

Fast and Interpretable Support Vector Classification Based on Truncated ANOVA Decompositions

Kseniya Akhalaya Franziska Nestler Daniel Potts

February 6, 2024

Support Vector Machines (SVMs) are an important tool for performing classification on scattered data, where one usually has to deal with many data points in high-dimensional spaces. We propose solving SVMs in primal form using feature maps based on trigonometric functions or wavelets. In small dimensional settings the Fast Fourier Transform (FFT) and related methods are a powerful tool in order to deal with the considered basis functions. For growing dimensions the classical FFT-based methods become inefficient due to the curse of dimensionality. Therefore, we restrict ourselves to multivariate basis functions, each one of them depends only on a small number of dimensions. This is motivated by the well-known sparsity of effects and recent results regarding the reconstruction of functions from scattered data in terms of truncated analysis of variance (ANOVA) decomposition, which makes the resulting model even interpretable in terms of importance of the features as well as their couplings. The usage of small superposition dimensions has the consequence that the computational effort no longer grows exponentially but only polynomially with respect to the dimension. In order to enforce sparsity regarding the basis coefficients, we use the frequently applied ℓ_2 -norm and, in addition, ℓ_1 -norm regularization. The found classifying function, which is the linear combination of basis functions, and its variance can then be analyzed in terms of the classical ANOVA decomposition of functions. Based on numerical examples we show that we are able to recover the signum of a function that perfectly fits our model assumptions. We obtain better results with ℓ_1 -norm regularization, both in terms of accuracy and clarity of interpretability.

Keywords: Support Vector Machines, Primal SVM, ℓ_1 -norm Regularization, Interpretability, ANOVA, FISTA, Gradient Descent, Trigonometric Feature Maps, Wavelets, Grouped Transformations

1. Introduction

In the following, we shortly review some basic work on Support Vector Machines (SVMs) for classification problems using the general references [42, 41], which we also recommend for all

further details and interesting theoretical results. Assume, we are given a finite training data set $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_M, y_M) \in \mathcal{X} \times \{-1, +1\}$. Here, \mathcal{X} is a nonempty finite set consisting of feature vectors $\mathbf{x}_j \in \mathbb{R}^d$, which we also call inputs or rather instances. The y_j are called labels or outputs. The classical SVM approach aims to construct a classifier $F : \mathcal{X} \rightarrow \{-1, +1\}$, also called decision function or predictor, of the form

$$F(\mathbf{x}) = \text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle + b), \quad \mathbf{w} \in \mathbb{R}^d, b \in \mathbb{R}, \quad (1.1)$$

where $\langle \cdot, \cdot \rangle : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is an inner product on \mathbb{R}^d with the agreement that $\text{sign}(0) := 0$. The main goal is to find a weight vector $\mathbf{w} \in \mathbb{R}^d$ and bias $b \in \mathbb{R}$ such that the prediction given by $F(\mathbf{x}_j)$ is correct for many or ideally all training data points (\mathbf{x}_j, y_j) , $j = 1, 2, \dots, M$. Let us assume that the hyperplane defined by

$$H(\mathbf{w}, b) := \left\{ \mathbf{x} \in \mathbb{R}^d : \langle \mathbf{w}, \mathbf{x} \rangle + b = 0 \right\}$$

with normal vector $\mathbf{w} / \|\mathbf{w}\|_2$ perfectly separates the training data set $\{(\mathbf{x}_j, y_j)\}_{j=1,2,\dots,M}$ into two classes with indices $\mathcal{I}_- := \{j \in 1, 2, \dots, M : y_j = -1\}$ and $\mathcal{I}_+ := \{j \in 1, 2, \dots, M : y_j = +1\}$. Thus, we want to ensure that feature vectors \mathbf{x}_j with negative corresponding labels y_j are on the negative side of the hyperplane $H(\mathbf{w}, b)$, i.e., $\langle \mathbf{w}, \mathbf{x}_j \rangle + b < 0$ for $j \in \mathcal{I}_-$ and feature vectors \mathbf{x}_j with positive corresponding labels y_j are on the positive side, i.e., $\langle \mathbf{w}, \mathbf{x}_j \rangle + b > 0$ for $j \in \mathcal{I}_+$. These two conditions are often presented in a single equation $y_j (\langle \mathbf{w}, \mathbf{x}_j \rangle + b) > 0$ and the data set is called linearly separable. For a linearly separable data set there may exist various separating hyperplanes. To find the best one, we first compute the distance between a point, described by input vector \mathbf{x}_j , and a separating hyperplane $H(\mathbf{w}, b)$ by

$$y_j \left(\left\langle \frac{\mathbf{w}}{\|\mathbf{w}\|_2}, \mathbf{x}_j \right\rangle + \frac{b}{\|\mathbf{w}\|_2} \right), \quad j = 1, 2, \dots, M$$

to derive the margin, which is defined as smallest distance,

$$\gamma := \min_{j=1,2,\dots,M} y_j \left(\left\langle \frac{\mathbf{w}}{\|\mathbf{w}\|_2}, \mathbf{x}_j \right\rangle + \frac{b}{\|\mathbf{w}\|_2} \right).$$

Further, we choose the separating hyperplane that maximizes the margin between these two classes of training data. The corresponding classifier is called maximal margin classifier or hard margin classifier, which was first introduced in [46], and constructs a required hyperplane under the constraint $\|\mathbf{w}\|_2 = 1$, i.e.,

$$\gamma \|\mathbf{w}\|_2 = \min_{j=1,2,\dots,M} y_j (\langle \mathbf{w}, \mathbf{x}_j \rangle + b).$$

Finally, by using an additional scaling assumption $y_j (\langle \mathbf{w}, \mathbf{x}_j \rangle + b) \geq 1$, we see that γ becomes maximal if and only if $\|\mathbf{w}\|_2$ becomes minimal. In summary, the hard margin classifier aims to find parameters $\mathbf{w} \in \mathbb{R}^d$ and $b \in \mathbb{R}$ solving the following quadratic optimization problem with linear constraints

$$\frac{1}{2} \|\mathbf{w}\|_2^2 \rightarrow \min_{\mathbf{w}, b} \quad \text{subject to} \quad y_j (\langle \mathbf{w}, \mathbf{x}_j \rangle + b) \geq 1, \quad j = 1, 2, \dots, M \quad (1.2)$$

and unique solution, see [45, Chapter 5]. As we can see, the formulation does not allow any violations of the margin condition, hence the expression "hard". There are many different methods for solving such optimization problems efficiently, see for example [28, Chapter 16].

To resolve the issue that a given data set is not linearly separable, we use a nonlinear SVM, first proposed in [9], which maps the feature vectors $\mathbf{x}_j \in \mathcal{X}$ into a possibly infinite-dimensional space H , the so-called feature space, by a typically non-linear map $\Phi : \mathbb{D}^d \rightarrow H$, defined on a finite domain \mathbb{D}^d , called feature map. In the presence of noise, it can happen that we need to misclassify some feature vectors \mathbf{x}_j in order to avoid overfitting, especially if $d \geq M$. Recall that in (1.2) the constraints force the hyperplanes to cause no error on the training data set. In contrast, the soft margin approach, suggested in [13], reduces these constraints by requiring only $y_j (\langle \mathbf{w}, \mathbf{x}_j \rangle + b) \geq 1 - \xi_j$ with slack variables $\xi_j \geq 0$, $j = 1, 2, \dots, M$.

In this paper we will focus on feature maps Φ given by

$$\Phi : \mathbb{D}^d \rightarrow \mathbb{R}^{|\mathcal{I}|}, \quad \mathbf{x} \mapsto \Phi(\mathbf{x}) := (\varphi_{\mathbf{k}}(\mathbf{x}))_{\mathbf{k} \in \mathcal{I}}, \quad (1.3)$$

where $\varphi_{\mathbf{k}}(\mathbf{x}) : \mathbb{D}^d \rightarrow \mathbb{R}$ are basis functions with certain multi indices $\mathbf{k} \in \mathcal{I} \subset \mathbb{Z}^d$. At this point, the question arises as to how this index set \mathcal{I} should be chosen, which is not trivial, especially in high-dimensional spaces. The choice of the index set \mathcal{I} is very important here and may have a major influence on the performance of the considered algorithms. For details concerning the choice of an index set \mathcal{I} we refer to Section 2.1. Using a function $f : \mathcal{X} \rightarrow \mathbb{R}$, we rewrite (1.1) as

$$F(\mathbf{x}) := \text{sign}(f(\mathbf{x})) = \text{sign} \left(\sum_{\mathbf{k} \in \mathcal{I}} \hat{f}_{\mathbf{k}} \varphi_{\mathbf{k}}(\mathbf{x}) \right), \quad \mathbf{x} \in \mathcal{X}, \mathcal{I} \subset \mathbb{Z}^d,$$

i.e., the affine linear function $\langle \mathbf{w}, \mathbf{x} \rangle + b$ is replaced by a linear combination of the basis functions $\varphi_{\mathbf{k}}$ with coefficients $\hat{f}_{\mathbf{k}} \in \mathbb{R}$ to be determined. Note that we omit the bias $b \in \mathbb{R}$ and include a constant function $\varphi_{\mathbf{0}}(\mathbf{x}) \equiv 1$ instead. In the following, we will call f the classifying function, which is again affine linear in the enlarged feature space $\{(\varphi_{\mathbf{k}}(\mathbf{x}))_{\mathbf{k} \in \mathcal{I}}, \mathbf{x} \in \mathbb{D}^d\} \subset \mathbb{R}^{|\mathcal{I}|}$, but non-linear in the original \mathbb{D}^d . Combining this idea with the slack variables ξ_j we introduced before, we obtain a quadratic optimization problem

$$\frac{1}{2} R(\hat{\mathbf{f}}) + C \sum_{j=1}^M \xi_j \rightarrow \min_{\hat{\mathbf{f}}, \boldsymbol{\xi}} \quad \text{subject to} \quad y_j \left(\langle \hat{\mathbf{f}}, \Phi(\mathbf{x}_j) \rangle \right) \geq 1 - \xi_j, \quad (1.4)$$

$$\xi_j \geq 0, \quad j = 1, 2, \dots, M$$

for the vector of basis coefficients $\hat{\mathbf{f}} = (\hat{f}_{\mathbf{k}})_{\mathbf{k} \in \mathcal{I}} \in \mathbb{R}^{|\mathcal{I}|}$, the vector of slack variables $\boldsymbol{\xi} = (\xi_j)_{j=1}^M \in \mathbb{R}^M$, regularization parameter $C > 0$ and regularization term $R(\hat{\mathbf{f}})$. In this paper, we will restrict our considerations to the cases $R(\hat{\mathbf{f}}) = \|\hat{\mathbf{f}}\|_2^2$ (ℓ_2 -norm regularization) and $R(\hat{\mathbf{f}}) = \|\hat{\mathbf{f}}\|_1$ (ℓ_1 -norm regularization). Since the bias b is now included in our feature map, also the constant part of the function f is penalized. This can make a small difference at the end, but is simpler in terms of implementation and is not uncommon when using non-linear regression, for example.

Further, we multiply the objective function in (1.4) by $2\lambda := \frac{1}{MC}$ and use a loss function $L : \{-1, +1\} \times \mathbb{R} \rightarrow [0, \infty)$ to reformulate the constrained problem as an unconstrained one

$$\lambda R(\hat{\mathbf{f}}) + \frac{1}{M} \sum_{j=1}^M L \left(y_j, \langle \hat{\mathbf{f}}, \Phi(\mathbf{x}_j) \rangle \right) \rightarrow \min_{\hat{\mathbf{f}}} \quad \hat{\mathbf{f}} \in \mathbb{R}^{|\mathcal{I}|} \quad (1.5)$$

with regularization parameter $\lambda > 0$. The loss function L is intended to penalize the presence of large slack variables ξ_j , i.e., it is ideally zero for all training data points (\mathbf{x}_j, y_j) , $j = 1, 2, \dots, M$.

The well-known hinge-loss is typically used in the context of SVMs. We will focus our attention on the squared hinge loss or rather truncated least squares loss function defined by

$$L\left(y_j, \langle \hat{\mathbf{f}}, \Phi(\mathbf{x}_j) \rangle\right) := \left(\max\left\{0, 1 - y_j \langle \hat{\mathbf{f}}, \Phi(\mathbf{x}_j) \rangle\right\}\right)^2,$$

since it is smooth. But also other loss functions, as for example least squares loss, are possible, see [42, Chapter 2]. At this point we would like to mention that well-known libraries for solving such SVMs are LIBSVM [11] and LIBLINEAR [17]. These are two open source machine learning libraries, with existing interfaces and extensions for many programming languages, e.g., Python or Julia. While LIBSVM incorporates algorithms for kernelized SVMs, LIBLINEAR implements linear SVMs. The mentioned references give an overview about theoretical convergence results, some implementation issues and performance comparisons to other state-of-the-art solvers.

As we can see, the standard form of SVM uses ℓ_2 -norm regularization, i.e., $R(\hat{\mathbf{f}}) = \|\hat{\mathbf{f}}\|_2^2$. By shrinking the magnitude of the coefficients, the ℓ_2 -norm penalty reduces the variance of the estimated coefficients, and thus can achieve better prediction accuracy. However, the ℓ_2 -norm penalty does in general not produce highly sparse coefficients and hence cannot automatically perform variable selection very well. This is the major limitation for applying support vector classification for high-dimensional data, where variable selection is essential for providing reasonable interpretations. Therefore, we use the expanded traditional SVM formulation, see [50], to enforce more sparsity and better variable selection, using a different regularization approach, e.g., ℓ_1 -norm regularization or rather a lasso penalty term $R(\hat{\mathbf{f}}) = \|\hat{\mathbf{f}}\|_1$ in (1.5), see [44].

