

A recursive procedure for density estimation on the binary hypercube*

Maxim Raginsky

*Department of Electrical and Computer Engineering
and Coordinated Science Laboratory
University of Illinois
Urbana, IL 61801
e-mail: maxim@illinois.edu*

Jorge G. Silva

*SAS Institute, Inc.
100 SAS Campus Drive
Cary, NC 27513
e-mail: jorge.silva@sas.com*

Svetlana Lazebnik

*Department of Computer Science
University of Illinois
Urbana, IL 61801
e-mail: slazebni@illinois.edu*

and

Rebecca Willett

*Department of Electrical and Computer Engineering
Duke University
Durham, NC 27708
e-mail: willett@duke.edu*

Abstract: This paper describes a recursive estimation procedure for multivariate binary densities (probability distributions of vectors of Bernoulli random variables) using orthogonal expansions. For d covariates, there are 2^d basis coefficients to estimate, which renders conventional approaches computationally prohibitive when d is large. However, for a wide class of densities that satisfy a certain sparsity condition, our estimator runs in probabilistic polynomial time and adapts to the unknown sparsity of the underlying density in two key ways: (1) it attains near-minimax mean-squared error for moderate sample sizes, and (2) the computational complexity is lower for sparser densities. Our method also allows for flexible control of the trade-off between mean-squared error and computational complexity.

AMS 2000 subject classifications: Primary 62G07; secondary 62G20, 62C20.

*This work was supported by NSF CAREER Award CCF-06-43947, DARPA Grant HR0011-07-1-003, and ARO Grant W911NF-09-1-0262.

Keywords and phrases: Minimax estimation, density estimation, adaptive estimation, binary hypercube, Walsh basis, sparsity.

Received December 2011.

Contents

1	Introduction	821
1.1	Organization of the paper	824
2	Preliminaries	824
2.1	Notation	824
2.2	The Walsh system	825
2.3	Sparsity and weak- ℓ^p balls	826
3	Density estimation via recursive Walsh thresholding	827
3.1	Analysis of performance	830
4	Proofs of the theorems	833
4.1	Preliminaries	834
4.1.1	Moment bound	834
4.1.2	Concentration bounds	834
4.1.3	Large separated subsets of $\mathcal{F}_d^{+,1}(p)$	835
4.2	Proof of Theorem 1	835
4.3	Proof of Theorem 2	838
4.4	Proof of Theorem 3	841
5	Simulations	843
5.1	Low-dimensional simulations	843
5.2	High-dimensional simulations	847
6	Summary and conclusion	848
A	Auxiliary proofs	850
A.1	Proof of Lemma 1	850
A.2	Proof of Lemma 2	851
A.3	Proof of Lemma 3	852
B	Empirical process representation	854
	Acknowledgements	855
	References	855

1. Introduction

This paper considers the problem of estimating a multivariate binary density from a number of independent observations. That is, we have n observations of the form $X_i \in \{0, 1\}^d$ which are independent and identically distributed (i.i.d.) samples from a population with a probability density f (with respect to the counting measure on the d -dimensional *binary hypercube* $\{0, 1\}^d$). We wish to estimate f on the basis of these observations. Multivariate binary data arise in a variety of applications:

- *Biostatistics.* Each X_i could represent a biochemical profile of a bacterial strain, where every component is a “yes-no” indicator of a presence of a particular biochemical marker [18, 38]. A recent paper [31] proposed a

methodology for representing gene expression data using binary vectors. More classical scenarios include recording the occurrence of a given symptom or a medical condition in a patient over time [12] or outcomes of a series of medical tests [1].

- *Quantitative methods in social sciences.* Each X_i could represent a respondent in a survey or a panel, where every component is a “yes-no” answer to a question [6], describe a voting record of a legislator, or correspond to co-occurrences of events in social networks [32].
- *Artificial intelligence.* Each X_i could represent a user query to a search engine or a database, where every component corresponds to the presence or absence of a particular keyword [13], or an image stored on a website like Flickr¹, where every component corresponds to a user-supplied tag from a given list.

Many situations involving multivariate binary data have the following features: (1) the number of covariates (or the dimension of the hypercube) d is such that the number of possible values each observation could take (2^d) is much larger than the sample size n ; (2) there is a “clustering effect” in the population, meaning that the shape of the underlying density is strongly influenced mainly by a small number of constellations of the d covariates. For example, a particular class of bacterial strains may be reliably identified by looking at a particular subset of the biomarkers; there may be several such classes in the population of interest, each associated with a distinct subset of biomarkers. Similarly, when working with panel data, it may be the case that the answers to some specific subset of questions are highly correlated among a particular group of the panel participants, and the responses of these participants to other questions are nearly random; moreover, there may be several such distinct groups in the panel.

These considerations call for a density estimation procedure that can effectively cope with “thin” samples (i.e., those samples for which $n < 2^d$) in terms of both estimation error and computational complexity, and at the same time automatically adapt to the possible clustering in the population, in the sense described above. We take the minimax point of view, where we assume that the unknown density f comes from a particular function class \mathcal{F} and seek an estimator that exactly or approximately attains the minimax mean-squared error

$$R_n^*(\mathcal{F}) = \inf_{\hat{f}} \sup_{f \in \mathcal{F}} \mathbb{E} \|\hat{f} - f\|_{L^2}^2,$$

where the infimum is over all estimators based on n i.i.d. samples from f . We will choose the class \mathcal{F} to model the “constellation effects” via a certain *sparsity* condition. Our choice of the L^2 risk, as opposed to other measures of risk more commonly used in density estimation, such as Hellinger, Kullback–Leibler or total variation risks, is dictated by the fact that the sparsity condition mentioned above is most naturally stated in a Hilbert space framework, which in turn facilitates the design of our estimation procedure, as well as the derivation of both upper and lower bounds on $R_n^*(f)$. We refer the reader to several other

¹<http://www.flickr.com>

works on density estimation that use L^2 risk [24, 7, 10, 40, 20, 19, 11]. We also note that, because the Euclidean L^1 norm (which for probability densities gives the total variation risk) dominates the Euclidean L^2 norm, and because the square of the Kullback–Leibler divergence dominates the total variation distance [8], lower bounds on the squared L^2 risk automatically translate into lower bounds on the squared L^1 (total variation) risk and on the Kullback–Leibler risk.

Because of the host of applications in which multivariate binary data naturally arise, several authors have investigated algorithms for estimation of their probability densities (see, e.g., [1, 28, 24, 7]). However, existing approaches either have very slow rates of error convergence or are computationally prohibitive when the number of covariates is very large. For example, the kernel density estimation scheme proposed by Aitchison and Aitken [1] has computational complexity $O(nd)$ (where n is the number of observations and d the number of covariates), yet its squared L^2 error decays at the rate $O(n^{-4/(4+d)})$ [33], which is disastrously slow for large d . In contrast, orthogonal series methods, which can potentially achieve near-minimax error decay rates, require the estimation of 2^d basis coefficients and do not easily admit computationally tractable estimation methods for very large d . For instance, using the Fast Walsh–Hadamard Transform to estimate the coefficients of a density in the Walsh basis (see below) using n samples requires $O(nd2^d)$ operations (see Appendix B in [28] and references therein).

In this paper we present a computationally tractable orthogonal series estimation method based on recursive block thresholding of empirical Walsh coefficients. In particular, the proposed method entails recursively examining empirical estimates of whole *blocks* of the 2^d different Walsh coefficients. At each stage of the algorithm, the overall weight of basis coefficients computed at previous stages is used to decide which remaining coefficients are most likely to be significant or insignificant, and computing resources are allocated accordingly. It is shown that this decision is accurate with high probability, so that insignificant coefficients are not estimated, while the significant coefficients are. This approach is similar in spirit to the algorithm of Goldreich and Levin [17], originally developed for applications to cryptography and later adopted by Kushilevitz and Mansour [22, 25] to the problem of learning Boolean functions using membership queries. Although there are significant differences between the problems of density estimation and function learning which are reflected in our estimation procedure, our algorithm inherits the computational tractability of the Goldreich–Levin scheme: in particular, it runs in probabilistic polynomial time.

The proposed estimator adapts to unknown sparsity of the underlying density in two distinct ways. First, it is near-minimax optimal for “moderate” sample sizes $d \leq n \leq 2^{2d/p}$, with an L^2 error decay rate of $O(2^{-d}(d/n)^{2r/(2r+1)})$, where $p \in (0, 1]$ is a measure of sparsity and $r = 1/2 - 1/p$. Moreover, the computational complexity of our algorithm is automatically lower for sparser densities. Sparsity has been recently recognized as a crucial enabler of accurate estimation in “big- d , small- n ” type problems [4, 5]. Specifically for densities on the binary hypercube, sparsity in the Walsh basis has a natural qualitative interpretation that the

shape of the density is influenced mainly by a small number of constellations of the covariates. For example, if the components of a multidimensional binary vector represent positive/negative outcomes in a series of medical tests, it is often the case that the outcomes of certain small constellations of tests play the determining role in the diagnosis.

There are several different series expansions on the binary hypercube presented in the literature, including the Rademacher–Walsh orthogonal series (see Appendix A in [28] and references therein) and the Bahadur expansion [2, 18]. We focus in this paper on the Walsh system, which is derived from Fourier analysis on finite groups (see, e.g., Chap. 4 of Tao and Vu [36]), for two reasons. First, the coefficients of a particular function in the Walsh system give us information about the influence of the various subsets of the d variables on the value of the function [34, 9]. Second, the Walsh functions of a length- d vector can be factorized into products of Walsh functions of multiple shorter vectors with lengths summing to d ; this is detailed in Section 2.2. This factorization is central to the efficiency of the proposed coefficient estimation method. The Walsh system is widely used in the context of learning Boolean functions [25], as well as in harmonic analysis of real-valued functions on the binary hypercube [34, 9].

1.1. Organization of the paper

The remainder of the paper is organized as follows. Section 2 contains the preliminaries on notation, the Walsh system, and sparsity classes on the binary hypercube. Next, in Section 3 we describe the motivation behind the thresholding methods in orthogonal series estimation on the binary hypercube, introduce our recursive thresholded estimator, and analyze its MSE and computational complexity. The theorems of Section 3 are proved in Section 4. Some illustrative simulation results are given in Section 5. The contributions of the paper are summarized in Section 6. Finally, some technical results are relegated to the appendices.

2. Preliminaries

2.1. Notation

The basic set $\{0, 1\}$ will be denoted by \mathbb{B} . For any integer $k > 1$, the components of binary strings $x \in \mathbb{B}^k$ will be denoted by $x^{(j)}$, $1 \leq j \leq k$: for any $x \in \mathbb{B}^k$, we have $x = (x^{(1)}, \dots, x^{(k)})$. We will use juxtaposition to denote concatenation of strings: if $y \in \mathbb{B}^k$ and $z \in \mathbb{B}^l$, then $yz \in \mathbb{B}^{k+l}$ is the string $x = (y^{(1)}, \dots, y^{(k)}, z^{(1)}, \dots, z^{(l)})$. For any $0 \leq k \leq d$, we will define the *prefix* mapping $\pi_k : \mathbb{B}^d \rightarrow \mathbb{B}^k$ and the *suffix* mapping $\sigma_k : \mathbb{B}^d \rightarrow \mathbb{B}^{d-k}$ by

$$\pi_k(x) \triangleq (x^{(1)}, \dots, x^{(k)}), \quad \sigma_k(x) \triangleq (x^{(k+1)}, \dots, x^{(d)}), \quad (1)$$

so that $x = \pi_k(x)\sigma_k(x)$ for any $x \in \mathbb{B}^d$ (note that both π_0 and σ_d return an empty string). Whenever we deal with vectors whose components are indexed

by the elements of \mathbb{B}^k for some k , we will always assume that the components are arranged according to the lexicographic ordering of the binary strings in \mathbb{B}^k . Given two real numbers a, b , we let $a \wedge b$ denote $\min\{a, b\}$. Also, throughout the paper, C is used to denote a generic constant whose value may change from line to line; specific absolute constants will be denoted by C_1, C_2 , etc.

Throughout the paper, we let $M \equiv 2^d$.

2.2. The Walsh system

For any integer $k \geq 1$, denote by μ_k the counting measure on \mathbb{B}^k and endow the space of functions $f : \mathbb{B}^k \rightarrow \mathbb{R}$ with the structure of the real Hilbert space $L^2(\mu_k)$ via the standard inner product

$$\langle f, g \rangle \triangleq \sum_{x \in \mathbb{B}^k} f(x)g(x).$$

The *Walsh system* (see references in the Introduction) in $L^2(\mu_k)$ is an orthonormal system $\Phi_k = \{\varphi_s : s \in \mathbb{B}^k\}$, defined by

$$\varphi_s(x) \triangleq \frac{1}{2^{k/2}}(-1)^{s \cdot x}, \quad \forall x \in \mathbb{B}^k \tag{2}$$

where $s \cdot x \triangleq \sum_{j=1}^k s^{(j)}x^{(j)}$. Hence, any $f \in L^2(\mu_k)$ has the Fourier–Walsh expansion

$$f = \sum_{s \in \mathbb{B}^k} \theta_s \varphi_s,$$

where $\theta_s \triangleq \langle f, \varphi_s \rangle$, $s \in \mathbb{B}^k$. To keep the notation simple, we will not explicitly mark the underlying dimension when working with the Walsh functions. When $k = d$, we will write Φ instead of Φ_d .

