

Minimax lower bounds for function estimation on graphs*

Alisa Kirichenko[†]

*Centrum Wiskunde en Informatica
P.O. Box 94079
1090 GB Amsterdam
The Netherlands
e-mail: alisa.kirichenko@cwi.nl*

and

Harry van Zanten

*Korteweg-de Vries Institute for Mathematics
University of Amsterdam
P.O. Box 94248
1090 GE Amsterdam
The Netherlands
e-mail: hvzanten@uva.nl*

Abstract: We study minimax lower bounds for function estimation problems on large graph when the target function is smoothly varying over the graph. We derive minimax rates in the context of regression and classification problems on graphs that satisfy an asymptotic shape assumption and with a smoothness condition on the target function, both formulated in terms of the graph Laplacian.

Keywords and phrases: Function estimation on graphs, minimax lower bounds.

Received September 2017.

1. Introduction

In recent years there has been substantial interest in high-dimensional estimation and prediction problems on large graphs. These can in many cases be seen as high-dimensional or nonparametric regression or classification problems in which the goal is to learn a “smooth” function on a given graph. Various methods have been proposed to deal with such problems, motivated by a variety of applications. Any sensible method employs some form of regularisation that takes the geometry of the graph into account. Examples of methods that have been considered include penalised least squares regression using a Laplacian-based penalty (e.g. Ando and Zhang (2007); Belkin et al. (2004); Kolaczyk (2009); Smola and Kondor (2003); Zhu and Hastie (2005)), penalisation using the total variation norm (e.g. Sadhanala et al. (2016)) and Bayesian regularisa-

*Research funded by the Netherlands Organisation for Scientific Research NWO.

[†]Research carried out while the author was at the University of Amsterdam.

tion (e.g. Hartog and van Zanten (2018), Bertozzi et al. (2017), Kirichenko and van Zanten (2017)).

There exist only a few papers that study theoretical aspects of the performance of nonparametric estimation procedures on graphs. Early references are Belkin et al. (2004), in which a theoretical analysis of a Tikhonov regularisation method is conducted in terms of algorithmic stability and Johnson and Zhang (2007) and Ando and Zhang (2007), who consider sub-sampling schemes for estimating a function on a graph. More recently, convergence rates have been obtained by Sadhanala et al. (2016) in the context of regression on a regular grid using total variation penalties and by Kirichenko and van Zanten (2017) for nonparametric Bayes procedures for regression and classification on more general graphs. The paper Sadhanala et al. (2016) also establishes minimax lower bounds for regression problems on grids.

In this paper we derive new minimax results for regression and binary classification on graphs, exhibiting the best possible rates that can be attained uniformly over certain classes of “smooth” functions on graphs. We consider simple undirected graphs that satisfy an assumption on their “asymptotic geometry”, formulated in terms of the graph Laplacian. This assumption, which is recalled in the next section, was introduced in Kirichenko and van Zanten (2017). The geometry assumption attaches a parameter r to a graph which essentially describes how the eigenvalues of its Laplacian behave. It was illustrated in Kirichenko and van Zanten (2017) that it is satisfied for many graphs of interest. Theoretically it can be shown to hold for instance for regular grids and tori of arbitrary dimensions. Moreover, for a given graph it can be verified empirically whether the assumption is reasonable and what the corresponding graph parameter r is. In Kirichenko and van Zanten (2017) this was done both for simulated “small world graphs” and for real protein-protein interaction graphs.

The geometry parameter r appears in the minimax lower bounds we derive. The other key ingredient is the regularity β of the function that is being estimated, defined in a suitable manner. We introduce a Sobolev-type smoothness condition on the target function using the graph Laplacian again to quantify smoothness. The geometry of the underlying graph and the smoothness of the target function together determine the minimax rate through the geometry parameter r and the smoothness parameter β . We have chosen our setup and normalisations in such a way that the optimal rates over balls of smooth functions that we obtain are of the usual form $n^{-\beta/(r+2\beta)}$. This shows that the geometry parameter r can be interpreted as some kind of “dimension” of the graph. In the regular grid case it is indeed precisely the dimension of the grid, as shown in Kirichenko and van Zanten (2017). However, the result holds for much more general graphs as well. In particular, the geometry parameter r does not need to be an integer.

For the sake of completeness we give two-sided results, that is, we also exhibit estimators that achieve the lower bounds, showing that the bounds are tight. However, these estimators are non-adaptive, in the sense that they depend on the smoothness parameter β , which will typically not be accessible in realistic settings. More interestingly perhaps, the lower bounds match the

upper bounds we obtained in Kirichenko and van Zanten (2017). This shows that the nonparametric Bayes procedures we proposed in the latter paper are smoothness-adaptive and rate-optimal. We note however that the procedure exhibited in Kirichenko and van Zanten (2017) that is adaptive on the whole range of regularities $\beta > 0$ is only rate-optimal up to a logarithmic factor. It might be of interest to study the possibility of procedures that achieve exactly the correct rate.

In the next section we introduce the general framework. Specifically, we define the geometry condition on the graph and the smoothness condition on the target function. In Section 3 we describe the regression and classification problems on a graph and present our results on the minimax rates for those problems. The mathematical proofs are given in Section 4.

2. Setting

Let G be a connected simple undirected graph with vertices labelled $\{1, \dots, n\}$. Let A be its adjacency matrix, i. e. A_{ij} is 1 or 0 according to whether or not there is an edge between vertices i and j . Let D be the diagonal matrix with element D_{ii} equal to the degree of vertex i . Let $L = D - A$ be the Laplacian of the graph.

