
INTERPRETABLE TRANSFORMED ANOVA APPROXIMATION ON THE EXAMPLE OF THE PREVENTION OF FOREST FIRES

A PREPRINT

📧 **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
potts@math.tu-chemnitz.de

November 7, 2021

ABSTRACT

The distribution of data points is a key component in machine learning. In most cases, one uses min-max normalization to obtain nodes in $[0, 1]$ or Z-score normalization for standard normal distributed data. In this paper, we apply transformation ideas in order to design a complete orthonormal system in the L_2 space of functions with the standard normal distribution as integration weight. Subsequently, we are able to apply the explainable ANOVA approximation for this basis and use Z-score transformed data in the method. We demonstrate the applicability of this procedure on the well-known forest fires data set from the UCI machine learning repository. The attribute ranking obtained from the ANOVA approximation provides us with crucial information about which variables in the data set are the most important for the detection of fires.

Keywords ANOVA · high-dimensional · approximation · interpretability · normal distribution

1 Introduction

In machine learning, the scale of our features is a key component in building models. When we work with data from applications we have to accept it as it is. In most cases, we cannot control where the nodes are lying. Let us, e.g., take recommendations in online shopping. We are only able analyze the customers that actually exist and what they bought in the shop. However, the features may lie on immensely different scales. If we measure, e.g., the time a customer spent in the shop in seconds as well as their age in years, the result will be a scale that contains values with thousands of seconds and a scale ranging from up to 90 years. Bringing those features on similar scales through normalization may significantly improve performance of our model.

Two common methods for data normalization are min-max-normalization and Z-score normalization, see e.g. [7]. The former method will yield data in the interval $[0, 1]$ and is especially useful if there is an intrinsic upper and lower bound for the values, e.g., when considering age. If we come back to our previous example, the time a customer spends in the shop would be less suitable since the values may have a wide range and we will probably have very few people with a significantly small or large time. In this case, the Z-score normalization makes much more sense. It tells us how many standard deviations our value lies away from the mean of the data resulting in a distribution with zero mean and variance one.

The explainable ANOVA approximation method introduced in [21, 23, 22] is based on the well-known multivariate analysis of variance (ANOVA) decomposition, see e.g. [3, 24, 14, 13, 8, 18], and relies on the existence of a complete orthonormal system in the space which is suitable for fast matrix-vector multiplication algorithms in grouped transformations, c.f. [2]. Until now, this method was always applied with min-max-normalization since it relied on the space $L_2([0, 1]^d)$ of square-integrable functions over the cube with the half-period cosine basis. It is our goal to modify the approach in order to create the possibility to work with standard normal distributed data, i.e., data that has been obtained through Z-score normalization.

We aim to achieve this by using the transformation ideas from [17] and [16] in order to construct a complete orthonormal system in the space

$$L_2(\mathbb{R}^d, \omega) := \left\{ f: \mathbb{R}^d \rightarrow \mathbb{R}: \|f\|_{L_2(\mathbb{R}^d)} := \sqrt{\int_{\mathbb{R}^d} |f(\mathbf{x})|^2 \omega(\mathbf{x}) d\mathbf{x}} < \infty \right\} \quad (1)$$

with the probability density of the standard normal distribution

$$\omega(\mathbf{x}) := \prod_{i=1}^d \frac{1}{\sqrt{2\pi}} e^{-x_i^2/2} = (2\pi)^{-d/2} e^{-\|\mathbf{x}\|_2^2/2}. \quad (2)$$

Combining this transformation with the half-period cosine basis allows for fast multiplications in the grouped transformations and makes the ANOVA approximation method applicable for Z-score normalized data.

As an example, we apply this approach to a data set about the detection of forest fires, see [1, 4]. Constructing a model with the capability of efficiently predicting the size of the fire in this data set may provide a way of predicting the occurrences of fires. This creates the possibility of efficiently implementing appropriate counter-measures. In our time of climate change with massive forest fires every year, e.g., in Australia or the USA, this is an extremely current topic. With the interpretation capabilities of the ANOVA method, cf. [22], we are additionally able to explain the importance of our features and give reasonable explanation for the predictions.