For any k , the Walsh system Φ_k is a tensor product basis induced by $\Phi_1 = \{\varphi_0, \varphi_1\}$, where

$$\varphi_0(x) = \frac{1}{\sqrt{2}} \quad \text{and} \quad \varphi_1(x) = \frac{1}{\sqrt{2}}(-1)^x$$

for any $x \in \mathbb{B}$. That is, for any $k \geq 1$, any $\varphi_s \in \Phi_k$ has the form

$$\varphi_s = \varphi_{s^{(1)}} \otimes \varphi_{s^{(2)}} \otimes \cdots \otimes \varphi_{s^{(k)}},$$

which means that

$$\varphi_s(x) = \prod_{i=1}^k \varphi_{s^{(i)}}(x^{(i)}), \quad \forall x \in \mathbb{B}^k.$$

This generalizes to the following useful factorization property of the Walsh functions: for any $k \geq 1$ and any $l \leq k$, we have

$$\varphi_s = \varphi_{\pi_{l,k}(s)} \otimes \varphi_{\sigma_{l,k}(s)}, \quad \forall s \in \mathbb{B}^k \tag{3}$$

where $\pi_{l,k}$ and $\sigma_{l,k}$ denote the prefix and the suffix mappings defined on \mathbb{B}^k analogously to (1). This means that, for products of functions on disjoint subsets of the d variables, the Fourier–Walsh coefficients also have the product form.

2.3. Sparsity and weak- ℓ^p balls

Our interest lies with functions whose Fourier–Walsh representations satisfy a certain sparsity constraint. Given a function f on \mathbb{B}^d , let $\theta(f)$ denote the vector of its Fourier–Walsh coefficients. We will assume that the components of $\theta(f)$ decay according to a power law. Formally, let $\theta_{(1)}, \dots, \theta_{(M)}$, where $M = 2^d$, be the components of $\theta(f)$ arranged in decreasing order of magnitude:

$$|\theta_{(1)}| \geq |\theta_{(2)}| \geq \dots \geq |\theta_{(M)}|.$$

Given some $0 < p < \infty$, we say that $\theta(f)$ belongs to the Marcinkiewicz, or weak- ℓ^p , ball of radius R [3, 21], and write $\theta(f) \in w\ell^p(R)$, if

$$|\theta_{(m)}| \leq R \cdot m^{-1/p}, \quad 1 \leq m \leq M. \quad (4)$$

It is not hard to show that the Fourier–Walsh coefficients of any probability density on \mathbb{B}^d are bounded by $1/\sqrt{M}$. With this in mind, let us define the function class

$$\mathcal{F}_d(p) \triangleq \left\{ f : \mathbb{B}^d \rightarrow \mathbb{R} : \theta(f) \in w\ell^p(1/\sqrt{M}) \right\} \quad (5)$$

We are particularly interested in the case $0 < p \leq 1$.

We will need approximation properties of weak- ℓ^p balls as listed, e.g., in [5]. The basic fact is that the power-law condition (4) particularized to the elements of $\mathcal{F}_d(p)$ is equivalent to the concentration estimate

$$|\{s \in \mathbb{B}^d : |\theta_s|^2 \geq \lambda\}| \leq \left(\frac{1}{M\lambda} \right)^{p/2} \quad (6)$$

valid for all $\lambda > 0$. Additionally, for any $1 \leq k \leq M$, let $\theta_k(f)$ denote the vector $\theta(f)$ with $\theta_{(k+1)}, \dots, \theta_{(M)}$ set to zero. Then it follows from (4) that

$$\|\theta(f) - \theta_k(f)\|_{\ell_M^2} \leq CM^{-1/2}k^{-r} \quad (7)$$

where $r \triangleq 1/p - 1/2$, and C is some constant that depends only on p . Hence, given any $f \in \mathcal{F}_d(p)$ and denoting by f_k the function obtained from it by truncating all but the k largest Fourier–Walsh coefficients, we get from Parseval’s identity that

$$\|f - f_k\|_{L^2(\mu_d)} \leq CM^{-1/2}k^{-r}. \quad (8)$$

Thus, the assumption that f belongs to the sparsity class $\mathcal{F}_d(p)$ for some p can be interpreted qualitatively as saying that the behavior of f is strongly influenced by a small number of subsets of the d covariates. The number of these influential subsets decreases as $p \rightarrow 0$. To make this qualitative statement somewhat more precise, we recall the notion of *influence* of a covariate $i \in \{1, \dots, d\}$ on the function f [9]:

$$\text{Inf}_i(f) \triangleq \sum_{s \in \mathbb{B}^d: s^{(i)}=1} \theta_s(f)^2.$$

The larger this value, the more effect the “bit flip” $x^{(i)} \mapsto x^{(i)} \oplus 1$ has on the value of f , as measured by the average

$$\frac{1}{M} \sum_{x \in \mathbb{B}^d} |f(x) - f(x \oplus e_i)|^2,$$

where $e_i \in \mathbb{B}^d$ has 1 in the i th coordinate and 0 everywhere else. By the concentration estimate (6), the number of “large” coefficients in the Fourier–Walsh expansion of any $f \in \mathcal{F}_d(p)$ decreases as $p \rightarrow 0$. At the same time, the squared norm $\|f\|_{L^2(\mu_d)}^2$ of any $p \in \mathcal{F}_d(p)$ also decreases as $p \rightarrow 0$. Consequently, the number of “influential” covariates [i.e., those $i \in \{1, \dots, d\}$ with sufficiently large value of $\text{Inf}_i(f)$] becomes smaller as $p \rightarrow 0$.

3. Density estimation via recursive Walsh thresholding

Let X_1, \dots, X_n be independent random variables in \mathbb{B}^d with common unknown density f . We wish to estimate f on the basis of this sample. For densities defined on the Euclidean space, nonparametric estimators based on hard or soft thresholding of empirically estimated coefficients of the target density in a suitably chosen basis (e.g., a wavelet basis) attain near-minimax rates of convergence of the squared-error risk over rich classes of densities [10, 19]. Thresholding is a means of controlling the bias-variance trade-off.

Several authors have investigated the use of term-by-term thresholding rules for density estimation on the binary hypercube. There, one begins by computing the empirical estimates

$$\widehat{\theta}_s = \frac{1}{n} \sum_{i=1}^n \varphi_s(X_i) \tag{9}$$

of the Fourier–Walsh coefficients of f , and then forming the thresholded estimator

$$\widehat{f} = \sum_{s \in \mathbb{B}^d} I_{\{T(\widehat{\theta}_s) \geq \lambda_n\}} \widehat{\theta}_s \varphi_s, \tag{10}$$

where $T(\cdot)$ is some real-valued statistic and $I_{\{\cdot\}}$ is the indicator function. Based on the observation that

$$\text{Var } \widehat{\theta}_s = \frac{1}{n} \left(\frac{1}{M} - \theta_s^2 \right), \tag{11}$$

while the squared bias incurred by omitting the term $\widehat{\theta}_s \varphi_s$ from the estimator (10) is θ_s^2 , Ott and Kronmal [28] considered the ideal thresholded estimator

$$\widehat{f}^* = \sum_{s \in \mathbb{B}^d} I_{\{\theta_s^2 > 1/M(n+1)\}} \widehat{\theta}_s \varphi_s. \tag{12}$$

The idea is to keep only those terms where the variance (11) is smaller than the bias θ_s^2 . Clearly, \widehat{f}^* is impractical because the thresholding criterion depends on

the unknown coefficients θ_s . Instead, Ott and Kronmal [28] suggested that one could mimic the ideal estimator (12) by replacing θ_s^2 in the thresholding criterion by the unbiased estimator $(n\hat{\theta}_s^2 - 1/M)/(n-1)$, leading to the practical estimator

$$\hat{f}_{\text{WT}} = \sum_{s \in \mathbb{B}^d} I_{\{\hat{\theta}_s^2 > 2/M(n+1)\}} \hat{\theta}_s \varphi_s, \quad (13)$$

where WT stands for “Walsh thresholding.” This estimator was further improved by Liang and Krishnaiah [24] and Chen, Krishnaiah and Liang [7], who replaced the hard thresholding rule in (13) with shrinkage rules.

The main disadvantage of such termwise thresholding is the need to compute empirical estimates of all $M = 2^d$ Fourier–Walsh coefficients. While this is not an issue when d is comparable to $\log n$, it is clearly impractical when $d \gg \log n$. In order to alleviate this difficulty, we will consider a recursive thresholding approach, which will allow us to reject whole *groups* of empirical coefficients based on efficiently implementable thresholding rules. The main idea behind this approach is motivated by the following argument.

Given some $1 \leq k \leq d$, we can represent any function $f \in L^2(\mu_d)$ with the Fourier–Walsh coefficients $\{\theta_s : s \in \mathbb{B}^d\}$ as

$$\begin{aligned} f &= \sum_{u \in \mathbb{B}^k} \sum_{v \in \mathbb{B}^{d-k}} \theta_{uv} \varphi_{uv} \\ &= \sum_{u \in \mathbb{B}^k} \left(\sum_{v \in \mathbb{B}^{d-k}} \theta_{uv} \varphi_v \right) \otimes \varphi_u \\ &\equiv \sum_{u \in \mathbb{B}^k} f_u \otimes \varphi_u, \end{aligned}$$

where, for each $u \in \mathbb{B}^k$, $f_u \triangleq \sum_{v \in \mathbb{B}^{d-k}} \theta_{uv} \varphi_v$ is a function in $L^2(\mu_{d-k})$. The Fourier–Walsh coefficients of f_u are precisely those coefficients of f that are indexed by $s \in \mathbb{B}^d$ with $\pi_k(s) = u$. By Parseval’s identity, we have

$$W_u \triangleq \|f_u\|_{L^2(\mu_{d-k})}^2 = \sum_{v \in \mathbb{B}^{d-k}} \theta_{uv}^2.$$

This leads to the following observation: for any $\lambda > 0$,

$$W_u < \lambda \text{ for some } u \in \mathbb{B}^k \quad \Rightarrow \quad \theta_{uv}^2 < \lambda \text{ for every } v \in \mathbb{B}^{d-k}.$$

The usefulness of this observation for our purposes comes from the fact that we can represent the strings $s \in \mathbb{B}^d$, and hence the elements of the Walsh system in $L^2(\mu_d)$, by the leaves of a complete binary tree of depth d . Suppose we wanted to pick out only those coefficients of f whose squared magnitude exceeds some threshold λ . If we knew that $W_u \leq \lambda$ for some $u \in \mathbb{B}^k$, then this would tell us that the square of every coefficient corresponding to a leaf descending from u does not exceed λ . Hence, we could start at the root of the tree and at each internal node u that has not yet been visited check whether $W_u \geq \lambda$; if not,

then we would delete u and all of its descendants from the tree without having to compute explicitly the corresponding coefficients. At the end of the process (i.e., when we get to the leaves), we will be left only with those $s \in \mathbb{B}^d$ for which $\theta_s^2 \geq \lambda$. If $f \in \mathcal{F}_d(p)$ for some p , then the resulting squared L^2 error will be

$$\sum_{s \in \mathbb{B}^d} I_{\{\theta_s^2 < \lambda\}} \theta_s^2 \leq CM^{-1}(M\lambda)^{-2r/(2r+1)},$$

where $r = 1/p - 1/2$, as before.

We will follow this reasoning in constructing our density estimator. We begin by developing a suitable estimator for W_u . To do that, we shall rely on the following lemma (see Appendix A.1 for the proof):

Lemma 1. *For any density f on \mathbb{B}^d , any $1 \leq k \leq d$, and any $u \in \mathbb{B}^k$, we have*

$$f_u(y) = \mathbb{E}_f \{ \varphi_u(\pi_k(X)) I_{\{\sigma_k(X)=y\}} \}, \forall y \in \mathbb{B}^{d-k}$$

and

$$W_u = \mathbb{E}_f \{ \varphi_u(\pi_k(X)) f_u(\sigma_k(X)) \}.$$

This lemma suggests that, for each $1 \leq k \leq d$ and each $u \in \mathbb{B}^k$, an empirical estimate of W_u can be obtained by

$$\begin{aligned} \widehat{W}_u &= \frac{1}{n} \sum_{i=1}^n \varphi_u(\pi_k(X_i)) \left[\frac{1}{n} \sum_{j=1}^n \varphi_u(\pi_k(X_j)) I_{\{\sigma_k(X_i)=\sigma_k(X_j)\}} \right] \\ &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \varphi_u(\pi_k(X_i)) \varphi_u(\pi_k(X_j)) I_{\{\sigma_k(X_i)=\sigma_k(X_j)\}}. \end{aligned} \tag{14}$$

Although this is a biased estimator, it has the following useful property (see Appendix A.2 for the proof):

Lemma 2. *For any $1 \leq k \leq d$ and any $u \in \mathbb{B}^k$,*

$$\widehat{W}_u = \sum_{v \in \mathbb{B}^{d-k}} \widehat{\theta}_{uv}^2, \tag{15}$$

where each $\widehat{\theta}_{uv}$ is an empirical estimate of θ_{uv} computed according to (9).

Another advantage of computing \widehat{W}_u indirectly via (14), rather than (15), is that, while the latter requires $O(2^{d-k}n)$ operations, the former requires only $O(n^2d)$ operations. This can amount to significant computational savings when $k < d - \log(nd)$. When $k \geq d - \log(nd)$, it becomes more efficient to use the direct estimator (15).

Now that we have a way of estimating W_u , we can define our density estimation procedure. Provided the threshold scales appropriately with the sample size, we will be able to achieve a good balance between the estimation error (variance) and the approximation error (squared bias) and attain near-minimax

rates of convergence. In our analysis, we shall actually consider a more flexible strategy: for every $1 \leq k \leq d$, we shall compare the estimate \widehat{W}_u of W_u to a threshold that depends not only on the sample size n , but also on k . More specifically, we will let

$$\lambda_{k,n} = \frac{\alpha_k}{n}, \quad 1 \leq k \leq d \quad (16)$$

where the sequence $\{\alpha_k\}_{k=1}^d$ satisfies $\alpha_1 \geq \alpha_k \geq \dots \geq \alpha_d > 0$. In particular, this set-up covers the following two extreme cases:

1. $\alpha_k = \text{const}$ for all k – this covers the standard case of always comparing against the same threshold (that depends on n)
2. $\alpha_k = \text{const} \cdot 2^{d-k}$ – this corresponds to thresholding not the sum of (a particular subset of) the coefficients, but their *average*.