It should also be noted that usually SVMs are solved in the dual form including a kernel matrix of size $M \times M$, cf. [37] and references therein. If the regularization $R(\hat{\mathbf{f}}) = \|\hat{\mathbf{f}}\|_2^2$ is used in (1.4), for instance, the dual optimization problem reads as

$$\frac{1}{2} \boldsymbol{\alpha}^\top Q \boldsymbol{\alpha} - \mathbf{1} \boldsymbol{\alpha} \rightarrow \min_{\boldsymbol{\alpha}} \quad \text{subject to} \quad \mathbf{y}^\top \boldsymbol{\alpha} = 0, \quad (1.6)$$

$$0 \geq \alpha_j \geq C, \quad j = 1, 2, \dots, M,$$

where $Q \in \mathbb{R}^{M \times M}$ is the so-called kernel matrix, which is typically positive (semi-)definite, and $Q_{ij} = y_i y_j K(\mathbf{x}_i, \mathbf{x}_j)$ for a given kernel function $K(\cdot, \cdot)$. A frequently used kernel function is the radial basis function (RBF) kernel or rather the Gaussian kernel, for instance. Kernel matrices are typically non-sparse and badly conditioned, which makes it hard to work with them in case of large data sets. Thus, it has become a common approach in the field of kernel methods to replace the kernel matrix Q by a low-rank approximation $Q \approx \Psi^\top \Psi$ in order to compute matrix-vector products more efficiently. This is the main idea of the well-known random Fourier features approach [34], for instance. A similar approach is applied in [27, 47], where the authors use an expansion of the kernel function K in terms of trigonometric basis functions and make use of nonuniform FFTs in order to compute the matrix-vector products in an efficient manner. In other words, in both settings we go back from the kernel setting to the feature-map formulation, where one can jump between primal and dual formulations.

This is the main motivation for the present work. Instead of taking a detour through the kernel formulation and solving the dual optimization problem, we directly work with the original or rather primal formulation combined with a finite dimensional feature map (1.3). Our feature map is built from basis functions, with which we can work efficiently for huge amount of data. If the optimization (1.5) is solved iterative, one has to compute matrix-vector

products of the form

$$\Phi \hat{\mathbf{f}} = \left(\left\langle \hat{\mathbf{f}}, \Phi(\mathbf{x}_j) \right\rangle \right)_{j=1}^M, \quad \Phi := (\varphi_{\mathbf{k}}(\mathbf{x}_j))_{j=1,2,\dots,M, \mathbf{k} \in \mathcal{I}} \in \mathbb{R}^{M \times |\mathcal{I}|}$$

in each iteration. In order to obtain an efficient algorithm, it is essential to replace the classic matrix-vector multiplication with a more favorable alternative. In case of trigonometric bases, for instance, we can efficiently approximate the required matrix-vector products using FFT-based methods for nonequispaced data [16, 6, 40, 30]. When making use of wavelets, the resulting matrices are sparse and, thus, the matrix-vector products can be computed efficiently, see [14, Chapter 3] or [48, Chapter 6].

In order to overcome the curse of dimensionality, we assume that high-dimensional interactions between features are negligible and restrict our feature map Φ to multivariate basis functions that only depend on a small number of dimensions each, which is equivalent to a corresponding restriction of the index set \mathcal{I} . The underlying rather simple structure of orthonormal systems makes it even possible to understand the importance of single features or couplings between them, based on the analysis of variance (ANOVA) decomposition. This work thus continues the previous works [31, 33, 24], which dealt exclusively with regression problems. Combining this idea with a customized SVM approach for classification problems, we are able to solve classification tasks in an interpretable fashion. In summary, our main contribution is the transfer of the ANOVA-based regression idea to classification scenarios by using the corresponding feature maps within SVMs.

This paper is organized as follows. In Section 2.2 we introduce the applied basis functions as well as the classical ANOVA decomposition of functions. Furthermore, we summarize certain properties of this decomposition and discuss its interpretability in terms of Sobol indices, [38, 39]. Finally, the efficient computation based on grouped transformations [2] is explained. We continue with the discussion of the implemented optimization algorithms in Section 3, where we distinguish between ℓ_2 -norm and ℓ_1 -norm regularization. Both penalty terms are considered in combination with the squared hinge loss, which is smooth. Thus, the ℓ_2 -norm regularized optimization problem is solved with a gradient descent method, whereas the ℓ_1 -norm approach is treated based on FISTA [5]. Finally, we present first numerical experiments in Section 4 and conclude with a short summary in Section 5.

2. ANOVA Decomposition for Interpretability

In this section we define the classical analysis of variance (ANOVA) decomposition of functions, cf. [10, 25, 23, 19], and summarize the main theoretical results presented in [31, 33, 24], where the ANOVA decomposition was studied in the context of series expansions involving trigonometric functions or wavelets, see Section 2.1. Furthermore, the interpretability of the ANOVA decomposition based on the Sobol indices or rather global sensitivity indices, cf. [38, 39], is discussed in Section 2.2. Finally, we take a closer look into the results presented in [2], namely the concept of grouped index sets and grouped transformations, which provide us an improvement in the complexity for the evaluation of the arising matrix-vector products in (1.5), see Section 2.3.

2.1. Function spaces, systems and frequency sets

In this Section we lay the foundation and present basic definitions, that are indispensable for the rest of this paper. In our setting, we consider square-integrable functions over a domain $\mathbb{D} \in \{\mathbb{T}, [0, 1/2]\}$ with the torus $\mathbb{T} := \mathbb{R}/\mathbb{Z} \simeq [-1/2, 1/2)$, i.e.,

$$f \in L_2(\mathbb{D}^d) := \left\{ f : \mathbb{D}^d \rightarrow \mathbb{C} : \|f\|_{L_2(\mathbb{D}^d)} := \sqrt{\int_{\mathbb{D}^d} |f(\mathbf{x})|^2 d\mathbf{x}} < \infty \right\}$$

with dimension $d \in \mathbb{N}$. Note that $L_2(\mathbb{D}^d)$ is a separable Hilbert space with inner product

$$\langle f, g \rangle_{L_2(\mathbb{D}^d)} = \frac{1}{\tau^d} \int_{\mathbb{D}^d} f(\mathbf{x}) \overline{g(\mathbf{x})} d\mathbf{x}, \quad f, g \in L_2(\mathbb{D}^d),$$

where we denote by $\tau \in \{1, 1/2\}$ the corresponding interval length. Using the property of separability, we know that there exists a countable complete orthonormal system or basis in the separable Hilbert space. In the case of tensor product spaces, we are able to construct a basis from the one-dimensional case as follows. We construct a basis of $L_2(\mathbb{D}^d)$ for $d > 1$ by defining the basis functions

$$\varphi_{\mathbf{k}}(\mathbf{x}) := \prod_{i=1}^d \varphi_{k_i}(x_i), \quad (2.1)$$

for $\mathbf{k} = (k_i)_{i=1}^d \in \mathbb{Z}^d$. Now we may write any $f \in L_2(\mathbb{D}^d)$ in terms of a basis expansion

$$f(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} c_{\mathbf{k}}(f) \varphi_{\mathbf{k}}(\mathbf{x})$$

with the basis coefficients $c_{\mathbf{k}}(f) = \langle f, \varphi_{\mathbf{k}} \rangle$, which are simply obtained as inner products of the single basis functions and the function f , due to the orthogonality. Note, that Parseval's identity

$$\|f\|_{L_2(\mathbb{D}^d)} = \sqrt{\sum_{\mathbf{k} \in \mathbb{Z}^d} |c_{\mathbf{k}}(f)|^2}$$

holds true, see [49, Chapter 5].

In general, a function $f \in L_2(\mathbb{D}^d)$ has an infinite number of basis coefficients $c_{\mathbf{k}}(f)$, $\mathbf{k} \in \mathbb{Z}^d$, of course. However, if the coefficients $c_{\mathbf{k}}(f)$ become negligibly small for large indices \mathbf{k} , which is typically the case for sufficiently smooth functions, we may truncate the infinite sum and replace \mathbb{Z}^d by a finite frequency set or index set, which we denote by $\mathcal{I} \subseteq \mathbb{Z}^d$, in order to get a good approximation of f in terms of a partial sum

$$f(\mathbf{x}) \approx \sum_{\mathbf{k} \in \mathcal{I}} c_{\mathbf{k}}(f) \varphi_{\mathbf{k}}(\mathbf{x}). \quad (2.2)$$

However, in practice the function f is not known and we may fit the basis coefficients in (2.2) such that the resulting partial sum describes our data in an appropriate manner. In case of binary classification, for instance, we may assume that our data is described by a function $f \in L_2(\mathbb{D}^d)$ but we only have information about the signum of f , i.e., we only know labels

$y_j = \text{sign}(f(\mathbf{x}_j))$ for the given feature vectors \mathbf{x}_j . Thus, by solving (1.5) we try to find a partial sum with coefficients $\hat{f}_{\mathbf{k}}$, $\mathbf{k} \in \mathcal{I}$, that satisfies

$$y_j \stackrel{!}{=} \text{sign} \left(\sum_{\mathbf{k} \in \mathcal{I}} \hat{f}_{\mathbf{k}} \varphi_{\mathbf{k}}(\mathbf{x}_j) \right).$$

Note, that this does not imply that we are approximating the coefficients of the original function f , i.e., in general we will compute coefficients $\hat{f}_{\mathbf{k}} \not\approx c_{\mathbf{k}}(f)$. In addition, also the index set \mathcal{I} is in general not known in practice. In other words, we also have to determine a reasonable index set \mathcal{I} , which is not straight forward in high dimensions.

In this paper, we make use of the fact that we are able to compute the finite sum in (2.2) for arbitrary feature vectors $\mathbf{x}_j \in \mathbb{D}^d$, $j = 1, 2, \dots, M$ in an efficient manner in case of small dimensions d , see Section 2.3. But first, let us take a closer look at the possible function systems and corresponding index sets that we will use in the upcoming sections concerning numerical experiments.

2.1.1. The cosine basis

The half-period cosine system with basis functions

$$\varphi_{\mathbf{k}}^{\text{cos}}(\mathbf{x}) := 2^{|\text{supp } \mathbf{k}|} \prod_{i=1}^d \cos(2\pi k_i x_i), \quad \mathbf{x} = (x_i)_{i=1}^d,$$

where $|\text{supp } \mathbf{k}| = |\{j \in \{1, 2, \dots, d\} : k_j \neq 0\}|$ denotes the number of non-zero entries in $\mathbf{k} \in \mathbb{N}_0^d$, is a well-known basis of $L_2([0, 1/2]^d)$, cf. [43, Chapter 9]. It is obtained by applying the tensor product ansatz (2.1) to the univariate cosine basis.

To obtain finite sums, we truncate the series expansion using a finite frequency or rather index set. For a given multi degree $\mathbf{N} = (N_i)_{i=1}^d \in (2\mathbb{N})^d$ we denote by

$$\mathcal{I}_{\mathbf{N}} := \{0, \dots, N_1 - 1\} \times \dots \times \{0, \dots, N_d - 1\} \quad (2.3)$$

the corresponding frequency index set of cardinality $\prod_{i=1}^d N_i$. Given such an index set $\mathcal{I}_{\mathbf{N}}$, we define the classifying function by

$$S(\mathcal{X}, \mathcal{I}_{\mathbf{N}})f(\mathbf{x}) := \sum_{\mathbf{k} \in \mathcal{I}_{\mathbf{N}}} \hat{f}_{\mathbf{k}} \varphi_{\mathbf{k}}^{\text{cos}}(\mathbf{x}), \quad (2.4)$$

where $\hat{\mathbf{f}} = (\hat{f}_{\mathbf{k}})_{\mathbf{k} \in \mathcal{I}_{\mathbf{N}}}$ solve (1.5). As explained above, the partial sum (2.4) or rather the included coefficients $\hat{f}_{\mathbf{k}}$ are computed by solving (1.5) in order to given $S(\mathcal{X}, \mathcal{I}_{\mathbf{N}})f(\mathbf{x}_j) = \text{sign}(f(\mathbf{x}_j))$ for the given training data set, where it is our model assumption that such a function f , which describes the data in this way, exists.

Finally, we create the corresponding feature map via (1.3) and define the associated feature matrix by

$$\Phi^{\text{cos}}(\mathcal{X}, \mathcal{I}_{\mathbf{N}}) := (\varphi_{\mathbf{k}}^{\text{cos}}(\mathbf{x}))_{\mathbf{x} \in \mathcal{X}, \mathbf{k} \in \mathcal{I}_{\mathbf{N}}} \in \mathbb{R}^{M \times |\mathcal{I}_{\mathbf{N}}|}. \quad (2.5)$$

Note that a straightforward matrix-vector multiplication requires $\mathcal{O}(M |\mathcal{I}_{\mathbf{N}}|)$ arithmetical operations, which can be improved by applying the fast cosine transform for nonequispaced data (NFCT), see [30, Chapter 7] or [22], yielding a complexity of $\mathcal{O}(|\mathcal{I}_{\mathbf{N}}| \log |\mathcal{I}_{\mathbf{N}}| + M)$. With growing dimension d the number of frequencies $|\mathcal{I}_{\mathbf{N}}|$ will grow exponentially in d and computing an NFCT would be far too expensive. However, in case of small dimensions $d \lesssim 4$ such algorithms are still very efficient.

2.1.2. The Chui-Wang basis

Next we present a system consisting of periodized, translated and dilated wavelets following [24], where the authors derived a lot of theory on hyperbolic wavelet regression, including considerations about the ANOVA decomposition of such functions. For more basic theory about wavelets we refer to the classical literature as for instance [12, 14, 48]. While the above discussed cosine system perfectly fits into the framework described at the beginning of Section 2.1, the notation becomes somewhat more complicated in case of wavelets. That is, for example, that we have to deal with two different indices \mathbf{j} and \mathbf{k} now. Moreover, wavelet systems do in general not form an orthonormal basis. Nevertheless, the essential principles remain the same.

