GROUPED TRANSFORMATIONS AND REGULARIZATION IN HIGH-DIMENSIONAL EXPLAINABLE ANOVA APPROXIMATION

A PREPRINT

Felix Bartel
Faculty of Mathematics
Chemnitz University of Technology
09107 Chemnitz
felix.bartel@math.tu-chemnitz.de

Daniel Potts
Faculty of Mathematics
Chemnitz University of Technology
09107 Chemnitz
potts@math.tu-chemnitz.de

Michael Schmischke
Faculty of Mathematics
Chemnitz University of Technology
09107 Chemnitz
michael.schmischke@math.tu-chemnitz.de

November 7, 2021

ABSTRACT

In this paper we propose a tool for high-dimensional approximation based on trigonometric polynomials where we allow only low-dimensional interactions of variables. In a general high-dimensional setting, it is already possible to deal with special sampling sets such as sparse grids or rank-1 lattices. This requires black-box access to the function, i.e., the ability to evaluate it at any point. Here, we focus on scattered data points and grouped frequency index sets along the dimensions. From there we propose a fast matrix-vector multiplication, the grouped Fourier transform, for high-dimensional grouped index sets. Those transformations can be used in the application of the previously introduced method of approximating functions with low superposition dimension based on the analysis of variance (ANOVA) decomposition where there is a one-to-one correspondence from the ANOVA terms to our proposed groups. The method is able to dynamically detected important sets of ANOVA terms in the approximation. In this paper, we consider the involved least-squares problem and add different forms of regularization: Classical Tikhonov-regularization, namely, regularized least squares and the technique of group lasso, which promotes sparsity in the groups. As for the latter, there are no explicit solution formulas which is why we applied the fast iterative shrinking-thresholding algorithm to obtain the minimizer. Moreover, we discuss the possibility of incorporating smoothness information into the least-squares problem. Numerical experiments in under-, overdetermined, and noisy settings indicate the applicability of our algorithms. While we consider periodic functions, the idea can be directly generalized to non-periodic functions as well.

Keywords analysis of variance · ANOVA · explainable approximation · fast iterative shrinkage-thresholding algorithm · FISTA · group lasso · high dimensional approximation · multivariate trigonometric polynomials · nonequispaced fast Fourier transform · NFFT · LSQR

1 Introduction

Discrete transformations like the discrete Fourier transform, the discrete cosine transform, and many others play an important role in a large variety of applications in applied mathematics and other sciences. They have since given rise to algorithms like the fast Fourier transform (FFT) and the fast cosine transform that allow for the fast multiplication of the associated matrices, cf. [34]. The FFT with its many applications belongs to the most important algorithms of our

time. Those concepts have been generalized to allow for nodes to be nonequispaced resulting in the nonequispaced fast Fourier transform (NFFT) and the nonequispaced fast cosine transform, cf. [21, 34, 20]. While those algorithms have a good complexity, it depends on the size of the frequency index set which grows exponentially in the dimension for many common examples.

For high dimensional approximation with trigonometric polynomials there already exist efficient methods for special sampling sets, such as sparse grids [14, 16, 9] or rank-1 lattice nodes [17, 18]. In [35] the multivariate analysis of variance (ANOVA) decomposition, cf. [5, 39, 27, 26, 15, 12], was related to a decomposition in the frequency domain that allowed for the Fourier coefficients to be considered in subsets of coordinate indices related to the ANOVA terms which we propose as groups. If one assumes sparsity in the index sets related to sparsity in the ANOVA decomposition, it is possible to decompose high-dimensional nonequispaced discrete transforms into multiple transforms associated with ANOVA terms. In this paper, we introduce the grouped Fourier transformation which is able to handle high-dimensional frequency index sets if they have sparsity in their support, i.e., the number of nonzero elements in the frequencies is limited by a so called superposition threshold, cf. [35].

In [35] a method was proposed for the approximation of periodic functions with a low superposition dimension, see e.g. [5, 26, 6, 32, 13], and a sparse ANOVA decomposition. In other words, we assume that the variables of a function only interact in low-dimensional terms. From a theoretical viewpoint, it has been proven in [35, Section 4] that periodic functions of specific dominating-mixed smoothness with product and order-dependent weights yield low superposition dimensions and can therefore be approximated well by the method. In applications, one may assume that we have sparsity-of-effects or that the Pareto principle holds, see e.g. [44, 13], meaning that many real world systems are dominated by only a small number of low complexity interactions. The frequency index sets used in the method have a grouped structure and we have a one-to-one relation to the truncated ANOVA decomposition. The central part of the approximation approach is given by a least-squares problem with a Fourier matrix of a grouped index set which can be handled by the proposed grouped Fourier transform. Moreover, sensitivity analysis, cf. [30], is used in order to rank the influence of the dimension interactions or ANOVA terms and dynamically detect important ones. This can be displayed or interpreted as an explainable ANOVA network and may lead to a further truncation of the ANOVA decomposition to an *active set* of terms. One main advantage of the Fourier approach is that we are able to immediately compute an approximation to the variance of the ANOVA terms. The idea is in principle related to the multivariate decomposition method, cf. [25, 8], for the approximation of high- to infinite-dimensional integrals as well as functions, see [43]. However, we are considering the approximation of functions and work with random or scattered data instead of quasi-Monte Carlo methods. Moreover, our detection of the active set is dynamic and comes from sensitivity analysis instead of a-priori information.

In this paper, we use the grouped Fourier transform to solve the appearing least-squares problem and discuss enhancements by combing it with regularization approaches: The first method uses least-squares minimization with the LSQR algorithm, see [33], as in [35]. However, we add a Tikhonov regularization term and also discuss the incorporation of smoothness where we use Sobolev weights to adapt our algorithm for different decay properties of the Fourier coefficients. Here, the sensitivity analysis on the approximation lets us identify important terms which lead to an active set. Solving the least-squares system with this active set yields the benefit of a simpler model function and improves approximation quality, cf. results in [35].

In the second approach we try to combine both steps, i.e., approximation and identification of an active set of terms by using the group lasso approach from [46] which immediately promotes sparsity in the groups. This approach has been used and adapted to several scenarios, see e.g. [28, 45, 40]. Since we set the groups as the ANOVA terms, it refers to sparsity in the ANOVA decomposition, i.e., the identification of an active set of terms. This leads to a non-linear minimization problem, which we solve by using the fast iterative shrinking-thresholding Algorithm (FISTA), cf. [3]. To our knowledge this regularization was only done for ℓ_2-norms of groups, but here we were able to incorporate smoothness information by using Sobolev-type norms.

The paper is organized as follows: In Section 2 we reiterate on the ANOVA decomposition for periodic functions as discussed in [35] and introduce grouped index sets. In Section 3 we introduce the grouped Fourier transform for frequency sets of a grouped structure and propose a fast algorithm to compute it. Furthermore, we discuss complexity as well as the error of this algorithm. We apply the grouped Fourier transform in Section 4 to the least-squares problem (14) of approximating high-dimensional functions with the method proposed in [35]. In the Section 4.1 we add a ℓ_2 regularization term and therein use information about the smoothness of the function. The solver in this case is the well-known LSQR method, cf. [33]. Section 4.2 contains the group lasso regularization idea using FISTA, see Algorithm 1, as a solver. Here, the idea to use smoothness information is considered as well. We provide numerical examples with a special test function for both methods. In addition, we show a possible extension to non-periodic functions with an example from applications in Section 5.

2 Classical ANOVA Decomposition

In this section we discuss the main properties of the *classical ANOVA decomposition*, see [5, 39, 27, 26, 15, 12], for periodic functions $f\colon \mathbb{T}^d \to \mathbb{C}$ defined over the d-dimensional torus $\mathbb{T} \coloneqq \mathbb{R}/\mathbb{Z}$. The decomposition has proven useful in understanding the reason behind the success of certain quadrature methods for high-dimensional integration [29, 4, 10] and also infinite-dimensional integration [1, 11, 25]. For a general description of multivariate decompositions, we refer to [26] and for the classical ANOVA decomposition of periodic functions to [35].

For a function $f \in \mathrm{L}_2(\mathbb{T}^d)$ with spatial dimension $d \in \mathbb{N}$, we introduce the integral projections

$$\mathrm{P}_{\boldsymbol{u}} f(\boldsymbol{x}) = \int_{\mathbb{T}^{d-|\boldsymbol{u}|}} f(\boldsymbol{x}) \mathrm{d}\boldsymbol{x}_{\boldsymbol{u}^{\mathrm{c}}} \tag{1}$$

for a subset of coordinate indices $\boldsymbol{u} \subseteq \mathcal{D} \coloneqq \{1, 2, \dots, d\}$ and its complement $\boldsymbol{u}^{\mathrm{c}} = \mathcal{D} \setminus \boldsymbol{u}$. Additionally, for vectors $\boldsymbol{x} \in \mathbb{C}^d$ indexed with a subset $\boldsymbol{u} \subseteq \mathcal{D}$ we define $\boldsymbol{x}_{\boldsymbol{u}} \coloneqq (x_i)_{i \in \boldsymbol{u}}$. This leads to the ANOVA terms

$$f_{\boldsymbol{u}}(\boldsymbol{x}) \coloneqq \mathrm{P}_{\boldsymbol{u}} f(\boldsymbol{x}) - \sum_{\boldsymbol{v} \subsetneq \boldsymbol{u}} f_{\boldsymbol{v}}(\boldsymbol{x}) = \sum_{\substack{\boldsymbol{k} \in \mathbb{Z}^d \\ \operatorname{supp} \boldsymbol{k} = \boldsymbol{u}}} \mathrm{c}_{\boldsymbol{k}}(f)\, \mathrm{e}^{2\pi \mathrm{i} \langle \boldsymbol{k}, \boldsymbol{x} \rangle}, \boldsymbol{u} \subseteq \mathcal{D}, \tag{2}$$

with $\operatorname{supp} \boldsymbol{k} = \{j \in \mathcal{D}\colon k_j \neq 0\}$, and Fourier coefficients

$$\mathrm{c}_{\boldsymbol{k}}(f) = \int_{\mathbb{T}^d} f(\boldsymbol{x})\, \mathrm{e}^{-2\pi \mathrm{i} \langle \boldsymbol{k}, \boldsymbol{x} \rangle}\, \mathrm{d}\boldsymbol{x}.$$

Then the classical ANOVA decomposition of f is given by

$$f = \sum_{\boldsymbol{u} \subseteq \mathcal{D}} f_{\boldsymbol{u}}.$$

The uniqueness of the decomposition as well as the relation to the series expansion (2) have been proven in [35]. Note that our specific choice of the integral projections (1) lead to the classical ANOVA decomposition. A different choice of projection is also possible, e.g., the anchored variant used in the multivariate decomposition method, see [25, 8].

In order to measure the importance of an ANOVA term $f_{\boldsymbol{u}}$ in relation to the function f, we use *global sensitivity indices* or *Sobol indices*

$$\varrho(\boldsymbol{u}, f) \coloneqq \frac{\sigma^2(f_{\boldsymbol{u}})}{\sigma^2(f)} \quad \text{with} \quad \sigma^2(f) \coloneqq \|f\|^2_{\mathrm{L}_2(\mathbb{T}^d)} - |\mathrm{c}_{\boldsymbol{0}}(f)|^2, \tag{3}$$

cf. [41, 42, 27]. This motivates the concept of effective dimensions. The superposition dimension is a particular notion of effective dimension and is given as

$$d^{(\mathrm{sp})} = \min \left\{ s \in \mathcal{D}\colon \sum_{\substack{\boldsymbol{u} \subseteq \mathcal{D} \\ |\boldsymbol{u}| \leq s}} \sigma^2(f_{\boldsymbol{u}}) \geq \alpha \sigma^2(f) \right\} \tag{4}$$

for an accuracy $\alpha \in [0, 1]$. In other words, it is the smallest integer $s \in \mathcal{D}$ such that the variance of the function can mostly be explained by up to s-dimensional ANOVA terms.

The number of ANOVA terms is 2^d and therefore grows exponentially in the dimension which reflects the well-known curse of dimensionality. In order to circumvent the curse, we use the idea of truncating the ANOVA decomposition and only taking a certain number of terms into account. A subset U of the power set $\mathcal{P}(\mathcal{D})$ is called *subset of ANOVA terms* if for every $\boldsymbol{u} \in U$, all subsets $\boldsymbol{v} \subseteq \boldsymbol{u}$ are also elements of U. A special truncation possibility is motivated by the superposition dimension $d^{(\mathrm{sp})} \in \mathcal{D}$. We may form a set of terms

$$U_{d_s} \coloneqq \{\boldsymbol{u} \subseteq \mathcal{D}\colon\ |\boldsymbol{u}| \leq d_s\} \tag{5}$$

where the order is limited by a superposition threshold $d_s \in \mathcal{D}$, see e.g. [15, 35]. The condition that a subset of ANOVA terms $U \subseteq \mathcal{P}(\mathcal{D})$ has to be downward closed can be relaxed of one assumes $f_{\boldsymbol{u}} \equiv 0$ for $\boldsymbol{u} \notin U$.

For a subset of ANOVA terms $U \subseteq \mathcal{P}(\mathcal{D})$, we define the truncation

$$\mathrm{T}_U f \coloneqq \sum_{\boldsymbol{u} \in U} f_{\boldsymbol{u}}. \tag{6}$$

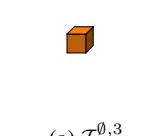

(a) $\mathcal{I}_6^{\emptyset,3}$

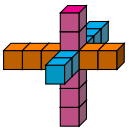

(b) magenta: $\mathcal{I}_6^{\{1\},3}$, orange: $\mathcal{I}_6^{\{2\},3}$, and blue: $\mathcal{I}_6^{\{3\},3}$

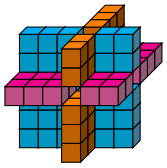

(c) magenta: $\mathcal{I}_6^{\{2,3\},3}$, orange: $\mathcal{I}_6^{\{1,3\},3}$, and blue: $\mathcal{I}_6^{\{1,2\},3}$

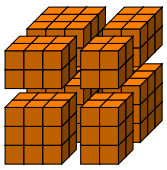

(d) $\mathcal{I}_6^{\{1,2,3\},3}$

Figure 1: Decomposition of $\mathcal{I}_6^3$ into its groups.

and the special case $\mathrm{T}_{d_s} f \coloneqq \mathrm{T}_{U_{d_s}} f$. In [35] it has been shown that functions of certain types of smoothness, e.g. dominating-mixed smoothness, have a low superposition dimension and can therefore be approximated well by a truncation with $\mathrm{T}_{d_s} f$, see also [32].