2 Transformed Half-Period Cosine

In this section, it is our goal to construct a complete orthonormal system in the space $L_2(\mathbb{R}^d, \omega)$ from (1) with the product density $\omega(\mathbf{x})$ from (2). This is the probability density function of the standard normal distribution, i.e., the normal distribution with zero mean and variance one. We have $\int_{\mathbb{R}^d} \omega(\mathbf{x}) d\mathbf{x} = 1$ as well as $\sup_{\mathbf{x} \in \mathbb{R}^d} \omega(\mathbf{x}) = (2\pi)^{-d/2}$ which implies $\omega \in L_\infty(\mathbb{R}^d)$.

We aim to construct the basis using transformation ideas from [16, 17] and the half-period cosine basis on $L_2([0, 1]^d)$. The orthonormal basis functions on $L_2([0, 1]^d)$ are given by

$$\phi_{\mathbf{k}}^{\cos}(\mathbf{x}) = \sqrt{2^{|\mathbf{k}|_0}} \prod_{i=1}^d \cos(\pi k_i x_i), \quad \mathbf{k} \in \mathbb{N}_0^d \quad (3)$$

with $\|\mathbf{k}\|_0 := |\text{supp } \mathbf{k}|$ and $\text{supp } \mathbf{k} := \{s \in \{1, 2, \dots, d\}: k_s \neq 0\}$. We start from a given function $f: \mathbb{R}^d \rightarrow \mathbb{R}$, $f \in L_2(\mathbb{R}^d, \omega)$, and aim to transform it onto the cube $[0, 1]^d$. As transformation we propose

$$\psi: [0, 1]^d \rightarrow \mathbb{R}^d, \quad \psi(\mathbf{x}) = \sqrt{2} \begin{pmatrix} \text{erf}^{-1}(2x_1 - 1) \\ \text{erf}^{-1}(2x_2 - 1) \\ \vdots \\ \text{erf}^{-1}(2x_d - 1) \end{pmatrix} \quad (4)$$

with the inverse transformation

$$\psi^{-1}: \mathbb{R}^d \rightarrow [0, 1]^d, \quad \psi^{-1}(\mathbf{x}) = \frac{1}{2} \begin{pmatrix} \text{erf}(x_1/\sqrt{2}) + 1 \\ \text{erf}(x_2/\sqrt{2}) + 1 \\ \vdots \\ \text{erf}(x_d/\sqrt{2}) + 1 \end{pmatrix}. \quad (5)$$

The error function is given by

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt.$$

As a result, we have the commutative diagram in Figure 1. This allows us to transform the half-period cosine to a complete orthonormal system on $L_2(\mathbb{R}^d, \omega)$ with the help of

Lemma 2.1. *Let $g, h \in L_2([0, 1]^d)$, $u, v \in L_2(\mathbb{R}^d, \omega)$ with probability density ω from (2), and transformation ψ, ψ^{-1} as in (4) and (5) respectively. Then*

$$\begin{aligned} \langle g \circ \psi^{-1}, h \circ \psi^{-1} \rangle_{L_2(\mathbb{R}^d, \omega)} &= \langle g, h \rangle_{L_2([0, 1]^d)} \\ \langle u, v \rangle_{L_2(\mathbb{R}^d, \omega)} &= \langle u \circ \psi, v \circ \psi \rangle_{L_2([0, 1]^d)} \end{aligned}$$

and subsequently $\|h\|_{L_2([0, 1]^d)} = \|h \circ \psi^{-1}\|_{L_2(\mathbb{R}^d, \omega)}$ and $\|u \circ \psi\|_{L_2([0, 1]^d)} = \|u\|_{L_2(\mathbb{R}^d, \omega)}$.




Figure 1: Commutative diagram of the function and the transformations.