As we shall see, this k -dependent scaling will allow us to flexibly trade off the expected L^2 error and the computational complexity of the resulting estimator. Now we describe our density estimator. Given the sequence $\boldsymbol{\lambda} = \{\lambda_{k,n}\}_{k=1}^d$, define the set

$$A(\boldsymbol{\lambda}) \triangleq \left\{ s \in \mathbb{B}^d : \widehat{W}_{\pi_k(s)} \geq \lambda_{k,n}, \forall 1 \leq k \leq d \right\} \quad (17)$$

and consider the density estimate

$$\widehat{f}_{\text{RWT}} \triangleq \sum_{s \in \mathbb{B}^d} I_{\{s \in A(\boldsymbol{\lambda})\}} \widehat{\theta}_s \varphi_s, \quad (18)$$

where RWT stands for “recursive Walsh thresholding.” To implement this estimator on a computer, we call the routine RECURSIVEWALSH, shown as Algorithm 1, with $u = \emptyset$ (the empty string, corresponding to the root of the tree) and with the desired threshold sequence $\boldsymbol{\lambda}$. The factors of $1/2$ in the updates for \widehat{W}_{u_0} and \widehat{W}_{u_1} arise because of the factorization property (3) of the Walsh basis functions: for any $k \geq 0$ and any $s, x, x' \in \mathbb{B}^{k+1}$ we have

$$\varphi_u(x) \varphi_u(x') = \varphi_{\pi_k(s)}(\pi_k(x)) \varphi_{\pi_k(s)}(\pi_k(x')) \left(-\frac{1}{\sqrt{2}} \right)^{s^{(k)}(x^{(k)} + x'^{(k)})}.$$

3.1. Analysis of performance

Let us denote by $\mathcal{F}_d^{+,1}(p)$ the set of all probability densities in $\mathcal{F}_d(p)$:

$$\mathcal{F}_d^{+,1}(p) \triangleq \left\{ f \in \mathcal{F}_d(p) : f \geq 0, \sum_{x \in \mathbb{B}^d} f(x) = 1 \right\}. \quad (19)$$

Our first main result is that, with appropriately tuned thresholds, the estimator (18) adapts to unknown sparsity of the Fourier–Walsh representation of f :

Algorithm 1 RECURSIVEWALSH(u, λ)

```

 $k \leftarrow \text{length}(u)$ 
if  $k = d$  then
     $\hat{\theta}_u \leftarrow \frac{1}{n} \sum_{i=1}^n \varphi_u(X_i)$ 
    if  $\hat{\theta}_u^2 \geq \lambda_{d,n}$  then
        output  $u, \hat{\theta}_u$ 
    end if
    return
end if
 $u_0 \leftarrow 0u$ 
 $u_1 \leftarrow 1u$ 
 $\widehat{W}_{u_0} \leftarrow \frac{1}{2n^2} \sum_{i=1}^n \sum_{j=1}^n \varphi_u(\pi_k(X_i)) \varphi_u(\pi_k(X_j)) I_{\{\sigma_{k+1}(X_i) = \sigma_{k+1}(X_j)\}}$ 
if  $\widehat{W}_{u_0} \leq \lambda_{k+1,n}$  then
    return
else
    RECURSIVEWALSH( $u_0, \lambda$ )
end if
 $\widehat{W}_{u_1} \leftarrow \frac{1}{2n^2} \sum_{i=1}^n \sum_{j=1}^n \varphi_u(\pi_k(X_i)) \varphi_u(\pi_k(X_j)) (-1)^{X_i^{(k+1)} + X_j^{(k+1)}} I_{\{\sigma_{k+1}(X_i) = \sigma_{k+1}(X_j)\}}$ 
if  $\widehat{W}_{u_1} \leq \lambda_{k+1,n}$  then
    return
else
    RECURSIVEWALSH( $u_1, \lambda$ )
end if

```

Theorem 1. *There exist absolute constants $C_1, C_2 > 0$, such that the following holds. Suppose the threshold sequence $\lambda = \{\lambda_{k,n}\}_{k=1}^d$, where $\lambda_{k,n} = \alpha_k/n$ for all k , is chosen in such a way that $\alpha_k \geq C_1 d/M = (C_1/M) \log M$ for all k , where $M \equiv 2^d$. Then the estimator (18) satisfies*

$$\sup_{f \in \mathcal{F}_d^{+1}(p)} \mathbb{E}_f \|f - \widehat{f}_{\text{RWT}}\|_{L^2(\mu_d)}^2 \leq \frac{C_2}{M} \left(\frac{\log M}{n} \right)^{2r/(2r+1)} \tag{20}$$

for all n and all $0 < p \leq 1$, where, as before, $r = 1/p - 1/2$. Moreover, the risk of (18) in the large-sample regime $n \geq M^{2r/(2r+1)}$ is bounded by

$$\sup_{f \in \mathcal{F}_d^{+1}(p)} \mathbb{E}_f \|f - \widehat{f}_{\text{RWT}}\|_{L^2(\mu_d)}^2 \leq \frac{C_2 \log M}{n}. \tag{21}$$

Remark 1. *Adaptation to p .* As the reader can easily see from the proof of the theorem in Section 4.2, the bounds (20) and (21) hold for all $0 < p \leq 2$. However, we were only able to derive lower bounds for the case of $p \leq 1$.

Remark 2. *Positivity and normalization issues.* As is the case with orthogonal series estimators, \widehat{f}_{RWT} may not necessarily be a bona fide density. In particular, there may be some $x \in \mathbb{B}^d$ such that $\widehat{f}_{\text{RWT}}(x) < 0$, and it may happen that $\int \widehat{f}_{\text{RWT}} d\mu_d \neq 1$. In principle, this can be handled by clipping the negative values

at zero and renormalizing; this procedure can only improve the expected L^2 error. In practice this may be computationally expensive when d is very large. If the estimate is suitably sparse, however, the renormalization can be carried out approximately using Monte-Carlo estimates of the appropriate sums. Moreover, in many applications the scaling of the density is not important. \square

Our second main result is a lower bound on the minimax L^2 risk attainable by any estimator over $\mathcal{F}_d^{+,1}(p)$. It shows, in particular, that our recursive estimator \hat{f}_{RWT} is minimax for “moderate” sample sizes $\log M \preceq n \preceq M^{2/p}$. For large sample sizes, $n \succeq M^{2/p}$, \hat{f}_{RWT} is no longer optimal — in particular, it is outperformed by both the simple histogram estimator

$$\hat{f}_{\text{hist}}(x) = \frac{1}{n} \sum_{i=1}^n I_{\{X_i=x\}}$$

and by the unthresholded orthogonal series estimator

$$\hat{f}(x) = \sum_{s \in \mathbb{B}^d} \hat{\theta}_s \varphi_s(x),$$

both of which attain the optimal $O(1/n)$ risk. Note, however, that for $n \geq M^{2r/(2r+1)} = M^{1-p/2}$, the risk of \hat{f}_{RWT} is within a factor of $\log M$ of the minimax rate. The precise statement is as follows:

Theorem 2. *Consider the problem of estimating an unknown $f \in \mathcal{F}_d^{+,1}(p)$, $0 < p \leq 1$, from n i.i.d. samples X_1, \dots, X_n . Then the following statements hold:*

1. *Suppose that $\log M \leq n \leq M^{2(1-\epsilon)/p}$ for some $\epsilon \in (0, 1)$. Then there exists a positive constant $C = C(p, \epsilon)$, such that*

$$\inf_{\hat{f}_n} \sup_{f \in \mathcal{F}_d^{+,1}(p)} \mathbb{E}_f \|\hat{f}_n - f\|_{L^2(\mu_d)}^2 \geq \frac{C}{M} \left(\frac{\log M}{n} \right)^{\frac{2r}{2r+1}}. \quad (22)$$

where, as before, $M = 2^d$.

2. *Suppose that $n \geq M^{2/p}$ and $M \geq 4$. Then there exists an absolute constant $C > 0$, such that*

$$\inf_{\hat{f}_n} \sup_{f \in \mathcal{F}_d^{+,1}(p)} \mathbb{E}_f \|\hat{f}_n - f\|_{L^2(\mu_d)}^2 \geq \frac{C}{n}. \quad (23)$$

Our third, and final, main result bounds the running time of the algorithm used for computing \hat{f}_{RWT} :

Theorem 3. *Fix any $f \in \mathcal{F}_d(p)$. Given any $\delta \in (0, 1)$, provided each α_k is chosen so that*

$$nC_1(2^k a_{k,n}^2 \wedge 2^{k/2} a_{k,n}) \geq \frac{\log(2^k d / \delta)}{\log e}, \quad (24)$$

where

$$a_{k,n} \triangleq \frac{1}{5} \sqrt{\frac{\alpha_k}{n}} - \sqrt{\frac{C_2^2}{2^k n}}$$

and $C_1, C_2 > 0$ are certain absolute constants, then Algorithm 1 runs in

$$O\left(n^2 d \left(\frac{n}{M}\right)^{p/2} K(\boldsymbol{\alpha}, p)\right) \tag{25}$$

time with probability at least $1 - \delta$, where $K(\boldsymbol{\alpha}, p) \triangleq \sum_{k=1}^d \alpha_k^{-p/2}$ and, as before, $M = 2^d$.

Remark 3. *Trade-off between time complexity and MSE.* By controlling the rate at which the sequence α_k decays with k , we can trade off MSE against complexity. Consider the following two extreme cases: (1) $\alpha_1 = \dots = \alpha_d \sim 1/M$ and (2) $\alpha_k \sim 2^{d-k}/M$. The first case, which reduces to the term-by-term thresholding, achieves the same bias-variance trade-off as the Ott–Kronmal estimator [28]. However, it has $K(\boldsymbol{\alpha}, p) = O(M^{p/2}d)$, resulting in $O(d^2 n^{2+p/2})$ complexity. The second case, which leads to a very severe estimator that will tend to reject a lot of coefficients, has MSE of $O(n^{-2r/(2r+1)} M^{-1/(2r+1)})$, but $K(\boldsymbol{\alpha}, p) = O(M^{p/2})$, leading to a considerably better $O(dn^{2+p/2})$ complexity. From the computational viewpoint, it is preferable to use rapidly decaying thresholds. However, this reduction in complexity will be offset by a corresponding increase in MSE. In fact, using RWT with exponentially decaying α_k 's in practice is not advisable as its low complexity is mainly due to the fact that it will tend to reject even the big coefficients very early on, especially when d is large. To achieve a good balance between complexity and MSE, a moderately decaying threshold sequence might be best, e.g., $\alpha_k \sim (d - k + 1)^m/M$ for some $m \geq 1$. As $p \rightarrow 0$, the effect of $\boldsymbol{\lambda}$ on complexity becomes negligible, and the complexity tends to $O(n^2 d)$. \square

Remark 4. *Incoherence.* Note that for any of the above choices of α_k , the proposed method requires $\text{polylog}(M)$ operations. One intuitive explanation for why such fast computation is possible is that the Walsh basis is “incoherent” (to use term common in compressed sensing and group testing literature) with the canonical basis of $L^2(\mu_d)$. Similar polylog computation results were achieved by Gilbert *et al.* in the context of fast sparse Fourier approximation [14, 15] and group testing [16]. Their strategies also had connections to the Goldreich–Levin algorithm [17], as well as to the work of Kushilevitz and Mansour on sparse Boolean function estimation [22, 25]. \square

4. Proofs of the theorems

In this section we prove our three main results. However, before proceeding to the proofs, we collect all the technical tools that we will be using: moment bounds, concentration inequalities, and an approximation-theoretic lemma pertaining to class $\mathcal{F}_d^{+,1}(p)$.

4.1. Preliminaries

4.1.1. Moment bound

We will need the following result of Rosenthal [30]. Let U_1, \dots, U_n be i.i.d. random variables with $\mathbb{E}U_i = 0$ and $\mathbb{E}U_i^2 \leq \sigma^2$. Then for any $m \geq 2$ there exists some c_m such that

$$\mathbb{E} \left| n^{-1} \sum_{i=1}^n U_i \right|^m \leq c_m \left(\frac{\sigma^m}{n^{m/2}} + \frac{\mathbb{E}|U_1|^m}{n^{m-1}} \right). \quad (26)$$

4.1.2. Concentration bounds

We will need the well-known Hoeffding inequality: if U_1, \dots, U_n are i.i.d. bounded random variables such that $\mathbb{E}U_i = 0$ and $|U_i| \leq b < \infty$ for all $1 \leq i \leq n$, then

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n U_i \right| > t \right) \leq 2 \exp \left(-\frac{nt^2}{2b^2} \right) \quad (27)$$

The following result is due to Talagrand [35]. Let U_1, \dots, U_n be i.i.d. random variables, let $\varepsilon_1, \dots, \varepsilon_n$ be independent Rademacher random variables [i.e., $\mathbb{P}(\varepsilon_i = -1) = \mathbb{P}(\varepsilon_i = 1) = 1/2$] also independent of U_1, \dots, U_n , and let \mathcal{F} be a class of functions uniformly bounded by $L > 0$. Then if there exist some $v, H > 0$ such that $\sup_{g \in \mathcal{F}} \text{Var } g(U) \leq v$ and

$$\mathbb{E} \left\{ \sup_{g \in \mathcal{F}} \sum_{i=1}^n \varepsilon_i g(U_i) \right\} \leq nH \quad (28)$$

for all n , then there are universal constants C_1 and C_2 such that, for every $\tau > 0$,

$$\mathbb{P} \left(\sup_{g \in \mathcal{F}} \nu_n(g) \geq \tau + C_2 H \right) \leq \exp \left\{ -nC_1 \left(\frac{\tau^2}{v} \wedge \frac{\tau}{L} \right) \right\}, \quad (29)$$

where

$$\nu_n(g) \triangleq \frac{1}{n} \sum_{i=1}^n g(U_i) - \mathbb{E}g(U), \quad \forall g \in \mathcal{F} \quad (30)$$

is the empirical process indexed by \mathcal{F} .