To obtain finite dimensions, we truncate the series expansion (2) to finite frequency sets which we introduce in the following. We start with the one-dimensional frequency set $\mathcal{I}_N \coloneqq \mathbb{Z} \cap [-N/2, N/2)$, $N \in 2\mathbb{N}$, with $2\mathbb{N}$ being the even natural numbers. For $N \in 2\mathbb{N}$ and $\boldsymbol{u} \subseteq \mathcal{D}$ we define

$$\mathcal{I}_N^{\boldsymbol{u},d} \coloneqq \left\{\boldsymbol{k} \in \mathbb{Z}^d \colon \operatorname{supp} \boldsymbol{k} = \boldsymbol{u} \text{ and } k_j \in \mathcal{I}_N \text{ for } j \in \boldsymbol{u}\right\}. \tag{7}$$

Remark 2.1. *For a bandwidth $N \in 2\mathbb{N}$ we have $\{\boldsymbol{\ell}_{\boldsymbol{u}} \colon \boldsymbol{\ell} \in \mathcal{I}_N^{\boldsymbol{u},d}\} = (\mathcal{I}_N \setminus \{0\})^{|\boldsymbol{u}|}$ and, thus, we obtain for the cardinality $|\mathcal{I}_N^{\boldsymbol{u},d}| = (N-1)^{|\boldsymbol{u}|}$.*

For a given subset of ANOVA terms $U \subseteq \mathcal{P}(\mathcal{D})$ and bandwidths $N_{\boldsymbol{u}} \in 2\mathbb{N}$, $\boldsymbol{u} \in U$, we define a *grouped index set* as the disjoint union

$$\mathcal{I}_{\boldsymbol{N}}(U) \coloneqq \bigcup_{\boldsymbol{u} \in U} \mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d} \tag{8}$$

with $\boldsymbol{N} \coloneqq (N_{\boldsymbol{u}})_{\boldsymbol{u} \in U}$. The corresponding Fourier partial sum is then given by

$$S_{\mathcal{I}_{\boldsymbol{N}}(U)} f(\boldsymbol{x}) \coloneqq \sum_{\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}(U)} c_{\boldsymbol{k}}(f)\, \mathrm{e}^{2\pi \mathrm{i} \langle \boldsymbol{k}, \boldsymbol{x} \rangle}. \tag{9}$$

Remark 2.2. *Setting $N_{\boldsymbol{u}} = N \in 2\mathbb{N}$ for a bandwidth $N \in 2\mathbb{N}$ and $\boldsymbol{u} \subseteq \mathcal{D}$, we have a disjoint union of the hypercube*

$$\mathcal{I}_N^d = \mathcal{I}_{(N,N,\dots,N)}(\mathcal{P}(\mathcal{D})) = \bigcup_{\boldsymbol{u} \subseteq \mathcal{D}} \mathcal{I}_N^{\boldsymbol{u},d}. \tag{10}$$

This identity is visualized in Figure 1.

3 Grouped Fourier Transform

The *nonequispaced discrete Fourier transformation* in $d \in \mathbb{N}$ dimensions is given by

$$f_{\boldsymbol{x}} = \sum_{\boldsymbol{k} \in \mathcal{I}_N^d} \hat{f}_{\boldsymbol{k}}\, \mathrm{e}^{2\pi \mathrm{i} \langle \boldsymbol{k}, \boldsymbol{x} \rangle}$$

for $\boldsymbol{x} \in \mathcal{X}$, where $\mathcal{X} \subseteq \mathbb{T}^d$ is an arbitrary finite set of nodes on the torus $\mathbb{T}$. Moreover, we have the index set $\mathcal{I}_N^d \subseteq \mathbb{Z}^d$ from (10), coefficients $\hat{f}_{\boldsymbol{k}} \in \mathbb{C}$, $\boldsymbol{k} \in \mathcal{I}_N^d$, and $\langle \cdot, \cdot \rangle$ the Euclidean inner product, see [34, Chapter 7]. The related Fourier matrix is given by

$$\boldsymbol{F}(\mathcal{X}, \mathcal{I}) = \left(\mathrm{e}^{2\pi \mathrm{i} \langle \boldsymbol{k}, \boldsymbol{x} \rangle}\right)_{\boldsymbol{x} \in \mathcal{X}, \boldsymbol{k} \in \mathcal{I}} \in \mathbb{C}^{|\mathcal{X}| \times |\mathcal{I}|}, \tag{11}$$

where a straightforward matrix-vector multiplication requires $\mathcal{O}(|\mathcal{X}|\,|\mathcal{I}|)$ arithmetical operations.

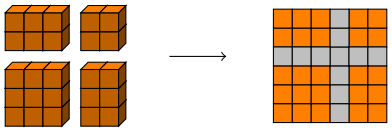

Figure 2: Visualization of the extension for $\mathcal{I}_6^{\{2,3\},3}$ to $\mathcal{I}_6^2$.

In the following, we introduce the *grouped Fourier transform* as a special case with the frequency index sets $\mathcal{I}_{\boldsymbol{N}}(U)$, $U \subseteq \mathcal{P}(\mathcal{D})$, $\boldsymbol{N} \in (2\mathbb{N})^{|U|}$, as in (8), i.e., they consist of a disjoint union of lower dimensional frequency index sets with zeros along some dimensions. Index sets of this type were also used in [35]. In this section, we discuss all necessities to carry out the computations in lower dimensions instead of the high spatial dimension. This will peak in the identity (13), which is then used to develop fast algorithms.

Definition 3.1. *Let U be a subset of $\mathcal{P}(\mathcal{D})$, where $\mathcal{P}$ stands for the powerset, and $\mathcal{I}(U)$ a grouped index set of structure* (8) *with $N_{\boldsymbol{u}} \in 2\mathbb{N}$, $\boldsymbol{u} \in U$, bandwidth parameters, and $\hat{f}_{\boldsymbol{k}} \in \mathbb{C}$ coefficients for $\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}(U)$. On an arbitrary set of nodes $\boldsymbol{x} \in \mathcal{X}$, the* grouped Fourier transform *is then given by*

$$f_{\boldsymbol{x}} = \sum_{\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}(U)} \hat{f}_{\boldsymbol{k}} \, \mathrm{e}^{2\pi \mathrm{i} \langle \boldsymbol{k}, \boldsymbol{x} \rangle} = \sum_{\boldsymbol{u} \in U} \sum_{\boldsymbol{k} \in \mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}} \hat{f}_{\boldsymbol{k}} \, \mathrm{e}^{2\pi \mathrm{i} \langle \boldsymbol{k}, \boldsymbol{x} \rangle}.$$

Computing the grouped Fourier transform from Definition 3.1 via the naive approach, i.e., setting up the Fourier matrix $\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))$ and performing a matrix-vector multiplication with $\hat{\boldsymbol{f}} = (\hat{\boldsymbol{f}}(\boldsymbol{u}))_{\boldsymbol{u} \in U}$ and $\hat{\boldsymbol{f}}(\boldsymbol{u}) = (\hat{f}_{\boldsymbol{k}})_{\boldsymbol{k} \in \mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}}$ results in the complexity class

$$\mathcal{O}\left(|\mathcal{X}| \sum_{\boldsymbol{u} \in U} |\mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}| \right) = \mathcal{O}\left(|\mathcal{X}| \sum_{\boldsymbol{u} \in U} (N_{\boldsymbol{u}} - 1)^{|\boldsymbol{u}|} \right),$$

cf. Remark 2.1

This is infeasible in high dimensions, but can be reduced as follows. First, we observe that we are working with a block matrix, i.e.,

$$\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U)) = \begin{pmatrix} \boldsymbol{F}(\mathcal{X}, \mathcal{I}_{N_{\boldsymbol{u}_1}}^{\boldsymbol{u}_1,d}) & \boldsymbol{F}(\mathcal{X}, \mathcal{I}_{N_{\boldsymbol{u}_2}}^{\boldsymbol{u}_2,d}) & \cdots & \boldsymbol{F}(\mathcal{X}, \mathcal{I}_{N_{\boldsymbol{u}_{|U|}}}^{\boldsymbol{u}_{|U|},d}) \end{pmatrix} \tag{12}$$

when introducing an order on the sets $\boldsymbol{u}$. Using the sparse structure of $\mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}$, we obtain

$$\begin{aligned} \boldsymbol{F}\left(\mathcal{X}, \mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}\right) &= \left(\mathrm{e}^{2\pi \mathrm{i} \langle \boldsymbol{k}, \boldsymbol{x} \rangle} \right)_{\boldsymbol{x} \in \mathcal{X}, \boldsymbol{k} \in \mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}} \\ &= \left(\mathrm{e}^{2\pi \mathrm{i} \langle \boldsymbol{\ell}, \boldsymbol{x}_{\boldsymbol{u}} \rangle} \right)_{\boldsymbol{x} \in \mathcal{X}, \boldsymbol{\ell} \in (\mathcal{I}_{N_{\boldsymbol{u}}} \setminus \{0\})^{|\boldsymbol{u}|}}. \end{aligned}$$

For $\hat{\boldsymbol{f}}(\boldsymbol{u})$ with frequency domain $\mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}$ we now introduce an extension such that the coefficients corresponding to zero-frequencies are filled with zeros and the frequency-dimension is reduced to $|\boldsymbol{u}|$. Thus, we have frequencies in $\mathcal{I}_{N_{\boldsymbol{u}}}^{|\boldsymbol{u}|}$. This extension is visualized in Figure 2 and can be achieved formally by defining a vector $\hat{\boldsymbol{g}}(\boldsymbol{u}) \in \mathbb{C}^{|\mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}|}$ with entries

$$(\hat{\boldsymbol{g}}(\boldsymbol{u}))_{\boldsymbol{\ell} \in \mathcal{I}_{N_{\boldsymbol{u}}}^{|\boldsymbol{u}|}} = \begin{cases} \hat{f}_{\boldsymbol{h}} & : \operatorname{supp} \boldsymbol{\ell} = \{1, 2, \ldots, |\boldsymbol{u}|\} \\ 0 & : \text{otherwise} \end{cases}$$

with $\boldsymbol{h} \in \mathbb{Z}^d$ the frequency such that $\boldsymbol{h}_{\boldsymbol{u}} = \boldsymbol{\ell}$ and $\boldsymbol{h}_{\boldsymbol{u}^{\mathrm{c}}} = \boldsymbol{0}$.

We obtain the identity

$$\boldsymbol{F}\left(\mathcal{X}, \mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}\right) \hat{\boldsymbol{f}}(\boldsymbol{u}) = \boldsymbol{F}\left(\mathcal{X}_{\boldsymbol{u}}, \mathcal{I}_{N_{\boldsymbol{u}}}^{|\boldsymbol{u}|}\right) \hat{\boldsymbol{g}}(\boldsymbol{u}), \tag{13}$$

with $\mathcal{X}_{\boldsymbol{u}} \coloneqq \{\boldsymbol{x}_{\boldsymbol{u}} \colon \boldsymbol{x} \in \mathcal{X}\}$ where the left-hand side operates in d dimensions and the right-hand side in $|\boldsymbol{u}|$ dimensions. Note that $\mathcal{X}_{\boldsymbol{u}}$ is a multiset, i.e., it may contain duplicates. Furthermore, the right-hand side of (13) can be computed via a nonequispaced fast Fourier transform (NFFT), see [21, 34, 20], in the complexity $\mathcal{O}(N_{\boldsymbol{u}}^{|\boldsymbol{u}|} \log N_{\boldsymbol{u}} + |\log \varepsilon|^{|\boldsymbol{u}|} |\mathcal{X}|)$ for precision ε. In summary, we have

$$\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U)) \hat{\boldsymbol{f}} = \sum_{\boldsymbol{u} \in U} \boldsymbol{F}\left(\mathcal{X}, \mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}\right) \hat{\boldsymbol{f}}(\boldsymbol{u}) = \sum_{\boldsymbol{u} \in U} \boldsymbol{F}\left(\mathcal{X}_{\boldsymbol{u}}, \mathcal{I}_{N_{\boldsymbol{u}}}^{|\boldsymbol{u}|}\right) \hat{\boldsymbol{g}}(\boldsymbol{u})$$

which means that we have reduced one d-dimensional NFFT to $|U|$ low-dimensional NFFTs. While the curse of dimensionality effectively hinders fast multiplication with $\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))\hat{\boldsymbol{f}}$ in d dimensions we only need up to d_s-dimensional NFFTs for $U \subseteq U_{d_s}$ and the number of terms in U_{d_s} grows only polynomial in the spatial dimension d for a fixed d_s, see (4). Moreover, the low-dimensional transformations can be computed independent of each other which allows for parallelization. Given enough resources, it would be possible to compute multi-threaded variants of each single transformation simultaneously which yields a high benefit in execution time. It follows that the multiplication with $\boldsymbol{F}$ can be accomplished in

$$\mathcal{O}\left(\sum_{\boldsymbol{u}\in U}\left(N_{\boldsymbol{u}}^{|\boldsymbol{u}|}\log N_{\boldsymbol{u}} + |\log\varepsilon|^{|\boldsymbol{u}|}|\mathcal{X}|\right)\right)$$

by the NFFT for precision ε.