Proof. Let $g, h \in L_2([0, 1]^d)$ and $u, v \in L_2(\mathbb{R}^d, \omega)$. Then we insert the definition and perform a change of variables as follows

$$\begin{aligned} \langle g \circ \psi^{-1}, h \circ \psi^{-1} \rangle_{L_2(\mathbb{R}^d, \omega)} &= \int_{\mathbb{R}^d} g(\psi^{-1}(\mathbf{x})) h(\psi^{-1}(\mathbf{x})) \omega(\mathbf{x}) \, d\mathbf{x} \\ &= \int_{\mathbb{R}^d} g(\mathbf{t}) h(\mathbf{t}) \omega(\psi(\mathbf{t})) \psi'(\mathbf{t}) \, d\mathbf{t}. \end{aligned}$$

As functional determinant we obtain $\psi'(\mathbf{t}) = \prod_{i=1}^d \sqrt{2\pi} e^{\text{erf}^{-2}(2x_i-1)}$ and subsequently

$$\omega(\psi(\mathbf{t})) \psi'(\mathbf{t}) = \prod_{i=1}^d \frac{1}{\sqrt{2\pi}} e^{-\text{erf}^{-2}(2x_i-1)} \cdot \sqrt{2\pi} e^{\text{erf}^{-2}(2x_i-1)} = 1.$$

This proves the first equality. For the second equality, we may use an analogous procedure. \square

Theorem 2.2. *The functions $(\phi_{\mathbf{k}}^{\text{trafo}})_{\mathbf{k} \in \mathbb{N}_0^d}$ with*

$$\phi_{\mathbf{k}}^{\text{trafo}}(\mathbf{x}) := (\phi_{\mathbf{k}}^{\text{cos}} \circ \psi^{-1})(\mathbf{x}) = \sqrt{2}^{\|\mathbf{k}\|_0} \prod_{i=1}^d \cos\left(\pi k_i \frac{\text{erf}(x_i/\sqrt{2}) + 1}{2}\right) \quad (6)$$

form a complete orthonormal system in $L_2(\mathbb{R}^d, \omega)$.

Proof. Clearly, we have $\phi_{\mathbf{k}}^{\text{trafo}} \in L_2(\mathbb{R}^d, \omega)$, $\mathbf{k} \in \mathbb{N}_0^d$ by Lemma 2.1. The orthonormality follows from Lemma 2.1 and the orthonormality of the half-period cosine basis, i.e., for $\mathbf{k}, \ell \in \mathbb{N}_0^d$ we have

$$\begin{aligned} \langle \phi_{\mathbf{k}}^{\text{trafo}}, \phi_{\ell}^{\text{trafo}} \rangle_{L_2(\mathbb{R}^d, \omega)} &= \langle \phi_{\mathbf{k}}^{\text{cos}} \circ \psi^{-1}, \phi_{\ell}^{\text{cos}} \circ \psi^{-1} \rangle_{L_2(\mathbb{R}^d, \omega)} \\ &= \langle \phi_{\mathbf{k}}^{\text{cos}}, \phi_{\ell}^{\text{cos}} \rangle_{L_2([0, 1]^d)} = \delta_{\mathbf{k}, \ell}. \end{aligned}$$

It remains to show that the system is complete, i.e., for every $f \in L_2(\mathbb{R}^d, \omega)$ we have

$$\lim_{n \rightarrow \infty} \left\| f - \sum_{i=1}^n \langle f, \phi_{\mathbf{k}_n}^{\text{trafo}} \rangle_{L_2(\mathbb{R}^d, \omega)} \phi_{\mathbf{k}_n}^{\text{trafo}} \right\|_{L_2(\mathbb{R}^d, \omega)} = 0$$

with \mathbf{k}_n , $n = 1, 2, \dots$, an order of $\mathbf{k} \in \mathbb{N}_0^d$. First of all, we have