The Chui-Wang basis functions are obtained as follows. First, we introduce the nested function spaces V_j for $j \in \mathbb{Z}$ by $V_j = \text{span}\{\phi(2^j \cdot -k) : k \in \mathbb{Z}\}$, where the function ϕ is called the scaling function. In the case of Chui-Wang wavelets, the scaling function is a cardinal B-spline $B_m : \mathbb{R} \rightarrow \mathbb{R}$ of order $m \in \mathbb{N}$, which are recursively defined via

$$B_1(x) := \begin{cases} 1, & -1/2 < x < 1/2 \\ 0, & \text{otherwise} \end{cases} \quad \text{and} \quad B_m(x) := \int_{x-1/2}^{x+1/2} B_{m-1}(y) dy.$$

As a consequence, we deduce

$$\cdots \subset V_{-1} \subset V_0 \subset V_1 \subset \cdots.$$

Moreover, the following properties, as $\bigcap_{j \in \mathbb{Z}} V_j = \{0\}$ and $\overline{\bigcup_{j \in \mathbb{Z}} V_j} = L_2(\mathbb{R})$ hold true for the spaces V_j . Now, the wavelet space W_j is defined as the orthogonal complement of a space V_j in V_{j+1} , i.e., $V_{j+1} = V_j \oplus W_j$ for $j \in \mathbb{N}_0$, and, consequently, we obtain an orthogonal decomposition of $L_2(\mathbb{R})$ by

$$V_{j+1} = V_j \oplus W_j, \quad j \in \mathbb{N}_0,$$

as well as an orthogonal decomposition of $L_2(\mathbb{R})$ by

$$L_2(\mathbb{R}) = V_0 \oplus \bigoplus_{j=0}^{\infty} W_j.$$

The wavelet spaces W_j , $j \in \mathbb{N}_0$, themselves are generated by a function φ , called the wavelet, in terms of $W_j = \text{span}\{\varphi_{j,k} : k \in \mathbb{Z}\}$, where

$$\varphi_{j,k}(x) = 2^{j/2} \varphi(2^j x - k), \quad k \in \mathbb{Z}$$

are scaled and shifted versions of the wavelet function φ . Using a cardinal B-spline of order m the corresponding wavelet function φ is the Chui-Wang wavelet, [12, Chapter 6]. Note that in the special case $m = 1$ we obtain the well-known Haar wavelets [20]. The Haar wavelet is one of the simplest wavelets, but it is not even continuous. However, this apparent weakness can become an advantage in the case one wants to approximate functions containing abrupt changes or jumps.

Now we define the one-periodic versions of the scaling function ϕ and the wavelet φ by

$$\phi_{j,k}^{\text{per}}(x) := \sum_{\ell \in \mathbb{Z}} \phi_{j,k}(x + \ell) \quad \text{and} \quad \varphi_{j,k}^{\text{per}}(x) := \sum_{\ell \in \mathbb{Z}} \varphi_{j,k}(x + \ell).$$

Motivated by the theoretical results presented in [24], we make use of the periodic Chui-Wang wavelet basis in this paper, although we deal with non-periodic examples later on. Especially, we make use of periodic Haar wavelets in our numerical examples since sharp jumps between classes have to be modeled. Please note that an appropriate extension of the wavelet basis to be used for non-periodic functions on the interval is by no means straight forward.

Note in addition that wavelets do in general not induce an orthonormal basis, but a Riesz-basis. Thus, in order to compute the wavelet coefficients one has to make use of the so-called dual basis $\varphi_{j,k}^*$ fulfilling $\langle \varphi_{j,k}, \varphi_{i,\ell}^* \rangle = \delta_{i,j} \delta_{k,\ell}$. Then, any function $f \in L_2(\mathbb{T})$ can be decomposed as

$$\begin{aligned} f &= \langle f, \varphi_{-1,0}^{\text{per}*} \rangle \varphi_{-1,0}^{\text{per}} + \sum_{j \geq 0} \sum_{k \in \mathbb{Z}} \langle f, \varphi_{j,k}^{\text{per}*} \rangle \varphi_{j,k}^{\text{per}} \\ &=: c_{-1,0}(f) \varphi_{-1,0}^{\text{per}} + \sum_{j \geq 0} \sum_{k \in \mathbb{Z}} c_{j,k}(f) \varphi_{j,k}^{\text{per}}. \end{aligned}$$

In the special case of the Haar wavelet system we have $\varphi_{j,k}^* = \varphi_{j,k}$, i.e., we are again in the setting of an orthonormal system.

Applying the tensor-product approach (2.1) to the wavelet setting, we analogously define the multi-variate periodic wavelets via

$$\varphi_{\mathbf{j},\mathbf{k}}^{\text{chui,per}}(\mathbf{x}) = \prod_{i=1}^d \varphi_{j_i,k_i}^{\text{chui,per}}(x_i)$$

with multi-indices $\mathbf{j} = (j_i)_{i=1}^d$, $j_i \in \{-1, 0, 2, \dots\}$, and $\mathbf{k} \in \mathcal{K}_{\mathbf{j}}$, where

$$\mathcal{K}_{\mathbf{j}} = \times_{i=1}^d \begin{cases} \{0\}, & j_i = -1 \\ \{0, 1, \dots, 2^{j_i} - 1\}, & j_i \geq 0. \end{cases}$$

In order to truncate the series appropriately we define the finite index set

$$\mathcal{J}_{\mathbf{N}} = \left\{ \mathbf{j} \in \mathbb{Z}^d : \mathbf{j} \geq -\mathbf{1}, \sum_{i,j_i \geq 0} \frac{j_i}{N_i} \leq 1 \right\} \quad (2.6)$$

with multi-degree $\mathbf{N} = (N_i)_{i=1}^d \in \mathbb{N}^d$ and introduce the partial sum

$$S(\mathcal{X}, \mathcal{J}_{\mathbf{N}})f(\mathbf{x}) := \sum_{\mathbf{j} \in \mathcal{J}_{\mathbf{N}}} \sum_{\mathbf{k} \in \mathcal{K}_{\mathbf{j}}} \hat{f}_{\mathbf{j},\mathbf{k}} \varphi_{\mathbf{j},\mathbf{k}}^{\text{chui,per}}(\mathbf{x}), \quad (2.7)$$

where, again, the unknown coefficients $\hat{\mathbf{f}} = (\hat{f}_{\mathbf{j},\mathbf{k}})_{\substack{\mathbf{j} \in \mathcal{J}_{\mathbf{N}} \\ \mathbf{k} \in \mathcal{K}_{\mathbf{j}}}}$ have to be learned via (1.5) based on the set of feature vectors \mathcal{X} in order to give $S(\mathcal{X}, \mathcal{J}_{\mathbf{N}})f(\mathbf{x}_j) = \text{sign}(f(\mathbf{x}_j))$ for a presumed and unknown function f .

From (2.6) we see that we do not use the full tensor product in the multi-variate setting, but rather a hyperbolic subset of it, see Figure 2.1, motivated by the theoretical results presented in [24]. In the special case that the same bandwidth N is applied in all dimensions, i.e., $\mathbf{N} = (N, N, \dots, N)^\top \in \mathbb{N}^d$, the last condition in (2.6) becomes

$$\sum_{i,j_i \geq 0} j_i \leq N.$$




Figure 2.1: Illustration, of 2-dimensional indices $\mathbf{k} = (k_1, k_2)$, represented by the gray squares, where $\mathbf{j} = (j_1, j_2) \in \mathcal{J}_{(3,3)}$ and $\mathbf{k} \in \mathcal{K}_{\mathbf{j}}$ for each \mathbf{j} . This graphic is taken from [24].

Building up the resulting feature map via (1.3) and defining the resulting feature matrix gives

$$\Phi^{\text{chui}}(\mathcal{X}, \mathcal{J}_{\mathbf{N}}) := \left(\varphi_{\mathbf{j}, \mathbf{k}}^{\text{chui, per}}(\mathbf{x}) \right)_{\mathbf{x} \in \mathcal{X}, \substack{\mathbf{j} \in \mathcal{J}_{\mathbf{N}} \\ \mathbf{k} \in \mathcal{K}_{\mathbf{j}}}} \in \mathbb{R}^{M \times \sum_{\mathbf{j} \in \mathcal{J}_{\mathbf{N}}} |\mathcal{K}_{\mathbf{j}}|}. \quad (2.8)$$

Note that the wavelets are localized functions and in case of the Haar system even compactly supported, which means that the matrix defined above is sparse. Thus, one can compute matrix vector products involving this matrix in an efficient manner, as long as the dimension d is small.

2.2. ANOVA decomposition and effective dimensions

The main idea of the classical ANOVA decomposition of functions is to decompose a d -variate function in 2^d so-called ANOVA terms $f_{\mathbf{u}}$, where each term represents a subset of coordinate indices $\mathbf{u} \subseteq [d] := \{1, 2, \dots, d\}$. Each single term depends only on the variables in the corresponding subset and the number of these variables $|\mathbf{u}|$ is called the order of the ANOVA term. At this point we would like to refer to the references, regarding the classical ANOVA decomposition and its use for the function approximation, that were already mentioned at the beginning of this chapter. In order to introduce the ANOVA decomposition of a function $f \in L_2(\mathbb{D}^d)$ we define the projection operators

$$P_{\mathbf{u}}f(\mathbf{x}) = \frac{1}{\tau^{d-|\mathbf{u}|}} \int_{\mathbb{D}^{d-|\mathbf{u}|}} f(\mathbf{x}) d\mathbf{x}_{\mathbf{u}^c}$$

for a subset of coordinate indices \mathbf{u} and its complementary set $\mathbf{u}^c := [d] \setminus \mathbf{u}$. Then, the ANOVA terms can be defined recursively via

$$f_{\mathbf{u}}(\mathbf{x}) := P_{\mathbf{u}}f(\mathbf{x}) - \sum_{\mathbf{v} \subsetneq \mathbf{u}} f_{\mathbf{v}}(\mathbf{x}), \quad (2.9)$$

where $\mathbf{x}_{\mathbf{u}} := (x_i)_{i \in \mathbf{u}} \in \mathbb{D}^{|\mathbf{u}|}$ denotes the restriction of \mathbf{x} to the variables present in the subset \mathbf{u} . In case of the cosine basis we obtain [31]

$$f_{\mathbf{u}}(\mathbf{x}) = \sum_{\substack{\mathbf{k} \in \mathbb{N}_0^d \\ \text{supp } \mathbf{k} = \mathbf{u}}} c_{\mathbf{k}}(f) \varphi_{\mathbf{k}}^{\text{cos}}(\mathbf{x}),$$

where $\text{supp } \mathbf{k} := \{s \in [d] : k_s \neq 0\}$ is the set of non-zero dimensions in a multi frequency \mathbf{k} . Thus, the ANOVA decomposition introduces a disjoint decomposition of the set of basis coefficients or rather frequencies, i.e., the ANOVA terms are orthogonal to each other with respect to the L_2 inner product. An ANOVA term $f_{\mathbf{u}}$ includes exactly those frequencies \mathbf{k} that are non-zero in the dimensions present in \mathbf{u} and equal to zero otherwise.

In a very similar fashion, we conclude for the wavelet basis [24]

$$f_{\mathbf{u}}(\mathbf{x}) = \sum_{\substack{\mathbf{j} \in \mathbb{Z}^d \\ \mathbf{j}_{\mathbf{u}} > -\mathbf{1}, \mathbf{j}_{\mathbf{u}^c} = -\mathbf{1}}} \sum_{\mathbf{k} \in \mathcal{K}_{\mathbf{j}}} c_{\mathbf{j}, \mathbf{k}}(f) \varphi_{\mathbf{j}, \mathbf{k}}^{\text{chui, per}}(\mathbf{x}).$$

The function f can then be uniquely decomposed as

$$f(\mathbf{x}) = f_{\emptyset} + \sum_{i=1}^d f_{\{i\}}(x_i) + \sum_{i=1}^{d-1} \sum_{j=i+1}^d f_{\{i,j\}}(\mathbf{x}_{\{i,j\}}) + \cdots + f_{[d]}(\mathbf{x}) = \sum_{\mathbf{u} \subseteq [d]} f_{\mathbf{u}}(\mathbf{x}) \quad (2.10)$$

into $|\mathcal{P}([d])| = 2^d$ ANOVA terms. The relation to the series expansion (2.9) have first been proven in [31]. But we still have the curse of dimensionality in (2.10), since the number of all ANOVA terms grows exponentially with the dimension d and we need to reduce this in an appropriate manner. To this end, we make use of truncated ANOVA decomposition by taking into account only variable interactions involving only a small number of dimensions. This is motivated by the well-known sparsity of effects stating that many phenomena are well described by few low-dimensional interactions, see [10, 23, 15, 32].

As an example, for small $d_s \in [d]$, $d_s \leq 3$, we could only make use of subsets \mathbf{u} that contain d_s or less variables, i.e., the restricted subset of ANOVA terms is

$$U_{d_s} := \{\mathbf{u} \subseteq [d] : |\mathbf{u}| \leq d_s\} \quad (2.11)$$

where d_s is called superposition threshold or rather the superposition dimension. The cardinality of the set U_{d_s} is given by $1 + \binom{d}{1} + \cdots + \binom{d}{d_s}$ and, thus, only grows polynomially in d for fixed superposition dimension d_s .