Remark 3.2. *For an error estimate we have for the NFFT*

$$\left|\sum_{\boldsymbol{k}\in\mathcal{I}_{N_{\boldsymbol{u}}}^{\boldsymbol{u},d}} \hat{f}_{\boldsymbol{k}} \mathrm{e}^{2\pi\mathrm{i}\langle\boldsymbol{k},\boldsymbol{x}\rangle} - \mathit{NFFT}\left(\hat{\boldsymbol{g}}(\boldsymbol{u})\right)\right| \leq C_{\boldsymbol{u}}(N_{\boldsymbol{u}})\left\|\hat{\boldsymbol{f}}(\boldsymbol{u})\right\|_1$$

where $C_{\boldsymbol{u}}(N_{\boldsymbol{u}})$ decays exponentialy in $N_{\boldsymbol{u}}$, cf. [34, Theorem 7.8] and references therein. Since the grouped Fourier transform is a sum of the above, we can bound the overall error by

$$\left(\max_{\boldsymbol{u}\in U} C_{\boldsymbol{u}}(N_{\boldsymbol{u}})\right)\left\|\hat{\boldsymbol{f}}\right\|_1.$$

The adjoint problem, i.e., the multiplication of $\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))^{\mathsf{H}}$ with a vector $\boldsymbol{f} \in \mathbb{C}^{|\mathcal{X}|}$ can be realized in the same fashion. We have

$$\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))^{\mathsf{H}}\boldsymbol{f} = \begin{pmatrix} \boldsymbol{F}\left(\mathcal{X}, \mathcal{I}_{N_{\boldsymbol{u}_1}}^{\boldsymbol{u}_1,d}\right)^{\mathsf{H}}\boldsymbol{f} \\ \vdots \\ \boldsymbol{F}\left(\mathcal{X}, \mathcal{I}_{N_{\boldsymbol{u}_{|U|}}}^{\boldsymbol{u}_{|U|},d}\right)^{\mathsf{H}}\boldsymbol{f} \end{pmatrix} = \begin{pmatrix} \boldsymbol{F}\left(\mathcal{X}_{\boldsymbol{u}_1}, \mathcal{I}_{N_{\boldsymbol{u}_1}}^{|\boldsymbol{u}_1|}\right)^{\mathsf{H}}\boldsymbol{f} \\ \vdots \\ \boldsymbol{F}\left(\mathcal{X}_{\boldsymbol{u}_{|U|}}, \mathcal{I}_{N_{\boldsymbol{u}_{|U|}}}^{|\boldsymbol{u}_{|U|}|}\right)^{\mathsf{H}}\boldsymbol{f} \end{pmatrix}$$

which yields the same benefits. We can decompose our d-dimensional adjoint transform in $|U|$ low-dimensional transforms and are able to compute all of them simultaneously by using parallelization.

Remark 3.3. *The considerations in this section are not limited to the exponential functions, but rather, can be extended to other complete orthonormal systems. In fact, for $\exp(2\pi\mathrm{i}\langle\boldsymbol{k},\cdot\rangle)$ and $\cos(\pi\langle\boldsymbol{k},\cdot\rangle)$ the presented ideas are already implemented in the julia package `GroupedTransforms` which can be found on GitHub, see `github.com/NFFT/GroupedTransforms`, where the nonequispaced fast cosine transform (NFCT) is used for performance in the latter. We refer to the references [37, 36].*

4 Approximation

In this section, we consider the problem of approximating periodic functions $f\colon \mathbb{T}^d \to \mathbb{C}$ with high spatial dimension $d \in \mathbb{N}$. In particular, we have a scattered data setting where the given data about the unknown function f are sampling values $\boldsymbol{y} = (f(\boldsymbol{x}) + \eta)_{\boldsymbol{x}\in\mathcal{X}}$ which may contain noise η on a finite set of nodes $\mathcal{X} \subseteq \mathbb{T}^d$. We are looking for an approximation with a Fourier partial sum (9) of the truncated ANOVA decomposition and want to incorporate the fast algorithms from Section 3. We assume that the function has a low superposition dimension (4). This can occur either because it belongs to a function class that has this property, see e.g. [35], or it stems from a real world application with sparsity-of-effects.

We use the method proposed in [35] for approximation, which suggests to choose a superposition threshold $d_s \in \mathcal{D}$. In particular, we will use the ANOVA terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} \in U_{d_s}$, from (5) with grouped frequency index set $\mathcal{I}_{\boldsymbol{N}}(U_{d_s})$, see (8), and bandwidth parameters $N_{\boldsymbol{u}} \in 2\mathbb{N}$. Since the Fourier coefficients $\mathrm{c}_{\boldsymbol{k}}(f) \in \mathbb{C}$, $\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}(U_{d_s})$, are not known, we are going to approximate them by coefficients $\hat{f}_{\boldsymbol{k}} \approx \mathrm{c}_{\boldsymbol{k}}(f)$ that we determine from the given data $\mathcal{X}$ and $\boldsymbol{y}$. We obtain an approximate Fourier partial sum

$$S_{\mathcal{I}_{\boldsymbol{N}}(U_{d_s})}^{\mathcal{X}} f(\boldsymbol{x}) \coloneqq \sum_{\boldsymbol{k}\in\mathcal{I}_{\boldsymbol{N}}(U_{d_s})} \hat{f}_{\boldsymbol{k}}\, \mathrm{e}^{2\pi\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}}$$

with $\hat{f}_{\boldsymbol{k}} \in \mathbb{C}$ such that $S_{\mathcal{I}_{\boldsymbol{N}}(U_{d_s})}^{\mathcal{X}} f(\boldsymbol{x}) \approx S_{\mathcal{I}_{\boldsymbol{N}}(U_{d_s})} f(\boldsymbol{x})$.

For any subset of ANOVA terms $U \subseteq U_{d_s}$ (with equality possible) computing the approximate coefficients $\hat{f}_{\boldsymbol{k}}$ is done via solving the least-squares problem

$$\min_{\hat{\boldsymbol{f}}} \frac{1}{2} \left\| \boldsymbol{y} - \boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U)) \hat{\boldsymbol{f}} \right\|_2^2 \tag{14}$$

with grouped index set $\mathcal{I}_{\boldsymbol{N}}(U)$, bandwidth parameters $N_{\boldsymbol{u}} \in 2\mathbb{N}$, $\boldsymbol{u} \in U$, and Fourier matrix (11). For $|\mathcal{X}| > |\mathcal{I}_{\boldsymbol{N}}(U)|$ and $\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))$ full rank, the system has a unique solution. Moreover, if $\mathcal{X}$ are uniformly distributed i.i.d. nodes and the oversampling factor is large enough, the matrix has full rank with high probability and we get an estimate for the approximation error, cf. [35, Theorem 6.7] and [19, Section 5.2].

In Section 4.1 we use the iterative LSQR method to solve (14). However, we are going to extend this by adding an ℓ_2 regularization term known from the ordinary Tikhonov regularization and simultaneously consider weights which incorporate information about the decay rate of the Fourier coefficients, i.e., the smoothness of the function. In particular, we use weights corresponding to functions from a Sobolev type space

$$\mathrm{H}^{\omega}(\mathbb{T}^d) \coloneqq \left\{ f \in \mathrm{L}_2(\mathbb{T}^d) \colon \|f\|_{\mathrm{H}^{\omega}(\mathbb{T}^d)} \coloneqq \sqrt{\sum_{\boldsymbol{k} \in \mathbb{Z}^d} \omega^2(\boldsymbol{k}) \left| \mathrm{c}_{\boldsymbol{k}}(f) \right|^2} < \infty \right\}$$

for a weight function $\omega \colon \mathbb{Z}^d \to [1, \infty)$, see e.g. [23, 35]. We use sensitivity analysis on $S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_{d_s})} f(\boldsymbol{x})$ to detect which of the ANOVA terms $f_{\boldsymbol{u}}$ for $\boldsymbol{u} \in U_{d_s}$ are of a high-importance to the function and determine an active set

$$U^* = U^{(\boldsymbol{\varepsilon})}_{\mathcal{X}, \boldsymbol{y}} \coloneqq \emptyset \cup \{ \boldsymbol{u} \subseteq \mathcal{D} \colon \varrho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_{d_s})} f) > \varepsilon_{|\boldsymbol{u}|} \} \tag{15}$$

with threshold vector $\boldsymbol{\varepsilon} \in (0,1)^{d_s}$. Solving the problem (14) with $\mathcal{I}_{\boldsymbol{N}}(U^*)$ and new bandwidth parameters $N_{\boldsymbol{u}} \in 2\mathbb{N}$, $\boldsymbol{u} \in U^*$, yields the approximation $S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U^*)} f(\boldsymbol{x})$.

Our goal is to compare how ℓ_2 regularization impacts the addition of noise which was not considered in [35] and has been formulated as an open problem. Moreover, we discuss the addition of known Sobolev smoothness information leading to a smaller *search space* which can be seen by using the equivalent Ivanov formulation of the regularization functional, cf. [31], where we minimize the data fitting term subject to a bound on the Sobolev norm. This allows for an oversampling factor $|X| \,/\, |\mathcal{I}_{\boldsymbol{N}}(U)|$ close to or smaller than one. In this case, we will even be in an undetermined setting, i.e., $|X| < |\mathcal{I}_{\boldsymbol{N}}(U)|$.

In Section 4.2 we use a regularization technique to introduce an approach that combines the two steps from before, i.e., approximation and active set detection. The method is a variation of the group lasso technique from [46] that promotes sparsity in defined groups. In our case, the groups will be the ANOVA terms or subsets of coordinate indices $\boldsymbol{u} \in U_{d_s}$. As in the prior case, we add Sobolev weights which will be used to incorporate smoothness information about the function.

Both approaches rely on iterative solvers which require fast multiplication with $\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))$ and its adjoint which is provided by the grouped Fourier transform from Section 3. As a quality measure we use the approximation error

$$\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U)) \coloneqq \frac{1}{\|f\|_{\mathrm{L}_2(\mathbb{T}^d)}} \left\| f - S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U)} f \right\|_{\mathrm{L}_2(\mathbb{T}^d)} \tag{16}$$

which can be computed if the norm and the Fourier coefficients of the function are known for the purpose of computing the error via Parseval's identity

$$\left\| f - S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U)} f \right\|^2_{\mathrm{L}_2(\mathbb{T}^d)} = \|f\|^2_{\mathrm{L}_2(\mathbb{T}^d)} + \sum_{\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}(U)} \left| \hat{f}_{\boldsymbol{k}} - \mathrm{c}_{\boldsymbol{k}}(f) \right|^2 - \sum_{\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}(U)} \left| \mathrm{c}_{\boldsymbol{k}}(f) \right|^2.$$

For comparability reasons, we work with the same test function throughout this section

$$f : \mathbb{T}^9 \to \mathbb{R},$$
$$\boldsymbol{x} \mapsto B_2(x_1) B_4(x_3) B_6(x_8) + B_2(x_2) B_4(x_5) B_6(x_6) + B_2(x_4) B_4(x_7) B_6(x_9), \tag{17}$$

where B_2, B_4 and B_6 are parts of univariate, shifted, scaled, and dilated B-splines of order 2, 4, and 6, respectively. Their Fourier series is given by

$$B_j(x) \coloneqq c_j \sum_{k \in \mathbb{Z}} \operatorname{sinc}^j \left(\frac{\pi \cdot k}{j} \right) \cos(\pi \cdot k) \, \mathrm{e}^{2\pi \mathrm{i} k \cdot x} \quad \text{for} \quad j = 2, 4, 6,$$

with $\operatorname{sinc}(x) \coloneqq \sin(x)/x$ and the three normalization constants $c_2 \coloneqq \sqrt{3/4}$, $c_4 \coloneqq \sqrt{315/604}$, $c_6 \coloneqq \sqrt{277200/655177}$ such that $\|B_j\|_{\mathrm{L}_2(\mathbb{T})} = 1$, $j = 2, 4, 6$.

The test function works well for computing the error (16) since we have the Fourier coefficients $\mathrm{c}_{\boldsymbol{k}}(f)$ and the norm $\|f\|_{\mathrm{L}_2(\mathbb{T}^d)}$ explicitly given. Moreover, The function f has superposition dimension $d^{(\mathrm{sp})} = 3$, see (4), for arbitrary high accuracy, i.e., it can be represented by at most three-dimensional ANOVA terms with $f = \mathrm{T}_3 f$. This leads to $d_s = 3$ being the optimal choice for the superposition threshold with no error caused by the ANOVA truncation. We have the active set of terms

$$\boldsymbol{u} \in U^* \coloneqq \mathcal{P}(\{1,3,8\}) \cup \mathcal{P}(\{2,5,6\}) \cup \mathcal{P}(\{4,7,9\})$$

with $f_{\boldsymbol{u}} = 0$ for $\boldsymbol{u} \notin U^*$. The function also has dominating-mixed smoothness of $3/2 - \varepsilon$ for every $\varepsilon > 0$, i.e., $f \in \mathrm{H}^{\omega_\varepsilon}(\mathbb{T}^9)$ with

$$\omega_\varepsilon(\boldsymbol{k}) = \prod_{j \in \operatorname{supp} \boldsymbol{k}} (1 + |k_j|)^{\frac{3}{2}-\varepsilon},$$

cf. [38].

The ANOVA terms $f_{\boldsymbol{u}}$ can be computed analytically such that we obtain

$$\begin{aligned} f_\emptyset &= 3 \prod_{j \in \{2,4,6\}} c_j \\ f_{\{i\}}(x_i) &= \frac{\prod_{j \in \{2,4,6\}} c_j}{c_{\mathrm{o}(i)}} \left(B_{\mathrm{o}(i)}(x_i) - c_{\mathrm{o}(i)}\right) \end{aligned} \tag{18}$$

for the constant and the one-dimensional terms with $i = 1, 2, \ldots, 9$ and

$$\mathrm{o}(i) \coloneqq \begin{cases} 2 & : i \in \{1,2,4\} \\ 4 & : i \in \{3,5,7\} \\ 6 & : i \in \{8,6,9\}. \end{cases}$$

We find for the two-dimensional terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} = \{i,j\}$, $i, j \in \{1, 2, \ldots, 9\}$, that $f_{\{i,j\}} \equiv 0$ for $\boldsymbol{u} \notin U^*$ and

$$f_{\{i,j\}}(x_{\{i,j\}}) = \frac{\prod_{j \in \{2,4,6\}} c_j}{c_{\mathrm{o}(i)} c_{\mathrm{o}(j)}} \left(B_{\mathrm{o}(i)}(x_i) - c_{\mathrm{o}(i)}\right) \left(B_{\mathrm{o}(j)}(x_j) - c_{\mathrm{o}(j)}\right) \tag{19}$$

for $\boldsymbol{u} \in U^*$. Finally, we get for the three-dimensional terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} = \{i,j,\ell\}$, $i, j, \ell \in \{1, 2, \ldots, 9\}$, that $f_{\boldsymbol{u}} \equiv 0$ again for $\boldsymbol{u} \notin U^*$ and

$$f_{\{i,j,\ell\}}(x_{\{i,j,\ell\}}) = \left(B_{\mathrm{o}(i)}(x_i) - c_{\mathrm{o}(i)}\right) \left(B_{\mathrm{o}(j)}(x_j) - c_{\mathrm{o}(j)}\right) \left(B_{\mathrm{o}(\ell)}(x_\ell) - c_{\mathrm{o}(\ell)}\right) \tag{20}$$

for $\boldsymbol{u} \in U^*$. However, in the approximation scenario we only have the data about the function f and not f itself. The norm and the exact Fourier coefficients are only used to compute the error $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))$ and compare the order of the sensitivity indices.

