}|,1} M_{|\mathbf{J}^c|,1} + 1) C_{\rho,r} \left(\frac{1}{3} D_{\min} \lambda s^{p+2}\right)^{r-\frac{1}{2}} n^{-r}. \end{aligned}$$

If $n^r s^{(2\rho+4+\theta)(\frac{1}{2}-r)+\rho+2+\theta} \geq \tilde{C}_{D,\theta}$, then $\|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}(\mathbf{S}_{\mathcal{Z},\mathbf{J}} + \lambda \mathbf{D}_n)^{-1} \mathbf{D}_n \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq \frac{1}{6} \|\nabla f_{\rho}\|_{\mathcal{K}}$ implies

$$\|\mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \mathbf{S}_{\mathcal{Z},\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq \frac{\|\nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} D_{\min}^2 \tilde{C}_{M,\theta}^2}{324 D_{\max} (\kappa^2 C_{\rho} M_{|\mathbf{J}|,1} M_{|\mathbf{J}^c|,1} + 1)} s^{\rho+2+\theta}.$$

Therefore,

$$\text{Prob}(\Omega_1^c | \mathcal{Z} \in \Omega_0 \cap \Omega_4) \leq \text{Prob}\{\|\mathbf{L}_{\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho} - \mathbf{S}_{\mathcal{Z},\mathbf{J}} \nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} \geq C_{D,\kappa,\rho} s^{\rho+2+\theta}\}, \tag{80}$$

where $C_{D,\kappa,\rho} = \frac{\|\nabla_{\mathbf{J}} f_{\rho}\|_{\mathcal{K}} D_{\min}^2 \tilde{C}_{M,\theta}^2}{324 D_{\max} (\kappa^2 C_{\rho} M_{|\mathbf{J}|,1} M_{|\mathbf{J}^c|,1} + 1)}$. Together with Lemma 11, we get the desired result by letting

$$C_{\Omega_1} = C_{D,\kappa,\rho}^2 / (8\kappa^2 (\text{Diam}(X))^2 \|\nabla f_{\rho}\|_{\infty} C_{D,\kappa,\rho} + 32\kappa^4 C_{\rho} M_{|\mathbf{J}|,4} M_{|\mathbf{J}^c|,0} \|\nabla f_{\rho}\|_{\infty}^2). \quad \square$$

Lemma 15 Let $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1}^n$ be i.i.d. draws from a probability distribution ρ on Z . Assume $\mathbf{f}_{\mathbf{J}} \in \mathbb{H}_{\mathcal{K}}^{|\mathbf{J}|}$ and independent of \mathcal{Z} . For any $\epsilon > 0$, we have

$$\text{Prob}\{\|\tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) - \tilde{Y}_{\mathbf{J}} - (\tilde{\mathbf{L}}_{\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) - \tilde{\mathbf{f}}_{\rho,s,\mathbf{J}})\|_{\mathcal{K}} \geq \epsilon\} \leq 2 \exp\left\{-\frac{\epsilon^2}{2(\widehat{M}_s \epsilon + \widehat{\sigma}_s^2)}\right\},$$

where

$$\widehat{M}_s = \frac{4(\kappa \text{Diam}(X))^2}{n} (2M + \|\mathbf{f}_{\mathbf{J}}\|_{\infty})$$

and

$$\widehat{\sigma}_s^2 = (\kappa^2 \sqrt{C_{\rho} M_{|\mathbf{J}|,2} M_{|\mathbf{J}^c|,2}} \|\mathbf{f}_{\mathbf{J}}\|_{\infty} + 2M\kappa \sqrt{C_{\rho} M_{|\mathbf{J}^c|,2}})^2 \frac{s^{\rho+4}}{n}.$$

Proof Let

$$F(\mathcal{Z}) = \tilde{\mathbf{S}}_{\mathcal{Z},\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) - \tilde{Y}_{\mathbf{J}} = \frac{1}{n(n-1)} \sum_{i,j=1}^n \omega_{i,j}^s (y_i - y_j + \mathbf{f}_{\mathbf{J}}(\mathbf{x}_i) \cdot (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}})) (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}}) \mathcal{K}_{\mathbf{x}_i}.$$