Remark 5. Typically, some additional regularity conditions on \mathcal{F} are needed to ensure measurability of the supremum $\sup_{g \in \mathcal{F}} \nu_n(g)$ of the empirical process (30). However, when U takes values in a finite set, as is the case in this paper, there is no need for such conditions because any uniformly bounded class of real-valued functions on a finite set is separable: it contains a countable subset \mathcal{F}_0 , such that for any $g \in \mathcal{F}$ there exists a sequence $g_1, g_2, \dots \in \mathcal{F}_0$ converging to g pointwise. Such a separability property ensures measurability of suprema over \mathcal{F} [37, p. 110].

4.1.3. Large separated subsets of $\mathcal{F}_d^{+,1}(p)$

In the sequel, we will be interested in large subsets of the class of densities $\mathcal{F}_d^{+,1}(p) \subset \mathcal{F}_d(p)$, whose elements are separated from one another by a given fixed amount, as measured by the norm $\|\cdot\|_{L^2(\mu_d)}$. The following lemma, whose proof is given in Appendix A, will be useful:

Lemma 3. *Let $r = 1/p - 1/2$. Let s_1, \dots, s_M , $M = 2^d$, be the lexicographic ordering of the elements of \mathbb{B}^d . Given a positive real parameter a and an integer $k \in \{1, \dots, M - 1\}$, let $\Theta(M, k, a) \subset \mathbb{R}^{M-1}$ consist of $(M - 1)$ -dimensional real vectors having exactly k nonzero components, each of which is equal to either a or $-a$. With this, define the set $\mathcal{F}(k, a) \subset L^2(\mu_d)$ by*

$$\mathcal{F}(k, a) = \left\{ f : \theta_{s_1}(f) = \frac{1}{\sqrt{M}}, (\theta_{s_j}(f))_{2 \leq j \leq M} \in \Theta(M, k, a) \right\}.$$

Suppose that k and a are such that

$$ka \leq \frac{1}{2\sqrt{M}} \quad \text{and} \quad a \leq \frac{1}{\sqrt{M}}(k + 1)^{-1/p}. \tag{31}$$

Then the following statements hold:

1. The set $\mathcal{F}(k, a)$ is contained in $\mathcal{F}_d^{+,1}(p)$.
2. For any two $f, f' \in \mathcal{F}(k, a)$ we have

$$D(f||f') \leq 2Mka^2, \tag{32}$$

where $D(\cdot||\cdot)$ is the Kullback–Leibler divergence (relative entropy) [8].

3. If $k = M - 1$, then there exists a set $\tilde{\mathcal{F}}(M - 1, a) \subset \mathcal{F}(M - 1, a)$ with the following properties:
 - $\|f - f'\|_{L^2(\mu_d)}^2 \geq (M - 1)a^2$ for all $f, f' \in \tilde{\mathcal{F}}(M - 1, a)$ with $f \neq f'$
 - $\log |\tilde{\mathcal{F}}(M - 1, a)| \geq (M - 1)/8$
4. If $M - 1 \geq 4k$, then there exists a set $\tilde{\mathcal{F}}(k, a) \subset \mathcal{F}(k, a)$ with the following properties:
 - $\|f - f'\|_{L^2(\mu_d)}^2 \geq ka^2$ for all $f, f' \in \tilde{\mathcal{F}}(k, a)$ with $f \neq f'$
 - $\log |\tilde{\mathcal{F}}(k, a)| \geq 0.233 k \left(\log \frac{M-1}{k} + 1\right)$

4.2. Proof of Theorem 1

Let us decompose the squared L^2 error as

$$\|f - \hat{f}_{\text{RWT}}\|_{L^2(\mu_d)}^2 = \sum_s I_{\{s \in A(\boldsymbol{\lambda})\}} (\theta_s - \hat{\theta}_s)^2 + \sum_s I_{\{s \in A(\boldsymbol{\lambda})^c\}} \theta_s^2 \equiv T_1 + T_2.$$

We start by observing that any $s \in A(\boldsymbol{\lambda})$ necessarily satisfies $\widehat{W}_s \equiv \hat{\theta}_s^2 \geq \lambda_{d,n}$, while for any $s \in A(\boldsymbol{\lambda})^c$ there exists some $1 \leq k \leq d$ such that $\widehat{W}_{\pi_k(s)} < \lambda_{k,n}$,

which implies, in turn, that $\widehat{\theta}_{\pi_k(s)t}^2 < \lambda_{k,n}$ for all $t \in \mathbb{B}^{d-k}$ and, in particular, that $\widehat{\theta}_s^2 < \lambda_{k,n} \leq \lambda_{1,n}$. Therefore, defining the sets

$$A_1 = \{s \in \mathbb{B}^d : \widehat{\theta}_s^2 \geq \lambda_{d,n}\} \quad \text{and} \quad A_2 = \{s \in \mathbb{B}^d : \widehat{\theta}_s^2 < \lambda_{1,n}\},$$

we can bound T_1 and T_2 as

$$T_1 \leq \sum_s I_{\{s \in A_1\}} (\theta_s - \widehat{\theta}_s)^2$$

and

$$T_2 \leq \sum_s I_{\{s \in A_2\}} \theta_s^2.$$

Further, defining

$$B = \left\{s \in \mathbb{B}^d : \theta_s^2 < \lambda_{d,n}/2\right\} \quad \text{and} \quad S = \left\{s \in \mathbb{B}^d : \theta_s^2 \geq 3\lambda_{1,n}/2\right\},$$

we can write

$$T_1 \leq \sum_s I_{\{s \in A_1 \cap B\}} (\theta_s - \widehat{\theta}_s)^2 + \sum_s I_{\{s \in A_1 \cap B^c\}} (\theta_s - \widehat{\theta}_s)^2 \equiv T_{11} + T_{12},$$

and

$$T_2 \leq \sum_s I_{\{s \in A_2 \cap S\}} \theta_s^2 + \sum_s I_{\{s \in A_2 \cap S^c\}} \theta_s^2 \equiv T_{21} + T_{22}.$$

Applying (6) and (11), we get

$$\begin{aligned} \mathbb{E}T_{12} &\leq \sum_{s \in B^c} \mathbb{E}(\theta_s - \widehat{\theta}_s)^2 \\ &\leq \frac{1}{Mn} |\{s : \theta_s^2 \geq \lambda_{d,n}/2\}| \\ &\leq \frac{1}{Mn} \left(\frac{2}{M\lambda_{d,n}}\right)^{p/2} \\ &= \frac{1}{Mn} \left(\frac{2n}{M\alpha_d}\right)^{p/2} \\ &= \frac{1}{M} n^{p/2-1} \underbrace{\left(\frac{2}{2^d \alpha_d}\right)^{p/2}}_{\leq 1} \\ &\leq \frac{1}{M} n^{-2r/(2r+1)}. \end{aligned} \tag{33}$$

To bound T_{22} we apply (8):

$$\begin{aligned} \mathbb{E}T_{22} &\leq \sum_{s \in \mathbb{B}^d} I_{\{\theta_s^2 < 3\alpha_1/2n\}} \theta_s^2 \leq \frac{C}{M} \left(\frac{M\alpha_1}{n}\right)^{2r/(2r+1)} \\ &\leq \frac{C}{M} \left(\frac{\log M}{n}\right)^{2r/(2r+1)}. \end{aligned} \tag{34}$$

In order to deal with the large-deviation terms T_{11} and T_{21} , we will need some moment and concentration bounds which are listed in Section 4.1. First, using Cauchy–Schwarz, we get

$$\mathbb{E}T_{11} \leq \sum_s \sqrt{\mathbb{E}(\theta_s - \widehat{\theta}_s)^4 \cdot \mathbb{P}(s \in A_1 \cap B)} \tag{35}$$

To estimate the fourth moment in (35), we apply the bound (26) to $U_i = \varphi_s(X_i) - \theta_s$, $1 \leq i \leq n$, and $m = 4$. Then

$$\mathbb{E}U_i^2 = \mathbb{E}(\varphi_s(X_i) - \theta_s)^2 = \frac{1}{M} - \theta_s^2 \leq \frac{1}{M}$$

and

$$\mathbb{E}|U_1|^4 \leq \frac{1}{M^2},$$

so that

$$\mathbb{E}(\theta_s - \widehat{\theta}_s)^4 \leq c_4 \left(\frac{1}{M^2 n^2} + \frac{1}{M^2 n^3} \right) \leq \frac{2c_4}{M^2 n^2}.$$

To handle the probability that $s \in A_1 \cap B$, we first estimate

$$|\widehat{\theta}_s - \theta_s|^2 = (\theta_s - \widehat{\theta}_s)^2 = \theta_s^2 - 2\theta_s \widehat{\theta}_s + \widehat{\theta}_s^2 \geq \theta_s^2 - 2|\theta_s \widehat{\theta}_s| + \widehat{\theta}_s^2 = (|\theta_s| - |\widehat{\theta}_s|)^2.$$

From this we conclude that $\widehat{\theta}_s^2 \geq \lambda_{d,n}$ and $\theta_s^2 < \lambda_{d,n}/2$ together imply

$$|\widehat{\theta}_s - \theta_s| \geq \frac{1}{5} \sqrt{\lambda_{d,n}} = \frac{1}{5} \sqrt{\frac{\alpha_d}{n}}$$

(the factor of $1/5$ is simply a lower bound on $1 - 1/\sqrt{2}$). Therefore,

$$\mathbb{P}(s \in A_1 \cap B) \leq \mathbb{P}\left(|\widehat{\theta}_s - \theta_s| \geq \frac{1}{5} \sqrt{\frac{\alpha_d}{n}}\right).$$

Applying Hoeffding’s inequality (27) to $U_i = \varphi_s(X_i) - \theta_s$, $1 \leq i \leq n$, with $b = 2/\sqrt{M}$ and using the fact that $\alpha_d \geq C_1 d/M$, we get

$$\mathbb{P}\left(|\widehat{\theta}_s - \theta_s| \geq \frac{1}{5} \sqrt{\frac{\alpha_d}{n}}\right) \leq 2 \exp(-Cd) \leq \frac{2}{MC}.$$

for some absolute constant $C > 0$. If C_1 is chosen so that $C \geq 2$, then we will have

$$\mathbb{E}T_{11} \leq \frac{2\sqrt{c_4}}{Mn}. \tag{36}$$

Finally,

$$\mathbb{E}T_{21} \leq \sum_s \mathbb{P}(s \in A_2 \cap S) \theta_s^2 \leq \sum_{s \in S} \mathbb{P}(s \in A_2) \theta_s^2.$$

Using (the one-sided version of) Hoeffding's inequality,

$$\begin{aligned} \mathbb{P}(s \in A_2) &= \mathbb{P}\left(\widehat{\theta}_s < \sqrt{\lambda_{1,n}}\right) \\ &= \mathbb{P}\left(\widehat{\theta}_s - \theta_s < \sqrt{\lambda_{1,n}} - \theta_s\right) \\ &\leq \exp\left[-C'n\left(\theta_s - \sqrt{\lambda_{1,n}}\right)^2\right] \end{aligned} \quad (37)$$

for some absolute constant $C' > 0$. By definition of the set S ,

$$S \subseteq \bigcup_{\ell=2}^{\infty} \left\{s \in \mathbb{B}^d : \frac{3(\ell-1)\lambda_{1,n}}{2} \leq \theta_s^2 < \frac{3\ell\lambda_{1,n}}{2}\right\}. \quad (38)$$

Therefore, combining (37) and (38) and using (6), we get

$$\begin{aligned} \mathbb{E}T_{21} &\leq \sum_{s \in S} \exp\left[-C'n\left(\theta_s - \sqrt{\lambda_{1,n}}\right)^2\right] \theta_s^2 \\ &\leq \frac{3\lambda_{1,n}}{2} \sum_{\ell=2}^{\infty} \exp(-C''\lambda_{1,n}n\ell) \cdot \ell \left| \left\{s \in \mathbb{B}^d : \frac{3(\ell-1)\lambda_{1,n}}{2} \leq \theta_s^2 < \frac{3\ell\lambda_{1,n}}{2}\right\} \right| \\ &\leq \frac{C}{M} \left(\frac{\log M}{n}\right)^{2r/(2r+1)} \sum_{\ell=2}^{\infty} \frac{1}{M^\ell} \ell^{2r/(2r+1)} \\ &\leq \frac{C}{M} \left(\frac{\log M}{n}\right)^{2r/(2r+1)}. \end{aligned} \quad (39)$$

Putting together Eqs. (33), (34), (36), and (39), we get the bound (20). To get (21), we use much cruder bounds:

$$\begin{aligned} \mathbb{E}T_1 &\leq \sum_s \mathbb{E}[(\theta_s - \widehat{\theta}_s)^2] \leq M \cdot \frac{CM}{n} = \frac{C}{n}; \\ \mathbb{E}T_{22} &\leq \sum_s I_{\{\theta_s^2 < 3\alpha_1/2n\}} \theta_s^2 \leq M \cdot \frac{C \log M}{Mn} = \frac{C \log M}{n}; \\ \mathbb{E}T_{21} &\leq \sum_{s \in S} \theta_s^2 \leq \frac{|S|}{M} \leq \frac{1}{M} \left(\frac{2}{3M\lambda_{1,n}}\right)^{p/2} \leq \frac{C}{M} \left(\frac{n}{\log M}\right)^{p/2} \leq \frac{Cn^{p/2-1}}{Mn} \leq \frac{C}{n}, \end{aligned}$$

where in the last bound we have used the assumption that $n \geq M^{2r/(2r+1)}$ (the bounds on $\mathbb{E}T_1$ and $\mathbb{E}T_{22}$ hold for all values of n). Combining these three bounds, we get (21).