4.1 LSQR

In this section, we discuss the addition of an ℓ_2 regularization the problem (14) and the incorporation of a-priori information about the decay of the Fourier coefficients. We characterize the smoothness of a function by the decay of the Fourier coefficients, i.e., a function $f \in \mathrm{H}^\omega(\mathbb{T}^d)$ has the smoothness defined by the weight $\omega \colon \mathbb{Z}^d \to [1, \infty)$. This directly implies for the decay $|\mathrm{c}_{\boldsymbol{k}}(f)| \in \mathrm{o}(\omega^{-1}(\boldsymbol{k}))$ since $\omega^2(\boldsymbol{k}) \, |\mathrm{c}_{\boldsymbol{k}}(f)|^2 \to 0$ for $\|\boldsymbol{k}\|_\infty \to \infty$ by definition. In order to incorporate this information we add a term to (14) and obtain the modified problem

$$\min_{\hat{\boldsymbol{f}}} \frac{1}{2} \left\| \boldsymbol{y} - \boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U)) \hat{\boldsymbol{f}} \right\|_2^2 + \lambda \left\| \hat{\boldsymbol{f}} \right\|_{\hat{\boldsymbol{W}}}^2 \tag{21}$$

with regularization parameter $\lambda > 0$, weighted norm $\|\hat{\boldsymbol{f}}\|_{\hat{\boldsymbol{W}}}^2 = \hat{\boldsymbol{f}}^{\mathsf{H}} \hat{\boldsymbol{W}} \hat{\boldsymbol{f}}$ and $\hat{\boldsymbol{W}} = \operatorname{diag}(\omega(\boldsymbol{k}))_{\boldsymbol{k} \in \mathcal{I}}$ a diagonal matrix. This can be the identity if no smoothness information is known and $\omega(\boldsymbol{k}) \equiv 1$, i.e., all frequencies are penalized equally. Problem (21) always has a unique solution in our setting which we aim to find by applying the iterative LSQR algorithm, see [33].

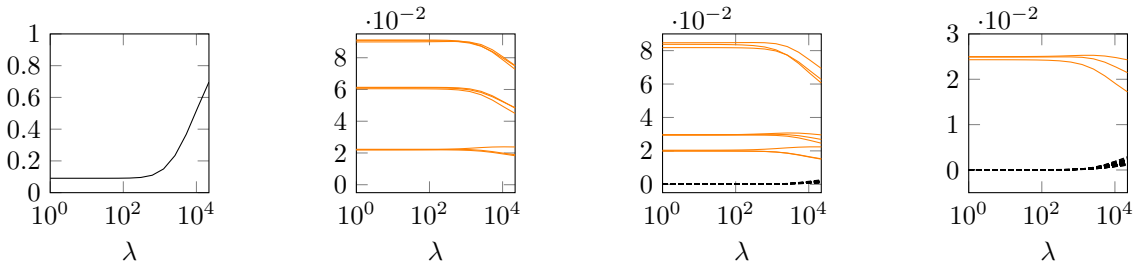

(a) $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U^*))$ (b) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 1$ (c) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 2$ (d) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 3$

Figure 3: LSQR with small bandwidths given in (23) on exact data with $s = 0$ (orange: active ANOVA terms $U^\star$, dashed: inactive ANOVA terms in the test function).

We apply the method with our newly obtained minimization problem (21) to the test function (17) which we sample at $\mathcal{X} \subseteq \mathbb{T}^d$, $|\mathcal{X}| = 10\,000$, uniform i.i.d. random nodes and use the Sobolev weights

$$\omega_s(\boldsymbol{k}) = \prod_{j=1}^{9} (1 + |k_j|)^s \tag{22}$$

for $s \geq 0$ in $\hat{\boldsymbol{W}}$. In the following numerical experiments we use different parameters s and $N_{\boldsymbol{u}}$ as well as $\lambda \in [\mathrm{e}^0, \mathrm{e}^{10}]$. The results are visualized with the L_2-error (16) for the approximation $S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U^*)} f$ and the global sensitivity indices $\varrho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$, cf. (3), for each $|\boldsymbol{u}| = 1, 2, 3$ separately. Note that the errors and global sensitivity indices have been averaged over 100 random draws of the node set $\mathcal{X}$.

(i) In the first setting we choose $\omega(\boldsymbol{k}) = 1$ for all $\boldsymbol{k} \in \mathcal{I}$, which penalizes all frequencies equally. The regularization term in (21) then becomes a non-weighted ℓ_2 regularization without smoothness information. To compensate for this lack of information, we choose small bandwidths

$$N_{\boldsymbol{u}} = \begin{cases} 26 & \text{for all } \boldsymbol{u} \text{ with } |\boldsymbol{u}| = 1 \\ 6 & \text{for all } \boldsymbol{u} \text{ with } |\boldsymbol{u}| = 2 \\ 4 & \text{for all } \boldsymbol{u} \text{ with } |\boldsymbol{u}| = 3, \end{cases} \tag{23}$$

which add up to 3 394 frequencies such that we are in an overdetermined setting. The results of solving the minimization is depicted in Figure 3. We are able to clearly detect the active set U^* and distinguish the global sensitivity indices of the contained terms from the rest. In (b) we observe that the global sensitivity indices $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ are larger than zero for all nine one-dimensional ANOVA terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} \in U^*$, $|\boldsymbol{u}| = 1$, occuring in the test function. We see that the sensitivity indices are on three different levels. They relate to the smoothness if we consider the analytical term (18), i.e., depending on which of the three splines is involved, the corresponding ANOVA term is more smooth if the spline is of a higher order. The nine two-dimensional terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} \in U^*$, $|\boldsymbol{u}| = 2$, are clearly separated from the inactive terms in $\mathcal{P}(\mathcal{D}) \setminus U^*$, see (c). The sensitivity indices here are also grouped around three levels since in (19) there is always a product of two different splines yielding $\binom{3}{2} = 3$ types of smoothness for those terms. The plot (d) shows that we can also distinguish the three active terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} \in U^*$, $|\boldsymbol{u}| = 3$, which are of the same smoothness since they are always a product of the three involved splines, see (20). The minimal relative L_2-error is $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U^*)) \approx 9.1 \cdot 10^{-2}$ when using the active set of terms U^* after sensitivity analysis. This experiment provides a comparison for the addition of noise and smoothness information.

(ii) In the next experiment we consider the effect of adding 10% Gaussian noise which results in Figure 4. We observe that the minimal L_2-error $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U^*)) \approx 0.165$ increases in comparison to the previous experiment as expected. We get a typical behaviour of under- and overfitting in the L_2-error which can be recognized in the global sensitivity indices as well: For small λ the ANOVA terms $f_{\boldsymbol{u}}$ not occurring in the test function are not zero in our approximation since it fits the noise. For large λ on the other hand we also penalize ANOVA terms $f_{\boldsymbol{u}}$ which actually occur in the test function as can be seen by the orange lines dropping. However, it is evident that we are able to achieve the goal of detecting the active ANOVA terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} \in U^*$, since their sensitivity indices are well-separated from the sensitivity indices of the inactive terms.

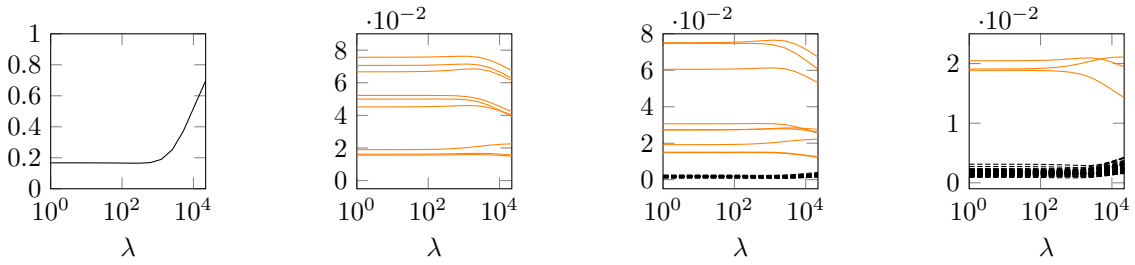

(a) $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U^*))$ (b) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 1$ (c) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 2$ (d) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 3$

Figure 4: LSQR with small bandwidths given in (23) on noisy data with $s = 0$ (orange: active ANOVA terms $U^\star$, dashed: inactive ANOVA terms in the test function).

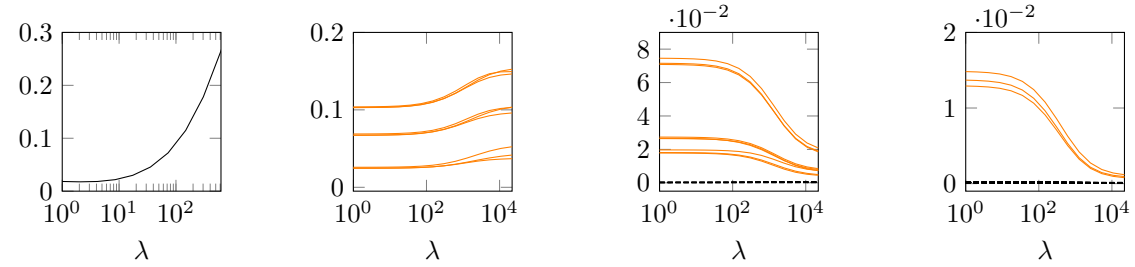

(a) $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U^*))$ (b) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 1$ (c) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 2$ (d) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 3$

Figure 5: LSQR with large bandwidths given in (24) on exact data with $s = 1.5$ (orange: active ANOVA terms $U^\star$, dashed: inactive ANOVA terms in the test function).

(iii) In this experiment, we are interested in the addition of smoothness information and how it influences the minimization. We choose large bandwidths

$$N_{\boldsymbol{u}} = \begin{cases} 352 & \text{for } |\boldsymbol{u}| = 1 \\ 20 & \text{for } |\boldsymbol{u}| = 2 \\ 8 & \text{for } |\boldsymbol{u}| = 3, \end{cases} \tag{24}$$

which results in $44\,968$ frequencies and an undetermined optimization problem (21) with $|\mathcal{I}_{\boldsymbol{N}}(U_3)| > |\mathcal{X}|$. The addition of the weighted ℓ_2 regularization term in (21) leads to a reduction of the *search space* to the functions in the Sobolev type space $\mathrm{H}^{\omega_s}(\mathbb{T}^d)$ by incorporating the weights ω_s from (22). We choose $s = 1.5$ in $\hat{\boldsymbol{W}}(\boldsymbol{u})$ which propagate functions of this dominating mixed smoothness since for the test function we have $f \in \mathrm{H}^{\omega_{1.5-\varepsilon}}(\mathbb{T}^d)$ for every $\varepsilon > 0$.

The results for the experiment without noise are shown in Figure 5. We see that in this undetermined setting, the incorporation of smoothness information allowed us to efficiently solve the problem and with $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U^*)) \approx 1.8 \cdot 10^{-2}$ even achieve a better L_2-error than in the previous experiments. The active ANOVA terms in U^* are also clearly separable from the inactive ANOVA terms. We also recognize the different smoothness levels of the ANOVA terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} \in U^*$, through the sensitivity indices as described in (i).

(iv) The same experiment with the addition of 10% Gaussian noise can be seen in Figure 6. In comparison to the noisy experiment with small bandwidths in (ii) we achieve a smaller minimal L_2-error of $\varepsilon(\mathcal{X}, \mathcal{I}(U^*)) \approx 0.189$. In addition, the different smoothness levels of the ANOVA terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} \in U^*$, are clearly better distinguishable than in experiment (ii).

Remark 4.1. *The optimal choice of λ in this paper is determined via the minimum L_2-error $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))$ from* (16)*. As mentioned before, this error can only be explicitly calculated if we generate the data as evaluations from a known synthetic test function. In the general case, one could use cross-validation techniques, e.g., with the methods in [2] which allow for fast computation of the cross-validation score. This score can then be used to choose a λ which will be close to the $\mathrm{L}_2(\mathbb{T}^d)$-optimal one, i.e., avoiding over- or underfitting.*

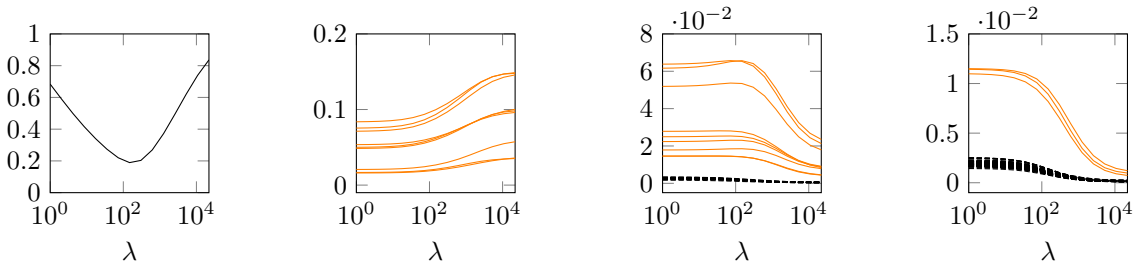

(a) $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U^*))$ (b) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 1$ (c) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 2$ (d) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 3$

Figure 6: LSQR with large bandwidths given in (24) on noisy data with $s = 1.5$ (orange: active ANOVA terms $U^\star$, dashed: inactive ANOVA terms in the test function).

4.2 Group Lasso

As it is promoted in Section 2, after choosing a superposition threshold d_s and the set of ANOVA terms U_{d_s}, we often have only a few of those ANOVA terms active. In other words, there is sparsity in the groups $\boldsymbol{u} \in U_{d_s}$. The ANOVA terms $f_{\boldsymbol{u}}$ corresponding to these active groups, however, do not inherit this sparse structure, i.e., there is no sparsity in their Fourier coefficients. Therefore, a method is needed which promotes sparsity on the scale of the groups $\boldsymbol{u} \in U_{d_s}$, but not within the coefficients of this group $\hat{\boldsymbol{f}}(\boldsymbol{u})$, cf. (13). Incorporating this type of sparsity in the regularization combines the previously required sensitivity analysis into an automated process during the minimization.