Moreover, using the ANOVA decomposition we are able to obtain information about the importance of each term $f_{\mathbf{u}}(\mathbf{x})$ with respect to the function f . In (2.9) we can see that the ANOVA terms have a nice interpretation in terms of the basis coefficients. The relative importance with respect to f can be measured in terms of global sensitivity indices (GSIs) or rather Sobol indices [39, 38], which measure the contribution of the single subsets \mathbf{u} to the overall variance of the function f . Combining (2.9) and Parseval's identity gives

$$\varrho(\mathbf{u}, f) := \begin{cases} \frac{\sum_{\text{supp}(\mathbf{k})=\mathbf{u}} |c_{\mathbf{k}}|^2}{\sum_{\mathbf{k} \neq \mathbf{0}} |c_{\mathbf{k}}|^2} & : \text{cosine basis,} \\ \frac{\sum_{\mathbf{j}_{\mathbf{u}} > -\mathbf{1}, \mathbf{j}_{\mathbf{u}^c} = -\mathbf{1}} \sum_{\mathbf{k} \in \mathcal{K}_{\mathbf{j}}} |c_{\mathbf{j}, \mathbf{k}}|^2}{\sum_{\mathbf{j} \neq -\mathbf{1}} \sum_{\mathbf{k} \in \mathcal{K}_{\mathbf{j}}} |c_{\mathbf{j}, \mathbf{k}}|^2} & : \text{wavelet basis,} \end{cases} \quad (2.12)$$

for which we easily conclude $\varrho(\mathbf{u}, f) \in [0, 1]$. If $\varrho(\mathbf{u}, f) \in [0, 1]$ is large then the ANOVA term $f_{\mathbf{u}}$ significantly contributes to the overall variance of the function f , i.e., making changes in that dimensions might cause a significant change of the function value. In contrast, variables or subsets \mathbf{u} might be negligible when $\varrho(\mathbf{u}, f) \approx 0$. This makes especially sense if one searches for functions for regression or classification with small variance.

For a computed classifying function of the form (2.4) or rather (2.7), depending on given training data set $\{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_M, y_M)\}$, we use the fitted coefficients $\hat{f}_{\mathbf{k}}$ or rather $\hat{f}_{j,\mathbf{k}}$ to compute the corresponding global sensitivity indices via (2.12). Furthermore, we are able to reduce the number of ANOVA terms such that we select only indices for more valuable ones. Based on the computed GSIs, we define an active set of ANOVA terms

$$U^\varepsilon := \emptyset \cup \{\mathbf{u} \in U : \varrho(\mathbf{u}, S(\mathcal{X}, \mathcal{I}_{\mathbf{N}})) > \varepsilon\} \subseteq U \quad (2.13)$$

or rather

$$U^\varepsilon := \emptyset \cup \{\mathbf{u} \in U : \varrho(\mathbf{u}, S(\mathcal{X}, \mathcal{J}_{\mathbf{N}})) > \varepsilon\} \subseteq U, \quad (2.14)$$

for a given tolerance $\varepsilon > 0$, in order to identify ANOVA terms with a small contribution to the overall variance of the computed classifying function. Having identified a suitable active set U^ε , we may re-build the feature map (1.3) and solve (1.5) again by only using the necessary groups of variables, where we probably now have a smaller number of coefficients to be calculated. Moreover, this will be hopefully a beneficial way to improve our prediction of the unknown classifying function. The same procedure has already been successfully applied in case of regression problems, see [32]. In order to be able to implement these theoretical considerations in an efficient manner, we introduce the concept of Grouped Transformations in the following section.

2.3. Grouped transformations

The main idea of grouped transformations is to omit high-dimensional interactions and consider the sum of low-dimensional transformations, which are efficient, see [2]. Later in our numerical experiments we make use of the *GroupedTransforms.jl* package [3], with which the below described grouped transformations can be computed efficiently. In case of the cosine basis, the corresponding matrix-vector products are computed based on the NFCT algorithm [21, 36]. In case of wavelets, the considered matrices are sparse, allowing for an efficient computation of the matrix-vector product, cf. [24, 3].

For a given subset of ANOVA terms $U \subseteq \mathcal{P}([d])$ we define a grouped index set as a disjoint union

$$\mathcal{I}_{\mathbf{N}}(U) = \bigcup_{\mathbf{u} \in U} \mathcal{I}_{\mathbf{N}^{\mathbf{u}}}^{\mathbf{u},d}$$

where $\mathbf{N} = (\mathbf{N}^{\mathbf{u}})_{\mathbf{u} \in U}$ is a $|U|$ -tuple with entries $\mathbf{N}^{\mathbf{u}} \in (2\mathbb{N})^{|\mathbf{u}|}$ and

$$\mathcal{I}_{\mathbf{N}^{\mathbf{u}}}^{\mathbf{u},d} = \{\mathbf{k} \in \mathbb{N}_0^d : \text{supp } \mathbf{k} = \mathbf{u}, (\mathbf{k}_{\mathbf{u}})_i \in \mathcal{I}_{N_i^{\mathbf{u}}}, i = 1, 2, \dots, |\mathbf{u}|\}$$

for $\mathbf{u} \in U$ are the corresponding frequency index sets. As an example, in a 4-dimensional setting let $\mathbf{u} := \{2, 3\}$ and $\mathbf{N}^{\mathbf{u}} := (6, 24) \in (2\mathbb{N})^2$. Then, the index set $\mathcal{I}_{\mathbf{N}^{\mathbf{u}}}^{\mathbf{u},d}$ contains indices $\mathbf{k} = (0, k_2, k_3, 0)$ with $k_2 \in \{1, 2, \dots, 5\}$ and $k_3 \in \{1, 2, 3\}$, see (2.3).

For the Chui-Wang basis some adjustments are necessary. Analogously to $\mathcal{I}_{\mathbf{N}}(U)$ we define the index sets

$$\mathcal{J}_{\mathbf{N}}(U) = \bigcup_{\mathbf{u} \in U} \mathcal{J}_{\mathbf{N}^{\mathbf{u}}}^{\mathbf{u},d}$$

with tuple $\mathbf{N} = (\mathbf{N}^{\mathbf{u}})_{\mathbf{u} \in U}$, $\mathbf{N}^{\mathbf{u}} \in \mathbb{N}^{|\mathbf{u}|}$ and for $\mathbf{u} \in U$ we get

$$\mathcal{J}_{\mathbf{N}^{\mathbf{u}}}^{\mathbf{u},d} = \left\{ \mathbf{j} \in \mathbb{Z}^d : \mathbf{j}_{\mathbf{u}} > -\mathbf{1}, \mathbf{j}_{\mathbf{u}^c} = -\mathbf{1}, \sum_{i, (j_{\mathbf{u}})_i \geq 0} \frac{(j_{\mathbf{u}})_i}{N_i^{\mathbf{u}}} \leq 1 \right\}. \quad (2.15)$$

In order to truncate the ANOVA decomposition of f , we assume that the function f has low superposition dimension $d_s \ll d$ and define $U_{d_s} := \{\mathbf{u} \subseteq [d] : |\mathbf{u}| \leq d_s\}$ in the following.

Making use of the grouped index sets introduced above, the feature map (2.5) or rather (2.8) has a block structure, i.e.,

$$\Phi^{\cos}(\mathcal{X}, \mathcal{I}_{\mathbf{N}}(U)) = \left(\Phi^{\cos}(\mathcal{X}, \mathcal{I}_{\mathbf{N}^{\mathbf{u}_1}}^{\mathbf{u}_1,d}) \quad \Phi^{\cos}(\mathcal{X}, \mathcal{I}_{\mathbf{N}^{\mathbf{u}_2}}^{\mathbf{u}_2,d}) \quad \dots \quad \Phi^{\cos}(\mathcal{X}, \mathcal{I}_{\mathbf{N}^{\mathbf{u}_{|U|}}}^{\mathbf{u}_{|U|},d}) \right)$$

and

$$\Phi^{\text{chui}}(\mathcal{X}, \mathcal{J}_{\mathbf{N}}(U)) = \left(\Phi^{\text{chui}}(\mathcal{X}, \mathcal{J}_{\mathbf{N}^{\mathbf{u}_1}}^{\mathbf{u}_1,d}) \quad \Phi^{\text{chui}}(\mathcal{X}, \mathcal{J}_{\mathbf{N}^{\mathbf{u}_2}}^{\mathbf{u}_2,d}) \quad \dots \quad \Phi^{\text{chui}}(\mathcal{X}, \mathcal{J}_{\mathbf{N}^{\mathbf{u}_{|U|}}}^{\mathbf{u}_{|U|},d}) \right).$$

In case of the cosine system, for instance, we see that a matrix-vector product $\Phi^{\cos}(\mathcal{X}, \mathcal{I}_{\mathbf{N}}(U))\hat{\mathbf{f}}$ is simply a sum of low-dimensional transformations

$$\Phi^{\cos}(\mathcal{X}, \mathcal{I}_{\mathbf{N}}(U))\hat{\mathbf{f}} = \sum_{\mathbf{u} \subseteq U} \Phi^{\cos}(\mathcal{X}, \mathcal{I}_{\mathbf{N}^{\mathbf{u}}}^{\mathbf{u},d}) \hat{\mathbf{f}}(\mathbf{u}), \quad (2.16)$$

for which the computational effort no longer grows exponentially for growing d and fixed group structure or superposition dimension d_s . By $\hat{\mathbf{f}}(\mathbf{u})$ we mean the vector of basis coefficients restricted to indices \mathbf{k} belonging to the ANOVA term $f_{\mathbf{u}}$. The same also applies to the wavelet case, see [24, Section 4.1].

3. Optimization Algorithms

In this chapter we focus on the problem of estimating the model parameters $\hat{\mathbf{f}}$ for the given training data set $\{(\mathbf{x}_j, y_j)\}_{j=1,2,\dots,M}$, i.e., solving the optimization problem (1.5). In order to obtain a classifier we need to construct some function f , such that $\text{sign}(f(\mathbf{x}_j)) = y_j$ for as many inputs \mathbf{x}_j , $j = 1, 2, \dots, M$, as possible. In this chapter we write Φ for given grouped transform, which are the feature matrices $\Phi^{\cos}(\mathcal{X}, \mathcal{I}_{\mathbf{N}}(U))$ or $\Phi^{\text{chui}}(\mathcal{X}, \mathcal{J}_{\mathbf{N}}(U))$, given in (2.5) and (2.8), where U is a given subset of ANOVA terms and \mathcal{X} is a given set of instances. The related fitted coefficients $\hat{\mathbf{f}}^{\cos}$ and $\hat{\mathbf{f}}^{\text{chui}}$ we denote in general by $\hat{\mathbf{f}}$ with index set \mathcal{I} .

The common way to solve the SVM problem (1.5) is solving the corresponding dual optimization problem based on kernel functions. However, the kernel matrices become too large and badly conditioned for huge amount of data. Well, if we use the approach, what we built in the previous chapter, we are able to solve our SVMs in primal form using grouped transformations as feature maps.

In this chapter we present two algorithms, which can be used for solving the optimization problem (1.5) with ℓ_2 -norm regularization, see Section 3.1, and ℓ_1 -norm regularization in Section 3.2.

3.1. ℓ_2 -norm regularization

First we discuss unconstrained optimization problem of the form

$$P(\hat{\mathbf{f}}) := \lambda \|\hat{\mathbf{f}}\|_2^2 + \frac{1}{M} \sum_{j=1}^M \left(\max \left\{ 0, 1 - y_j \langle \hat{\mathbf{f}}, \Phi(\mathbf{x}_j) \rangle \right\} \right)^2 \rightarrow \min_{\hat{\mathbf{f}}} \quad \hat{\mathbf{f}} \in \mathbb{R}^{|\mathcal{I}|} \quad (3.1)$$

using regularization term $\|\hat{\mathbf{f}}\|_2^2$, which enforces the small variance of the $\Phi \hat{\mathbf{f}}$. To solve this problem numerically, we apply a gradient descent method, which uses the negative gradient direction $d^k := -\nabla P(\hat{\mathbf{f}}^{(k)})$ in each iteration and improves the objective in that direction. In order to compute the gradient of $P(\hat{\mathbf{f}})$ with respect to the vector of basis coefficients $\hat{\mathbf{f}}$, we first introduce the following function $h : \mathbb{R} \rightarrow \mathbb{R}$ and its derivative, given by

$$h(t) := \max\{0, t\}^2, \quad h'(t) = \begin{cases} 2t & : t \geq 0, \\ 0 & : t < 0. \end{cases}$$

Consequently we have

$$\nabla P(\hat{\mathbf{f}}) = 2\lambda \hat{\mathbf{f}} - \frac{1}{M} \sum_{j=1}^M y_j \cdot \Phi(\mathbf{x}_j) \cdot h'(1 - y_j \langle \hat{\mathbf{f}}, \Phi(\mathbf{x}_j) \rangle). \quad (3.2)$$

Thinking in terms of matrix-vector multiplications, the approximation of the objective function (3.1) and its gradient (3.2) is realized as follows. The matrix-vector product $\Phi \hat{\mathbf{f}}$ is a vector of length M containing the vector products $\langle \hat{\mathbf{f}}, \Phi(\mathbf{x}_j) \rangle$. Due to the ANOVA structure of the considered index sets, the matrix-vector product $\Phi \hat{\mathbf{f}}$ can be approximated efficiently in terms of a grouped transform, as explained above. In case of the cosine system, for instance, a sum of NFCTs in small dimensions is computed. Having computed $\Phi \hat{\mathbf{f}}$, we just need pointwise operations in order to evaluate $P(\hat{\mathbf{f}}) = \lambda \|\hat{\mathbf{f}}\|_2^2 + \frac{1}{M} \text{sum}(\max.(\mathbf{0}, \mathbf{1} - \mathbf{y} \cdot *(\Phi \hat{\mathbf{f}})))$, where a dot indicates a pointwise operation.

In order to compute $\nabla P(\hat{\mathbf{f}})$ we also need to multiply with the transposed matrix Φ^\top in an efficient way. This is possible in a very similar fashion. In case of the cosine basis, the resulting algorithm is known as transposed NFCT, see [30, Chapter 7]. Finally, we obtain $\nabla P(\hat{\mathbf{f}}) = 2\lambda \hat{\mathbf{f}} - \frac{1}{M} \text{sum}((\Phi^\top \mathbf{y}) \cdot *h' \cdot (\mathbf{1} - \mathbf{y} \cdot *(\Phi \hat{\mathbf{f}})))$.