4.3. Proof of Theorem 2

The proof of the first part uses a popular information-theoretic technique due to Yang and Barron [39, 40]; we only outline the main steps. The first step is to

lower-bound the minimax risk by the minimum probability of error in a multiple hypothesis test. Let \mathcal{F}_0 be an arbitrary subset of $\mathcal{F}_d^{+,1}(p)$. Then

$$\inf_{\hat{f}} \sup_{f \in \mathcal{F}_d^{+,1}(p)} \mathbb{E}_f \|\hat{f} - f\|_{L^2(\mu_d)}^2 \geq \inf_{\hat{f}} \sup_{f \in \mathcal{F}_0} \mathbb{E}_f \|\hat{f} - f\|_{L^2(\mu_d)}^2.$$

In particular, suppose that the set \mathcal{F}_0 is finite, $\mathcal{F}_0 = \{f^{(1)}, \dots, f^{(N)}\}$, and δ -separated in $L^2(\mu_d)$, i.e.,

$$\|f^{(i)} - f^{(j)}\|_{L^2(\mu_d)} \geq \delta, \quad \forall i, j \in \{1, \dots, N\}; i \neq j \tag{40}$$

Then a standard argument [40] gives

$$\inf_{\hat{f}} \sup_{f \in \mathcal{F}_0} \mathbb{E}_f \|\hat{f} - f\|_{L^2(\mu_d)}^2 \geq \frac{\delta^2}{4} \min_{\tilde{f}} \mathbb{P}(\tilde{f} \neq f^{(Z)}), \tag{41}$$

where the random variable Z is uniformly distributed over the set $\{1, \dots, N\}$, and the minimum is over all estimators \tilde{f} based on X^n that take values in the packing set $\{f^{(1)}, \dots, f^{(N)}\}$. Applying Fano's inequality [8], we can write

$$\min_{\tilde{f}} \mathbb{P}(\tilde{f} \neq f^{(Z)}) \geq 1 - \frac{I(Z; X^n) + \log 2}{\log N}, \tag{42}$$

where $I(Z; X^n)$ is the Shannon mutual information [8] between the random index $Z \in \{1, \dots, N\}$ and the observations $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} f^{(Z)}$. In this particular case, we have

$$\begin{aligned} I(Z; X^n) &= \frac{n}{N} \sum_{k=1}^N \sum_{x \in \mathbb{B}^d} f^{(k)}(x) \log \frac{f^{(k)}(x)}{N^{-1} \sum_{\ell=1}^N f^{(\ell)}(x)} \\ &= \frac{n}{N} \sum_{k=1}^N D(f^{(k)} \| \bar{f}), \end{aligned} \tag{43}$$

where \bar{f} denotes the mixture density $N^{-1} \sum_{\ell=1}^N f^{(\ell)}$. The next step consists in upper-bounding this mutual information. To that end, suppose that there exists some $\Delta > 0$, such that

$$D(f \| f') \leq \Delta, \quad \forall f, f' \in \mathcal{F}_0. \tag{44}$$

Using convexity of the relative entropy and (44), for every $k \in \{1, \dots, N\}$ we have

$$D(f^{(k)} \| \bar{f}) \leq \frac{1}{N} \sum_{\ell=1}^N D(f^{(k)} \| f^{(\ell)}) \leq \Delta.$$

Substituting this into (43), we see that $I(Z; X^n) \leq n\Delta$. Combining this bound with (41), we get

$$\inf_{\hat{f}} \sup_{f \in \mathcal{F}_0} \mathbb{E}_f \|\hat{f} - f\|_{L^2(\mu_d)}^2 \geq \frac{\delta^2}{4} \left[1 - \frac{n\Delta + \log 2}{\log N} \right] \tag{45}$$

In particular, let $k \in \{1, \dots, M - 1\}$ and $a > 0$ satisfy the conditions (31) of Lemma 3, as well as

$$a^2 \leq \frac{C}{Mnk} \log |\tilde{\mathcal{F}}(k, a)| \tag{46}$$

for a suitable constant $C > 0$, where $\tilde{\mathcal{F}}(k, a)$ are the subsets of $\mathcal{F}_d^{+,1}(p)$ described in Lemma 3. If we let $\mathcal{F}_0 = \tilde{\mathcal{F}}(k, a)$, then, by Lemma 3, (44) holds with $\Delta = \frac{1}{2n} \log \frac{|\tilde{\mathcal{F}}(k, a)|}{4}$. This, in conjunction with (45), gives

$$\inf_{\hat{f}} \sup_{f \in \mathcal{F}_0} \mathbb{E}_f \|\hat{f} - f\|_{L^2(\mu_d)}^2 \geq C\delta^2(k, a), \tag{47}$$

where

$$\delta(k, a) \triangleq \min \left\{ \|f - f'\|_{L^2(\mu_d)} : f, f' \in \tilde{\mathcal{F}}(k, a); f \neq f' \right\}$$

is the minimal $L^2(\mu_d)$ -separation between any two distinct elements of $\tilde{\mathcal{F}}(k, a)$. We can now consider the following cases:

1. Suppose that $M^{2/p} \leq n$. Then we take $k = M - 1$ and $a^2 = \frac{C}{Mn}$. Because in this case

$$k \vee (k + 1)^{1/p} = M^{1/p} \leq \sqrt{n} \leq \frac{C}{a\sqrt{M}},$$

and $\log |\tilde{\mathcal{F}}(M - 1, a)| \geq (M - 1)/8$, the conditions (31) and (46) will be satisfied for a suitable choice of C . Moreover, by Lemma 3, we have

$$\delta^2(k, a) = \delta^2(M - 1, a) \geq (M - 1)a^2 \geq \frac{C}{n}.$$

Substituting this into (45), we obtain (23).

2. Suppose that $n \leq M^{2(1-\epsilon)/p}$ for some $\epsilon \in (0, 1)$. Let $k = C \left(\frac{n}{\log M}\right)^{p/2}$ and $a = C' \left(\frac{\log M}{Mn}\right)^{1/2}$. Then

$$\begin{aligned} k \vee (k + 1)^{1/p} &\leq (2k)^{1/p} \\ &= (2C)^{1/p} \left(\frac{n}{\log M}\right)^{1/2} \\ &= \frac{(2C)^{1/p} C'}{a\sqrt{M}}. \end{aligned}$$

If we choose C and C' in such a way that $(2C)^{1/p} C' \leq 1/2$, then (31) will be satisfied. Moreover, we must have

$$M - 1 \geq 4k = 4C \left(\frac{n}{\log M}\right)^{p/2}. \tag{48}$$

With our assumptions on n and M , this will hold for all sufficiently large M . Next, we check that (46) is satisfied. Assuming that (48) holds, Lemma 3 implies that

$$\begin{aligned} \frac{1}{Mnk} \log |\tilde{\mathcal{F}}(k, a)| &\geq \frac{C}{Mn} \log \frac{M-1}{k} \\ &\geq \frac{C}{Mn} \log \frac{(M-1)^{2/p} \log M}{n}. \end{aligned}$$

Again, using our assumption on n and M , as well as the fact that $a^2 = \frac{C \log M}{Mn}$, we can guarantee that (46) holds, with an appropriate choice of $C = C(p, \epsilon)$. By Lemma 3, we will have

$$\delta^2(k, a) \geq ka^2 \geq \frac{C}{M} \left(\frac{\log M}{n} \right)^{1-\frac{p}{2}} = \frac{C}{M} \left(\frac{\log M}{n} \right)^{\frac{2r}{2r+1}},$$

and we obtain (22).

4.4. Proof of Theorem 3

The time complexity of the algorithm is determined by the number of recursive calls made to RECURSIVEWALSH. Recall that, for each $1 \leq k \leq d$, a recursive call to RECURSIVEWALSH is made for every $u \in \mathbb{B}^k$ for which $\widehat{W}_u \geq \lambda_{k,n}$. Let us say that a recursive call to RECURSIVEWALSH(u, λ) is *correct* if $W_u \geq \lambda_{k,n}/2$. We will show that, with high probability, only the correct recursive calls are made at every $1 \leq k \leq d$. The probability of making at least one incorrect recursive call is given by

$$\begin{aligned} &\mathbb{P} \left(\bigcup_{k=1}^d \bigcup_{u \in \mathbb{B}^k} \{ \widehat{W}_u \geq \lambda_{k,n}, W_u < \lambda_{k,n}/2 \} \right) \\ &\leq \sum_{k=1}^d \sum_{u \in \mathbb{B}^k} \mathbb{P} \left(\widehat{W}_u \geq \lambda_{k,n}, W_u < \lambda_{k,n}/2 \right). \end{aligned}$$

For a given $u \in \mathbb{B}^k$, let

$$\widehat{f}_u \triangleq \sum_{v \in \mathbb{B}^{d-k}} \widehat{\theta}_{uv} \varphi_v.$$

Then $\widehat{W}_u \geq \lambda_{k,n}$ and $W_u < \lambda_{k,n}/2$ together imply that

$$\begin{aligned} \|f_u - \widehat{f}_u\|_{L^2(\mu_{d-k})}^2 &= \sum_{v \in \mathbb{B}^{d-k}} (\widehat{\theta}_{uv} - \theta_{uv})^2 \\ &= \widehat{W}_u - 2 \sum_{v \in \mathbb{B}^{d-k}} \widehat{\theta}_{uv} \theta_{uv} + W_u \end{aligned}$$

$$\begin{aligned}
&\geq \widehat{W}_u - 2\sqrt{\widehat{W}_u W_u} + W_u \\
&= \left(\sqrt{\widehat{W}_u} - \sqrt{W_u}\right)^2 \\
&\geq \left(1 - \frac{1}{\sqrt{2}}\right)^2 \lambda_{k,n} \\
&\geq \lambda_{k,n}/25.
\end{aligned}$$

Now, as shown in Appendix B, for each $u \in \mathbb{B}^k$, the norm $\|f_u - \widehat{f}_u\|_{L^2(\mu_{d-k})}$ can be expressed as a supremum of an empirical process over a suitable function class, to which we can then apply Talagrand's bound (29) with $L = 1/\sqrt{2^k}$, $v = 1/2^k$, and $H = 1/\sqrt{2^k n}$. Hence,

$$\begin{aligned}
\mathbb{P}(\widehat{W}_u \geq \lambda_{k,n}, W_u < \lambda_{k,n}/2) &\leq \mathbb{P}(\|f_u - \widehat{f}_u\|_{L^2(\mu_{d-k})}^2 \geq \sqrt{\lambda_{k,n}}/5) \\
&\leq \exp\left\{-nC_1(2^k a_{k,n}^2 \wedge 2^{k/2} a_{k,n})\right\},
\end{aligned}$$

where for each $1 \leq k \leq d$,

$$a_{k,n} = \frac{1}{5}\sqrt{\frac{\alpha_k}{n}} - \sqrt{\frac{C_2^2}{2^k n}}.$$

Here, C_1 and C_2 are the absolute constants in Talagrand's bound (29). Given $\delta > 0$, if we choose α_k according to (24), then

$$\mathbb{P}(\widehat{W}_u \geq \lambda_{k,n}, W_u < \lambda_{k,n}/2) \leq \frac{\delta}{d^{2k}}, \quad \forall u \in \mathbb{B}^k.$$

Summing over $1 \leq k \leq d$ and $u \in \mathbb{B}^k$, we see that, with probability at least $1 - \delta$, only the correct recursive calls will be made.

Next, we give an upper bound on the number of the correct recursive calls. For each $1 \leq k \leq d$, $W_u \geq \lambda_{k,n}/2$ implies that there exists at least one $v \in \mathbb{B}^{d-k}$ such that $\theta_{uv}^2 \geq \lambda_{k,n}/2$. Since for every $1 \leq k \leq d$ each θ_s contributes to exactly one W_u , we have by the pigeonhole principle that

$$\begin{aligned}
|\{u \in \mathbb{B}^k : W_u \geq \lambda_{k,n}/2\}| &\leq |\{s \in \mathbb{B}^d : \theta_s^2 \geq \lambda_{k,n}/2\}| \\
&\leq \left(\frac{2}{M\lambda_{k,n}}\right)^{p/2},
\end{aligned}$$

where in the second line we used (6). Hence, the number of the correct recursive calls in Algorithm 1 is bounded by

$$N = \sum_{k=1}^d \left(\frac{2}{M\lambda_{k,n}}\right)^{p/2} = \left(\frac{2n}{M}\right)^{p/2} \sum_{k=1}^d \alpha_k^{-p/2}.$$

At each recursive call, we compute an estimate of the corresponding W_{u_0} and W_{u_1} , which requires $O(n^2 d)$ operations. Therefore, with probability at least $1 - \delta$, the time complexity of the algorithm is given by (25).

5. Simulations

Although an extensive empirical evaluation is outside the scope of this paper, we have implemented the proposed estimator, and now present some simulation results to demonstrate its small-sample performance on synthetic data in low- and high-dimensional regimes. In the low-dimensional regime, it is feasible to obtain the “ground truth” by exhaustively computing all the 2^d Walsh coefficients and to compare it with our estimate. In the high-dimensional regime, our comparison is based on the density values at randomly generated samples. Additionally, we present a number of computational strategies that greatly enhance computational efficiency in the high-dimensional regime.

5.1. Low-dimensional simulations

We generated synthetic observations from a mixture density f on a 15-dimensional binary hypercube. The mixture has 10 components, where each component is a product density with 12 randomly chosen covariates having Bernoulli(1/2) distributions, and the other three having Bernoulli(0.9) distributions. For $d = 15$, it is still feasible to quickly compute the ground truth, consisting of 32768 values of f and its Walsh coefficients. These values are shown in Fig. 1 (left). As can be seen from the coefficient profile in the bottom of the figure, this density is clearly sparse. Fig. 1 also shows the estimated probabilities and the Walsh coefficients for sample sizes $n = 5000$ (middle) and $n = 10000$ (right).