A function f on the (vertices of the) graph is simply a function $f : \{1, \dots, n\} \rightarrow \mathbb{R}$. We measure distances and norms of functions using the norm $\|\cdot\|_n$ defined by $\|f\|_n^2 = n^{-1} \sum_{i=1}^n f^2(i)$. The corresponding inner product of two functions f and g is denoted by

$$\langle f, g \rangle_n = \frac{1}{n} \sum_{i=1}^n f(i)g(i).$$

The Laplacian is nonnegative definite (Cvetković et al. (2010)). Hence we can order the Laplacian eigenvalues by magnitude and denote them by

$$0 = \lambda_{n,0} < \lambda_{n,1} \leq \lambda_{n,2} \leq \dots \leq \lambda_{n,n-1}.$$

(The smallest one always equals 0 and since the graph is connected, the second one is positive, see for instance Cvetković et al. (2010)). We fix a corresponding sequence of eigenfunctions ψ_i , orthonormal with respect to the inner product $\langle \cdot, \cdot \rangle_n$.

We derive our results under an asymptotic geometry assumption on the graph, first introduced in Kirichenko and van Zanten (2017), formulated in terms of the Laplacian eigenvalues.

Assumption. *We say that the geometry assumption is satisfied with parameter $r \geq 1$ if there exist $i_0 \in \mathbb{N}$, $\kappa \in (0, 1]$ and $C_1, C_2 > 0$ such that for all n large enough,*

$$C_1 \left(\frac{i}{n}\right)^{2/r} \leq \lambda_i \leq C_2 \left(\frac{i}{n}\right)^{2/r}, \quad \text{for all } i \in \{i_0, \dots, \kappa n\}.$$

Very roughly speaking the condition means that asymptotically, or from “far away”, the graph looks like an r -dimensional grid with n vertices. From Kirichenko and van Zanten (2017) we know the assumption is satisfied for d -dimensional grids with r equal to the dimension d , hence our results on the minimax rates include the usual statements for regression and classification with regular, fixed design. We stress however that the constant r does not need to be a natural number. For given graphs the parameter r can be calculated numerically. For example, in Kirichenko and van Zanten (2017) a Watts-Strogatz “small world” graph is considered which satisfies the condition with r equal to 1.4. Observe that we do not assume the existence of a “limiting manifold” for the graph as $n \rightarrow \infty$. See Kirichenko and van Zanten (2017) for more discussion of the geometry assumption and more examples.

We describe the smoothness of the function of interest by assuming it belongs to a Sobolev-type ball of the form

$$H^\beta(Q) = \left\{ f : \left\langle f, (I + (n^{\frac{2}{r}}L)^\beta)f \right\rangle_n \leq Q^2 \right\} \quad (2.1)$$

for some $\beta, Q > 0$ (independent of n). This should be viewed as the natural discrete graph version of the usual notion of a Sobolev ball of functions on $[0, 1]^r$. (This is most easily seen in the case of the path graph and $r = 1$, as illustrated in Example 3.1 of Kirichenko and van Zanten (2017).) The particular normalisation, which depends on the geometry parameter r , ensures non-trivial asymptotics. Again, we stress that we do not assume that the functions on the graph are discretised versions of certain continuous objects on a “limiting manifold”.

3. Main results

Now that we have introduced ways to quantify the graph geometry and the regularity of the target function, we can formulate our minimax results for regression and classification. In both cases, $G = G_n$ will be a connected simple undirected graph with vertices $1, \dots, n$, satisfying the geometry assumption for $r \geq 1$. The target function will be a regression function that is observed with additive Gaussian noise in the regression case and a binary regression function in the classification case.

In the regression case we assume that we have observations $Y = (Y_1, \dots, Y_n)$ at the vertices of the graph satisfying

$$Y_i = f(i) + \sigma \xi_i, \quad i = 1, \dots, n, \quad (3.1)$$

where the ξ_i are independent standard Gaussians, $\sigma > 0$ and $f : \{1, \dots, n\} \rightarrow \mathbb{R}$ is the unknown function of interest. We denote the corresponding distribution of Y by \mathbb{P}_f and the associated expectation by \mathbb{E}_f . Our main result in this setting is the following.¹

¹We write $a_n \asymp b_n$ if $0 < \liminf a_n/b_n \leq \limsup a_n/b_n < \infty$.

Theorem 3.1 (Regression). *Suppose that the graph satisfies the geometry assumption for $r \geq 1$. Then for all $\beta, Q > 0$*

$$\inf_{\hat{f}} \sup_{f \in H^\beta(Q)} \mathbb{E}_f \|\hat{f} - f\|_n^2 \asymp n^{-2\beta/(2\beta+r)}, \tag{3.2}$$

where the infimum is taken over all estimators $\hat{f} = \hat{f}(Y_1, \dots, Y_n)$.

The theorem shows that the minimax rate for the regression problem on the graph is equal to $n^{-\beta/(2\beta+r)}$. We obtain the upper bound on the rate by constructing a projection estimator \tilde{f} for which

$$\sup_{f \in H^\beta(Q)} \mathbb{E}_f \|\tilde{f} - f\|_n^2 \lesssim n^{-2\beta/(2\beta+r)}.$$

The proof shows that this estimator depends on the regularity level of the target function and therefore is not adaptive. An adaptive (Bayesian) procedure is exhibited in Kirichenko and van Zanten (2017).