We summarize the the gradient descent method in Algorithm 1, where the choice of the step size is realized by the well-known Armijo backtracking strategy [1], see Algorithm 2. Using stopping criteria, which ensures that the gradient $\nabla P(\hat{\mathbf{f}})$ is sufficiently small, we can also stop earlier.

3.2. ℓ_1 -norm regularization

In this section we discuss the optimization method for the following unconstrained convex optimization problem

$$P(\hat{\mathbf{f}}) := \lambda \|\hat{\mathbf{f}}\|_1 + \frac{1}{M} \sum_{j=1}^M \left(\max \left\{ 0, 1 - y_j \langle \hat{\mathbf{f}}, \Phi(\mathbf{x}_j) \rangle \right\} \right)^2 \rightarrow \min_{\hat{\mathbf{f}}} \quad \hat{\mathbf{f}} \in \mathbb{R}^{|\mathcal{I}|}, \quad (3.3)$$

for given $\Phi \in \mathbb{R}^{M \times |\mathcal{I}|}$, which will either be equal to $\Phi^{\text{cos}}(\mathcal{X}, \mathcal{I}_N(U))$, as defined in (2.5), or $\Phi^{\text{chui}}(\mathcal{X}, \mathcal{I}_N(U))$, cf. (2.8). Solving such a problem is more challenging than the one discussed

Algorithm 1 Gradient descent algorithm for the solution of (3.1)

Input: set of instance-label pairs $\{(\mathbf{x}_j, y_j) : j = 1, 2, \dots, M\} \in \mathcal{X} \times \{-1, +1\}$, $\mathcal{X} \subseteq \mathbb{D}^d$, grouped transform
initial guess $\hat{\mathbf{f}}^{(0)}$, regularization parameter $\lambda > 0$, maximal number of iterations $K \in \mathbb{N}$

1: **for** $k = 1, 2, \dots, K$ **do**
2: $d^k = -\nabla P(\hat{\mathbf{f}}^{(k)})$
3: **if** $\|\nabla P(\hat{\mathbf{f}}^{(k)})\|_2 < 10^{-8}$ **then**
4: **break**
5: **end if**
6: compute a step size s_k via Algorithm 2
7: $\hat{\mathbf{f}}^{(k+1)} \leftarrow \hat{\mathbf{f}}^{(k)} + s_k d^k$
8: **end for**

Output: $\hat{\mathbf{f}}^{(K)}$ minimizer of (3.1)

Algorithm 2 Armijo backtracking rule

Input: current iterate $\hat{\mathbf{f}}^{(k)}$, search direction d^k , initial step size $s_k^{(0)}$, $\sigma \in (0, 1)$, $\xi \in (0, 1)$,
precomputed function value $P(\hat{\mathbf{f}}^{(k)})$ and gradient $\nabla P(\hat{\mathbf{f}}^{(k)})$

1: set $n := 0$
2: **while** $P(\hat{\mathbf{f}}^{(k)} + s_k^{(n)} d^k) > P(\hat{\mathbf{f}}^{(k)}) + \sigma s_k^{(n)} \|\nabla P(\hat{\mathbf{f}}^{(k)})\|_2^2$ **do**
3: $s_k^{(n+1)} \leftarrow \xi s_k^{(n)}$
4: $n \leftarrow n + 1$
5: **end while**

Output: $s_k^{(n)}$ step size

in the previous section, because $\|\hat{\mathbf{f}}\|_1$ is non-smooth, whereas the second part

$$q(\hat{\mathbf{f}}) := \frac{1}{M} \sum_{j=1}^M \left(\max \left\{ 0, 1 - y_j \langle \hat{\mathbf{f}}, \Phi(\mathbf{x}_j) \rangle \right\} \right)^2$$

has a Lipschitz continuous gradient with respect to $\hat{\mathbf{f}}$. In order to solve (3.3) we will make use of the well-known fast iterative shrinkage thresholding algorithm (FISTA), see [5], which became popular during the past decades. In addition, we would like to mention that FISTA has also been used in [2] for high-dimensional explainable ANOVA approximation, where regression problems have been considered using a group Lasso regularization term.

Let us briefly review the basic approximation model, which we are going to use. For any $L > 0$, consider the following quadratic approximation of $P(\hat{\mathbf{f}})$ at a given point $\hat{\mathbf{h}}$

$$Q_L(\hat{\mathbf{f}}, \hat{\mathbf{h}}) := q(\hat{\mathbf{h}}) + \langle \hat{\mathbf{f}} - \hat{\mathbf{h}}, \nabla q(\hat{\mathbf{h}}) \rangle + \frac{L}{2} \|\hat{\mathbf{f}} - \hat{\mathbf{h}}\|_2^2 + \lambda \|\hat{\mathbf{f}}\|_1$$

which admits a unique minimizer

$$p_L(\hat{\mathbf{h}}) = \arg \min_{\hat{\mathbf{f}}} \left\{ \lambda \|\hat{\mathbf{f}}\|_1 + \frac{L}{2} \left\| \hat{\mathbf{f}} - \underbrace{\left(\hat{\mathbf{h}} - \frac{1}{L} \nabla q(\hat{\mathbf{h}}) \right)}_{=: \mathbf{a}} \right\|_2^2 \right\}, \quad (3.4)$$

where L plays the role of a step size. Using the well-known definition of proximal operator, see [4, Chapter 6], [29], given by

$$\text{prox}_{\lambda\|\cdot\|_1}(\mathbf{a}) = \arg \min_{\hat{\mathbf{f}}} \left(\lambda \|\hat{\mathbf{f}}\|_1 + \frac{1}{2} \|\hat{\mathbf{f}} - \mathbf{a}\|_2^2 \right),$$

we get for $i = 1, 2, \dots, |\hat{\mathbf{f}}|$ following solution of (3.4)

$$p_L(\hat{h}_i) = \text{prox}_{\lambda\|\cdot\|_1}(a_i) = \text{sign}(a_i) \max(0, |a_i| - \lambda). \quad (3.5)$$

We summarize the FISTA algorithm with backtracking step size rule in Algorithm 3.

Algorithm 3 FISTA with backtracking for the solution of (3.3)

Input: set of instance-label pairs $\{(\mathbf{x}_j, y_j) : j = 1, 2, \dots, M\} \in \mathcal{X} \times \{-1, +1\}$, $\mathcal{X} \subseteq \mathbb{D}^d$, grouped transform initial guess $\hat{\mathbf{f}}^{(0)}$, regularization parameter $\lambda > 0$, maximal number of iterations $K \in \mathbb{N}$, $L_0 > 0$ a

```

1:  $\hat{\mathbf{h}}^{(1)} \leftarrow \hat{\mathbf{f}}^{(0)}, t_1 \leftarrow 1$ 
2: for  $k = 1, 2, \dots, K$  do
3:    $L_k \leftarrow L_{k-1}$ 
4:    $\hat{\mathbf{f}}^{(k)} \leftarrow p_{L_k}(\hat{\mathbf{h}}^{(k)})$ 
5:   if  $|P(\hat{\mathbf{f}}^{(k)}) - P(\hat{\mathbf{f}}^{(k-1)})| < 10^{-8}$  or  $\|\hat{\mathbf{f}}^{(k-1)} - \hat{\mathbf{f}}^{(k)}\|_2 < 1F^{-8}$  then
6:     break
7:   end if
8:   while  $q(\hat{\mathbf{f}}^{(k)}) < q(\hat{\mathbf{h}}^{(k)}) + \langle \hat{\mathbf{f}}^{(k)} - \hat{\mathbf{h}}^{(k)}, \nabla q(\hat{\mathbf{h}}^{(k)}) \rangle + \frac{L_k}{2} \|\hat{\mathbf{f}}^{(k)} - \hat{\mathbf{h}}^{(k)}\|_2^2$  do
9:      $L_k \leftarrow \theta L_k$ 
10:     $\hat{\mathbf{f}}^{(k)} \leftarrow p_{L_k}(\hat{\mathbf{h}}^{(k)})$ 
11:   end while
12:    $t_{k+1} \leftarrow (1 + \sqrt{1 + 4t_k^2})/2$ 
13:    $\hat{\mathbf{h}}^{(k+1)} \leftarrow \hat{\mathbf{f}}^{(k)} + (t_{k+1})(\hat{\mathbf{f}}^{(k)} - \hat{\mathbf{f}}^{(k-1)})$ 
14: end for
Output:  $\hat{\mathbf{f}}^{(K)}$  minimizer of (3.3)

```

4. Numerical Experiments

In this chapter we present results of some numerical experiments. We consider one-dimensional examples, see Section 4.1, as well as examples in higher dimensions in Section 4.2 and Section 4.3. The available set of instances $\mathcal{X} \subset \mathbb{R}^d$ is consisting of M randomly chosen and uniquely distributed feature vectors $\mathbf{x}_j \in [0, 1/2]^d$ or rather $\mathbf{x}_j \in [-1/2, 1/2]^d$, $j = 1, 2, \dots, M$, depending on the applied system of basis functions.

Using a known test function f we obtain the corresponding labels $y_j \in \{-1, +1\}$ by computing $y_j = \text{sign}f(\mathbf{x}_j)$. Our goal is not to recover function f , but rather to construct a classifying function that matches the given training data points (\mathbf{x}_j, y_j) , $j = 1, 2, \dots, M$. This means that we determine basis coefficients $\hat{\mathbf{f}}$ on a presumed grouped index set \mathcal{I} with presumed active set U , as explained before. In order to measure the performance of our prediction on new data

points from test data set, we define the classification accuracy (CA) by

$$\text{CA} = \frac{\text{number of correct predictions}}{\text{total number of predictions}} \cdot 100\%.$$

In our numerical examples we build the test function in such a way that it perfectly fits to our feature map, in order to ensure that a perfectly classifying functions exists and should be found by our algorithm. This is a fairly simple test scenario but certainly the first to take in order to verify the functionality of our method.

For more general functions or real-world data sets, we oftentimes do not know anything about the active set of dimensions or the bandwidths that will be required. However, the testing of more complex test functions and the application of our algorithms to real world data will be subject of future work.

We implemented the presented optimization algorithms in the programming language Julia, where we make use of the publicly available *ANOVAapprox.jl*, [35], and *GroupedTransforms.jl* Julia packages, [3].

4.1. One-dimensional example

The one-dimensional example is theoretically very simple, but it offers a good possibility of visualization. We generate our training data set by using a zero-mean function f defined via

$$f^{\text{cos}}(x) := \sum_{k=1}^5 c_k \varphi_k^{\text{cos}}(x), \quad c_k = \frac{k+1}{4}$$

or rather

$$f^{\text{chui}}(x) := \sum_{j=0}^2 \sum_{k \in \mathcal{K}_j} c_{j,k} \varphi_{j,k}^{\text{chui,per}}(x), \quad c_{j,k} = \frac{k+1}{4}.$$

As we can see, we construct the test function in such a way that we have $|\mathcal{I}_{(6)}(\{1\})| = 5$ non-zero basis coefficients in case of the cosine basis, and $\sum_{j \in \mathcal{J}_{(2)}(\{1\})} |\mathcal{K}_j| = 7$ non-zero wavelet coefficients, otherwise. For the entire experiment, we set $M = 50$ and use 100 randomly chosen instances and their corresponding labels for testing. We determine the coefficients by solving the optimization problems (3.1) or (3.3), depending on the applied regularization.

As the presumed active set we use $U = \{\emptyset, \{1\}\}$, which is not surprising in a one-dimensional setting, and exactly incorporate those basis functions in the feature map, which have non-zero coefficients in the above defined test functions. In the following we denote by \hat{f}^{cos} and \hat{f}^{chui} the fitted basis coefficients of the cosine and the wavelet basis functions, respectively.

As regularization parameter we use, without further considerations, $\lambda = 0.01$. Note that we consider different regularization parameters in the multi-variate example later on. In Figure 4.1 we plot the obtained classifying functions for the ℓ_2 -norm regularization approach (orange) and when applying ℓ_1 -norm regularization (blue). In addition, we plot the sign of the test functions in black and the randomly generated training data by the little black dots. In order to obtain a more representative result, we performed 100 runs in total, resulting in an average test accuracy of 94.93% for the approach with cosine basis with ℓ_1 -norm regularization and 94.87% if we apply ℓ_2 -norm regularization. In case of using wavelets we achieve an average accuracy of 100% for both regularization approaches.

Note that the zeros of the classifying wavelet function have fixed positions, making it fairly easy to achieve this good result. In addition, our model fits the test function perfectly. Although the perfect ansatz is also used in case of the cosine system, we face some problems if training data points are missing close to the zero points of the test function. Then, test data points close to the jumps may be misclassified. So, even in such a constructed test scenario, it is not easy to achieve an accuracy of 100%.