To study the trade-off between MSE and complexity, we implemented three different thresholding schemes: (1) constant, $\lambda_{k,n} = 2/(2^d n)$, (2) logarithmic, $\lambda_{k,n} = 2 \log(d - k + 2)/(2^d n)$, and (3) linear, $\lambda_{k,n} = 2(d - k + 1)/(2^d n)$. The thresholds at $k = d$ are set to twice the variance of the empirical estimate of any coefficient whose value is zero; this forces the estimator to reject empirical coefficients whose values cannot be reliably distinguished from zero. Occasionally, spurious coefficients get retained, as can be seen in Fig. 1 (middle) for the estimate for $n = 5000$. Fig. 2 shows the performance of \hat{f}_{RWT} . Fig. 2(a) is a plot of MSE vs. sample size. In agreement with the theory, MSE is the smallest for the constant thresholding scheme [which is simply an efficient recursive implementation of a term-by-term thresholding estimator with $\lambda_n \sim 1/Mn$], and then it increases for the logarithmic and for the linear schemes. Fig. 2(b,c) shows the running time (in seconds) and the number of recursive calls made to RECURSIVEWALSH vs. sample size. The number of recursive calls is a platform-independent way of gauging the computational complexity of the algorithm, although it should be kept in mind that each recursive call has $O(n^2 d)$ overhead. Also, the number of recursive calls depends on whether a binary or N -ary tree is utilized. The N -ary tree scheme is explained in detail below, in Section 5.2. We have used $N = 256$ in the simulations, as this setting leads to much reduced computations times vs. a binary tree.

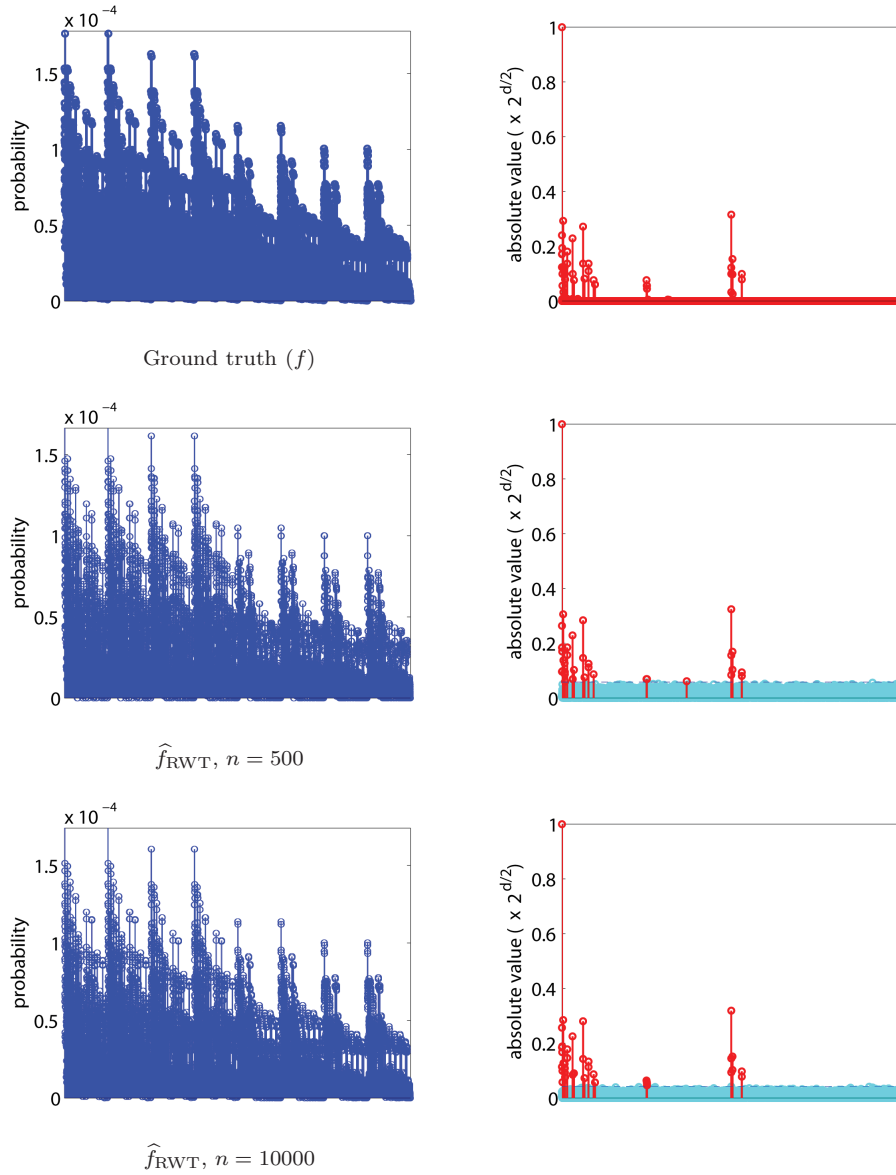


FIG 1. Ground truth (top) and estimated density for $n = 5000$ (middle) and $n = 10000$ (bottom) with constant thresholding. Left column: true and estimated probabilities (clipped at zero and renormalized) arranged in lexicographic order. Right column: absolute values of true and estimated Walsh coefficients arranged in lexicographic order. For the estimated densities, the coefficient plots also show the threshold level (dotted line) and absolute values of the rejected coefficients (lighter color).

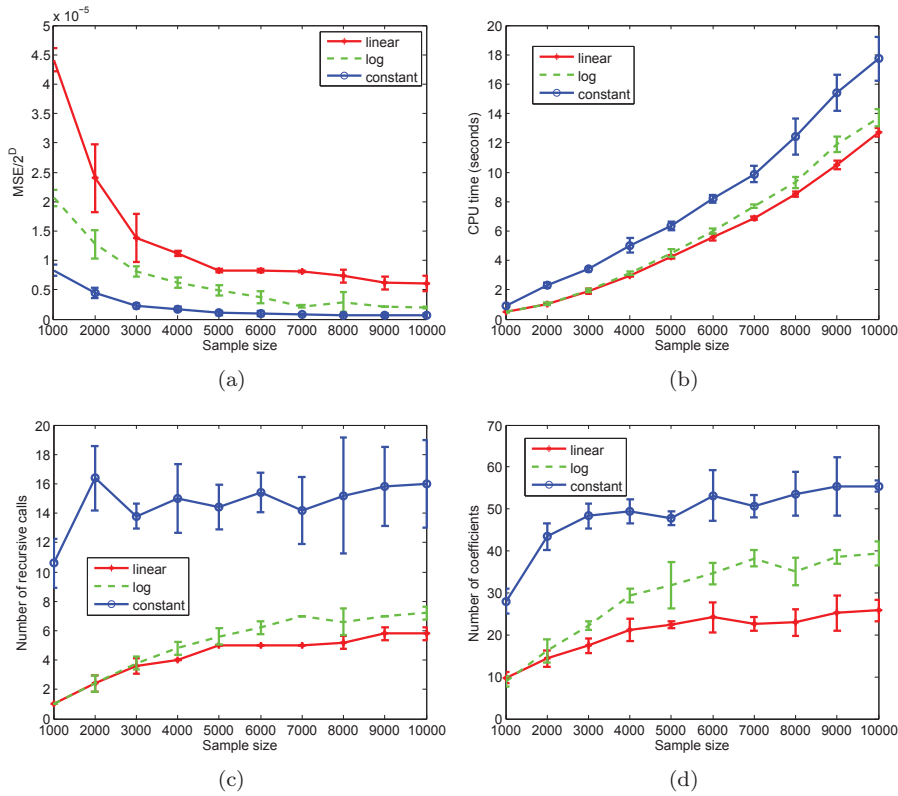


FIG 2. Small-sample performance of \hat{f}_{FRWT} in estimating f with three different thresholding schemes: (a) MSE; (b) running time (in seconds); (c) number of recursive calls; (d) number of coefficients retained by the algorithm. All results are averaged over five independent runs for each sample size (the error bars show the standard deviations).

The running time increases polynomially with n , and is the largest for the constant scheme, followed by the logarithmic and the linear schemes. We see that, while the MSE of the logarithmic scheme is fairly close to that of the constant scheme, its complexity is considerably lower, in terms of both the number of recursive calls and the running time. In all three cases, the number of recursive calls decreases with n due to the fact that weight estimates become increasingly accurate with n , which causes the expected number of false discoveries (i.e., making a recursive call at an internal node of the tree only to reject its descendants later) to decrease. Finally, Fig. 2(d) shows the number of coefficients retained in the estimate. This number grows with n as a consequence of the fact that the threshold decreases with n , while the number of accurately estimated coefficients increases.

Additionally, we have performed comparisons between our proposed method and two alternatives: the Ott and Kronmal thresholding estimator [28]; and

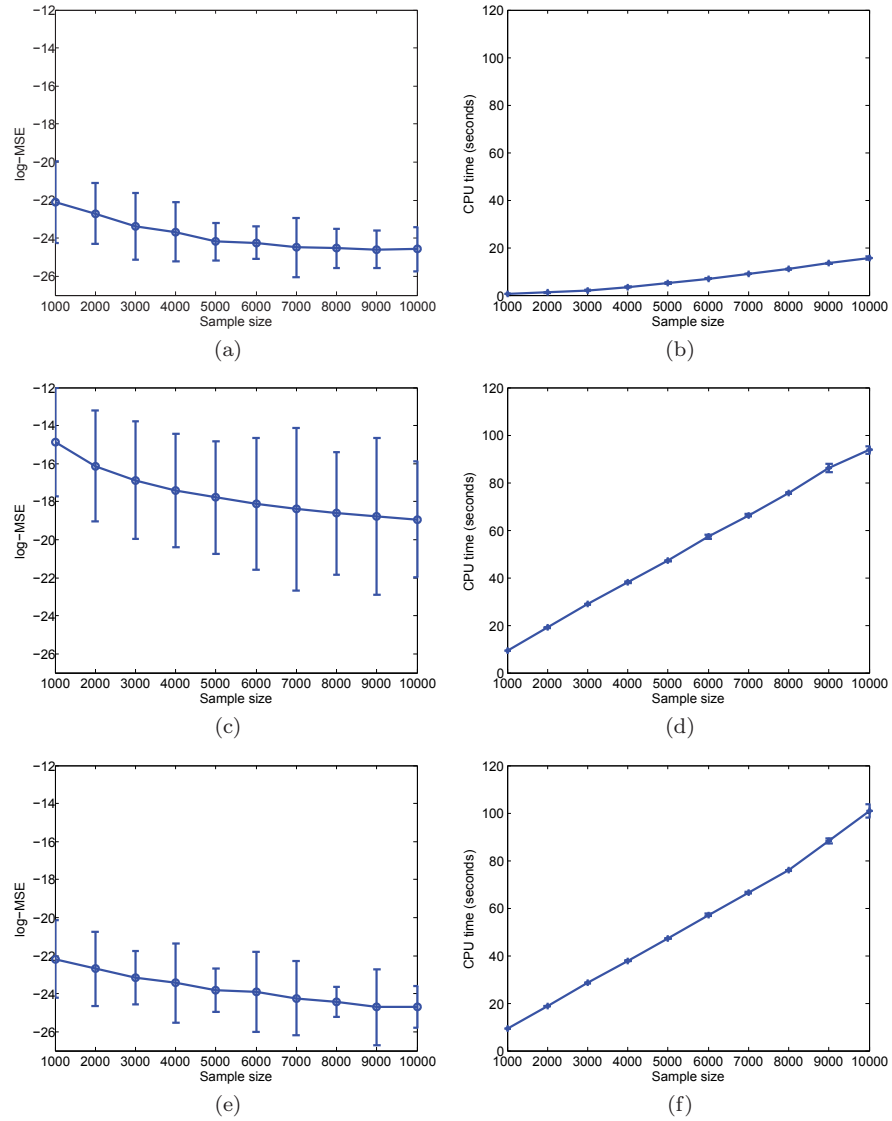


FIG 3. Comparison of \hat{f}_{RWT} with the Ott and Kronmal ($O\&K$) estimator [28] and with an exhaustive search for the best MSE. The plots show: (a) \hat{f}_{RWT} MSE; (b) \hat{f}_{RWT} times; (c) $O\&K$ MSE; (d) $O\&K$ times; (e) Exhaustive search MSE; (f) Exhaustive search times. All results are averaged over ten independent runs for each sample size (the error bars show the standard deviations).

an exhaustive search over all possible thresholds for the best MSE. As seen in Fig. 3, our thresholding estimator provides close to the best possible MSE with far lower computational cost than the alternatives.

5.2. High-dimensional simulations

Although our algorithm has been implemented in MATLAB and therefore can be much further optimized for speed, we have devised several strategies for traversing the coefficient tree efficiently and circumventing computer-architecture related challenges for high-dimensional problems. We describe those strategies and demonstrate them using the same experimental set-up as above, but with much larger dimensionality.

- Direct computation of the coefficients near the leaves of the tree.** As discussed in 2.3, the direct estimator (15) for \widehat{W}_u becomes computationally more efficient than the indirect estimator (14) for $k \geq d - \log(nd)$. Hence, near the bottom of the tree, instead of continuing to traverse the remaining levels based on the weights \widehat{W}_u , we simply compute all coefficients $\widehat{\theta}_u$ at the leaves of the corresponding subtree. This is always the most sensible course of action, given the fact that (15) requires the $\widehat{\theta}_u$ anyway.
- N -ary tree.** The number of levels to be traversed can be reduced by a factor of $\log N$ by considering an N -ary, rather than a binary, tree, at the cost of an increased number (N) of branches per level. We have found this trade-off to be worthwhile in many cases due to the possibility of vectorizing the computation of \widehat{W}_u for all the branches in each level and taking advantage of optimized routines for matrix algebra.
- Open-node queue.** While our algorithm lends itself to recursive implementation, many computer/operating system architectures impose a hard limit on the recursion level due to stack-size restrictions (recursive function calls typically use stack memory). This becomes a problem when the dimension d is high and the tree is accordingly very deep. We have circumvented the issue by implementing a queue system where the so-called open nodes (the tree branches awaiting processing) are sorted according to some criterion. This amounts to transferring the “stack” to user memory, where the only limit on the number of nodes is the free memory size. Possible criteria for sorting the nodes include depth-first, breadth-first and highest weight \widehat{W}_u . We have used the latter criterion in our high-dimensional simulations.
- Pruning high-frequency Walsh coefficients.** For a given Walsh function φ_s , the Hamming weight (i.e., number of “on” bits) of s is a measure of frequency; higher-frequency coefficients have a higher proportion of ones in s . Because in many problems it is appropriate to assume that the signals of interest have low frequency, we have included the ability to impose a limit on the order of the Walsh coefficients by ignoring any branches with more than m “on” bits. The choice of m depends on the problem context and on computational resources.
- Weight-adaptive thresholding.** For some datasets, significant gains can be achieved by varying the thresholds $\lambda_{k,n}$ in a data-driven manner at each level of the tree; as an alternative to the preset schedules α_k in


(16), it is possible to take the weights \widehat{W}_u for each branch at level k , and then expand only the top q branches with the highest \widehat{W}_u . This is equivalent to making α_k not only level-dependent, but also dependent on the sequence of weights at level k . The value of q controls the trade-off between computation speed and accuracy.