A method which accomplishes this mixture of sparsity and non-sparsity constraints is called group lasso and was introduced in [46]. With group lasso, one seeks for the solution $\hat{\boldsymbol{f}}^\star$ of

$$\min_{\hat{\boldsymbol{f}}} \frac{1}{2} \|\boldsymbol{y} - \boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))\hat{\boldsymbol{f}}\|_2^2 + \lambda \sum_{\boldsymbol{u} \in U} \|\hat{\boldsymbol{f}}(\boldsymbol{u})\|_{\hat{\boldsymbol{W}}(\boldsymbol{u})} \tag{25}$$

with grouped index set $\mathcal{I}_{\boldsymbol{N}}(U)$, $\boldsymbol{N} = (N_{\boldsymbol{u}})_{\boldsymbol{u} \in U}$, see (8), where we use the weighted norm $\|\hat{\boldsymbol{f}}(\boldsymbol{u})\|^2_{\hat{\boldsymbol{W}}(\boldsymbol{u})} = \hat{\boldsymbol{f}}(\boldsymbol{u})^{\mathsf{H}} \hat{\boldsymbol{W}}(\boldsymbol{u}) \hat{\boldsymbol{f}}(\boldsymbol{u})$ and $\hat{\boldsymbol{W}}(\boldsymbol{u}) = \operatorname{diag}(\omega(\boldsymbol{k}))_{\boldsymbol{k} \in \mathcal{I}^{\boldsymbol{u},d}_{N_{\boldsymbol{u}}}}$ diagonal matrices with weight functions $\omega : \mathbb{Z}^d \to [1, \infty)$ for $\boldsymbol{u} \in U$. Notice the lack of the square in the regularization term. By incorporating squares in (25) we would attain the same functional as in Section 4.1. However, omitting them allows for an ℓ_1-norm structure around the groups for a sparsifying effect similar to the basic lasso approach itself. In this group case, it pushes towards a small number of active ANOVA terms.

Remark 4.2. *With differently sized groups, i.e., a different amount of Fourier coefficients in each group $\mathcal{I}^{\boldsymbol{u},d}_{N_{\boldsymbol{u}}}$ of the index set, see* (7)*, lasso would set smaller groups with fewer frequencies to zero for sparsity as data fitting is easier with groups of more degrees of freedom. This could be tackled by introducing an additional weight for $\boldsymbol{u} \in U_{d_s}$ counterbalancing the group size. We work with groups of the same amount of frequencies throughout this section making this additional weight obsolete.*

The missing squares bring another difficulty with them, namely, the lack of differentiability in the regularization term. Thus (25) is a non-smooth convex minimization problem, which we tackle by the use of a proximal gradient method. Entranced by the computational simplicity and quadratic rate of convergence, we use the fast iterative shrinking-thresholding algorithm (FISTA), cf. [3], as our algorithm of choice.

In order to formulate the algorithm, we have to introduce the projection p_L. For a vector $\hat{\boldsymbol{h}}$ of Fourier coefficients, $p_L(\hat{\boldsymbol{h}})$ is given by the minimizer of

$$\frac{1}{2} \left\| \hat{\boldsymbol{f}} - \left(\hat{\boldsymbol{h}} - \frac{1}{L} \boldsymbol{F}^{\mathsf{H}}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U)) \left(\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U)) \hat{\boldsymbol{h}} - \boldsymbol{y} \right) \right) \right\|_2^2 + \frac{\lambda}{L} \sum_{\boldsymbol{u} \in U} \|\hat{\boldsymbol{f}}_{\boldsymbol{u}}\|_{\hat{\boldsymbol{W}}(\boldsymbol{u})}. \tag{26}$$

With that, we are able to write the FISTA algorithm adopted to our setting:

Remark 4.3. *The computational most expensive part is the projection p_L which we will reduce to one multiplication with $\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))$ and one with its adjoint which we realize by the grouped transformation from Section 3. Note that there is also a variant of the FISTA algorithm which uses constant step size L_k which makes the inner loop obsolete. For that, the complexity of the FISTA step boils down to the complexity of computing p_L. The quadratic convergence rate assures us, that a small number of iterations is sufficient.*

Algorithm 1 FISTA for group lasso

Input: initial guess for the Fourier coefficents $\hat{\boldsymbol{f}}_0$,
stepsize parameters $L_0 > 0$, $\eta > 1$ and
maximal iteration count $K \in \mathbb{N}$

Output: $\hat{\boldsymbol{f}}^\star = \hat{\boldsymbol{f}}^{(K)}$ minimizer of (25)

$\hat{\boldsymbol{h}}^{(1)} \leftarrow \hat{\boldsymbol{f}}^{(0)}$, $t_1 \leftarrow 1$
for $k = 1, \dots, K$ **do**
  $L_k \leftarrow L_{k-1}$
  $\hat{\boldsymbol{f}}^{(k)} \leftarrow p_{L_k}(\hat{\boldsymbol{h}}^{(k)})$
  while $\|\boldsymbol{y} - \boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))\hat{\boldsymbol{f}}^{(k)}\|_2^2 - \|\boldsymbol{y} - \boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))\hat{\boldsymbol{h}}^{(k)}\|_2^2$
  $> 2\langle \hat{\boldsymbol{f}}^{(k)} - \hat{\boldsymbol{h}}^{(k)}, \boldsymbol{F}^{\mathsf{H}}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))(\boldsymbol{F}(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U))\hat{\boldsymbol{h}}^{(k)} - \boldsymbol{y})\rangle + L_k \|\hat{\boldsymbol{f}}^{(k)} - \hat{\boldsymbol{h}}^{(k)}\|_2^2$ **do**
    $L_k \leftarrow \eta L_k$
    $\hat{\boldsymbol{f}}^{(k)} \leftarrow p_{L_k}(\hat{\boldsymbol{h}}^{(k)})$
  end while
  $t_{k+1} \leftarrow (1 + \sqrt{1 + 4t_k^2})/2$
  $\hat{\boldsymbol{h}}^{(k+1)} \leftarrow \hat{\boldsymbol{f}}^{(k)} + (t_k - 1)/(t_{k+1})(\hat{\boldsymbol{f}}^{(k)} - \hat{\boldsymbol{f}}^{(k-1)})$
end for

To compute the projection p_L we develop an explicit formula which will be done in the next theorem with the use of the following lemma.

Lemma 4.4. *Let $\mathcal{I} \subseteq \mathbb{Z}^d$ be a finite frequency index set, $\omega : \mathbb{Z}^d \to [0, \infty)$ a weight function, and $\hat{\boldsymbol{W}} = \operatorname{diag}(\omega(\boldsymbol{k}))_{\boldsymbol{k} \in \mathcal{I}}$. Further, let $\boldsymbol{x}_1^\star(\lambda)$ and $\boldsymbol{x}_2^\star(\xi)$ be the minimizers of*

$$\frac{1}{2}\|\boldsymbol{x} - \boldsymbol{y}\|_2^2 + \lambda \|\boldsymbol{x}\|_{\hat{\boldsymbol{W}}} \quad \textit{and} \quad \frac{1}{2}\|\boldsymbol{x} - \boldsymbol{y}\|_2^2 + \xi \|\boldsymbol{x}\|_{\hat{\boldsymbol{W}}}^2,$$

respectively. Then for $\xi > 0$ we have

(i)

$$\boldsymbol{x}_2^\star(\xi) = \operatorname{diag}\left(\frac{1}{1 + 2\xi\omega(\boldsymbol{k})}\right)_{\boldsymbol{k} \in \mathcal{I}} \boldsymbol{y},$$

(ii)

$$\boldsymbol{x}_1^\star(\lambda) = \boldsymbol{x}_2^\star(\xi) \quad \textit{for} \quad \lambda = \left\| \left(\frac{y_{\boldsymbol{k}}}{1/(2\xi) + \omega(\boldsymbol{k})} \right)_{\boldsymbol{k} \in \mathcal{I}} \right\|_{\hat{\boldsymbol{W}}},$$

(iii) and

$$\boldsymbol{x}_1^\star(\lambda) = \boldsymbol{0} \quad \textit{for} \quad \lambda \geq \left\| \left(\frac{y_{\boldsymbol{k}}}{\omega(\boldsymbol{k})} \right)_{\boldsymbol{k} \in \mathcal{I}} \right\|_{\hat{\boldsymbol{W}}}.$$

Proof. (i) We simply calculate the roots of the gradient

$$\boldsymbol{x} - \boldsymbol{y} + 2\xi \hat{\boldsymbol{W}} \boldsymbol{x} \stackrel{!}{=} \boldsymbol{0} \quad \Leftrightarrow \quad \boldsymbol{x}_2^\star(\xi) = \operatorname{diag}\left(\frac{1}{1 + 2\xi\omega(\boldsymbol{k})}\right)_{\boldsymbol{k} \in \mathcal{I}} \boldsymbol{y}.$$

The convexity of the function assures the minimizing property.

(ii) Similarly to (i), computing $\boldsymbol{x}_1^\star(\lambda)$ can be done by root-finding of

$$\nabla_{\boldsymbol{x}} \left(\frac{1}{2}\|\boldsymbol{x} - \boldsymbol{y}\|_2^2 + \lambda \|\boldsymbol{x}\|_{\hat{\boldsymbol{W}}} \right) = \boldsymbol{x} - \boldsymbol{y} + \frac{\lambda}{\|\boldsymbol{x}\|_{\hat{\boldsymbol{W}}}} \operatorname{diag}(\omega(\boldsymbol{k}))_{\boldsymbol{k} \in \mathcal{I}} \boldsymbol{x}$$

since the function in question is strictly convex. Using the ansatz $\boldsymbol{x} = \boldsymbol{x}_2^\star(\xi)$ and (i) leads to

$$\left(\operatorname{diag}\left(\frac{1}{1 + 2\xi\omega(\boldsymbol{k})}\right)_{\boldsymbol{k} \in \mathcal{I}} - \boldsymbol{I} + \frac{\lambda \operatorname{diag}\left(\frac{\omega(\boldsymbol{k})}{1 + 2\xi\omega(\boldsymbol{k})}\right)_{\boldsymbol{k} \in \mathcal{I}}}{\left\| \operatorname{diag}\left(\frac{1}{1 + 2\xi\omega(\boldsymbol{k})}\right)_{\boldsymbol{k} \in \mathcal{I}} \boldsymbol{y} \right\|_{\hat{\boldsymbol{W}}}} \right) \boldsymbol{y} \stackrel{!}{=} \boldsymbol{0}.$$

This is certainly fulfilled if the diagonal matrix is zero in each component, i.e.,

$$\frac{1}{1+2\xi\omega(\boldsymbol{k})} - 1 + \frac{\lambda}{\left\|\operatorname{diag}\left(\frac{1}{1+2\xi\omega(\boldsymbol{k})}\right)_{\boldsymbol{k}\in\mathcal{I}}\boldsymbol{y}\right\|_{\hat{\boldsymbol{W}}}} \frac{\omega(\boldsymbol{k})}{1+2\xi\omega(\boldsymbol{k})} \stackrel{!}{=} 0,$$

or equivalently,

$$\lambda \stackrel{!}{=} \left\|\operatorname{diag}\left(\frac{1}{1/(2\xi)+\omega(\boldsymbol{k})}\right)_{\boldsymbol{k}\in\mathcal{I}}\boldsymbol{y}\right\|_{\hat{\boldsymbol{W}}}.$$

(iii) We want to minimize

$$\frac{1}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_2^2 + \lambda\|\boldsymbol{x}\|_{\hat{\boldsymbol{W}}}.$$

Using polar coordinates we see that varying the arguments in $\boldsymbol{x}$ does not change the second summand. Hence, $\boldsymbol{x}$ and $\boldsymbol{y}$ have to have the same arguments in each component. Without loss of generality, we restrict to the moduli of these numbers, i.e., positive numbers.

Then

$$\begin{aligned}\frac{1}{2}\|\mathbf{0}-\boldsymbol{y}\|_2^2 + \lambda\|\mathbf{0}\|_{\hat{\boldsymbol{W}}} &= \frac{1}{2}\|-\boldsymbol{x}+\boldsymbol{x}-\boldsymbol{y}\|_2^2 \\ &\leq \frac{1}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_2^2 + \left\langle \operatorname{diag}\left(\sqrt{\omega(\boldsymbol{k})}\right)_{\boldsymbol{k}\in\mathcal{I}}\boldsymbol{x},\ \operatorname{diag}\left(\frac{1}{\sqrt{\omega(\boldsymbol{k})}}\right)_{\boldsymbol{k}\in\mathcal{I}}\boldsymbol{y}\right\rangle.\end{aligned}$$

By the Cauchy-Schwarz inequality, we obtain

$$\frac{1}{2}\|\mathbf{0}-\boldsymbol{y}\|_2^2 + \lambda\|\mathbf{0}\|_2 \leq \frac{1}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_2^2 + \|\boldsymbol{x}\|_{\hat{\boldsymbol{W}}}\left\|\operatorname{diag}\left(1/\sqrt{\omega(\boldsymbol{k})}\right)_{\boldsymbol{k}\in\mathcal{I}}\boldsymbol{y}\right\|_2.$$

Now making use of $\|\operatorname{diag}(1/\sqrt{\omega(\boldsymbol{k})})_{\boldsymbol{k}\in\mathcal{I}}\boldsymbol{y}\|_2 < \lambda$ we end up with

$$\frac{1}{2}\|\mathbf{0}-\boldsymbol{y}\|_2^2 + \lambda\|\mathbf{0}\|_2 \leq \frac{1}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_2^2 + \lambda\|\boldsymbol{x}\|_{\hat{\boldsymbol{W}}}$$

which holds for all $\boldsymbol{x}$ and, hence, $\mathbf{0}$ is the minimizer for this case.

□

Now, using Lemma 4.4, we prove an explicit formula for the projection p_L defined via the minimizer of (26).