In the binary classification case we assume that the data Y_1, \dots, Y_n are independent $\{0, 1\}$ -valued variables, observed at the vertices of the graph. In this case the goal is to estimate the binary regression function ρ , or “soft label function” on the graph, defined by

$$\rho(i) = \mathbb{P}(Y_i = 1).$$

The function ρ , of course, determines the distribution of the data, which we therefore denote by \mathbb{P}_ρ . Again, the associated expectations are denoted by \mathbb{E}_ρ .

Technically the classification case is slightly more demanding. Different from the regression case we also have to impose conditions on the Laplacian eigenfunctions ψ_j in this case. Moreover, we impose the regularity condition not directly on the binary regression function ρ , but on a suitably transformed version of it, so that it maps into \mathbb{R} instead of $(0, 1)$. Concretely, we fix a differentiable link function $\Psi : \mathbb{R} \rightarrow (0, 1)$ such that $\Psi'/(\Psi(1 - \Psi))$ is uniformly bounded, and $\Psi' > 0$ everywhere. Note that for instance the sigmoid, or logistic link $\Psi(f) = 1/(1 + \exp(-f))$ satisfies this condition. Under these conditions the inverse $\Psi^{-1} : (0, 1) \rightarrow \mathbb{R}$ is well defined. In the classification setting the regularity condition will be formulated in terms of $\Psi^{-1}(\rho)$.

Recall from Section 2 that ψ_j is a sequence of eigenfunctions (or vectors) of the graph Laplacian L that is orthonormal with respect to the norm $\|\cdot\|_n$. In particular, the eigenfunctions ψ_j are normalised such that

$$\frac{1}{n} \sum_{i=1}^n \psi_j^2(i) = 1$$

for every j . In the following theorem we assume in addition that they are uniformly bounded by a common constant $C > 0$, which is independent of n , i.e.

that $|\psi_j(i)| \leq C$ for every i and j . We need this technical assumption in the proof of the lower bound.

Theorem 3.2 (Classification). *Suppose that the graph satisfies the geometry assumption for $r \geq 1$ and that the Laplacian eigenfunctions are uniformly bounded by a common constant $C > 0$, independent of n . Let $\Psi : \mathbb{R} \rightarrow (0, 1)$ be a differentiable link function as above with inverse Ψ^{-1} . Then for $\beta \geq r/2$ and $Q > 0$*

$$\inf_{\hat{\rho}} \sup_{\rho: \Psi^{-1}(\rho) \in H^\beta(Q)} \mathbb{E}_\rho \|\hat{\rho} - \rho\|_n^2 \asymp n^{-2\beta/(2\beta+r)},$$

where the infimum is taken over all estimators $\hat{\rho} = \hat{\rho}(Y_1, \dots, Y_n)$.

Also in this case the proof of the theorem provides an estimator that achieves the lower bound, but that estimator depends on β and hence is non-adaptive. Adaptive, rate-optimal (Bayes) procedures have been exhibited for this classification setting as well in Kirichenko and van Zanten (2017). Note that compared to the regression case, there is an extra technical requirement $\beta \geq r/2$. Additionally, we assume boundedness of the Laplacian eigenfunctions. In principle, for a specific case this can be verified numerically. For regular grids of arbitrary dimensions it is straightforward to see that this condition is fulfilled.

Indeed, by Merris (1998), for instance, the (unnormalised) eigenvectors of the graph Laplacian of the path graph are given by

$$\tilde{\psi}_j(i) = \cos(\pi i j/n - \pi j/2n).$$

For the $\|\cdot\|_n$ -norm of the j th eigenvector we then have

$$\begin{aligned} \|\tilde{\psi}_j\|_n^2 &= \frac{1}{n} \sum_{i=1}^n \cos^2(\pi i j/n - \pi j/2n) = \\ &= \frac{1}{2} + \frac{1}{4n \sin(\pi j/n)} \sum_{i=1}^n 2 \sin(\pi j/n) \cos((2i-1)\pi j/n). \end{aligned}$$

By well known trigonometric identities we have, for any $x \in \mathbb{R}$,

$$\sum_{i=1}^n 2 \sin x \cos(2ix - x) = \sum_{i=1}^n (\sin 2ix - \sin(2ix - 2x)) = \sin 2nx.$$

It follows that for any $j = 1, \dots, n-1$,

$$\|\tilde{\psi}_j\|_n^2 = \frac{1}{2} + \frac{1}{4n} \frac{\sin 2\pi j}{\sin \pi j/n} = \frac{1}{2}.$$

Notice that $\|\tilde{\psi}_0\|_n^2 = 1$. So we see that indeed, the normalised eigenvectors ψ_j of the Laplacian of the path graph are uniformly bounded by a common constant. The eigenvectors of a Cartesian product of two graphs are equal to the Kronecker products of pairs of eigenvectors associated with the Laplacians of those graphs. Since the grids of higher dimensions are products of path graphs, they satisfy the condition as well.

4. Proofs

4.1. Proof of Theorem 3.1

In the regression case we first expand the observations in the eigenbasis of the graph Laplacian, which brings the problem into the setting of the sequence formulation of the white noise model. Then we adapt techniques from the proof of Pinsker’s theorem as can it is given, for instance, in Tsybakov (2009).