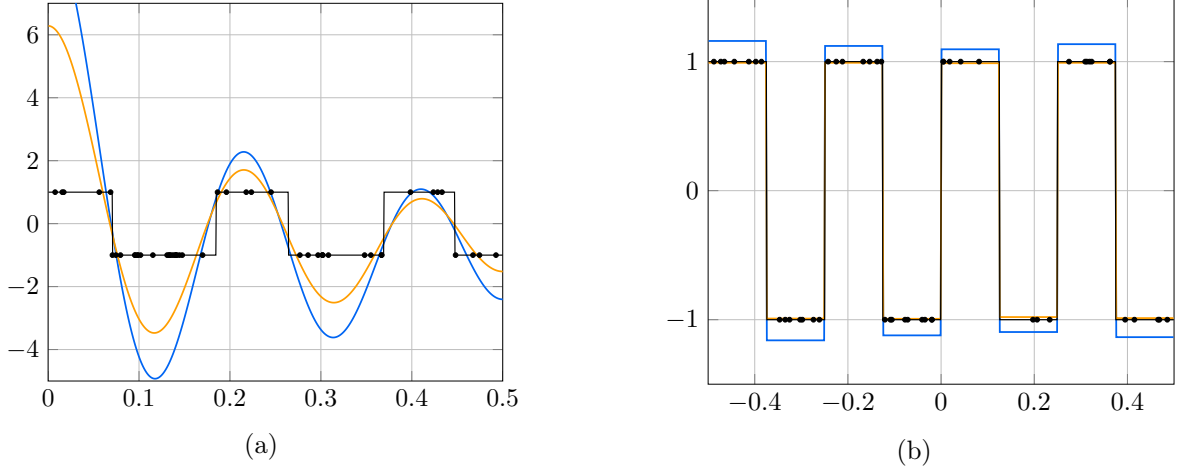


Figure 4.1: The sign of the test functions and the training data points (x_j, y_j) , $j = 1, 2, \dots, M$ are visualized in black. The obtained classifying functions are depicted in orange (ℓ_2 -norm regularization) and in blue (ℓ_1 -norm regularization). The regularization parameter is set to $\lambda = 0.01$. The classifying function $S(\mathcal{X}, \mathcal{I}_N)f(\mathbf{x})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\text{cos}}$ is visualized in (a) and $S(\mathcal{X}, \mathcal{J}_N)f(\mathbf{x})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\text{chui}}$ in (b).

4.2. Six-dimensional example

In this example we consider test functions $f : \mathbb{R}^6 \rightarrow \mathbb{R}$. Similar to the previous section, we construct these functions as linear combinations of cosine or wavelet basis functions, respectively. As the true active set we use $U^* := (\{4\}, \{6\}, \{2, 3\})$ and define

$$f^{\text{cos}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathcal{I}_{((6),(4),(2,4))}(U^*)} 1 \cdot \varphi_{\mathbf{k}}^{\text{cos}}(\mathbf{x})$$

and

$$f^{\text{chui}}(\mathbf{x}) = \sum_{\mathbf{j} \in \mathcal{J}_{((3),(3),(3,3))}(U^*)} \sum_{\mathbf{k} \in \mathcal{K}_{\mathbf{j}}} 1 \cdot \varphi_{\mathbf{j}, \mathbf{k}}^{\text{chui, per}}(\mathbf{x})$$

depending on the considered system of basis functions. With the applied band widths we have 19 non-zero basis coefficients in case of the cosine system and 71 non-zero wavelet coefficients. We generate our training data points (\mathbf{x}_j, y_j) , $j = 1, 2, \dots, M$ using test functions f^{cos} or rather f^{chui} .

As the presumed active set we use U_{d_s} with $d_s = 2$, i.e., we incorporate all subsets of cardinality ≤ 2 into our feature map. The feature map or rather the corresponding data

matrices are built by using the following band widths $\mathbf{N} = (\mathbf{N}^{\mathbf{u}})_{\mathbf{u} \in U}$ with $\mathbf{N}^{\mathbf{u}} \in \mathbb{N}^{|\mathbf{u}|}$,

$$\hat{\Phi}^{\text{cos}} := \Phi^{\text{cos}}(\mathcal{X}, \mathcal{I}_{\mathbf{N}}(U_2)), \quad \mathbf{N}^{\mathbf{u}} = \begin{cases} (6) & : |\mathbf{u}| = 1, \\ (4, 4) & : |\mathbf{u}| = 2 \end{cases} \quad (4.1)$$

and

$$\hat{\Phi}^{\text{chui}} := \Phi^{\text{chui}}(\mathcal{X}, \mathcal{J}_{\mathbf{N}}(U_2)), \quad \mathbf{N}^{\mathbf{u}} = \begin{cases} (3) & : |\mathbf{u}| = 1, \\ (3, 3) & : |\mathbf{u}| = 2. \end{cases} \quad (4.2)$$

So, we incorporate all multi-indices that are non-zero in the constructed test functions, but we also make use of a lot of basis functions that do not have any influence. In case of the cosine system we have $\hat{\mathbf{f}}^{\text{cos}} \in \mathbb{R}^{166}$, whereas only 19 are present in f^{cos} . For the wavelet system we have $\hat{\mathbf{f}}^{\text{chui}} \in \mathbb{R}^{876}$ and only 71 active coefficients in f^{chui} .

Note, that we need much more data points as in the previous example, since we are now in a six-dimensional setting. We use $M = 1000$ in case of the cosine basis and $M = 5000$ for the wavelet basis approach. The same number of randomly generated points are used as the test data set. Please note that we cannot directly compare both test cases, since the test functions are completely different and have different numbers of non-zero basis coefficients. Since the number of non-zero coefficients is roughly 5 times larger in the wavelet setting, we choose a 5 times larger training data set for that setting.

The choice of the regularization parameter is not negligible and, thus, we solve the considered optimization problems with different regularization parameters λ . In Figure 4.2 we plot the achieved accuracy on the test data set with respect to the applied λ . Again, all accuracies have to be understood as average values achieved over 100 runs. Using the ℓ_1 -norm regularization approach as an example, we see that there is an optimal value for λ . For larger or smaller values the classification rate becomes worse.

Although we used a lot of basis functions for fitting the model that have not been present in the original test function, we already achieve a good classification rate.

Now, of course, the question remains as to whether we can reconstruct the true active set with the help of the global sensitivity indices of the obtained classifying functions. To this end, we compute the global sensitivity indices of the classifying functions, which are plotted in Figure 4.3. It can be seen that the global sensitivity indices are significantly larger than zero only for subsets $\mathbf{u} \in U^*$, which is a very nice results. Almost all the other subsets have global sensitivity indices smaller or equal to 10^{-2} in case of ℓ_2 -norm regularization and even 10^{-6} when ℓ_1 -norm regularization is applied. In summary, with both optimization approaches we were able to detect the original active set U^* . The result is even clearer in the case of ℓ_1 -norm regularization.

In addition, we applied the threshold parameter $\varepsilon = 0.01$ in order to compute the corresponding active set $U^{0.01}$, see (2.13) or (2.14), and fitted the classifying functions again, only using this active set, applying the same bandwidths as defined above. By that we are able to improve the classification performance further, see Figure 4.4.




Figure 4.2: Average of mean CA over 100 runs using ℓ_1 -norm regularization (blue) and ℓ_2 -norm regularization (orange). In (a) we visualize results achieved by classifying functions $S(\mathcal{X}, \mathcal{I}_{\mathcal{N}}(U_2))f^{\text{cos}}(\mathbf{x})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\text{cos}}$, where the number of generated training and test data points have been set to 1000. In (b), we see results using classifying functions $S(\mathcal{X}, \mathcal{J}_{\mathcal{N}}(U_2))f(\mathbf{x})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\text{chui}}$ and 5000 generated training and test data points.

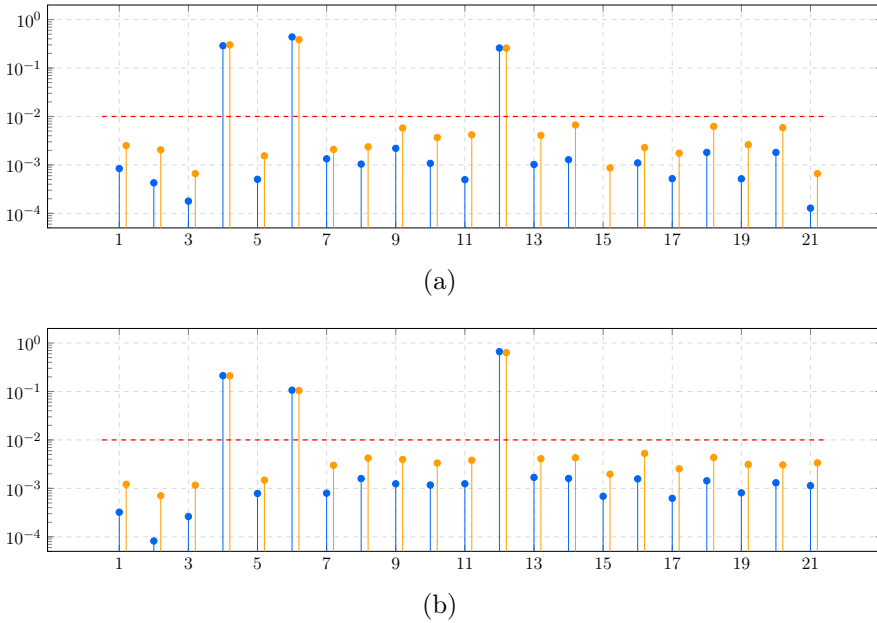


Figure 4.3: Global sensitivity indices $\rho(\mathbf{u}, S(\mathcal{X}, \mathcal{I}_{\mathcal{N}}(U_2))f^{\text{cos}})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\text{cos}}$ are visualized in (a) and $\rho(\mathbf{u}, S(\mathcal{X}, \mathcal{J}_{\mathcal{N}}(U_2))f^{\text{chui}})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\text{chui}}$ in (b). The results using ℓ_1 -norm regularization are shown in blue and for ℓ_2 -norm regularization in orange for $\mathbf{u} \in U_2$, where $|U_2| = 21$.




Figure 4.4: Average of mean CA over 100 runs using ℓ_1 -norm regularization (blue) and ℓ_2 -norm regularization (orange). In (a) we visualize the results achieved by classifying functions $S(\mathcal{X}, \mathcal{I}_N(U^{0.01}))f^{\cos}(\mathbf{x})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\cos}$, where the number of generated training and test data points have been set to 1000. In (b), we see the results obtained with classifying functions $S(\mathcal{X}, \mathcal{J}_N(U^{0.01}))f^{\text{chui}}(\mathbf{x})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\text{chui}}$ and 5000 generated training and test data points.

4.3. Ten-dimensional example

For our third numerical example we introduce one of the so-called Friedman functions [18], which were used as benchmark example in [26] and have since become an often used example in approximation of functions with scattered data, see [7, 8]. Moreover, in [31] the authors achieved very good results regarding the reconstruction of Friedman functions. We focus on the Friedman 1 function

$$F^1 : [0, 1]^{10} \rightarrow \mathbb{R}, F^1(\mathbf{x}) = 10 \sin(\pi x_1 x_2) + 20 \left(x_3 - \frac{1}{2}\right)^2 + 10x_4 + 5x_5, \quad (4.3)$$

where only five of the ten variables have any influence on the function. Further, the mean value of the function is given by

$$M(F^1) = \int_{[0,1]^d} F^1(\mathbf{x}) d\mathbf{x} = 10 \int_0^1 \frac{1 - \cos(\pi x_1)}{\pi x_1} dx_1 + \frac{55}{6} \approx 14.4133, \quad (4.4)$$

where the included integral can be expressed in terms of the well-known incomplete Gamma function and has to be computed numerically. For our synthetic generated data set $\{(\mathbf{x}_j, y_j) : j = 1, 2, \dots, M\}$ we generate M random uniformly distributed feature vectors $\mathbf{x}_j \in [0, 1/2]^{10}$ or rather $\mathbf{x}_j \in [-1/2, 1/2]^{10}$ and obtain the corresponding labels, similar to the previous examples, by $y_j = \text{sign}(f^{\cos}(\mathbf{x}_j))$ or rather $y_j = \text{sign}(f^{\text{chui}}(\mathbf{x}_j))$, where we define the zero-mean functions

$$f^{\cos} : [0, \frac{1}{2}]^{10} \rightarrow \mathbb{R}, \quad f^{\cos}(\mathbf{x}) = F^1(2\mathbf{x}) - M(F^1)$$

and

$$f^{\text{chui}} : [-\frac{1}{2}, \frac{1}{2}]^{10} \rightarrow \mathbb{R}, \quad f^{\text{chui}}(\mathbf{x}) = F^1(\mathbf{x} + \frac{1}{2}) - M(F^1),$$

for the two considered bases.

Note that the ANOVA terms of the Friedman 1 function can be stated explicitly, see Appendix A.1, for an active set $U^* = \{\emptyset, \{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{1, 2\}\}$, where only one two-dimensional term is relevant. Furthermore, we numerically verified that all present ANOVA terms significantly contribute to $\text{sign}(F^1(\mathbf{x}) - M(F^1))$, as presented in Appendix A.2. Our goal will be to determine the true active set by analyzing the global sensitivity indices of the classifying function, similarly to the previous examples.

We begin by setting the superposition threshold to $d_s = 2$. Since the bandwidth $\mathbf{N} = (\mathbf{N}^{\mathbf{u}})_{\mathbf{u} \in U}$ with $\mathbf{N}^{\mathbf{u}} \in \mathbb{N}^{|\mathbf{u}|}$ is unknown in this setting, we consider different bandwidths in order to generate the feature map. To this end, we use the following notation in order to define the corresponding feature map and feature matrices.

$$\Phi^{\text{cos}} := \Phi^{\text{cos}}(\mathcal{X}, \mathcal{I}_{\mathbf{N}}(U_2)), \quad \mathbf{N}^{\mathbf{u}} = \begin{cases} (N_1) & : |\mathbf{u}| = 1, \\ (N_2, N_2) & : |\mathbf{u}| = 2 \end{cases} \quad (4.5)$$

and

$$\Phi^{\text{chui}} := \Phi^{\text{chui}}(\mathcal{X}, \mathcal{J}_{\mathbf{N}}(U_2)), \quad \mathbf{N}^{\mathbf{u}} = \begin{cases} (N_1) & : |\mathbf{u}| = 1, \\ (N_2, N_2) & : |\mathbf{u}| = 2. \end{cases} \quad (4.6)$$

Let us first analyze the results obtained by using the cosine basis, visualized in Figure 4.5. We generate $M = 1000$ training data points (\mathbf{x}_j, y_j) , $j = 1, 2, \dots, M$ and solve the considered optimization problems with different regularization parameters λ . Depending on the applied bandwidths \mathbf{N} from (4.5) with $N_1 = 6$ and $N_2 \in \{2, 4, 6\}$, we have $\hat{\mathbf{f}}^{\text{cos}} \in \mathbb{R}^{96}$ if $N_2 = 2$, $\hat{\mathbf{f}}^{\text{cos}} \in \mathbb{R}^{456}$ if $N_2 = 4$ and $\hat{\mathbf{f}}^{\text{cos}} \in \mathbb{R}^{1176}$ for $N_2 = 6$. Afterwards we compute the classification accuracy using the corresponding classification functions $S(\mathcal{X}, \mathcal{I}_{\mathbf{N}}(U_2))f^{\text{cos}}(\mathbf{x})$ and the generated test data set, containing M test data points. All accuracies have to be understood as average values achieved over 10 runs.