The first three strategies are important modifications to a naive implementation of our algorithm, but in no way impact the MSE. The latter two techniques, however, provide an approximation to the estimator proposed and analyzed in this paper; for appropriate values of m and q , they yield significant computational savings for a modest increase in MSE.


In Figure 4, we present plots of the MSE and computation time for simulated data with $d = 50$ and multiple sample sizes (n), using the aforementioned optimizations. As before, the data were generated from a Bernoulli mixture density, similar to the one used for Figure 2 but using ten 50-dimensional mixture components, where each component has 47 covariates with a Bernoulli(1/2) distribution and three covariates with a Bernoulli(0.9) distribution. The results are averaged over ten independent runs. We have limited the number of “on” bits in the Walsh binary strings to three and eight (*i.e.*, $m = 3$ and $m = 8$ respectively), expanded only the 16 subtrees with highest \widehat{W}_u at each level (*i.e.*, $q = 16$) and used an N -ary tree with $N = 256$. The subtrees in the open-node queue were sorted by decreasing \widehat{W}_u . Even in this high-dimensional regime, we achieve steadily decreasing MSE as a function of n , as well as approximately linear scaling in computation time. It is also apparent that setting $m = 3$ achieves essentially the same MSE but with an order-of-magnitude reduction in the computational effort.

6. Summary and conclusion


We have presented a computationally efficient adaptive procedure for estimating a multivariate binary density in the “big d , small n ” regime, which essentially forces a “nonparametric” approach. Many problems of current practical interest that involve multivariate binary data seem to pertain to populations with certain “constellation” effects among the d covariates. We have formalized this observation by focusing on a class of densities whose Walsh representations exhibit a certain power-law behavior. For moderate sample sizes, our estimator attains nearly minimax rates of MSE convergence over this class and runs in polynomial time with high probability. Moreover, the complexity improves for sparser densities. We have also reported the results of simulations, which show that our implemented estimator behaves in accordance with the theory even in the small-sample regime. In the future, we plan to test our method on real high-dimensional data sets. Another promising future direction is to investigate the relationship between various smoothness classes of densities on the binary hypercube defined in terms of their Fourier–Walsh representations and probability densities of binary Markov random fields [23]. Such a density will have



(a)



(b)



(c)

FIG 4. Performance for a large-dimensional problem ($d = 50$): (a) Log-MSE; (b)–(c) running time (in seconds) for $m = 3$ and $m = 8$, respectively. All results are averaged over ten independent runs for each sample size (the error bars show the standard deviations).

the form

$$f(x) = \frac{1}{Z} \prod_{i=1}^d e^{-h_i(x^{(i)}, x^{N_i})},$$

where for each $i \in \{1, \dots, d\}$ we have a neighborhood $N_i \subseteq \{1, \dots, d\} \setminus \{i\}$, the corresponding “local energy” function $h_i(\cdot)$ depends only on $x^{(i)}$ and on $x^{N_i} = (x^{(j)} : j \in N_i)$, and Z is the normalization constant known as the *partition function*. It is reasonable to assume that if most of the neighborhoods N_i are small, then most of the Fourier–Walsh coefficients of f will be small as well. Assuming specific bounds on the decay of the Fourier–Walsh coefficients of f amounts to assuming something about the decay of correlations in the Markov random field governed by f . It remains to be seen whether sparsity classes of the type investigated in this paper can serve as a good model of binary Markov random fields with polynomial decay of correlations, or whether one would need to introduce a binary analog of something like (weak) Besov bodies [3] in order to account for localization both in space (small number of i ’s with large neighborhoods) and in frequency (small number of large Fourier–Walsh coefficients). In either case, it would be worthwhile to investigate the use of recursive thresholding estimators of the type introduced in this paper to estimate Markov graphical models on the binary hypercube.

Appendix A: Auxiliary proofs

A.1. Proof of Lemma 1

Using the appropriate definitions, as well as the factorization property (3) of the Walsh functions, we have

$$\begin{aligned} \mathbb{E}_f \{ \varphi_u(\pi_k(X)) I_{\{\sigma_k(X)=y\}} \} &= \sum_{x \in \mathbb{B}^d} f(x) \varphi_u(\pi_k(x)) I_{\{\sigma_k(x)=y\}} \\ &= \sum_{z \in \mathbb{B}^k} f(zy) \varphi_u(z) \\ &= \sum_{z \in \mathbb{B}^k} \left(\sum_{(v,w) \in \mathbb{B}^k \times \mathbb{B}^{d-k}} \theta_{vw} \varphi_{vw}(zy) \right) \varphi_u(z) \\ &= \sum_{z \in \mathbb{B}^k} \left(\sum_{(v,w) \in \mathbb{B}^k \times \mathbb{B}^{d-k}} \theta_{vw} \varphi_v(z) \varphi_w(y) \right) \varphi_u(z) \\ &= \sum_{(v,w) \in \mathbb{B}^k \times \mathbb{B}^{d-k}} \theta_{vw} \varphi_w(y) \langle \varphi_v, \varphi_u \rangle \\ &= \sum_{w \in \mathbb{B}^{d-k}} \theta_{uw} \varphi_w(y) \\ &\equiv f_u(y). \end{aligned}$$

Similarly,

$$\begin{aligned} W_u &= \sum_{y \in \mathbb{B}^{d-k}} f_u^2(y) \\ &= \sum_{y \in \mathbb{B}^{d-k}} \mathbb{E} \{ \varphi_u(\pi_k(X)) I_{\{\sigma_k(X)=y\}} \} f_u(y) \\ &= \sum_{x \in \mathbb{B}^d} \sum_{y \in \mathbb{B}^{d-k}} f(x) \varphi_u(\pi_k(x)) f_u(y) I_{\{\sigma_j(x)=y\}} \\ &= \sum_{x \in \mathbb{B}^d} f(x) \varphi_u(\pi_k(x)) f_u(\sigma_k(x)) \\ &\equiv \mathbb{E}_f \{ \varphi_u(\pi_k(X)) f_u(\sigma_k(X)) \}, \end{aligned}$$

and the lemma is proved.

A.2. Proof of Lemma 2

We begin by showing that, for any $1 \leq k < d$ and any $u \in \mathbb{B}^k$,

$$\widehat{W}_u = \widehat{W}_{u0} + \widehat{W}_{u1}. \tag{A.1}$$

From the factorization properties of the Walsh functions, we have

$$\begin{aligned} \widehat{W}_{u0} &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \varphi_{u0}(\pi_{k+1}(X_i)) \varphi_{u0}(\pi_{k+1}(X_j)) I_{\{\sigma_{k+1}(X_i)=\sigma_{k+1}(X_j)\}} \\ &= \frac{1}{2n^2} \sum_{i=1}^n \sum_{j=1}^n \varphi_u(\pi_k(X_i)) \varphi_u(\pi_k(X_j)) I_{\{\sigma_{k+1}(X_i)=\sigma_{k+1}(X_j)\}} \end{aligned}$$

and

$$\begin{aligned} \widehat{W}_{u1} &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \varphi_{u1}(\pi_{k+1}(X_i)) \varphi_{u1}(\pi_{k+1}(X_j)) I_{\{\sigma_{k+1}(X_i)=\sigma_{k+1}(X_j)\}} \\ &= \frac{1}{2n^2} \sum_{i=1}^n \sum_{j=1}^n \varphi_u(\pi_k(X_i)) \varphi_u(\pi_k(X_j)) (-1)^{X_i^{(k+1)}} \\ &\quad \times (-1)^{X_j^{(k+1)}} I_{\{\sigma_{k+1}(X_i)=\sigma_{k+1}(X_j)\}}. \end{aligned}$$

Adding these two expressions and using the fact that

$$1 + (-1)^{X_i^{(k+1)}} (-1)^{X_j^{(k+1)}} = 2 I_{\{X_i^{(k+1)}=X_j^{(k+1)}\}},$$

we get

$$\begin{aligned} \widehat{W}_{u0} + \widehat{W}_{u1} &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \varphi_u(\pi_k(X_i)) \varphi_u(\pi_k(X_j)) \\ &\quad \times I_{\{\sigma_{k+1}(X_i)=\sigma_{k+1}(X_j), X_i^{(k+1)}=X_j^{(k+1)}\}} \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \varphi_u(\pi_k(X_i)) \varphi_u(\pi_k(X_j)) I_{\{\sigma_k(X_i) = \sigma_k(X_j)\}} \\
&\equiv \widehat{W}_u.
\end{aligned}$$

This proves (A.1). By induction, we have

$$\widehat{W}_u = \sum_{s \in \mathcal{L}(u)} \widehat{W}_s,$$

where $\mathcal{L}(u)$ denotes the set of all leaves descending from u . Since $\mathcal{L}(u) = \{uv \in \mathbb{B}^d : v \in \mathbb{B}^{d-k}\}$ and since

$$\widehat{W}_s = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \varphi_s(X_i) \varphi_s(X_j) \equiv \widehat{\theta}_s^2,$$

the lemma is proved.

A.3. Proof of Lemma 3

We first consider Item 1. By construction, for any $f \in \mathcal{F}(k, a)$ we have

$$f(x) \leq \frac{1}{M} + \sum_{j=2}^M |\theta_{s_j}(f)| \cdot \|\varphi_{s_j}\|_\infty = \frac{1}{M} + \frac{ka}{\sqrt{M}}$$

and

$$f(x) \geq \frac{1}{M} - \sum_{j=2}^M |\theta_{s_j}(f)| \cdot \|\varphi_{s_j}\|_\infty = \frac{1}{M} - \frac{ka}{\sqrt{M}}$$

Thus, if k and a satisfy the first condition in (31), then all $f \in \mathcal{F}(k, a)$ will be bounded between $1/2M$ and $3/2M$. To see that the second condition in (31) implies $\mathcal{F}(k, a) \subset \mathcal{F}_d(p)$, observe that in that case the Fourier–Walsh coefficients of f ordered according to decreasing magnitude are

$$|\theta_{(m)}(f)| = \begin{cases} \frac{1}{\sqrt{M}}, & m = 1 \\ a, & m = 2, \dots, k+1 \\ 0, & m = k+2, \dots, M \end{cases}$$

from which it follows that $f \in \mathcal{F}_d(p)$. Finally, any $f \in \mathcal{F}(k, a)$ is also a probability density because it is nonnegative and because

$$\sum_{x \in \mathbb{B}^d} f(x) = \langle f, 1 \rangle = \sqrt{M} \langle f, \varphi_0 \rangle = \sqrt{M} \theta_0(f) = 1.$$

This shows that $\mathcal{F}(k, a) \subset \mathcal{F}_d^{+,1}(p)$, as claimed.

We now move on to Item 2. To prove that the bound (32) holds for any two densities $f, f' \in \mathcal{F}(k, a)$, we use the fact that the Kullback–Leibler divergence $D(f\|f')$ is bounded above by the chi-square distance:

$$D(f\|f') \leq \chi^2(f, f') \triangleq \sum_{x \in \mathbb{B}^d} \frac{|f(x) - f'(x)|^2}{f'(x)}. \tag{A.2}$$

For any $x \in \mathbb{B}^d$ and $f, f' \in \mathcal{F}(k, a)$, we have

$$|f(x) - f'(x)|^2 = \left| \sum_{j=1}^M (\theta_{s_j}(f) - \theta_{s_j}(f')) \varphi_{s_j}(x) \right|^2 \leq \frac{2ka^2}{M}.$$

Using this together with the fact that $f, f' \geq 1/2M$ in (A.2), we arrive at (32).

For Item 3, we need the following well-known combinatorial result from coding theory, the so-called Varshamov–Gilbert bound (see, e.g., Lemma 4.7 in [26]): For any $m \in \mathbb{N}$, there exists a subset \mathcal{K}_m of \mathbb{B}^m with the following properties:

- for any $u, v \in \mathcal{K}_m$ with $u \neq v$,

$$\sum_{i=1}^m I_{\{u^{(i)} \neq v^{(i)}\}} \geq m/4$$

- $\log |\mathcal{K}_m| \geq m/8$

There is a one-to-one correspondence between $\mathcal{F}(M - 1, a)$ and the binary hypercube \mathbb{B}^{M-1} , under which each $f \in \mathcal{F}(M - 1, a)$ is mapped to $u_f \in \mathbb{B}^{M-1}$ with

$$u_f^{(j)} = I_{\{\theta_{s_{j+1}}(f)=a\}}, \quad j = 1, \dots, M - 1.$$

Given the Varshamov–Gilbert set $\mathcal{K}_{M-1} \subset \mathbb{B}^{M-1}$, let

$$\tilde{\mathcal{F}}(M - 1, a) \triangleq \left\{ f \in \mathcal{F}(M - 1, a) : u_f \in \mathcal{K}_{M-1} \right\}.$$

Then for any two $f, f' \in \tilde{\mathcal{F}}(M - 1, a)$, we have

$$\|f - f'\|_{L^2(\mu_d)}^2 = 4a^2 \sum_{j=1}^{M-1} I_{\{u_f^{(j)} \neq u_{f'}^{(j)}\}} \geq (M - 1)a^2, \tag{A.3}$$

and $\log |\mathcal{F}(M - 1, a)| = \log |\mathcal{K}_{M-1}| \geq (M - 1)/8$. This takes care of Item 3.