4.1.1. Preliminaries

Let $Y = (Y_1, \dots, Y_n)$, $\xi = (\xi_1, \dots, \xi_n)$, and let ψ_i be the orthonormal eigenfunctions of the graph Laplacian. Denote $\tilde{\xi}_i = \langle \xi, \psi_i \rangle_n$ and observe that the $\tilde{\xi}_i$ are centred Gaussian with

$$\mathbb{E} \tilde{\xi}_i \tilde{\xi}_j = \frac{1}{n} \delta_{ij}.$$

The inner products $Z_i = \langle Y, \psi_i \rangle_n$ satisfy the following relation for $i = 0, \dots, n-1$

$$Z_i = \langle Y, \psi_i \rangle_n = f_i + \sigma \tilde{\xi}_i,$$

where f_i are coefficients in the decomposition of the target function $f_0 = \sum_{i=0}^{n-1} f_i \psi_i$. Additionally, consider the decomposition of an estimator $\hat{f} = \sum_{i=0}^{n-1} \hat{f}_i \psi_i$. Then

$$\|\hat{f} - f_0\|_n^2 = \left\langle \sum_{i=0}^{n-1} (\hat{f}_i - f_i) \psi_i, \sum_{i=0}^{n-1} (\hat{f}_i - f_i) \psi_i \right\rangle_n = \sum_{i=0}^{n-1} (\hat{f}_i - f_i)^2.$$

Hence, the minimax rates for the original problem are of the same order as the minimax rates for the problem of recovering $f = (f_0, \dots, f_{n-1})$, given the observations

$$Z_i = f_i + \varepsilon \zeta_i, \tag{4.1}$$

where ζ_i are independent standard Gaussian and $\varepsilon = \frac{\sigma}{\sqrt{n}}$. To avoid confusion we define general ellipsoids on the space of coefficients for an arbitrary sequence $a_j > 0$ and some finite constant $Q > 0$

$$B_n(Q) = \{f \in \mathbb{R}^n : \sum_{j=0}^{n-1} a_j^2 f_j^2 \leq Q\}. \tag{4.2}$$

For a function f in the Sobolev-type ball $H^\beta(Q)$ its vector of coefficients belongs to $B_n(Q)$ with

$$a_j^2 = 1 + \lambda_j^{2\beta/r} n^{2\beta/r}, \quad j = 0, \dots, n-1.$$

In order to prove the theorem it is sufficient to show that

$$\inf_{\hat{f}} \sup_{f \in B_n(Q)} \mathbb{E}_f \left(\sum_{i=0}^{n-1} (\hat{f}_i - f_i)^2 \right) \asymp n^{-2\beta/(2\beta+r)}. \tag{4.3}$$

We are going to follow the proof of Pinsker's theorem (see for example Tsybakov (2009)) which studies a similar case in the setting of the Gaussian white noise model on the interval $[0, 1]$. The proof requires some modifications arising from the nature of our problem. The main differences with the usual lower bound result over Sobolev balls in the infinite sequence model are that we only have n observations and that our ellipsoids $B_n(Q)$ have a special form.

In order to proceed we first consider the problem of obtaining minimax rates in the class of linear estimators. We introduce Pinsker's estimator and recall the linear minimax lemma showing that Pinsker's estimator is optimal in the class of linear estimators. The risk of a linear estimator $\hat{f}(l) = (l_0 Z_0, \dots, l_{n-1} Z_{n-1})$ with $l = (l_0, \dots, l_{n-1}) \in \mathbb{R}^n$ is given by

$$R(l, f) = \mathbb{E}_f \sum_{j=0}^{n-1} (\hat{f}_j - f_j)^2 = \sum_{j=0}^{n-1} ((1 - l_j)^2 f_j^2 + \varepsilon^2 l_j^2).$$

For large n we introduce the following equation with respect to the variable x

$$\frac{\varepsilon^2}{x} \sum_{j=0}^{n-1} a_j (1 - x a_j)_+ = Q. \quad (4.4)$$

Suppose, there exists a unique solution x of (4.4). For such a solution, define a vector of coefficients l' consisting of entries

$$l'_j = (1 - x a_j)_+. \quad (4.5)$$

The linear estimator $\tilde{f} = \tilde{f}(l')$ is called the Pinsker estimator for the general ellipsoid $B_n(Q)$. The following lemma, which appears as Lemma 3.2 in Tsybakov (2009), shows that the Pinsker estimator is a linear minimax estimator.

Lemma 4.1 (Linear minimax lemma). *Suppose that $B_n(Q)$ is a general ellipsoid defined by (4.2) with $Q > 0$ and a positive set of coefficients $\{a_j\}_{j=0}^{n-1}$. Suppose there exists a unique solution x of (4.4) and suppose that the associated coefficients l'_j defined by (4.5) satisfy*

$$S = \varepsilon^2 \sum_{j=0}^{n-1} l'_j < \infty. \quad (4.6)$$

Then the linear minimax risk satisfies

$$\inf_{l \in \mathbb{R}^n} \sup_{f \in B_n(Q)} R(l, f) = \sup_{f \in B_n(Q)} R(l', f) = S. \quad (4.7)$$

In order to be able to apply Lemma 4.1 in our graph setting we need the following technical lemma.