By independence, the expected value of $F(\mathcal{Z})$ equals

$$\begin{aligned} & \frac{1}{n(n-1)} \sum_{i=1}^n E_{\mathbf{z}_i} \sum_{i \neq j} E_{\mathbf{z}_j} \{ \omega_{i,j}^s (y_i - y_j + \mathbf{f}_{\mathbf{J}}(\mathbf{x}_i) \cdot (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}})) (\mathbf{x}_{j,\mathbf{J}^c} - \mathbf{x}_{i,\mathbf{J}^c}) \mathcal{K}_{\mathbf{x}_i} \} \\ &= \frac{1}{n} \sum_{i=1}^n E_{\mathbf{z}_i} \int_Z \omega^s(\mathbf{x}_i - \mathbf{u}) (y_i - v + \mathbf{f}_{\mathbf{J}}(\mathbf{x}_i) \cdot (\mathbf{u}_{\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}})) (\mathbf{u}_{\mathbf{J}^c} - \mathbf{x}_{i,\mathbf{J}^c}) \mathcal{K}_{\mathbf{x}_i} d\rho(\mathbf{u}, v). \end{aligned}$$

It follows that $E_{\mathcal{Z}} F(\mathcal{Z}) = \tilde{\mathbf{L}}_{\mathbf{J}}(\mathbf{f}_{\mathbf{J}}) - \tilde{\mathbf{f}}_{\rho,s,\mathbf{J}}$. Now we would apply Proposition 12 to the function $F(\mathcal{Z})$ to get our error bound on $F(\mathcal{Z}) - E_{\mathcal{Z}} F(\mathcal{Z})$. Let $i \in \{1, \dots, n\}$, we know that

$$\begin{aligned} & F(\mathcal{Z}) - E_{\mathbf{z}_i} F(\mathcal{Z}) \\ &= \frac{1}{n(n-1)} \sum_{j \neq i} \omega_{i,j}^s (y_i - y_j + \mathbf{f}_{\mathbf{J}}(\mathbf{x}_i) \cdot (\mathbf{x}_{j,\mathbf{J}} - \mathbf{x}_{i,\mathbf{J}})) (\mathbf{x}_{j,\mathbf{J}^c} - \mathbf{x}_{i,\mathbf{J}^c}) \mathcal{K}_{\mathbf{x}_i} \end{aligned}$$

$$\begin{aligned}
 &+ \frac{1}{n(n-1)} \sum_{j \neq i} \omega_{i,j}^s (y_i - y_j + \mathbf{f}_J(\mathbf{x}_j) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_{i,J})) (\mathbf{x}_{j,J^c} - \mathbf{x}_{i,J^c}) \mathcal{K}_{\mathbf{x}_j} \\
 &- \frac{1}{n(n-1)} \sum_{j \neq i} \int_Z e^{-\frac{|\mathbf{x}-\mathbf{x}_j|^2}{2s^2}} (y - y_j + \mathbf{f}_J(\mathbf{x}) \cdot (\mathbf{x}_{j,J} - \mathbf{x}_J)) (\mathbf{x}_{j,J^c} - \mathbf{x}_{J^c}) \mathcal{K}_x d\rho(\mathbf{x}, y) \\
 &- \frac{1}{n(n-1)} \sum_{j \neq i} \int_Z e^{-\frac{|\mathbf{x}-\mathbf{x}_j|^2}{2s^2}} (y - y_j + \mathbf{f}_J(\mathbf{x}_j) \cdot (\mathbf{x}_J - \mathbf{x}_{j,J})) (\mathbf{x}_{J^c} - \mathbf{x}_{i,J^c}) \mathcal{K}_{\mathbf{x}_j} d\rho(\mathbf{x}, y).
 \end{aligned}$$

Note that $\text{Diam}(X) = \sup_{\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2 \in X} |\tilde{\mathbf{x}}_1 - \tilde{\mathbf{x}}_2|$. Therefore, $\max\{|\mathbf{x}_J - \mathbf{x}_{j,J}|, |\mathbf{x}_{J^c} - \mathbf{x}_{j,J^c}|\} \leq |\mathbf{x} - \mathbf{x}_j| \leq \text{Diam}(X)$. For any $\mathbf{x} \in X$, we see that

$$\|F(\mathcal{Z}) - E_{z_i}(F(\mathcal{Z}))\|_{\mathcal{K}} \leq \widehat{M}_s = \frac{4(\kappa \text{Diam}(X))^2}{n} (2M + \|\mathbf{f}_J\|_{\infty}).$$