Theorem 4.5. *Let U be a subset of $\mathcal{P}(\{1,\dots,d\})$ and diagonal matrices $\hat{\boldsymbol{W}}(\boldsymbol{u}) = \operatorname{diag}(\omega(\boldsymbol{k}))_{\boldsymbol{k}\in\mathcal{I}^{\boldsymbol{u},d}_{N_{\boldsymbol{u}}}}$ with $I^{\boldsymbol{u},d}_{N_{\boldsymbol{u}}}$ from (7) and strictly positive weights for $\boldsymbol{u}\in U$. For any $\boldsymbol{y}\in\mathbb{C}^N$, $N=\sum_{\boldsymbol{u}\in U}(1-N_{\boldsymbol{u}})^{|\boldsymbol{u}|}$, and $\lambda>0$, the minimizer of*

$$\frac{1}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_2^2 + \lambda\sum_{\boldsymbol{u}\in U}\|\boldsymbol{x}_{\boldsymbol{u}}\|_{\hat{\boldsymbol{W}}(\boldsymbol{u})}$$

is given by

$$\boldsymbol{x}^\star_{\boldsymbol{u}} = \begin{cases}\operatorname{diag}\left(\frac{1}{1+\xi(\boldsymbol{u})\omega(\boldsymbol{k})}\right)_{\boldsymbol{k}\in\mathcal{I}^{\boldsymbol{u},d}_{N_{\boldsymbol{u}}}}\boldsymbol{y} & \text{for } \lambda \leq \left\|\left(\frac{y_{\boldsymbol{k}}}{\omega(\boldsymbol{k})}\right)_{\boldsymbol{k}\in\mathcal{I}^{\boldsymbol{u},d}_{N_{\boldsymbol{u}}}}\right\|_{\hat{\boldsymbol{W}}(\boldsymbol{u})} \\ \mathbf{0} & \text{otherwise}\end{cases}$$

for $\boldsymbol{u}\in U$ with $\xi(\boldsymbol{u})$ fulfilling

$$\left\|\left(\frac{y_{\boldsymbol{k}}}{1/\xi(\boldsymbol{u})+\omega(\boldsymbol{k})}\right)_{\boldsymbol{k}\in\mathcal{I}^{\boldsymbol{u},d}_{N_{\boldsymbol{u}}}}\right\|_{\hat{\boldsymbol{W}}(\boldsymbol{u})} = \lambda.$$

Proof. Because of the structure of the objective function, we can minimize for every $\boldsymbol{u}\in U$ separately. Lemma 4.4 (i)-(iii) then results the assertion. □

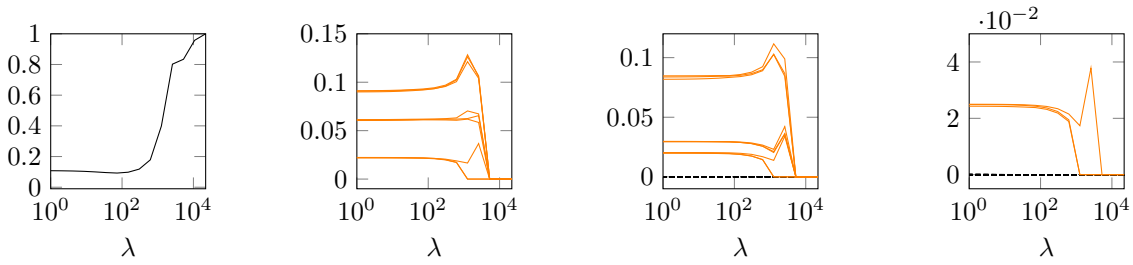

(a) $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U_3))$ (b) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 1$ (c) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 2$ (d) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 3$

Figure 7: FISTA with small bandwidths given in (23) on exact data with $s = 0$ (orange: active ANOVA terms $U^\star$, dashed: inactive ANOVA terms in the test function).

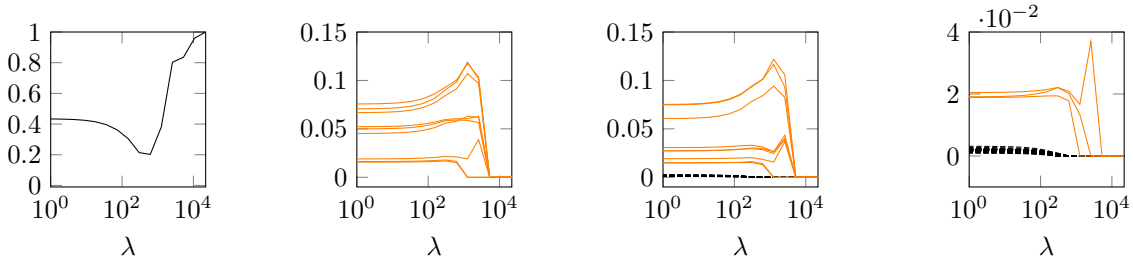

(a) $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U_3))$ (b) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 1$ (c) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 2$ (d) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 3$

Figure 8: FISTA with small bandwidths given in (23) on noisy data with $s = 0$ (orange: active ANOVA terms $U^\star$, dashed: inactive ANOVA terms in the test function).

Remark 4.6. *In Theorem 4.5 we still have to find $\xi(\boldsymbol{u})$ such that*

$$t(\xi) := \left\| \left(\frac{y_{\boldsymbol{k}}}{1/\xi(\boldsymbol{u}) + \omega(\boldsymbol{k})} \right)_{\boldsymbol{k} \in \mathcal{I}^{\boldsymbol{u},d}_{N_{\boldsymbol{u}}}} \right\|^2_{\hat{\boldsymbol{W}}(\boldsymbol{u})} = \lambda^2.$$

We do that in two steps:

(i) First, we determine whether $t(1) \geq \lambda^2$ or not. If so, ξ has to be smaller than one because of the monotonicity of the objective function. Then we use the bisection method on $t(\cdot)$ with the initial interval $[0, 1]$. Otherwise we do the same for the function $t(1/\cdot)$.

(ii) To obtain a more precise result we use some iterations with Newton's method afterwards.

Now for the numerical experiments. As in Section 4.1, we use the nine-dimensional function (17) which we sample at 10 000 uniform i.i.d. random nodes $\mathcal{X} \subseteq \mathbb{T}^9$. For four different settings we launched Algorithm 1 with $\lambda \in [\mathrm{e}^0, \mathrm{e}^{10}]$. We started with the initial guess $\hat{\boldsymbol{f}}_0 = \boldsymbol{0}$ and the biggest $\lambda = \mathrm{e}^{10}$ with the minimizer $\hat{\boldsymbol{f}}^\star(\lambda)$ which is presumably close to $\boldsymbol{0}$ for this case. For the remaining values of λ we used the last solution as new initial guess since $\hat{\boldsymbol{f}}^\star(\lambda)$ depends continuously on λ. For evaluation, we plotted the L_2-error and the global sensitivity indices (3) for each $|\boldsymbol{u}| = 1, 2, 3$ separately. Note that the errors and global sensitivity indices have been averaged over 100 random draws of the node set $\mathcal{X}$.

(i) In the first setting, we choose $\omega(\boldsymbol{k}) = 1$ for all $\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}(U)$, which penalizes all frequencies equally, cf. Section 4.1. To compensate for the lack of information we use the small bandwidths (23) such that we are in an overdetermined setting.

The results of this first experiment without noise can be seen in Figure 7. The minimal relative L_2-error is $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U_3)) \approx 9.2 \cdot 10^{-2}$ which is similar to the observations in Section 4.1. We are able to distinguish the global sensitivity indices of the active ANOVA terms very clearly from the ones not occurring in the test function. In (b) we observe that the global sensitivity indices $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ are larger than zero for all one-dimensional

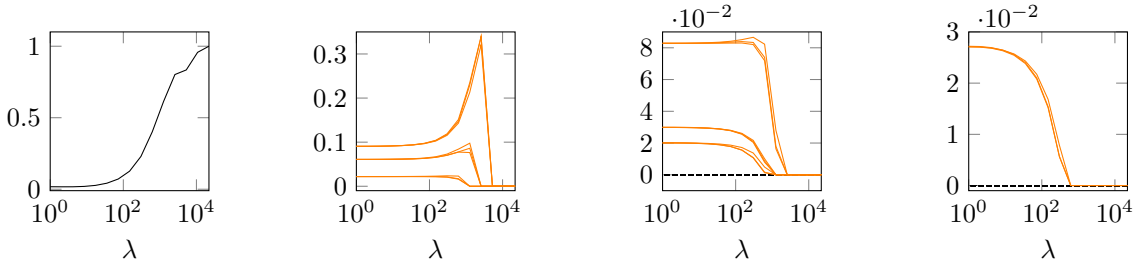

(a) $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U_3))$ (b) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 1$ (c) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 2$ (d) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 3$

Figure 9: FISTA with large bandwidths given in (24) on exact data with $s = 1.5$ (orange: active ANOVA terms $U^\star$, dashed: inactive ANOVA terms in the test function).

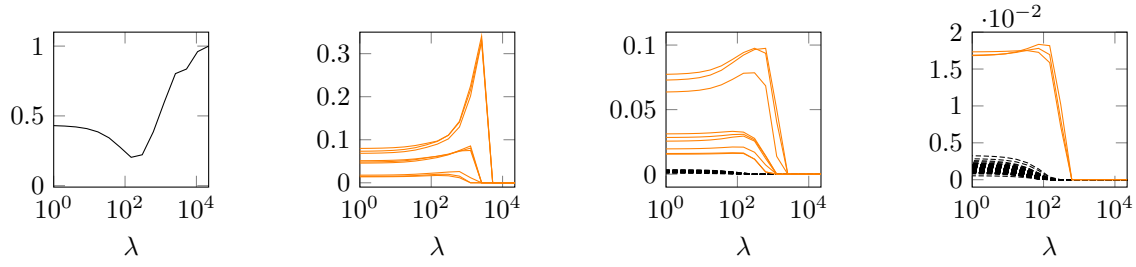

(a) $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U_3))$ (b) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 1$ (c) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 2$ (d) $\rho(\boldsymbol{u}, S^{\mathcal{X}}_{\mathcal{I}_{\boldsymbol{N}}(U_3)} f)$ with $|\boldsymbol{u}| = 3$

Figure 10: FISTA with large bandwidths given in (24) on noisy data with $s = 1.5$ (orange: active ANOVA terms $U^\star$, dashed: inactive ANOVA terms in the test function).

ANOVA terms $f_{\boldsymbol{u}}$ occuring in the test function. In (c) and (d) we have the expected sparsifying behaviour from the ℓ_1-regularization, as many groups are penalized to be exactly zero, for larger λ even the ones which occur in the test function. In summary, the group lasso regularization is able to detect the correct active terms. We also recognize the different smoothness of the ANOVA terms as mentioned in experiment (i) from Section 4.1.

(ii) In the next experiment we added 10% Gaussian noise which results in Figure 8. The minimal L_2-error has increased to $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U_3)) \approx 0.202$. We observe the typical behaviour of under- and overfitting in the L_2-error, see (a). This can also be recognized in the global sensitivity indices: For small λ the ANOVA terms $f_{\boldsymbol{u}}$ not occurring in the test function are not zero in our approximation since it fits the noise. For large λ on the other hand we also penalize active ANOVA terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} \in U^*$ which can be seen by the orange lines dropping to zero. Overall, the active ANOVA terms are distinguishable from the inactive terms and the group lasso is able to correctly detect the active set.

(iii) Now, we are interested in incorporating smoothness information about the function f: We choose the large bandwidths (24) which results in $53\,806$ frequencies such that we are underdetermined. With that many degrees of freedom we have to reduce our search space by using the Sobolev weights ω_s from (22) for $s = 1.5$ in $\hat{\boldsymbol{W}}(\boldsymbol{u})$. This corresponds to the smoothness of the function since $f \in \mathrm{H}^{\omega_{1.5-\varepsilon}}(\mathbb{T}^d)$.

The results without noise are depicted in Figure 9. The behaviour looks very similar to experiment (i) with smaller bandwidths, but the minimal L_2-error is smaller: $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U_3)) \approx 1.7 \cdot 10^{-2}$. Furthermore, under- and overfitting does not occur as the L_2-error curve is strictly increasing. Moreover, the active set of terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} \in U^*$, is correctly recognized by group lasso.

(iv) The same experiment with additional 10% Gaussian noise can be seen in Figure 10. In comparison to experiment (ii) with noise and small bandwidths we achieve a smaller minimal L_2-error $\varepsilon(\mathcal{X}, \mathcal{I}_{\boldsymbol{N}}(U_3)) \approx 0.203$. In addition, the ANOVA terms $f_{\boldsymbol{u}}$, $\boldsymbol{u} \in U^*$ are now better distinguishable as the smoothness weights $\omega_{1.5}$ filter out the non-smooth noise.

5 Application Data Example

In this section, we aim to apply our method to the well-known *adult census data set* from the UCI repository [7]. The data was extracted from the 1994 US census database. The goal is to use 14 attributes about a person or group of persons to predict whether they have an income of more than $50\,000$ dollars a year. This represents one of the most popular data sets from the UCI repository.

Since this data is not periodic, we use the generalization hinted at in Remark 3.3. This requires a complete orthonormal system of product structure in a space of non-periodic L_2 functions. We consider the space of square-integrable $\mathrm{L}_2([0,1]^d)$ over the cube $[0,1]^d$ with the half-period cosine basis, i.e., we have functions

$$\phi_{\boldsymbol{k}}(\boldsymbol{x}) = \sqrt{2}^{|\operatorname{supp} \boldsymbol{k}|} \prod_{j=1}^{d} \cos(\pi k_j x_j)$$

for frequencies $\boldsymbol{k} \in \mathbb{N}_0^d$ where $\mathbb{N}_0 = \{0, 1, \dots\}$. It is a well-known fact that $(\phi_{\boldsymbol{k}})_{\boldsymbol{k} \in \mathbb{N}_0^d}$ is a complete orthonormal system in $\mathrm{L}_2([0,1]^d)$, see e.g. [24]. Without going into explicit detail, it is evident that the considerations from Section 2 and Section 3 generalize to this system. However, the low dimensional transforms (13) will be realized trough a non-equispaced fast cosine transform (NFCT), see [34, Section 7.4], instead of the NFFT.

The data set requires a few simple pre-processing steps in order to be used with our approach. First of all, we remove the attribute *education* since this is already encoded in *education.num* and the attribute *fnlwgt* since it gives a weight to entries which is not normalized across federal states. We also transform categorical attributes to numerical attributes and remove data points with missing values. Moreover, we perform a min-max-normalization such that we have data points in the cube $[0,1]^{12}$. After that pre-processing we have a $45\,199$ data points which we split into a training node set $\mathcal{X}_{\mathrm{train}} = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \dots, \boldsymbol{x}_{M_{\mathrm{train}}}\} \subseteq [0,1]^{12}$ with $M_{\mathrm{train}} = 36\,160$ nodes and a test node set $\mathcal{X}_{\mathrm{test}} = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \dots, \boldsymbol{x}_{M_{\mathrm{test}}}\} \subseteq [0,1]^{12}$ with $M_{\mathrm{test}} = 9\,039$ nodes. Therefore, we have 80% of the overall data to obtain the approximation and 20% to verify the performance. Whether a person earns more than $50\,000$ dollars or not is a binary question, i.e., we have a vector $\boldsymbol{y}_{\mathrm{train}} \in \{0,1\}^{M_{\mathrm{train}}}$ such that for a person $\boldsymbol{x}_i \in \mathcal{X}_{\mathrm{train}}$ in the training set we know that they earn $50\,000$ dollar or more if $y_i = 1$ and they do not earn as much if $y_i = 0$.