Lemma 4.2. *Consider the ellipsoid $B_n(Q)$ defined by (4.2) with $Q > 0$ and*

$$a_j^2 = 1 + \lambda_j^{2\beta/r} n^{2\beta/r}, \quad j = 0, \dots, n-1.$$

Then, as $n \rightarrow \infty$, we have the following

(i) There exists a solution x of (4.4) which is unique and satisfies

$$x \asymp n^{-\beta/(2\beta+r)}.$$

(ii) The associated sum (4.6) of the coefficients of the Pinsker's estimator satisfies

$$S \asymp n^{-2\beta/(2\beta+r)}.$$

(iii) For $\varepsilon = \frac{\sigma}{\sqrt{n}}$ define $v_j = \frac{\varepsilon^2(1-xa_j)_+}{xa_j}$. Then

$$\max_{j=0, \dots, n-1} v_j^2 a_j^2 = O\left(n^{-r/(2\beta+r)}\right).$$

Proof. (i) According to Lemma 3.1 from Tsybakov (2009) for large enough n and for an increasing positive sequence a_j with $a_n \rightarrow +\infty$, as $n \rightarrow \infty$, there exists a unique solution of (4.4) given by

$$x = \frac{\varepsilon^2 \sum_{j=0}^{N-1} a_j}{Q + \varepsilon^2 \sum_{j=0}^{N-1} a_j^2},$$

where

$$N = \max \left\{ m : \varepsilon^2 \sum_{j=0}^{m-1} a_j (a_m - a_j) < Q \right\} < +\infty.$$

Consider N defined above. Denote

$$\begin{aligned} A_m &= \varepsilon^2 \sum_{j=0}^{m-1} a_j (a_m - a_j) = \\ &= n^{\beta/r-1} \sigma^2 \sum_{j=0}^{m-1} \sqrt{\left(1 + \lambda_j^{2\beta/r} n^{2\beta/r}\right) \left(\lambda_m^{2\beta/r} n^{2\beta/r} - \lambda_j^{2\beta/r} n^{2\beta/r}\right)}. \end{aligned}$$

Since the geometry condition on the graph is satisfied, for $j = i_0, \dots, \kappa n$ the eigenvalues of the graph Laplacian can be bounded in a following way

$$C_1 \left(\frac{j}{n}\right)^{2/r} \leq \lambda_j \leq C_2 \left(\frac{j}{n}\right)^{2/r}.$$

If $m < \kappa n$, then for some $K_1 > 0$

$$A_m \geq n^{-1} \sigma^2 C_1 \sum_{j=i_0}^{m-1} j^{\beta/r} \left(C_1 m^{\beta/r} - C_2 j^{\beta/r}\right) \geq K_1 n^{-1} m^{\beta/r} \sum_{j=i_0}^{m-1} j^{\beta/r}.$$

Since A_m is an increasing function of m , there exists $K_2 > 0$ such that for all

$$m > K_2 n^{r/(2\beta+r)}$$

it holds that $A_m > Q$. In a similar manner we can show that there exists $K_3 > 0$ such that for all

$$m < K_3 n^{r/(2\beta+r)}$$

it holds that $A_m < Q$.

This leads us to the conclusion that $N \asymp n^{r/(2\beta+r)}$. Then equation (4.4) has a unique solution that satisfies

$$x = \frac{\sigma^2}{n \left(Q + \frac{\sigma^2}{n} \sum_{j=0}^{N-1} a_j^2 \right)} \sum_{j=0}^{N-1} a_j \asymp \frac{1}{n} N^{1+\beta/r} \asymp n^{-\beta/(2\beta+r)}.$$

(ii) Since graph G satisfies the geometry assumption, we deduce from (i) that for some $K_4 > 0$

$$l'_j \asymp \left(1 - K_4 n^{-\beta/(2\beta+r)} j^{\beta/r} \right)_+, \text{ for } j = i_0, \dots, N.$$

For $j = 0, \dots, i_0 - 1$ we bound l'_j from above by 1. Then

$$S \lesssim n^{-1} i_0 + n^{-1} \sum_{j=i_0}^{N-1} \left(1 - K_4 n^{-\beta/(2\beta+r)} j^{\beta/r} \right)_+ \lesssim n^{-1} N \lesssim n^{-2\beta/(2\beta+r)}.$$

On the other hand,

$$S \gtrsim n^{-1} \sum_{j=i_0}^{N-1} \left(1 - K_4 n^{-\beta/(2\beta+r)} j^{\beta/r} \right)_+ \gtrsim n^{-2\beta/(2\beta+r)}.$$

(iii) Note that for $j > N$ we have $v_j^2 = 0$. We also know that $a_N < \frac{1}{x}$. Then

$$v_j^2 a_j^2 = \frac{\sigma^2 a_j (1 - x a_j)_+}{n x} \leq \frac{\sigma^2 a_N}{n x} \leq \frac{\sigma^2}{n x^2}.$$

Hence, as $n \rightarrow \infty$,

$$\max_{j=0, \dots, n-1} v_j^2 a_j^2 = O \left(n^{-r/(2\beta+r)} \right).$$

This finishes the proof of the lemma. □

4.1.2. Proof of the upper bound on the risk

Recall that we only need to provide the upper bound in (4.3). Consider the Pinsker estimator $\tilde{f} = (l'_0 Z_0, \dots, l'_{n-1} Z_{n-1})$ with

$$l'_j = (1 - x a_j)_+,$$

where $a_j^2 = 1 + \lambda_j^{2\beta/r} n^{2\beta/r}$ and x is a unique solution of (4.4). Using Lemma 4.1 and Lemma 4.2 we conclude that the Pinsker's estimator satisfies

$$\sup_{f \in B_n(Q)} \mathbb{E}_f \left(\sum_{i=0}^{n-1} (\tilde{f}_i - f_i)^2 \right) \lesssim n^{-2\beta/(2\beta+r)}.$$