Furthermore,

$$\begin{aligned}
 &(E_{z_i} \|F(\mathcal{Z}) - E_{z_i} F(\mathcal{Z})\|_{\mathcal{K}}^2)^{1/2} \\
 &\leq \frac{4\kappa}{n^2} \sum_{j \neq i} \left(\int_X \left(e^{-\frac{|\mathbf{x}-\mathbf{x}_j|^2}{2s^2}} (2M + \kappa \|\mathbf{f}_J\|_{\infty} |\mathbf{x}_{j,J} - \mathbf{x}_J|) |\mathbf{x}_{j,J^c} - \mathbf{x}_{J^c}| \right)^2 d\rho_X(\mathbf{x}) \right)^{\frac{1}{2}} \\
 &\leq (\kappa^2 \sqrt{C_{\rho} M_{|J|,2} M_{|J^c|,2}} \|\mathbf{f}_J\|_{\infty} + 2M\kappa \sqrt{C_{\rho} M_{|J^c|,2}}) \frac{s^{\frac{p}{2}+1}}{n}.
 \end{aligned}$$

Therefore,

$$\begin{aligned}
 \sigma^2 &= \sum_{i=1}^n \sup_{\mathcal{Z} \setminus \{z_i\} \in \mathcal{Z}^{n-1}} E_{z_i} \|F(\mathcal{Z}) - E_{z_i}(F(\mathcal{Z}))\|_{\mathcal{K}}^2 \\
 &\leq \widehat{\sigma}_s^2 := \left(\kappa^2 \sqrt{C_{\rho} M_{|J|,2} M_{|J^c|,2}} \|\mathbf{f}_J\|_{\infty} + 2M\kappa \sqrt{C_{\rho} M_{|J^c|,2}} \right)^2 \frac{s^{p+2}}{n}.
 \end{aligned}$$

The lemma follows directly from Proposition 12. □

Proof of Proposition 10 According to the definition of $\tilde{\mathbf{f}}_{\rho,s,J}$ in (68) and $\tilde{\mathbf{L}}_J$ in (70), we have

$$\begin{aligned}
 &\tilde{\mathbf{f}}_{\rho,s,J} - \tilde{\mathbf{L}}_J \nabla_J f_{\rho} \\
 &= \int_X \int_X w(\mathbf{x} - \mathbf{u}) (f_{\rho}(\mathbf{u}) - f_{\rho}(\mathbf{x}) - \nabla_J f_{\rho}(\mathbf{x}) \cdot (\mathbf{u}_J - \mathbf{x}_J)) (\mathbf{u}_{J^c} - \mathbf{x}_{J^c}) \mathcal{K}_x d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}).
 \end{aligned}$$

Assumption 1 implies

$$|f_{\rho}(\mathbf{u}) - f_{\rho}(\mathbf{x}) - \nabla_J f_{\rho}(\mathbf{x}) \cdot (\mathbf{u}_J - \mathbf{x}_J)| \leq C'_{\rho} |\mathbf{u} - \mathbf{x}|^2, \quad \forall \mathbf{u}, \mathbf{x} \in X.$$

Hence,

$$\begin{aligned}
 \|\tilde{\mathbf{f}}_{\rho,s,J} - \tilde{\mathbf{L}}_J \nabla_J f_{\rho}\| &\leq C'_{\rho} \int_X \int_X \omega^s(\mathbf{x} - \mathbf{u}) |\mathbf{u}_{J^c} - \mathbf{x}_{J^c}| |\mathbf{u} - \mathbf{x}|^2 \|\mathcal{K}_x\|_{\mathcal{K}} d\rho_X(\mathbf{x}) d\rho_X(\mathbf{u}) \\
 &\leq \kappa C'_{\rho} C_{\rho} \int_X \int_X e^{-\frac{|\mathbf{x}-\mathbf{u}|^2}{2s^2}} |\mathbf{u}_{J^c} - \mathbf{x}_{J^c}| |\mathbf{u} - \mathbf{x}|^2 d\mathbf{x} d\rho_X(\mathbf{u}) \\
 &\leq s^{p+3} \kappa C'_{\rho} C_{\rho} M_{|J|,2} M_{|J^c|,3}.
 \end{aligned}$$