As we can see, we get different CA performance depending on the regularization. We achieve the highest mean CA with $N_1 = 6$, $N_2 = 4$ and $\lambda = 2^{-4}$, in case of ℓ_1 -norm regularization. For the best results provided by ℓ_2 -norm regularization we need to choose $N_1 = 6$, $N_2 = 2$ and $\lambda = 2^{-7}$.

In Figure 4.6 we plot the resulting global sensitivity indices, which clearly show that the variables x_6 to x_{10} are significantly less important than the others. By applying the threshold parameter $\varepsilon = 0.01$ we are able to recover the actual active set

$$U^{0.01} = \{\emptyset, \{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{1, 2\}\} = U^*.$$

Next, let us consider the wavelet basis approach and test if we are able to get the same result. In this setting, we generate $M = 5000$ random feature vectors $\mathbf{x}_j \in [-1/2, 1/2]^{10}$ and compute the corresponding labels by $y_j = \text{sign}(f^{\text{chui}})$ to train our prediction models. Similarly to the cosine basis approach, we solve the considered optimization problems with different regularization parameters λ . Depending on different bandwidths \mathbf{N} from (4.6) with $N_1 = 3$ and $N_2 \in \{1, 2, 3\}$, we have $\hat{\mathbf{f}}^{\text{chui}} \in \mathbb{R}^{376}$ if $N_2 = 1$, $\hat{\mathbf{f}}^{\text{chui}} \in \mathbb{R}^{916}$ if $N_2 = 2$ and $\hat{\mathbf{f}}^{\text{chui}} \in \mathbb{R}^{2356}$ for $N_2 = 3$. We compute the classification accuracy on the test data set, containing M random test data points. Again, all accuracies have to be understood as average values achieved over 10 runs.

As before, we get different CA performance depending on the regularization. We achieve the highest mean CA with $N_1 = 3$, $N_2 = 2$ and $\lambda = 2^{-6}$ for the ℓ_1 -norm regularization. For

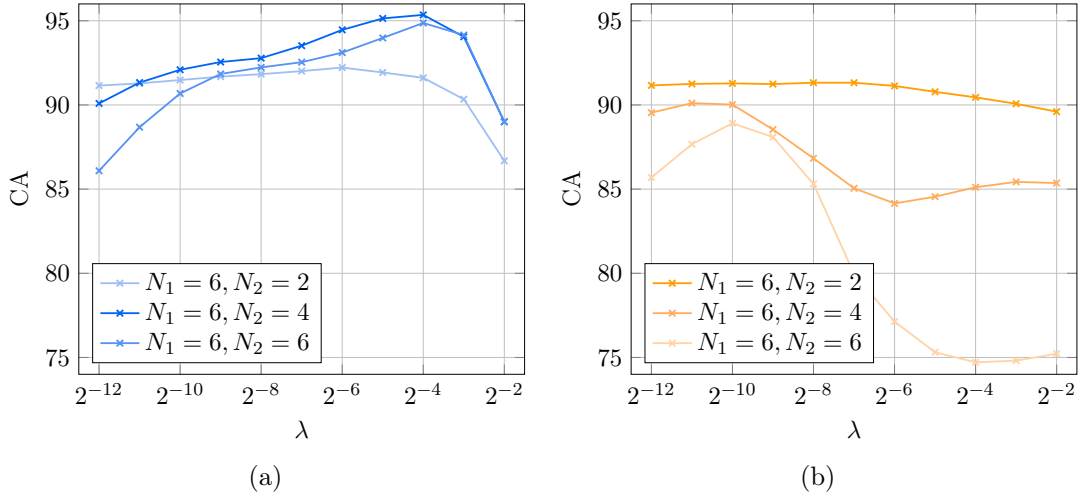


Figure 4.5: Visualization of mean CA over 10 runs achieved by the classifying functions $S(\mathcal{X}, \mathcal{I}_N(U_2))f^{\cos}(\mathbf{x})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\cos}$, using different regularization parameters λ and bandwidths \mathbf{N} , cf. (4.5). The results achieved with ℓ_1 -norm regularization are depicted in (a) and results achieved with ℓ_2 -norm regularization are shown in (b). The number of generated training and test data points have been set to 1000.

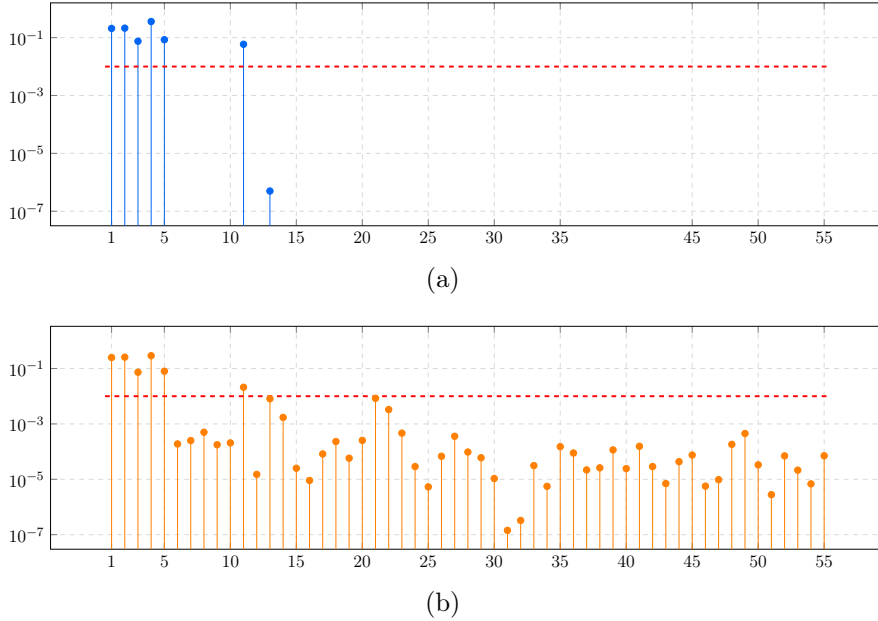


Figure 4.6: Global sensitivity indices $\rho(\mathbf{u}, S(\mathcal{X}, \mathcal{I}_N(U_2))f^{\cos})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\cos}$ using ℓ_1 -norm regularization (a), ℓ_2 -norm regularization (b) and corresponding best parameter choices, see Figure 4.5, for $\mathbf{u} \in U_2$, where $|U_2| = 56$. The resulting active set is $U^{0.01} = \{\{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{1, 2\}\}$.

the best results provided by ℓ_2 -norm regularization we need to choose $N_1 = 3$, $N_2 = 1$ and $\lambda = 2^{-8}$.

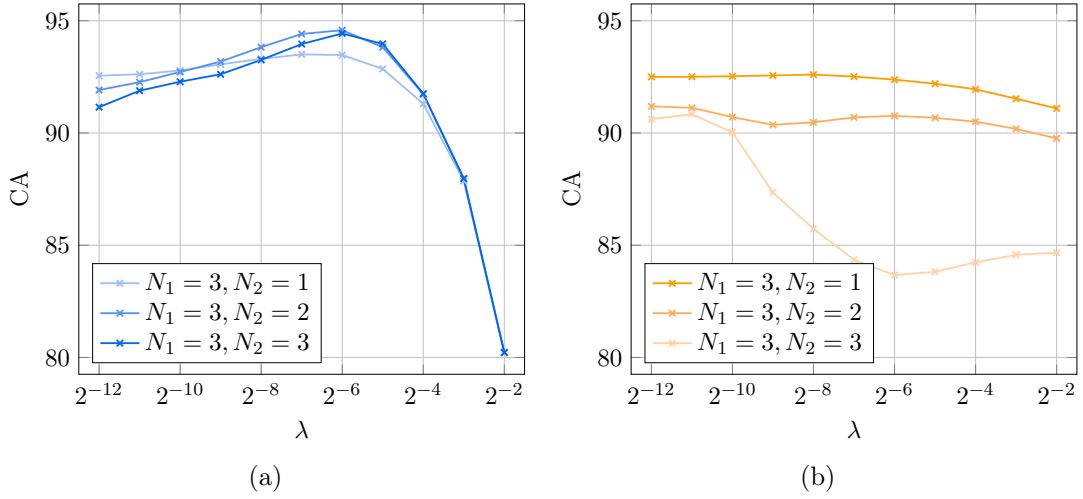



Figure 4.7: Visualization of mean CA over 10 runs achieved by the classifying functions $S(\mathcal{X}, \mathcal{J}_{\mathbf{N}}(U_2))f^{\text{chui}}(\mathbf{x})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\text{chui}}$, using different regularization parameters λ and bandwidths parameters \mathbf{N} , cf. (4.6). The number of generated training and test data points have been set to 5000. The results obtained with ℓ_1 -norm regularization are shown in (a) and results obtained with ℓ_2 -norm regularization are shown in (b).


In Figure 4.9 we have visualized the global sensitivity indices and get the same desired active set $U^{0.01}$. In addition, we use the active set $U^{0.01}$ to fit the classifying functions by using only this active set. Note that by using the bandwidths \mathbf{N} for cosine approach with $N_1 = 6$ we get $\hat{\mathbf{f}}^{\text{cos}} \in \mathbb{R}^{26}$ in case if $N_2 = 2$, and $\hat{\mathbf{f}}^{\text{cos}} \in \mathbb{R}^{34}$ for $N_2 = 4$. In case of wavelet approach we get $\hat{\mathbf{f}}^{\text{chui}} \in \mathbb{R}^{80}$ for $N_1 = 3$ and $N_2 = 1$, $\hat{\mathbf{f}}^{\text{chui}} \in \mathbb{R}^{92}$ for $N_1 = 3, N_2 = 2$. Moreover, we can even choose higher bandwidth, in order to get better results. For example using $N_1 = 4$ and $N_2 = 3$ we get $\hat{\mathbf{f}}^{\text{chui}} \in \mathbb{R}^{204}$, which improves the mean CA results. But, for instance, in case of cosine basis approach we couldn't find any improvement by using higher bandwidths.

5. Summary and outlook

Truncated ANOVA decomposition in the form of Fourier partial sums, and the like, have previously been used primarily for regression tasks. The goal of this work was to use such approaches in the form of a feature map within Support Vector Machines. The advantage here is that we can calculate the necessary matrix-vector products efficiently with grouped transformations, since we only restrict to couplings of small groups of variables. Furthermore, the models can be interpreted well in terms of the global sensitivity indices. Our numerical examples, in which we restricted our considerations to very simple test functions, show that we can achieve good classification results and that these can be interpreted as well. The next step would be more theoretical considerations, as well as the application of our methods to more complex test functions and real data sets.



(a)



(b)

Figure 4.8: Global sensitivity indices $\rho(\mathbf{u}, S(\mathcal{X}, \mathcal{J}_N(U_2))f^{\text{chui}})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\text{chui}}$ using ℓ_1 -norm regularization (a), ℓ_2 -norm regularization (b) and corresponding best parameter choices, see Figure 4.7, for $\mathbf{u} \in U_2$, where $|U_2| = 56$. The resulting active set is $U^{0.01} = \{\{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{1, 2\}\}$.

Acknowledgments

The authors gratefully acknowledge their support by the BMBF grant 01|S20053A (project SALE) and thank Felix Bartel, Michael Schmischke and Pascal Schröter for fruitful discussions.

A. The Friedman Function

In this section we discuss the Friedman 1 function (4.3) in more detail. First, we compute the active ANOVA terms in Section A.1. Based on a numerical test we verify how significant the single ANOVA terms contribute to the signum of the zero-mean function $F^1 - M(F^1)$, see Section A.2.

A.1. Computing the ANOVA-Decomposition of

For the Friedman 1 function as defined in (4.3) it is easy to see that the active set of subsets \mathbf{u} will be included in $U^* = \{\emptyset, \{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{1, 2\}\}$. In the following we state the ANOVA terms $F_{\mathbf{u}}^1(\mathbf{x}_{\mathbf{u}})$ for all $\mathbf{u} \in U^*$.