4.1.3. Proof of the lower bound on the risk

We follow the usual steps of the general reduction scheme for obtaining minimax rates (see e.g. Chapter 3 of Tsybakov (2009) for details). First, we note that it is sufficient to only take into account the first N coefficients in the decomposition of the target function, where

$$N = \max \left\{ m : \varepsilon^2 \sum_{j=0}^{m-1} a_j (a_m - a_j) < Q \right\}$$

with $a_j^2 = 1 + \lambda_j^{2\beta/r} n^{2\beta/r}$. Indeed, if we denote the minimax risk by R_n , i.e.

$$R_n = \inf_{\hat{f}} \sup_{f \in B_n(Q)} \mathbb{E}_f \left(\sum_{i=0}^{n-1} (\hat{f}_i - f_i)^2 \right),$$

and define

$$B_n(Q, N) = \{ f^{(N)} = (f_0, \dots, f_{N-1}, 0, \dots, 0) \in \mathbb{R}^n : \sum_{j=0}^{N-1} a_j^2 f_j^2 \leq Q \},$$

then we have $B_n(Q, N) \subseteq B_n(Q)$ and hence

$$R_n \geq \inf_{\hat{f}^{(N)} \in B_n(Q, N)} \sup_{f^{(N)} \in B_n(Q, N)} \mathbb{E}_f \sum_{j=0}^{N-1} (\hat{f}_j - f_j)^2. \tag{4.8}$$

Next we follow the usual step of bounding the minimax risk by a Bayes risk. Consider the density $\mu(f^{(N)}) = \prod_{j=0}^{N-1} \mu_{s_j}(f_j)$ with respect to the Lebesgue measure on \mathbb{R}^N . Here $s_j = (1 - \delta)v_j^2$ for some $\delta \in (0, 1)$ and μ_σ denotes the density of the Gaussian distribution with mean 0 and variance σ^2 . By (4.8) we can bound the minimax risk from below by the Bayes risk

$$R_n \geq \inf_{\hat{f}^{(N)} \in B_n(Q, N)} \sum_{j=0}^{N-1} \int_{B_n(Q, N)} \mathbb{E}_f (\hat{f}_j - f_j)^2 \mu(f^{(N)}) df^{(N)} \geq I^* - r^*, \tag{4.9}$$

where

$$I^* = \inf_{\hat{f}^{(N)} \in B_n(Q, N)} \sum_{j=0}^{N-1} \int_{\mathbb{R}^N} \mathbb{E}_f (\hat{f}_j - f_j)^2 \mu(f^{(N)}) df^{(N)};$$

$$r^* = \sup_{\hat{f}^{(N)} \in B_n(Q, N)} \sum_{j=0}^{N-1} \int_{B_n(Q, N)^c} \mathbb{E}_f (\hat{f}_j - f_j)^2 \mu(f^{(N)}) df^{(N)}$$

with $B(Q, N)^c = \mathbb{R}^N \setminus B_n(Q, N)$. From the proof of Pinsker's theorem, see Tsybakov (2009), we get the following bounds

$$I^* \gtrsim S,$$

$$r^* \lesssim \exp \left\{ -K \left(\max_{j=0, \dots, n-1} v_j^2 a_j^2 \right)^{-1} \right\}$$

for some $K > 0$. Using the results of Lemma 4.2 we conclude that $R_n \gtrsim n^{-2\beta/(2\beta+r)}$.

4.2. Proof of Theorem 3.2

In order to prove the result in the classification case we use Fano's lemma and the usual general scheme for reducing a minimax estimation problem to a minimax testing problem. (see for instance Tsybakov (2009) again).

4.2.1. Proof of the upper bound on the risk

We define the estimator that gives us an upper bound on the minimax risk based on the estimator \tilde{f} , which has been introduced in subsection 4.1.2 of the proof of Theorem 3.1. Consider the estimator

$$\tilde{\rho} = \Psi \left(\sum_{i=0}^{n-1} \tilde{f}_i \psi_i \right).$$

By the reasoning given in the aforementioned subsection and using the properties of the link function Ψ , we can see that

$$\begin{aligned} \sup_{\rho \in \{\rho: \Psi^{-1}(\rho) \in H^\beta(Q)\}} \mathbb{E}_\rho \|\tilde{\rho} - \rho\|_n^2 &\lesssim \\ &\lesssim \sup_{f \in B_n(Q)} \mathbb{E}_f \left(\sum_{i=0}^{n-1} (\tilde{f}_i - f_i)^2 \right) \lesssim n^{-2\beta/(2\beta+r)}. \end{aligned}$$

4.2.2. Proof of the lower bound on the risk

The proof of the lower bound on the risk is based on a corollary of Fano's lemma, see, for instance, Corollary 2.6 in Tsybakov (2009). Observe that by Markov's inequality for any soft label functions ρ_0, \dots, ρ_M with $M \in \mathbb{N}$ there exists $C > 0$ such that

$$\begin{aligned} \inf_{\hat{\rho}} \sup_{\rho \in \{\rho: \Psi^{-1}(\rho) \in H^\beta(Q)\}} \mathbb{E}_\rho n^{2\beta/(2\beta+r)} \|\hat{\rho} - \rho\|_n^2 &\gtrsim \\ &\gtrsim \inf_{\hat{\rho}} \max_{\rho \in \{\rho_0, \dots, \rho_M\}} \mathbb{P}_\rho \left(\|\hat{\rho} - \rho\|_n^2 \geq n^{-2\beta/(2\beta+r)} \right). \end{aligned}$$