Note that $\|\tilde{Y}_J - \tilde{S}_{Z,J} \nabla_J f_\rho\| \leq \|\tilde{f}_{\rho,s,J} - \tilde{L}_J \nabla_J f_\rho\| + \|\tilde{Y}_J - \tilde{S}_{Z,J} \nabla_J f_\rho - (\tilde{f}_{\rho,s,J} - \tilde{L}_J \nabla_J f_\rho)\|$. Choose $\lambda = \tilde{C}_{M,\theta} s^{p+2+\theta}$, we have

$$\text{Prob}(\Omega_2^c) \leq \text{Prob}\{\|\tilde{Y}_J - \tilde{S}_{Z,J} \nabla_J f_\rho - (\tilde{f}_{\rho,s,J} - \tilde{L}_J \nabla_J f_\rho)\|_{\mathcal{K}} \geq s^{p+2+\theta} \kappa C'_\rho C_\rho M_{|J|,2} M_{|J^c|,3}\} \tag{81}$$

The desired result follows directly from Lemma 15. □

Proof of Proposition 11 On the event Ω_0 , we have $\frac{1}{3} D_{\min} \mathbf{I} \leq \mathbf{D}_n \leq 3 D_{\max} \mathbf{I}$. Therefore,

$$\|(\mathbf{S}_{Z,J} + \lambda \mathbf{D}_n)^{-1} (Y_J - \mathbf{S}_{Z,J} \nabla f_\rho)\|_{\mathcal{K}} \leq \frac{3}{D_{\min} \lambda} \|Y_J - \mathbf{S}_{Z,J} \nabla f_\rho\|_{\mathcal{K}}.$$

On the event $\Omega_0 \cap \Omega_4$, we have

$$\begin{aligned} & \|\tilde{S}_{Z,J} (\mathbf{S}_{Z,J} + \lambda \mathbf{D}_n)^{-1} (Y_J - \mathbf{S}_{Z,J} \nabla f_\rho)\|_{\mathcal{K}} \\ & \leq \|\tilde{S}_{Z,J}\|_{\mathcal{K}} \|(\mathbf{S}_{Z,J} + \lambda \mathbf{D}_n)^{-1} (Y_J - \mathbf{S}_{Z,J} \nabla f_\rho)\|_{\mathcal{K}} \\ & \leq \frac{3(\kappa^2 C_\rho M_{|J|,1} M_{|J^c|,1} + 1)}{D_{\min} \lambda} s^{p+2+\theta} \|Y_J - \mathbf{S}_{Z,J} \nabla f_\rho\|_{\mathcal{K}}. \end{aligned}$$

Note that $\|Y_J - \mathbf{S}_{Z,J} \nabla f_\rho\|_{\mathcal{K}} \leq \|Y_J - \mathbf{S}_{Z,J} \nabla f_\rho + \mathbf{L}_J \nabla_J f_\rho - \mathbf{f}_{\rho,s,J}\|_{\mathcal{K}} + \|\mathbf{L}_J \nabla f_\rho - \mathbf{f}_{\rho,s,J}\|_{\mathcal{K}}$. Choosing $\lambda = \tilde{C}_{M,\theta} s^{p+2+\theta}$ with $\tilde{C}_{M,\theta}$ in (35) and using Lemma 10, we have

$$\begin{aligned} & \text{Prob}(\Omega_3^c | Z \in \Omega_0 \cap \Omega_4) \\ & \leq \text{Prob}\{\|Y_J - \mathbf{S}_{Z,J} \nabla f_\rho + \mathbf{L}_J \nabla_J f_\rho - \mathbf{f}_{\rho,s,J}\|_{\mathcal{K}} \geq s^{p+2+\theta}\} \\ & \leq \text{Prob}\left\{\|Y_J - \mathbf{f}_{\rho,s,J}\|_{\mathcal{K}} \geq \frac{s^{p+2+\theta}}{2}\right\} + \text{Prob}\left\{\|\mathbf{L}_J \nabla_J f_\rho - \mathbf{S}_{Z,J} \nabla f_\rho\|_{\mathcal{K}} \geq \frac{s^{p+2+\theta}}{2}\right\}. \end{aligned}$$

Together with Lemmas 11 and 12, we get the desired result. □

Acknowledgements This work was partially supported by a grant from National Science Foundation grant and a grant from University of California.

References

Aronszajn, N. (1950). Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68, 337–404.