The constant part of the function is given by

$$F_{\emptyset}^1 = P_{\emptyset}F^1 = M(F^1) \approx 14.4133,$$

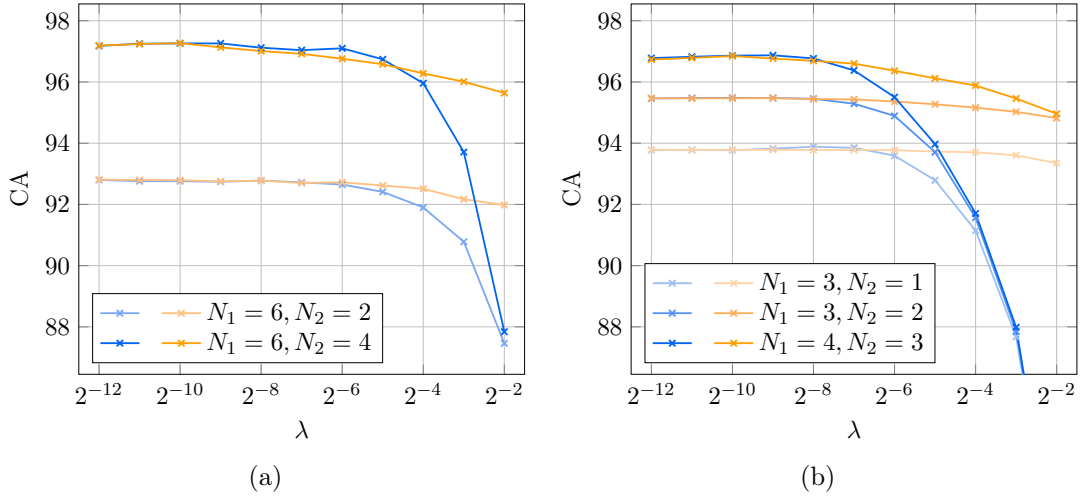


Figure 4.9: Visualization of improved mean CA over 10 runs depending on the regularization parameter λ and different bandwidths \mathbf{N} . In (a) we see the results achieved by using the classifying functions $S(\mathcal{X}, \mathcal{I}_{\mathbf{N}}(U^{0.01}))f^{\cos}(\mathbf{x})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\cos}$. Results using the classifying functions $S(\mathcal{X}, \mathcal{J}_{\mathbf{N}}(U^{0.01}))f^{\text{chui}}(\mathbf{x})$ with $\hat{\mathbf{f}} = \hat{\mathbf{f}}^{\text{chui}}$ are shown in (b). Depending on the regularization we used shades of blue to visualize the results obtained by ℓ_1 -norm regularization and shades of orange in case of ℓ_2 -norm regularization. We used 1000 training and test data points for the cosine basis approach and 5000 in case of wavelets.

as already stated above, cf. (4.4). Next, we compute for the subsets $\mathbf{u} \in U^*$ with $|\mathbf{u}| = 1$

$$\begin{aligned}
F_{\{1\}}^1(x_1) &= P_{\{1\}}F^1(x_1) - M(F^1) = \frac{10 - 10 \cos(\pi x_1)}{\pi x_1} + \frac{55}{6} - M(F^1), \\
F_{\{2\}}^1(x_2) &= P_{\{2\}}F^1(x_2) - M(F^1) = \frac{10 - 10 \cos(\pi x_2)}{\pi x_2} + \frac{55}{6} - M(F^1), \\
F_{\{3\}}^1(x_3) &= P_{\{3\}}F^1(x_3) - M(F^1) = 10 \int \frac{1 - \cos(\pi t)}{\pi t} dt + 20 \left(x_3 - \frac{1}{2}\right)^2 + \frac{15}{2} - M(F^1), \\
F_{\{4\}}^1(x_4) &= P_{\{4\}}F^1(x_4) - M(F^1) = 10 \int \frac{1 - \cos(\pi t)}{\pi t} dt + 10x_4 + \frac{25}{6} - M(F^1), \\
F_{\{5\}}^1(x_5) &= P_{\{5\}}F^1(x_5) - M(F^1) = 10 \int \frac{1 - \cos(\pi t)}{\pi t} dt + 5x_5 + \frac{20}{3} - M(F^1).
\end{aligned}$$

The value of the remaining integral

$$10 \int \frac{1 - \cos(\pi t)}{\pi t} dt \approx 5.2466$$

has to be computed numerically. Finally, the ANOVA term of the highest order is given by

$$\begin{aligned}
F_{\{1,2\}}^1(x_1, x_2) &= P_{\{1,2\}}F^1(x_1, x_2) - F_{\{1\}}^1(x_1) - F_{\{2\}}^1(x_2) - M(F^1) \\
&= 10 \sin(\pi x_1 x_2) + \frac{55}{6} - F_{\{1\}}^1(x_1) - F_{\{2\}}^1(x_2) - M(F^1) \\
&= 10 \sin(\pi x_1 x_2) + \frac{10 \cos(\pi x_1) - 10}{\pi x_1} + \frac{10 \cos(\pi x_2) - 10}{\pi x_2} - \frac{55}{6} + M(F^1).
\end{aligned}$$

It can easily be checked that the sum of all ANOVA terms gives in fact the function F^1 . The computed ANOVA terms are all non-zero and, thus, the active set is indeed given by U^* as stated above.

A.2. Importance of ANOVA Terms to the Signum

Based on the following numerical test we check the importance of the single ANOVA terms to the signum of the zero-mean function $F^1(\mathbf{x}) - M(F^1)$ in $d = 5$ dimensions, i.e., we restrict the Friedman-1 function to the important variables only.

We generated $N = 10^6$ randomly and uniformly distributed points $\mathbf{z}_j \in [0, 1]^5$, $j = 1, \dots, 5$. By $y_j = \text{sign}(F^1(\mathbf{z}_j) - M(F^1))$ we obtain a well-balanced data set, see Table A.1. Then, we checked whether leaving out one of the ANOVA terms will change the signum of the function values. It can be seen that all ANOVA terms contribute significantly to $\text{sign}(F^1 - M(F^1))$, see Table A.2. As an example, when omitting the ANOVA term $F_{\{4\}}^1$, the sign of the function value will change for approximately 20% of the data points.

Total number of points \mathbf{z}_j	10^6
Points with $\text{sign}(F^1(\mathbf{z}_j) - M(F^1)) = +1$	≈ 499600
Points with $\text{sign}(F^1(\mathbf{z}_j) - M(F^1)) = -1$	≈ 500400

Table A.1: By subtracting the mean value of the function we obtain two classes of approximately the same size.

subset \mathbf{u}	$\frac{\text{sign}(F^1(\mathbf{z}_j) - M(F^1)) \neq \text{sign}(F^1(\mathbf{z}_j) - M(F^1) - F_{\mathbf{u}}^1(\mathbf{z}_j, \mathbf{u}))}{\text{sign}(F^1(\mathbf{z}_j) - M(F^1) - F_{\mathbf{u}}^1(\mathbf{z}_j, \mathbf{u}))}$
\emptyset	$\approx 50.05\%$
$\{1\}$	$\approx 15.60\%$
$\{2\}$	$\approx 15.58\%$
$\{3\}$	$\approx 10.15\%$
$\{4\}$	$\approx 20.10\%$
$\{5\}$	$\approx 9.89\%$
$\{1, 2\}$	$\approx 8.05\%$

Table A.2: For each subset \mathbf{u} or rather ANOVA term $F_{\mathbf{u}}^1$ we state the percentage of the randomly chosen points \mathbf{z}_j , $j = 1, \dots, 10^6$, for which the sign changes if the ANOVA term is omitted.

References

- [1] L. Armijo. Minimization of functions having Lipschitz continuous first partial derivatives. *Pac. J. Math.*, 16(1):1–3, 1966.
- [2] F. Bartel, D. Potts, and M. Schmischke. Grouped transformations in high-dimensional explainable ANOVA approximation. *SIAM J. Sci. Comput.*, 44(3):1606–1631, 2022.
- [3] F. Bartel and M. Schmischke. GroupedTransforms.jl Julia Package. Contributor: P. Schröter, L. Weidensager.

- [4] A. Beck. *First-Order Methods in Optimization*. SIAM, Philadelphia, PA, 2017.
- [5] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.*, 2(1):183–202, 2009.
- [6] G. Beylkin. On the fast Fourier transform of functions with singularities. *Appl. Comput. Harmon. Anal.*, 2:363–381, 1995.
- [7] G. Beylkin, J. Garcke, and M. Mohlenkamp. Multivariate regression and machine learning with sums of separable functions. *SIAM J. Sci. Comput.*, 31:1840–1857, 2009.
- [8] P. Binev, W. Dahmen, and P. Lamby. Fast high-dimensional approximation with sparse occupancy trees. *J. Comput. Appl. Math.*, 235(8):2063 – 2076, 2011.
- [9] B. E. Boser, I. M. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In *Proc. 5th Annual Workshop on Computational Learning Theory*, pages 144–152, Pittsburgh, USA, 1992. Association for Computing Machinery (ACM).
- [10] R. Caffisch, W. Morokoff, and A. Owen. Valuation of mortgage-backed securities using Brownian bridges to reduce effective dimension. *J. Comput. Finance.*, 1:27–46, 1997.
- [11] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [12] C. K. Chui. *An Introduction to Wavelets*. Wavelet Analysis and its Applications. Academic Press, San Diego, USA, 1992.
- [13] C. Cortes and V. Vapnik. Support-vector networks. *Mach. Learn.*, 20:273–297, 1995.
- [14] I. Daubechies. *Ten Lectures on Wavelets*. SIAM, Philadelphia, USA, 1992.
- [15] R. DeVore, G. Petrova, and P. Wojtaszczyk. Approximation of functions of few variables in high dimensions. *Constr. Approx.*, 33:125–143, 2011.
- [16] A. Dutt and V. Rokhlin. Fast Fourier transforms for nonequispaced data. *SIAM J. Sci. Stat. Comput.*, 14:1368–1393, 1993.
- [17] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008. Software available at <http://www.csie.ntu.edu.tw/~cjlin/liblinear>.
- [18] J. H. Friedman, E. Grosse, and W. Stuetzle. Multidimensional Additive Spline Approximation. *SIAM Journal on Scientific and Statistical Computing*, 4(2):291–301, 1983.
- [19] C. Gu. *Smoothing Spline ANOVA Models*. Springer Series in Statistics. Springer Science+Business Media, New York, USA, 2013.
- [20] A. Haar. Zur Theorie der orthogonalen Funktionensysteme. *Math. Ann.*, 69:331–371, 1910.
- [21] J. Keiner, S. Kunis, and D. Potts. NFFT 3.5, C subroutine library. Contributors: F. Bartel, M. Fenn, T. Görner, M. Kircheis, T. Knopp, M. Quellmalz, M. Schmischke, T. Volkmer, A. Vollrath.

- [22] J. Keiner, S. Kunis, and D. Potts. Using NFFT3 - a software library for various nonequispaced fast Fourier transforms. *ACM Trans. Math. Software*, 36:Article 19, 1–30, 2009.
- [23] F. Y. Kuo, I. H. Sloan, G. W. Wasilkowski, and H. Woźniakowski. On decompositions of multivariate functions. *Math. Comput.*, 79(270):953–966, 2009.
- [24] L. Lippert, D. Potts, and T. Ullrich. Fast hyperbolic wavelet regression meets ANOVA. *Numer. Math.*, 154:155–207, 2023.
- [25] R. Liu and A. B. Owen. Estimating mean dimensionality of ANOVA decompositions. *J. Am. Stat. Assoc.*, 101(474):712–721, 2006.
- [26] D. Meyer, F. Leisch, and K. Hornik. The support vector machine under test. *Neurocomputing*, 55(1):169 – 186, 2003.
- [27] F. Nestler, M. Stoll, and T. Wagner. Learning in high-dimensional feature spaces using ANOVA-based fast matrix-vector multiplication. *Foundations of Data Science*, 4(3):423–440, 2022.
- [28] J. Nocedal and S. J. Wright. *Numerical Optimization*. Operations Research and Financial Engineering. Springer, New York, USA, 2006.
- [29] N. Parikh and S. Boyd. Proximal Algorithms. *Found. Trends Mach. Learn.*, 1(3):127–239, 2014.
- [30] G. Plonka, D. Potts, G. Steidl, and M. Tasche. *Numerical Fourier Analysis*. Applied and Numerical Harmonic Analysis. Birkhäuser, Cham, CH, 2018.
- [31] D. Potts and M. Schmischke. Approximation of high-dimensional periodic functions with Fourier-based methods. *SIAM J. Numer. Anal.*, 59(5):2393–2429, 2021.
- [32] D. Potts and M. Schmischke. Interpretable approximation of high-dimensional data. *SIAM J. Math. Data Sci.*, 3(4):1301–1323, 2021.
- [33] D. Potts and M. Schmischke. Learning multivariate functions with low-dimensional structures using polynomial bases. *J. Comput. Appl. Math.*, 403:113821, 2022.
- [34] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- [35] M. Schmischke. ANOVAapprox.jl Julia Package. Contributors: P. Schröter, F. Bartel, L. Weidensager, F. Nestler.
- [36] M. Schmischke. NFFT3.jl Julia Package. Contributors: F. Bartel, T. Riemer.
- [37] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. The MIT Press, 06 2018.
- [38] I. M. Sobol. On sensitivity estimation for nonlinear mathematical models. *Matem. Mod.*, 2(1):112–118, 1990.

- [39] I. M. Sobol. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Math. Comput. Simul.*, 55(1-3):271–280, 2001.
- [40] G. Steidl. A note on fast Fourier transforms for nonequispaced grids. *Adv. Comput. Math.*, 9:337–353, 1998.
- [41] G. Steidl. Supervised Learning by Support Vector Machines. In *Handbook of Mathematical Methods in Imaging*, page 959–1013. Springer, New York, USA, 2011.
- [42] I. Steinwart and A. Christmann. *Support Vector Machines*. Information Science and Statistics. Springer Science+Business Media, New York, USA, 2008.
- [43] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, Cambridge, USA, 2018.
- [44] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. R. Statist. Soc.*, 58(1):267–288, 1996.
- [45] V. N. Vapnik. *The Nature of Statistical Learning Theory*. Springer, 1995.
- [46] V. N. Vapnik and A. Lerner. Pattern recognition using generalized portrait method. *Autom. Remote Control*, 24(6):774–780, 1963.
- [47] T. Wagner, J. W. Pearson, and M. Stoll. A Preconditioned Interior Point Method for Support Vector Machines Using an ANOVA-Decomposition and NFFT-Based Matrix-Vector Products. *arXiv: 2312.00538*, 2023.
- [48] D. F. Walnut. *An Introduction to Wavelet Analysis*. Applied and Numerical Harmonic Analysis. Birkhäuser, New York, USA, 2004.
- [49] D. Werner. *Funktionalanalysis*. Springer, Berlin, Heidelberg, DE, 8 edition, 2018.
- [50] J. Zhu, S. Rosset, R. Tibshirani, and T. Hastie. 1-norm support vector machines. In *Proc. 16th International Conference on Neural Information Processing Systems*, pages 49–56, Whistler, CAN, 2003. MIT Press.