Consider probability measures P_0, P_1, \dots, P_M corresponding to the soft label functions ρ_0, \dots, ρ_M . For a test $\phi : \mathbb{R}^n \rightarrow \{0, 1, \dots, M\}$ define the average probability of error by

$$\bar{p}_M(\phi) = \frac{1}{M+1} \sum_{j=0}^M P_j(\phi \neq j).$$

Additionally, let

$$\bar{p}_M = \inf_{\phi} \bar{p}_M(\phi).$$

From the general scheme for obtaining lower bounds for minimax risk (for more detail see Chapter 2 of Tsybakov (2009)) we know that if ρ_0, \dots, ρ_M are such that for any pair $i, j \in \{0, \dots, M\}$

$$\|\rho_i - \rho_j\|_n \gtrsim n^{-2\beta/(2\beta+r)}, \text{ when } i \neq j, \tag{4.10}$$

then

$$\inf_{\hat{\rho}} \max_{\rho \in \{\rho_0, \dots, \rho_M\}} \mathbb{P}_{\rho} \left(\|\hat{\rho} - \rho\|_n^2 \geq Cn^{-2\beta/(2\beta+r)} \right) \gtrsim \bar{p}_M.$$

Also, Corollary 2.6 in Tsybakov (2009) states that if P_0, P_1, \dots, P_M satisfy

$$\frac{1}{M+1} \sum_{j=1}^M K(P_j, P_0) \leq \alpha \log M, \tag{4.11}$$

for some $0 < \alpha < 1$, then

$$\bar{p}_M \geq \frac{\log(M+1) - \log 2}{\log M} - \alpha.$$

Here $K(\cdot, \cdot)$ is the Kullback–Leibler divergence. Hence, if we construct the probability measures P_0, P_1, \dots, P_M corresponding to some soft label functions ρ_0, \dots, ρ_M for which (4.10) and (4.11) hold, we will have

$$\inf_{\hat{\rho}} \sup_{\rho \in \{\rho : \Psi^{-1}(\rho) \in H^{\beta}(Q)\}} \mathbb{E}_{\rho} n^{2\beta/(2\beta+r)} \|\hat{\rho} - \rho\|_n^2 \gtrsim \frac{\log(M+1) - \log 2}{\log M} - \alpha.$$

If $M \rightarrow \infty$, as $n \rightarrow \infty$, the required result follows.

Let $N = n^{r/(2\beta+r)}$ and let $\psi_0, \dots, \psi_{N-1}$ to be an orthonormal eigenbasis of the graph Laplacian L with respect to the $\|\cdot\|_n$ -norm. For $\delta > 0$ and $\theta = (\theta_0, \dots, \theta_{N-1}) \in \{\pm 1\}^N$ define

$$f_{\theta} = \delta N^{-(2\beta+r)/2r} \sum_{j=0}^{N-1} \theta_j \psi_j.$$

We will select M vectors of coefficients $\theta^{(j)}$ such that the probability measures corresponding to $\rho_j = \Psi(f_{\theta^{(j)}})$ will satisfy (4.11), where Ψ is the link function.

Observe that for small enough $\delta > 0$ functions f_θ belong to the class $H^\beta(Q)$. Indeed, using the geometry assumption we obtain

$$\begin{aligned} \left\langle f_\theta, (I + (n^{\frac{2}{r}}L)^\beta)f_\theta \right\rangle_n &= \delta^2 N^{-(2\beta+r)/r} \sum_{j=0}^{N-1} (1 + n^{2\beta/r} \lambda_j^\beta) \leq \\ &\leq \delta^2 N^{-(2\beta+r)/r} \left(N + C_2 i_0^{(2\beta+r)/r} + C_2 \sum_{j=i_0}^{N-1} j^{2\beta/r} \right) \leq K_1 \delta^2 \end{aligned}$$

for some constant $K_1 > 0$.

We pick a subset $\{\theta^{(1)}, \dots, \theta^{(M)}\}$ of $\{\pm 1\}^N$ such that for any pair $i, j \in \{1, \dots, M\}$ such that $i \neq j$ the vectors from the subset were sufficiently distant from each other

$$d_h(\theta^{(i)}, \theta^{(j)}) \gtrsim N, \tag{4.12}$$

where $d_h(\theta, \theta') = \sum_{j=0}^{N-1} \mathbb{1}_{\theta_j \neq \theta'_j}$ is the Hamming distance. By the Varshamov–Gilbert bound (see for example Lemma 2.9 in Tsybakov (2009)) we know that there exist such a subset $\{\theta^{(1)}, \dots, \theta^{(M)}\}$, and the size M of this subset satisfies

$$M \geq b^N$$

for some $1 < b < 2$. Let $\theta^{(0)} = (0, \dots, 0) \in \mathbb{R}^N$. We define a set of probability measures $\{P_0, \dots, P_M\}$ by setting $P_j = P_{\rho_j}$, where $\rho_j = \Psi(f_{\theta^{(j)}})$.