Finally, we consider Item 4. To that end, we will need a refinement of the Varshamov–Gilbert bound due to Reynaud-Bouret [29], which for our purposes can be stated as follows (see Lemma 4.10 in [26]): For any $m, k \in \mathbb{N}$ with $m \geq k$, let \mathbb{B}_k^m denote the subset of \mathbb{B}^m that consists of all $u \in \mathbb{B}^m$ with $\sum_{i=1}^m I_{\{u^{(i)}=1\}} = k$. If $m \geq 4k$, then there exists a set $\mathcal{K}_{m,k} \subset \mathbb{B}_k^m$ with the following properties:

- for any $u, v \in \mathcal{K}_{m,k}$ with $u \neq v$,

$$\sum_{i=1}^m I_{\{u^{(i)} \neq v^{(i)}\}} \geq k/2$$

- $\log |\mathcal{K}_{m,k}| \geq 0.233k \log(m/k)$

Now assume that $M - 1 \geq 4k$ and consider the corresponding set $\mathcal{K}_{M-1,k}$. To each $u \in \mathcal{K}_{M-1,k}$ we can associate 2^k elements of $\mathcal{F}(M - 1, k)$, say $\{f_{u,v} : v \in \mathbb{B}^k\}$, such that u determines the locations of the k nonzero Walsh–Fourier coefficients taking values in $\{-a, a\}$, while the choice of $v \in \mathbb{B}^k$ determines the signs of these coefficients. Thus, let

$$\tilde{\mathcal{F}}(k, a) \triangleq \left\{ f_{u,v} : u \in \mathcal{K}_{m,k}, v \in \mathbb{B}^k \right\}.$$

Consider now any two $f^{u,v}, f^{u',v'} \in \tilde{\mathcal{F}}(k, a)$, such that $(u, v) \neq (u', v')$. For any $j \in 1, \dots, M - 1$ such that $u^{(j)} \neq u'^{(j)}$, either $u^{(j)} = 0$ or $u'^{(j)} = 0$. Suppose the latter. Then $|\theta_{s_{j+1}}(f_{u,v})| = a$, while $|\theta_{s_{j+1}}(f_{u',v'})| = 0$. Consequently,

$$\|f_{u,v} - f_{u',v'}\|_{L^2(\mu_d)}^2 = a^2 \sum_{j=1}^{M-1} I_{\{u^{(j)} \neq u'^{(j)}\}} \geq ka^2$$

and

$$\begin{aligned} \log |\tilde{\mathcal{F}}(k, a)| &= k \log 2 + \log |\mathcal{K}_{M-1,k}| \\ &\geq k \log 2 + 0.233k \log \frac{M-1}{k} \\ &\geq 0.233k \left(\log \frac{M-1}{k} + 1 \right). \end{aligned}$$

This proves Item 4.

Appendix B: Empirical process representation

In this appendix, we show that for each $u \in \mathbb{B}^k$, the norm $\|f_u - \widehat{f}_u\|_{L^2(\mu_{d-k})}$ can be expressed as a supremum of an empirical process over a suitable function class; this was a key element of the proof of Theorem 3.

First, we show that $\|f_u - \widehat{f}_u\|_{L^2(\mu_{d-k})}$ can be expressed as an empirical process of the form (30) indexed by a suitable function class. To this end, define

$$\mathcal{F} \triangleq \left\{ \sum_{v \in \mathbb{B}^{d-k}} \xi(v) \varphi_{uv} : \xi \in L^2(\mu_{d-k}), \|\xi\|_{L^2(\mu_{d-k})} \leq 1 \right\}.$$

Then

$$\|f_u - \widehat{f}_u\|_{L^2(\mu_{d-k})} = \sup_{g \in \mathcal{F}} \nu_n(g). \tag{B.1}$$

Indeed, let X_1, \dots, X_n be n i.i.d. copies of $X \sim f$. Then

$$\begin{aligned} \nu_n(g) &= \frac{1}{n} \sum_{i=1}^n g(X_i) - \mathbb{E}g(X) \\ &= \frac{1}{n} \sum_{i=1}^n \sum_{v \in \mathbb{B}^{d-k}} \xi(v) \varphi_{uv}(X_i) - \sum_{v \in \mathbb{B}^{d-k}} \xi(v) \theta_{uv} \\ &= \sum_{v \in \mathbb{B}^{d-k}} (\hat{\theta}_{uv} - \theta_{uv}) \xi(v) \\ &\leq \|\xi\|_{L^2(\mu_{d-k})} \|f_u - \hat{f}_u\|_{L^2(\mu_{d-k})}, \end{aligned}$$

where in the last line we used Cauchy–Schwarz. This proves (B.1). Next, we determine the constants L , v and H that are needed to apply Talagrand’s bound (29). For any $\xi \in L^2(\mu_{d-k})$ with unit norm, we have

$$\begin{aligned} \sup_{x \in \mathbb{B}^d} \left| \sum_{v \in \mathbb{B}^{d-k}} \xi(v) \varphi_{uv}(x) \right| &\leq \|\xi\|_{L^2(\mu_{d-k})} \sup_{x \in \mathbb{B}^d} \sqrt{\sum_{v \in \mathbb{B}^{d-k}} \varphi_{uv}^2(x)} \\ &\leq \sqrt{2^{d-k} 2^{-d}} \\ &= 2^{-k/2}. \end{aligned}$$

Hence, any $g \in \mathcal{F}$ is bounded by $L \equiv 2^{-k/2}$. From this, we also get the bound $\text{Var } g \leq v$ with $v = L^2 = 2^{-k}$. Finally, to bound the Rademacher average (28), we note that \mathcal{F} is the unit ball in the RKHS induced by the kernel

$$K_u(x, y) \triangleq \sum_{v \in \mathbb{B}^{d-k}} \varphi_{uv}(x) \varphi_{uv}(y), \quad \forall x, y \in \mathbb{B}^d.$$

Then standard arguments (see, e.g., Section 2.4.2 in [27]) lead to the bound

$$\mathbb{E} \left\{ \sup_{g \in \mathcal{F}} \sum_{i=1}^n \varepsilon_i g(U_i) \right\} \leq \sqrt{\frac{n}{2^k}},$$

which gives $H = 1/\sqrt{2^k n}$.

Acknowledgements

The authors wish to thank the Associate Editor and the anonymous referees for their meticulous reading of the manuscript and for making numerous valuable comments and suggestions that helped greatly improve the paper.

References

- [1] J. AITCHISON AND C. G. G. AITKEN. Multivariate binary discrimination by the kernel method. *Biometrika*, 63(3):413–420, 1976. [MR0443222](#)

- [2] R. R. BAHADUR. A representation of the joint distribution of n dichotomous items. In H. Solomon, editor, *Studies in Item Analysis and Prediction*, pages 169–176. Stanford Univ. Press, 1961. [MR0121894](#)
- [3] J. BERGH AND J. LÖFSTRÖM. *Interpolation Spaces: An Introduction*. Springer-Verlag, 1976. [MR0482275](#)
- [4] E. CANDÈS. Modern statistical estimation via oracle inequalities. *Acta Numerica*, 15:257–325, 2006. [MR2269743](#)
- [5] E. J. CANDÈS AND T. TAO. Near-optimal signal recovery from random projections: universal encoding strategies? *IEEE Trans. Inform. Theory*, 52(12):5406–5425, December 2006. [MR2300700](#)
- [6] J. M. CARRO. Estimating dynamic panel data discrete choice models with fixed effects. *J. Econometrics*, 140:503–528, 2007. [MR2408917](#)
- [7] X. R. CHEN, P. R. KRISHNAIAH, AND W. W. LIANG. Estimation of multivariate binary density using orthogonal functions. *J. Multivariate Anal.*, 31:178–186, 1989. [MR1021829](#)
- [8] T. M. COVER AND J. A. THOMAS. *Elements of Information Theory*. Wiley, New York, 2nd edition, 2006. [MR2239987](#)
- [9] I. DINUR, E. FRIEDGUT, G. KINDLER, AND R. O’DONNELL. On the Fourier tails of bounded functions over the discrete cube. *Israel J. Math.*, 160(389-412), 2007. [MR2342503](#)
- [10] D. L. DONOHO, I. M. JOHNSTONE, G. KERKYACHARIAN, AND D. PICARD. Density estimation by wavelet thresholding. *Ann. Statist.*, 24(2):508–539, 1996. [MR1394974](#)
- [11] S. EFROMOVICH. *Nonparametric Curve Estimation*. Springer, 1999. [MR1705298](#)
- [12] M. J. GARCÍA-ZATTERA, A. JARA, E. LESAFFRE, AND D. DECLERCK. Conditional independence of multivariate binary data with an application in caries research. *Computational Statistics and Data Analysis*, 51:3223–3234, 2007. [MR2345637](#)
- [13] Z. GHAHRAMANI AND K. HELLER. Bayesian sets. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 435–442. MIT Press, Cambridge, MA, 2006. [MR2441315](#)
- [14] A. C. GILBERT, S. GUHA, P. INDYK, S. MUTHUKRISHNAN, AND M. STRAUSS. Near-optimal sparse Fourier representations via sampling. In *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing*, 2002. [MR2121138](#)
- [15] A. C. GILBERT, S. MUTHUKRISHNAN, AND M. J. STRAUSS. Improved time bounds for near-optimal sparse Fourier representation via sampling. In *Proc. SPIE Wavelets XI*, San Diego, CA, 2005.
- [16] A. C. GILBERT AND M. J. STRAUSS. Group testing in statistical signal recovery. Preprint, 2006.
- [17] O. GOLDBREICH AND L. LEVIN. A hard-core predicate for all one-way functions. In *Proc. 21st ACM Symp. on Theory of Computing*, pages 25–32, 1989.

- [18] M. GYLLENBERG AND T. KOSKI. Probabilistic models for bacterial taxonomy. *International Statistical Review*, 69(2):249–276, August 2001.
- [19] P. HALL, G. KERKYACHARIAN, AND D. PICARD. Block threshold rules for curve estimation using kernel and wavelet methods. *Ann. Statist.*, 26(3):922–942, 1998. [MR1635418](#)
- [20] P. HALL, S. PENEV, G. KERKYACHARIAN, AND D. PICARD. Numerical performance of block thresholded wavelet estimators. *Statistics and Computing*, 7:115–124, 1997.
- [21] I. M. JOHNSTONE. Minimax Bayes, asymptotic minimax and sparse wavelet priors. In S. S. Gupta and J. O. Berger, editors, *Statistical Decision Theory and Related Topics V*, pages 303–326. Springer, 1994. [MR1286310](#)
- [22] E. KUSHILEVITZ AND Y. MANSOUR. Learning decision trees using the Fourier spectrum. *SIAM J. Comput.*, 22(6):1331–1348, 1993. [MR1247193](#)
- [23] S. L. LAURITZEN. *Graphical Models*. Clarendon Press, Oxford, 1996. [MR1419991](#)
- [24] W.-Q. LIANG AND P. R. KRISHNAIAH. Nonparametric iterative estimation of multivariate binary density. *J. Multivariate Anal.*, 16:162–172, 1985. [MR0778495](#)
- [25] Y. MANSOUR. Learning Boolean functions via the Fourier transform. In V. P. Roychodhury, K.-Y. Siu, and A. Orlitsky, editors, *Theoretical Advances in Neural Computation and Learning*, pages 391–424. Kluwer, 1994.
- [26] P. MASSART. *Concentration Inequalities and Model Selection*. Springer, 2007. [MR2319879](#)
- [27] S. MENDELSON. A few notes on statistical learning theory. In S. Mendelson and A. J. Smola, editors, *Advanced Lectures in Machine Learning*, volume 2600 of *Lecture Notes in Computer Science*. Springer, 2003.
- [28] J. OTT AND R. A. KRONMAL. Some classification procedures for multivariate binary data using orthogonal functions. *J. Amer. Stat. Assoc.*, 71(354):391–399, June 1976. [MR0403084](#)
- [29] P. REYNAUD-BOURET. Adaptive estimation of the intensity of inhomogeneous Poisson processes via concentration inequalities. *Probab. Th. Rel. Fields*, 126:103–153, 2003. [MR1981635](#)
- [30] H. P. ROSENTHAL. On the span in l_p of sequences of independent random variables. *Israel J. Math.*, 8:273–303, 1972. [MR0271721](#)
- [31] I. SHMULEVICH AND W. ZHANG. Binary analysis and optimization-based normalization of gene expression data. *Bioinformatics*, 18(4):555–565, 2002.
- [32] J. SILVA AND R. WILLETT. Hypergraph-based detection of anomalous high-dimensional co-occurrences. *IEEE Trans. Pattern Anal. Mach. Intel.*, 31(3):563–569, 2009.
- [33] J. S. SIMONOFF. Smoothing categorical data. *J. Statist. Planning and Inference*, 47:41–60, 1995. [MR1360958](#)
- [34] M. TALAGRAND. On Russo’s approximate zero-one law. *Ann. Probab.*, 22(3):1576–1587, 1994. [MR1303654](#)
- [35] M. TALAGRAND. Sharper bounds for Gaussian and empirical processes. *Ann. Probab.*, 22:28–76, 1994. [MR1258865](#)

- [36] T. TAO AND V. H. VU. *Additive Combinatorics*. Cambridge Univ. Press, 2006. [MR2289012](#)
- [37] A. W. VAN DER VAART AND J. A. WELLNER. *Weak Convergence and Empirical Processes*. Springer, 1996. [MR1385671](#)
- [38] J. D. WILBUR, J. K. GHOSH, C. H. NAKATSU, S. M. BROUDER, AND R. W. DOERGE. Variable selection in high-dimensional multivariate binary data with applications to the analysis of microbial community DNA fingerprints. *Biometrics*, 58:378–386, June 2002. [MR1908178](#)
- [39] Y. YANG AND A. BARRON. Information-theoretic determination of minimax rates of convergence. Technical Report 28, Department of Statistics, Iowa State University, 1997.
- [40] Y. YANG AND A. BARRON. Information-theoretic determination of minimax rates of convergence. *Ann. Statist.*, 27(5):1564–1599, 1999. [MR1742500](#)