Bach, F. R. (2008). Consistency of the group Lasso and multiple kernel learning. *The Journal of Machine Learning Research*, 9, 1179–1225.

Belkin, M., & Niyogi, P. (2003). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6), 1373–1396.

Belkin, M., Niyogi, P., & Sindhwani, V. (2006). Manifold regularization: a geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7, 2434.

Bertin, K., & Lecué, K. (2008). Selection of variables and dimension reduction in high-dimensional non-parametric regression. *Electronic Journal of Statistics*, 2, 1224–1241.

Bickel, P., & Li, B. (2007). Local polynomial regression on unknown manifolds. In *IMS lecture notes-monograph series: Vol. 54. Complex datasets and inverse problems: tomography, networks and beyond*, (177–186).

Cai, J. F., Chan, R. H., & Shen, Z. (2008). A framelet-based image inpainting algorithm. *Applied and Computational Harmonic Analysis*, 24(2), 131–149.

- Combettes, P. L., & Wajs, V. R. (2005). Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4), 1168–1200 (electronic). doi:[10.1137/050626090](https://doi.org/10.1137/050626090).
- Cook, R. D. & Yin, X. (2001). Dimension reduction and visualization in discriminant analysis. *Australian & New Zealand Journal of Statistics*, 43(2), 147–199. doi:[10.1111/1467-842X.00164](https://doi.org/10.1111/1467-842X.00164) With a discussion by A. H. Welsh, Trevor Hastie, Mu Zhu, S. J. Sheather, J. W. McKean, Xuming He and Wing-Kam Fung and a rejoinder by the authors.
- Cucker, F., & Zhou, D. X. (2007). *Learning theory: an approximation theory viewpoint* (Vol. 24). Cambridge: Cambridge Univ Press.
- Daubechies, I., Defrise, M., & De Mol, C. (2004). An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11), 1413–1457. doi:[10.1002/cpa.20042](https://doi.org/10.1002/cpa.20042)
- Dhillon, I. S., Mallela, S., & Kumar, R. (2003). A divisive information theoretic feature clustering algorithm for text classification. *Journal of Machine Learning Research*, 3, 1265–1287.
- Do Carmo, M., & Flaherty, F. (1992). *Riemannian geometry*. Basel: Birkhauser.
- Donoho, D. L. (1995). De-noising by soft-thresholding. *IEEE Transactions on Information Theory*, 41(3), 613–627.
- Donoho, D. L. (2006). Compressed sensing. *IEEE Transactions on Information Theory*, 52(4), 1289–1306.
- Donoho, D., & Grimes, C. (2003). Hessian eigenmaps: locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10), 5591–5596.
- Efron, B., Hastie, T., Johnstone, I., & Tibshirani, R. (2004). Least angle regression. *Annals of Statistics*, 32(2), 407–499 with discussion, and a rejoinder by the authors.
- Fukumizu, K., Bach, F. R., & Jordan, M. I. (2009). Kernel dimension reduction in regression. *Annals of Statistics*, 37(4), 1871–1905. doi:[10.1214/08-AOS637](https://doi.org/10.1214/08-AOS637)
- Golub, G. H., & Van Loan, C. F. (1989). *Matrix computations*. Baltimore: Johns Hopkins University Press.
- Golub, T. R., Slonim, D. K., Tamayo, P., Huard, C., Gaasenbeek, M., Mesirov, J. P., Coller, H., Loh, M. L., Downing, J. R., Caligiuri, M. A., Bloomfield, C. D., & Lander, E. S. (1999). Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science*, 286(5439), 531–537.
- Guyon, I., Weston, J., Barnhill, S., & Vapnik, V. (2002). Gene selection for cancer classification using support vector machines. *Machine Learning*, 46(1), 389–422.
- Guyon, I., & Elisseeff, A. (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3, 1157–1182.
- Hiriart-Urruty, J., & Lemaréchal, C. (1993). *Convex analysis and minimization algorithms*. Berlin: Springer.
- Hristache, M., Juditsky, A., & Spokoiny, V. (2001). Structure adaptive approach for dimension reduction. *Annals of Statistics*, 29(6), 1537–1566.
- Lafferty, J., & Wasserman, L. (2008). Rodeo: sparse, greedy nonparametric regression. *Annals of Statistics*, 36(1), 28–63. doi:[10.1214/009053607000000811](https://doi.org/10.1214/009053607000000811)
- Langford, J., Li, L., & Zhang, T. (2009). Sparse online learning via truncated gradient. In D. Koller, D. Schuurmans, Y. Bengio, L. Bottou (Eds.) *Advances in neural information processing systems* (Vol. 21, pp. 905–912). Cambridge: MIT
- Li, K. C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414), 316–342 with discussion and a rejoinder by the author.
- Li, K. C. (1992). On principal Hessian directions for data visualization and dimension reduction: another application of Stein’s lemma. *Journal of the American Statistical Association*, 87(420), 1025–1039.
- Li, B., Zha, H., & Chiaromonte, F. (2005) Contour regression: a general approach to dimension reduction. *Ann Statist* pp 1580–1616.
- Lin, Y., & Zhang, H. H. (2006). Component selection and smoothing in multivariate nonparametric regression. *Annals of Statistics*, 34(5), 2272–2297. doi:[10.1214/009053606000000722](https://doi.org/10.1214/009053606000000722)
- Mackey, L. (2009). Deflation methods for sparse PCA. *Advances in Neural Information Processing Systems*, 21, 1017–1024.
- McDiarmid, C. (1989). On the method of bounded differences. *Surveys in Combinatorics*, 141, 148–188.
- Micchelli, C. A., & Pontil, M. (2005). On learning vector-valued functions. *Neural Computation*, 17(1), 177–204. doi:[10.1162/0899766052530802](https://doi.org/10.1162/0899766052530802)
- Micchelli, C. A., & Pontil, M. (2007). Feature space perspectives for learning the kernel. *Machine Learning*, 66, 297–319.
- Micchelli, C. A., Morales, J. M., & Pontil, M. (2010). A family of penalty functions for structured sparsity. *Advances in Neural Information Processing Systems*, 23, 1612–1623.
- Mukherjee, S. Zhou, D.X. (2006). Learning coordinate covariances via gradients. *Journal of Machine Learning Research*, 7, 519–549.
- Mukherjee, S., & Wu, Q. (2006). Estimation of gradients and coordinate covariation in classification. *Journal of Machine Learning Research*, 7, 2481–2514.