It is our goal to answer the same question for every person from the test set $\mathcal{X}_{\mathrm{test}}$. The general idea is to use an approximate partial sum

$$S^{\mathcal{X}_{\mathrm{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U)} f(\boldsymbol{x}) \coloneqq \sum_{\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}(U)} \hat{f}_{\boldsymbol{k}}\, \phi_{\boldsymbol{k}}(\boldsymbol{x})$$

of an unknown function $f \in \mathrm{L}_2([0,1]^d)$ with coefficients $\hat{f}_{\boldsymbol{k}}$ determined from the methods proposed in Section 4, U a subset of ANOVA terms and a grouped index set $\mathcal{I}_{\boldsymbol{N}}(U) \subseteq \mathbb{N}_0^d$. Now, we use as a model function

$$\tilde{f}^{\mathcal{X}_{\mathrm{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U)}(\boldsymbol{x}) \coloneqq \begin{cases} 0 & : S^{\mathcal{X}_{\mathrm{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U)} f(\boldsymbol{x}) < 0.5 \\ 1 & : S^{\mathcal{X}_{\mathrm{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U)} f(\boldsymbol{x}) \geq 0.5 \end{cases}$$

by deciding with a threshold of 0.5. The data set provides us with a way of validating our result by providing the correct answer $g(\boldsymbol{x}) \in \{0,1\}$ for every $\boldsymbol{x} \in \mathcal{X}_{\mathrm{test}}$ which has not been used in obtaining the model function. As a quality measure (or error), we use the percentage of correctly classified people

$$p(\mathcal{I}_{\boldsymbol{N}}(U), \mathcal{X}_{\mathrm{train}}, \mathcal{X}_{\mathrm{test}}) \coloneqq 100 \left(1 - \frac{1}{M_{\mathrm{test}}} \sum_{\boldsymbol{x} \in \mathcal{X}_{\mathrm{test}}} \left| \tilde{f}^{\mathcal{X}_{\mathrm{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U)}(\boldsymbol{x}) - g(\boldsymbol{x}) \right| \right) \in [0, 100].$$

Our first step is to assume a model with superposition threshold $d_s = 2$, i.e., we use the set of ANOVA terms $U_2 = \{\boldsymbol{u} \subseteq \{1, 2, \dots, 12\} \colon |\boldsymbol{u}| \leq 2\}$, cf. (5). In order to get the grouped index set $\mathcal{I}_{\boldsymbol{N}}(U_2)$ we choose parameters

$$N_{\boldsymbol{u}} = \begin{cases} 82 & : |\boldsymbol{u}| = 1 \\ 10 & : |\boldsymbol{u}| = 2 \end{cases} \tag{27}$$

for the related groups of frequencies

$$\mathcal{I}^{\boldsymbol{u},d}_{N_{\boldsymbol{u}}} \coloneqq \left\{ \boldsymbol{k} \in \mathbb{N}_0^d \colon \operatorname{supp} \boldsymbol{k} = \boldsymbol{u} \text{ and } 1 \leq k_j \leq N_{\boldsymbol{u}} - 1 \text{ for } j \in \boldsymbol{u} \right\}, \boldsymbol{u} \in U_2,$$

such that $|\mathcal{I}_{\boldsymbol{N}}(U_2)| = 6\,319$. Note that all of the following results have been obtained by using 10-fold cross-validation as in [22], i.e., the split of our data into training nodes $\mathcal{X}_{\mathrm{train}}$ and test nodes $\mathcal{X}_{\mathrm{test}}$ has been performed ten times and

the results were averaged. This averaging was also performed for the global sensitivity indices. We notice an optimal percentage of correctly classified people $p(\mathcal{I}_{\boldsymbol{N}}(U_2), \mathcal{X}_{\text{train}}, \mathcal{X}_{\text{test}}) = 86.15\%$ for the ℓ_2 regularization from Section 4.1 and $p(\mathcal{I}_{\boldsymbol{N}}(U_2), \mathcal{X}_{\text{train}}, \mathcal{X}_{\text{test}}) = 86.08\%$ for the group lasso approach from Section 4.2

In the case of ℓ_2 regularization, we try to improve the error by determining an active set of ANOVA terms. However, the situation is not as clear as for the synthetic example in Section 4. Trough testing, we determined that it delivers the best improvement if we choose $\boldsymbol{\varepsilon} = (0.1, 0.1) \in \mathbb{R}^2$ such that terms $f_{\boldsymbol{u}}$ with global sensitivity indices $\varrho(\boldsymbol{u}, S^{\mathcal{X}_{\text{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U_2)} f) < 0.1$ are removed for the approximation $S^{\mathcal{X}_{\text{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U^{(\boldsymbol{\varepsilon})}_{\mathcal{X}_{\text{train}},\boldsymbol{y}})}$. We may also increase the bandwidth for the one-dimensional terms such that we have

$$N_{\boldsymbol{u}} = \begin{cases} 300 & : |\boldsymbol{u}| = 1 \\ 10 & : |\boldsymbol{u}| = 2. \end{cases} \tag{28}$$

This reduces the number of terms from $|U_2| = 79$ to $|U^{(\boldsymbol{\varepsilon})}_{\mathcal{X}_{\text{train}},\boldsymbol{y}}| = 35$, cf. (15), and it also increases the number of correctly classified people to a best result of $p(\mathcal{I}_{\boldsymbol{N}}(U^{(\boldsymbol{\varepsilon})}_{\mathcal{X}_{\text{train}},\boldsymbol{y}}), \mathcal{X}_{\text{train}}, \mathcal{X}_{\text{test}}) = 86.30\%$.

In Figure 11 we have visualized our results. From Figure 11a it is evident that the group lasso approach does not identify important terms correctly for this data set which yields a lower number of correctly classified people with increasing regularization parameter. An explaination is the lack of sparsity in the underlying function. The FISTA algorithm still eliminates many ANOVA terms while sacrificing data fitting quality. Moreover, for the ℓ_2 regularization we see that the performance can be improved trough sensitivity analysis and reducing the number of terms. The (averaged) global sensitivity indices of the 12 one-dimensional ANOVA terms are depicted in Figure 11b for the ℓ_2 regularization and in Figure 11c for the group lasso. Moreover, we have the sensitivity indices for the 66 two-dimensional ANOVA terms are depicted in Figure 11d and Figure 11e. We can see that the group lasso does not identify the same terms as the sensitivity analysis. However, trough the validation on our test set, we observe that the sensitivity analysis yields an active set of terms $U^{(\boldsymbol{\varepsilon})}_{\mathcal{X}_{\text{train}},\boldsymbol{y}}$ which results in a larger number of correctly classified people.

In Figure 11b we observe that one term explains around 25% of the variance of our approximation. This attribute is the *capital loss* and it therefore is the most important attribute according to our model obtained by ℓ_2 regularization. The same term appears with about 35% of the variance in Figure 11c obtained by group lasso which means that even tough their performance differs, both models identify the same variable as vital to the problem. However, the group lasso approach starts to penalize this term with increasing regularization parameter which may be one reason the performance of this model is dropping.

Note that in order to compare our experiments to other known results, we need to be in a similar setting. We compare to the results from [22] and therefore use 10-fold cross-validation to verify our result. Moreover, it is in general possible to perform exploratory data analysis, e.g., remove outliers, which was not done in [22] and we also did not perform here. Table 1 shows the comparison. The ratio of training and test data is with 80% to 20% the same. Our approach yields a better classification result than all other approaches from [22] which include a support vector machine (SVM), a C4.5 decision-tree induction, a Naive-Bayes approach, and NBtree which is a hybrid between decision-tree and Naive-Bayesian classifiers.

method	correctly classified persons
SVM	85.03%
C4.5	84.46%
Naive Bayes	83.88%
NBtree	85.90%
ANOVAapprox (ℓ_2)	**86.30%**
ANOVAapprox (group lasso)	86.08%

Table 1: Comparison of the performance of different classification approaches from [22] on the adult census data set to methods proposed in this paper.

6 Conclusion

In the context of Fourier analysis, we introduced grouped frequency index sets $\mathcal{I}_{\boldsymbol{N}}(U)$ for a subset of ANOVA terms $U \subseteq \mathcal{P}(\mathcal{D})$, see (8), consisting of groups $\mathcal{I}^{\boldsymbol{u},d}_{N}$, see (7), along lower-dimensional cubes in the frequency domain. These can then be used to move the computations from the original d dimensions to the dimension of the groups by exploiting

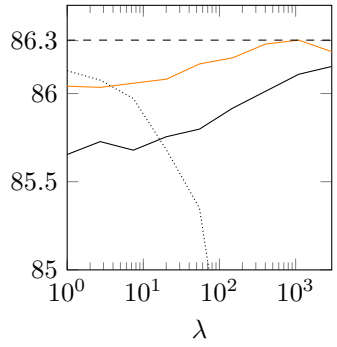

(a) $p(\mathcal{I}_{\boldsymbol{N}}(U_2), \mathcal{X}_{\text{train}}, \mathcal{X}_{\text{test}})$ for ℓ_2 regularization (black) and for group lasso (dotted). $p(\mathcal{I}_{\boldsymbol{N}}(U^{(\boldsymbol{\varepsilon})}_{\mathcal{X}_{\text{train}},\boldsymbol{y}}), \mathcal{X}_{\text{train}}, \mathcal{X}_{\text{test}})$ for ℓ_2 regularization after sensitivity analysis (orange).

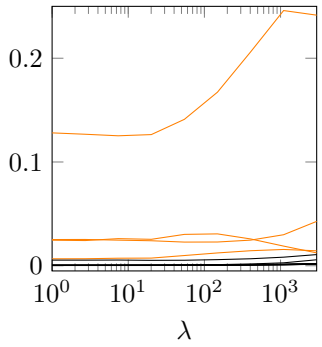

(b) $\varrho(\boldsymbol{u}, S^{\mathcal{X}_{\text{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U_2)} f)$, $|\boldsymbol{u}| = 1$, for ℓ_2 regularization and LSQR solver.

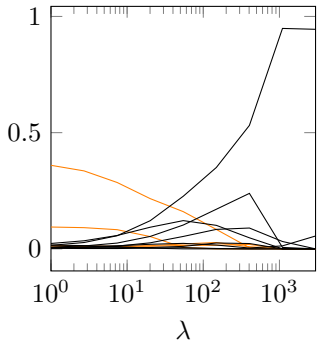

(c) $\varrho(\boldsymbol{u}, S^{\mathcal{X}_{\text{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U_2)} f)$, $|\boldsymbol{u}| = 1$, for group lasso and FISTA solver.

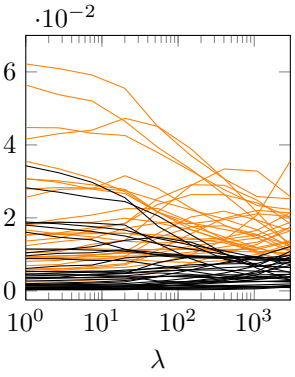

(d) $\varrho(\boldsymbol{u}, S^{\mathcal{X}_{\text{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U_2)} f)$, $|\boldsymbol{u}| = 2$, for ℓ_2 regularization and LSQR solver.

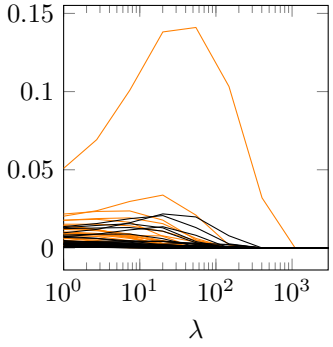

(e) $\varrho(\boldsymbol{u}, S^{\mathcal{X}_{\text{train}}}_{\mathcal{I}_{\boldsymbol{N}}(U_2)} f)$, $|\boldsymbol{u}| = 2$, for group lasso and FISTA solver.

Figure 11: Numerical experiments with the adult census data set using bandwidths (27) in $\mathcal{I}_{\boldsymbol{N}}(U_2)$ and bandwidths (28) in $\mathcal{I}_{\boldsymbol{N}}(U^{(\boldsymbol{\varepsilon})}_{\mathcal{X}_{\text{train}},\boldsymbol{y}})$.

the block structure (12) in the matrix. The resulting transformations are only of dimension up to a superposition threshold $d_s < d$. We presented a fast algorithm — the grouped Fourier transform — based on the identity (13). We were able to carry out every matrix-vector multiplication via the nonequispaced fast Fourier transformation (NFFT) to obtain the typical FFT-like complexity. Moreover, the transformation can be computed simultaneously and independent of the other transformations allowing for parallelization of the algorithm.

Even though we assumed a very special structure, there are proven applications: We have the one-to-one correspondence to the ANOVA decomposition, which we discussed in Section 2. As proven in [35], limiting the variable interactions to a superposition threshold $d_s < d$ works well for functions with a low superposition dimension. Moreover, certain function classes naturally lead to a low superposition dimension for high accuracy. When assuming sparsity-of-effects or the Pareto principle, one may also find that this idea works well for real world problems since those systems are dominated by a few low complexity interactions.

In Section 4 we presented two approaches to modify the least-squares problem (14) occurring in the approximation of the truncated ANOVA decomposition $\mathrm{T}_{d_s} f$, cf. (6) and the ideas in [35].

(i) The first approach follows the concept of minimizing the Tikhonov-functional (21). This extensions allows for a modification which is able to handle noise and incorporate smoothness information which reduces the *search space* of the minimization to functions of the smoothness class. This enables us to use significantly larger groups $\mathcal{I}_N^{\boldsymbol{u},d}$ reducing the cut-off error from considering the Fourier partial sum (9), cf. [35, Section 6].

(ii) The second approach uses the group lasso idea, where we proposed to use the ANOVA terms for the groups, see (25). In contrast to the classical Tikhonov-regularization this promotes sparsity in the ANOVA decomposition directly trough the regularization which eliminates the need for sensitivity analysis and therefore only one minimization needs to be solved. As there is no explicit formula for the solution of this non-smooth minimization problem, we used the FISTA algorithm to tackle it.

Both approaches performed well in the numerical experiments from Section 4 which cover the approximation of a synthetic test function in under-, overdertermined, and also noisy settings. It was successful to implement the sensitivity analysis into the regularization via a group lasso approach. Moreover, reducing the search space by considering functions of a certain smoothness also worked well. However, the situation is different for real data, i.e., the adult census data set, in Section 5. Here, we observe that group lasso has trouble identifying important terms and does not perform as well as a manual sensitivity analysis and ℓ_2 regularization. In summary, we were able to outperform known results from a SVM and other classical machine learning methods on the data set.