In order to show that the P_j satisfy (4.11) we present a technical lemma. In the classification setting the Kullback–Leibler divergence $K(\cdot, \cdot)$ satisfies

$$K(P_\rho, P_{\rho'}) = \sum_{i=1}^n \left(\rho(i) \log \frac{\rho(i)}{\rho'(i)} - (1 - \rho(i)) \log \frac{1 - \rho(i)}{1 - \rho'(i)} \right).$$

Lemma 4.3. *If $\frac{\Psi'}{\Psi(1-\Psi)}$ is bounded, then there exists $c > 0$ such that for any $v_1, v_2 \in \mathbb{R}^n$ we have*

$$K(P_{\Psi(v_1)}, P_{\Psi(v_2)}) \leq nc \|v_1 - v_2\|_n^2.$$

Proof. For every $x \in \mathbb{R}$ consider the function $g_x : \mathbb{R} \rightarrow \mathbb{R}$ defined as

$$g_x(y) = \Psi(x) \log \frac{\Psi(x)}{\Psi(y)} + (1 - \Psi(x)) \log \frac{1 - \Psi(x)}{1 - \Psi(y)}.$$

We see that $g'_x(y) = \frac{\Psi'(y)}{\Psi(y)(1-\Psi(y))} (\Psi(y) - \Psi(x))$. Then by Taylor’s theorem we can see that

$$|g_x(y)| \leq \sup_{v \in [x, y] \cup [y, x]} \left| \frac{\Psi'(v)}{\Psi(v)(1-\Psi(v))} \right| \sup_{v \in [x, y] \cup [y, x]} |\Psi'(v)| (x - y)^2.$$

The statement of the lemma follows. □

By Lemma 4.3, we obtain for some $K_2 > 0$

$$K(P_j, P_0) \leq K_2 n \|f_{\theta^{(j)}} - 0\|_n^2 = 4K_2 \delta^2 n N^{-2\beta/r},$$

since

$$\|f_\theta - f_{\theta'}\|_n^2 = 4\delta^2 N^{-(2\beta+r)/r} d_n(\theta, \theta'). \tag{4.13}$$

Observe that this bound does not depend on j . Hence,

$$\frac{1}{M+1} \sum_{j=1}^M K(P_j, P_0) \leq K_2 \delta^2 n N^{-2\beta/r} = K_2 \delta^2 \log M.$$

We can choose $\delta > 0$ to be small enough such that the condition (4.11) is satisfied with some $0 < \alpha < 1$.

To finish the proof of the theorem we need to show that ρ_0, \dots, ρ_M satisfy (4.10). From (4.12) and (4.13) we get

$$\|f_{\theta^{(i)}} - f_{\theta^{(j)}}\|_n \gtrsim n^{-\beta/(2\beta+r)}.$$

Moreover, by the assumption of the theorem we have for any $j = 1, \dots, M$

$$\max_{i=1, \dots, n} |f_{\theta^{(j)}}(i)| \lesssim N^{1-(2\beta+r)/2r} \max_{i=1, \dots, n} |\psi_j(i)| \lesssim N^{(r-2\beta)/2r}.$$

For $\beta \geq r/2$ the norm is then bounded by some constant, which does not depend on n or j . Hence, there exists $K_3 \geq 0$ such that for every $i = 1, \dots, n$ and every $j = 1, \dots, M$

$$|f_{\theta^{(j)}}(i)| \leq K_3.$$

Observe that since $\Psi'(x) \neq 0$ for any $x \in \mathbb{R}$, there exists $K_4 > 0$ such that for any $x, y \in [-K_3, K_3]$

$$|\Psi(x) - \Psi(y)| \geq K_4 |x - y|.$$

Thus, for any pair $i, j \in \{1, \dots, M\}$ such that $i \neq j$

$$\|\rho_i - \rho_j\|_n = \|\Psi(f_{\theta^{(i)}}) - \Psi(f_{\theta^{(j)}})\|_n \gtrsim n^{-\beta/(2\beta+r)}.$$

This completes the proof of the theorem.

References

- Ando, R. K. and Zhang, T. (2007). Learning on graph with laplacian regularization. *Advances in neural information processing systems* **19**, 25.
- Belkin, M., Matveeva, I. and Niyogi, P. (2004). Regularization and semi-supervised learning on large graphs. In *COLT*, volume 3120, pp. 624–638. Springer. [MR2177939](#)
- Bertozzi, A. L., Luo, X., Stuart, A. M. and Zygalakis, K. C. (2017). Uncertainty Quantification in the Classification of High Dimensional Data. *ArXiv e-prints*.

- Cvetković, D., Rowlinson, P. and Simić, S. (2010). An introduction to the theory of graph spectra. *Cambridge-New York*. [MR2571608](#)
- Hartog, J. and van Zanten, J. H. (2018). Nonparametric Bayesian label prediction on a graph. *Computational Statistics & Data Analysis* **120**, 111–131. [MR3742212](#)
- Johnson, R. and Zhang, T. (2007). On the effectiveness of Laplacian normalization for graph semi-supervised learning. *Journal of Machine Learning Research* **8**(4). [MR2332439](#)
- Kirichenko, A. and van Zanten, J. H. (2017). Estimating a smooth function on a large graph by Bayesian Laplacian regularisation. *Electron. J. Statist.* **11**(1), 891–915. [MR3629018](#)
- Kolaczyk, E. D. (2009). *Statistical analysis of network data*. Springer. [MR2724362](#)
- Merris, R. (1998). Laplacian graph eigenvectors. *Linear Algebra Appl.* **278**(1–3), 221–236. [MR1637359](#)
- Sadhanala, V., Wang, Y.-X. and Tibshirani, R. J. (2016). Total variation classes beyond 1d: Minimax rates, and the limitations of linear smoothers. In *Advances in Neural Information Processing Systems*, pp. 3513–3521.
- Smola, A. J. and Kondor, R. (2003). Kernels and regularization on graphs. In *Learning theory and kernel machines*, pp. 144–158. Springer.
- Tsybakov, A. B. (2009). *Introduction to nonparametric estimation*. Springer Series in Statistics. Springer, New York. [MR2724359](#)
- Zhu, J. and Hastie, T. (2005). Kernel logistic regression and the import vector machine. *Journal of Computational and Graphical Statistics* **14**(1), 185–205.