- Mukherjee, S., Wu, Q., & Zhou, D. (2010). Learning gradients on manifolds. *Bernoulli*, *16*(1), 181–207.
- Roweis, S., & Saul, L. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, *290*(5500), 2323–2326.
- Ruppert, D., & Ward, M. P. (1994). Multivariate locally weighted least squares regression. *Annals of Statistics*, *22*(3), 1346–1370.
- Samarov, A. M. (1993). Exploring regression structure using nonparametric functional estimation. *Journal of the American Statistical Association*, *88*(423), 836–847.
- Schölkopf, B., & Smola, A. (2002). *Learning with kernels: Support vector machines, regularization, optimization, and beyond*. Cambridge: MIT.
- Tenenbaum, J., Silva, V., Langford, J. (2000) A global geometric framework for nonlinear dimensionality reduction. *Science* *290*(5500): 2319–2323
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, *58*(1), 267–288.
- van der Vaart, A. W., & Wellner, J. A. (1996). *Springer Series in Statistics. Weak convergence and empirical processes*. New York: Springer. With applications to statistics.
- Vapnik, V. N. (1998). *Statistical learning theory. Adaptive and learning systems for signal processing, communications, and control*. New York: Wiley.
- Weston, J., Elisseeff, A., Schölkopf, B., & Tipping, M. (2003). Use of the zero norm with linear models and kernel methods. *Journal of Machine Learning Research*, *3*, 1439–1461.
- Xia, Y., Tong, H., Li, W. K., & Zhu, L. X. (2002). An adaptive estimation of dimension reduction space. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, *64*(3), 363–410. [10.1111/1467-9868.03411](https://doi.org/10.1111/1467-9868.03411).
- Ye, G. B. & Zhou, D. X. (2008). Learning and approximation by Gaussians on Riemannian manifolds. *Advances in Computational Mathematics*, *29*(3), 291–310.
- Zhang, T. (2004). Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*, *32*(1), 56–85. [doi:10.1214/aos/1079120130](https://doi.org/10.1214/aos/1079120130)
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, *67*(2), 301–320.
- Zou, H., Hastie, T., & Tibshirani, R. (2006). Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, *15*(2), 265–286.