The code is available in the Julia packages `ANOVAapprox` and `GroupedTransforms` on GitHub, see `https://github.com/NFFT/ANOVAapprox` and `https://github.com/NFFT/GroupedTransforms`. As shown in [35], it is also possible to work with a huge amount of data, namely multiple million nodes, using these methods. This would of course be a severely overdetermined setting.

Since we end up with a Fourier representation of the ANOVA decomposition, we can visualize it in the form of an explainable ANOVA network, see Figure 12. This representation immediately shows the couplings of the variables or the ANOVA terms and, with help of the global sensitivity indices (3), depicts the importance of individual ANOVA terms or, in other words, which terms are active and which are not.

Acknowledgments

Felix Bartel acknowledges funding by the European Social Fund (ESF), Project ID 100367298. Daniel Potts acknowledges funding by the German Research Foundation (Deutsche Forschungsgemeinschaft) - Project ID 416228727 - SFB 1410. Michael Schmischke is supported by the BMBF grant 01|S20053A.

References

[1] J. BALDEAUX AND M. GNEWUCH, *Optimal randomized multilevel algorithms for infinite-dimensional integration on function spaces with ANOVA-type decomposition*, SIAM J. Numer. Anal., 52 (2014), pp. 1128–1155, `https://doi.org/10.1137/120896001`.

[2] F. BARTEL, R. HIELSCHER, AND D. POTTS, *Fast cross-validation in harmonic approximation*, Appl. Comput. Harmon. Anal., 49 (2020), pp. 415–437, `https://doi.org/10.1016/j.acha.2020.05.002`.

[3] A. BECK AND M. TEBOULLE, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM J. Imaging Sci., 2 (2009), pp. 183–202, `https://doi.org/10.1137/080716542`.

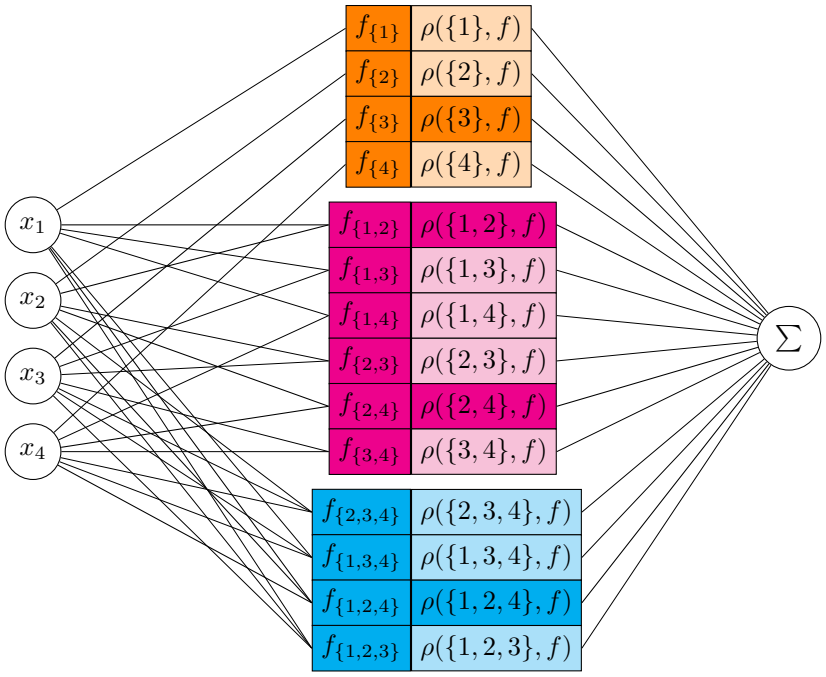

Figure 12: Explainable ANOVA network of a function f for $d_s = 3$. We visualize the different ANOVA terms $f_{\boldsymbol{u}}$ for $|\boldsymbol{u}| = 1$ in orange, $|\boldsymbol{u}| = 2$ in magenta and for $|\boldsymbol{u}| = 3$ in blue. The related global sensitivity indices $\rho(\boldsymbol{u}, f)$ serve as interpretable quantitiy for the ANOVA term.

[4] H.-J. BUNGARTZ AND M. GRIEBEL, *Sparse grids*, Acta Numer., 13 (2004), pp. 147–269, https://doi.org/10.1017/S0962492904000182.

[5] R. CAFLISCH, W. MOROKOFF, AND A. OWEN, *Valuation of mortgage-backed securities using Brownian bridges to reduce effective dimension*, J. Comput. Finance, 1 (1997), pp. 27–46, https://doi.org/10.21314/jcf.1997.005.

[6] R. DEVORE, G. PETROVA, AND P. WOJTASZCZYK, *Approximation of functions of few variables in high dimensions*, Constr. Approx., 33 (2010), pp. 125–143, https://doi.org/10.1007/s00365-010-9105-8.

[7] D. DUA AND C. GRAFF, *UCI machine learning repository*, 2017, http://archive.ics.uci.edu/ml.

[8] A. D. GILBERT, F. Y. KUO, D. NUYENS, AND G. W. WASILKOWSKI, *Efficient implementations of the multivariate decomposition method for approximating infinite-variate integrals*, SIAM J. Sci. Comput., 40 (2018), pp. A3240–A3266, https://doi.org/10.1137/17m1161890.

[9] M. GRIEBEL AND J. HAMAEKERS, *Fast discrete Fourier transform on generalized sparse grids*, in Sparse Grids and Applications - Munich 2012, J. Garcke and D. Pflüger, eds., vol. 97 of Lect. Notes Comput. Sci. Eng., Springer International Publishing, 2014, pp. 75–107, https://doi.org/10.1007/978-3-319-04537-5_4.

[10] M. GRIEBEL AND M. HOLTZ, *Dimension-wise integration of high-dimensional functions with applications to finance*, J. Complexity, 26 (2010), pp. 455–489, https://doi.org/10.1016/j.jco.2010.06.001.

[11] M. GRIEBEL, F. Y. KUO, AND I. H. SLOAN, *The ANOVA decomposition of a non-smooth function of infinitely many variables can have every term smooth*, Math. Comp., 86 (2016), pp. 1855–1876, https://doi.org/10.1090/mcom/3171.

[12] C. GU, *Smoothing Spline ANOVA Models*, Springer New York, 2013, https://doi.org/10.1007/978-1-4614-5369-7.

[13] A. HASHEMI, H. SCHAEFFER, R. SHI, U. TOPCU, G. TRAN, AND R. WARD, *Function approximation via sparse random features*, ArXiv e-prints 2103.03191, (2021).

[14] M. HEGLAND, *Adaptive sparse grids*, in ANZIAM J, 2001, pp. 335–353, https://doi.org/10.21914/anziamj.v44i0.685.

[15] M. HOLTZ, *Sparse grid quadrature in high dimensions with applications in finance and insurance*, vol. 77 of Lecture Notes in Computational Science and Engineering, Springer-Verlag, Berlin, 2011, https://doi.org/10.1007/978-3-642-16004-2.

[16] Y. JIANG AND Y. XU, *Fast discrete algorithms for sparse Fourier expansions of high dimensional functions*, J. Complexity, 26 (2010), pp. 51–81.

[17] L. KÄMMERER, *Reconstructing hyperbolic cross trigonometric polynomials by sampling along rank-1 lattices*, SIAM J. Numer. Anal., 51 (2013), pp. 2773–2796, http://dx.doi.org/10.1137/120871183.

[18] L. KÄMMERER, D. POTTS, AND T. VOLKMER, *Approximation of multivariate periodic functions by trigonometric polynomials based on rank-1 lattice sampling*, J. Complexity, 31 (2015), pp. 543–576, https://doi.org/10.1016/j.jco.2015.02.004.

[19] L. KÄMMERER, T. ULLRICH, AND T. VOLKMER, *Worst case recovery guarantees for least squares approximation using random samples*, Constr. Approx., 54 (2021), pp. 295–352, https://doi.org/10.1007/s00365-021-09555-0.

[20] J. KEINER, S. KUNIS, AND D. POTTS, *NFFT 3.5, C subroutine library*. http://www.tu-chemnitz.de/~potts/nfft. Contributors: F. Bartel, M. Fenn, T. Görner, M. Kircheis, T. Knopp, M. Quellmalz, M. Schmischke, T. Volkmer, A. Vollrath.

[21] J. KEINER, S. KUNIS, AND D. POTTS, *Using NFFT3 - a software library for various nonequispaced fast Fourier transforms*, ACM Trans. Math. Software, 36 (2009), pp. Article 19, 1–30, https://doi.org/10.1145/1555386.1555388.

[22] R. KOHAVI, *Scaling up the accuracy of naive-Bayes classifiers: a decision-tree hybrid*, in Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, 1996.

[23] T. KÜHN, S. MAYER, AND T. ULLRICH, *Counting via entropy: New preasymptotics for the approximation numbers of Sobolev embeddings*, SIAM J. Numer. Anal., 54 (2016), pp. 3625–3647, https://doi.org/10.1137/16m106580x.

[24] F. Y. KUO, G. MIGLIORATI, F. NOBILE, AND D. NUYENS, *Function integration, reconstruction and approximation using rank-1 lattices*, Math. Comp., 90 (2021), pp. 1861–1897, https://doi.org/10.1090/mcom/3595.

[25] F. Y. KUO, D. NUYENS, L. PLASKOTA, I. H. SLOAN, AND G. W. WASILKOWSKI, *Infinite-dimensional integration and the multivariate decomposition method*, J. Comput. Appl. Math., 326 (2017), pp. 217–234, https://doi.org/10.1016/j.cam.2017.05.031.

[26] F. Y. KUO, I. H. SLOAN, G. W. WASILKOWSKI, AND H. WOŹNIAKOWSKI, *On decompositions of multivariate functions*, Math. Comp., 79 (2009), pp. 953–966, https://doi.org/10.1090/s0025-5718-09-02319-9.

[27] R. LIU AND A. B. OWEN, *Estimating mean dimensionality of analysis of variance decompositions*, J. Amer. Statist. Assoc., 101 (2006), pp. 712–721, https://doi.org/10.1198/016214505000001410.

[28] L. MEIER, S. V. D. GEER, AND P. BÜHLMANN, *The group lasso for logistic regression*, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 70 (2008), pp. 53–71, https://doi.org/10.1111/j.1467-9868.2007.00627.x.

[29] H. NIEDERREITER, *Random Number Generation and Quasi-Monte Carlo Methods*, CBMS-NSF Regional Conference Series in Applied Mathematics, Society for Industrial and Applied Mathematics, 1992, https://doi.org/10.1137/1.9781611970081.

[30] J. E. OAKLEY AND A. O'HAGAN, *Probabilistic sensitivity analysis of complex models: a bayesian approach*, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 66 (2004), pp. 751–769, https://doi.org/10.1111/j.1467-9868.2004.05304.x.

[31] L. ONETO, S. RIDELLA, AND D. ANGUITA, *Tikhonov, Ivanov and Morozov regularization for support vector machine learning*, Mach. Learn., 103 (2016), pp. 103–136, https://doi.org/10.1007/s10994-015-5540-x.

[32] A. OWEN, *Effective dimension of some weighted pre-Sobolev spaces with dominating mixed partial derivatives*, SIAM J. Numer. Anal., 57 (2019), pp. 547–562, https://doi.org/10.1137/17m1158975.

[33] C. C. PAIGE AND M. A. SAUNDERS, *LSQR: An algorithm for sparse linear equations and sparse least squares*, ACM Trans. Math. Software, 8 (1982), pp. 43–71, https://doi.org/10.1145/355984.355989.

[34] G. PLONKA, D. POTTS, G. STEIDL, AND M. TASCHE, *Numerical Fourier Analysis*, Applied and Numerical Harmonic Analysis, Birkhäuser, 2018, https://doi.org/10.1007/978-3-030-04306-3.

[35] D. POTTS AND M. SCHMISCHKE, *Approximation of high-dimensional periodic functions with fourier-based methods*, SIAM J. Numer. Anal., 59 (2021), pp. 2393–2429, https://doi.org/10.1137/20m1354921.

[36] D. POTTS AND M. SCHMISCHKE, *Interpretable approximation of high-dimensional data*, SIAM J. Math. Data Sci. (accepted), (2021), https://arxiv.org/abs/2103.13787.

[37] D. POTTS AND M. SCHMISCHKE, *Learning multivariate functions with low-dimensional structures using polynomial bases*, J. Comput. Appl. Math., 403 (2022), p. 113821, https://doi.org/10.1016/j.cam.2021.113821.

[38] D. POTTS AND T. VOLKMER, *Sparse high-dimensional FFT based on rank-1 lattice sampling*, Appl. Comput. Harmon. Anal., 41 (2016), pp. 713–748.

[39] H. RABITZ AND O. F. ALIS, *General foundations of high dimensional model representations*, J. Math. Chem., 25 (1999), pp. 197–233, https://doi.org/10.1023/A:1019188517934.

[40] N. SIMON, J. FRIEDMAN, T. HASTIE, AND R. TIBSHIRANI, *A sparse-group lasso*, Journal of Computational and Graphical Statistics, 22 (2013), pp. 231–245, https://doi.org/10.1080/10618600.2012.681250.

[41] I. M. SOBOL, *On sensitivity estimation for nonlinear mathematical models*, Keldysh AppliedMathematics Institute, 1 (1990), pp. 112–118.

[42] I. M. SOBOL, *Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates*, Math. Comput. Simulation, 55 (2001), pp. 271–280, https://doi.org/10.1016/s0378-4754(00)00270-6.

[43] G. WASILKOWSKI AND H. WOŹNIAKOWSKI, *Liberating the dimension for function approximation*, J. Complexity, 27 (2011), pp. 86–110, https://doi.org/https://doi.org/10.1016/j.jco.2010.08.004.

[44] C. F. J. WU AND M. S. HAMADA, *Experiments - Planning, Analysis, and Optimization*, John Wiley & Sons, New York, 2011.

[45] H. YANG, Z. XU, I. KING, AND M. R. LYU, *Online learning for group lasso*, in ICML, 2010, pp. 1191–1198, https://icml.cc/Conferences/2010/papers/473.pdf.

[46] M. YUAN AND Y. LIN, *Model selection and estimation in regression with grouped variables*, J. R. Stat. Soc. Ser. B Stat. Methodol., 68 (2006), pp. 49–67, https://doi.org/10.1111/j.1467-9868.2005.00532.x.