$$\hat{c}_{\mathbf{k}} := \langle f, \phi_{\mathbf{k}}^{\text{trafo}} \rangle_{L_2(\mathbb{R}^d, \omega)} = \langle f \circ \psi, \phi_{\mathbf{k}}^{\text{cos}} \rangle_{L_2([0, 1]^d)}$$

by Lemma 2.1. We apply the norm equality from Lemma 2.1 to obtain

$$\lim_{n \rightarrow \infty} \left\| f - \sum_{i=1}^n \hat{c}_{\mathbf{k}_n} \phi_{\mathbf{k}_n}^{\text{trafo}} \right\|_{L_2(\mathbb{R}^d, \omega)} = \lim_{n \rightarrow \infty} \left\| f \circ \psi - \sum_{i=1}^n \hat{c}_{\mathbf{k}_n} \phi_{\mathbf{k}_n}^{\text{cos}} \right\|_{L_2([0, 1]^d)}$$

and since $(\phi_{\mathbf{k}}^{\text{cos}})_{\mathbf{k} \in \mathbb{N}_0^d}$ is complete in $L_2([0, 1]^d)$ and $f \circ \psi \in L_2([0, 1]^d)$, the limit is zero and our statement is proven. \square

In summary, we have constructed a complete orthonormal system $(\phi_{\mathbf{k}}^{\text{trafo}})_{\mathbf{k} \in \mathbb{N}_0^d}$ on the weighted space $L_2(\mathbb{R}^d, \omega)$ using transformation ideas from [16] and the well-known half-period cosine basis $(\phi_{\mathbf{k}}^{\text{cos}})_{\mathbf{k} \in \mathbb{N}_0^d}$ on $L_2([0, 1]^d)$. The transformation ψ in one dimension and the corresponding basis functions $\phi_{\mathbf{k}}^{\text{trafo}}$ are visualized in Figure 2.

(a) Transformation ψ from (4) in one dimension(b) Basis functions ϕ_k^{trafo} from (6) in one dimension for $k = 1$ (solid), $k = 2$ (dotted), and $k = 3$ (dashed)Figure 2: Transformation ψ and transformed basis functions ϕ_k^{trafo} in one dimension

3 Interpretable ANOVA Approximation

In this section, we briefly summarize the interpretable ANOVA (analysis of variance) approximation method and the idea of grouped transformations, see [21, 2]. The approach was considered for periodic functions, but has been expended to non-periodic functions in [23, 22]. In this paper, we focus on functions $f: \mathbb{R}^d \rightarrow \mathbb{R}$ from $L_2(\mathbb{R}^d, \omega)$ with probability density ω from (2). Since ω is the standard normal distribution, this function space is of a high relevance. It allows us e.g. to work with data from applications that has been Z-transformed, i.e., data with zero mean and variance one, see e.g. [7]. Since the transformed half-period cosine $(\phi_{\mathbf{k}}^{\text{trafo}})_{\mathbf{k} \in \mathbb{N}_0^d}$, see Theorem 2.2, is a complete orthonormal system in the space $L_2(\mathbb{R}^d, \omega)$, we have

$$f(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} c_{\mathbf{k}}(f) \phi_{\mathbf{k}}^{\text{trafo}}(\mathbf{x}), \quad c_{\mathbf{k}}(f) = \langle f, \phi_{\mathbf{k}}^{\text{trafo}} \rangle_{L_2(\mathbb{R}^d, \omega)}, \quad (7)$$

and through Parseval's identity $\|f\|_{L_2(\mathbb{R}^d, \omega)}^2 = \sum_{\mathbf{k} \in \mathbb{Z}^d} |c_{\mathbf{k}}(f)|^2$.

The classical ANOVA decomposition, cf. [3, 24, 14, 8], provides us with a unique decomposition in the frequency domain as shown in [21]. We denote the coordinate indices with $[d] = \{1, 2, \dots, d\}$ and subsets as bold small letters, e.g., $\mathbf{u} \subseteq [d]$. The ANOVA terms are defined as

$$f_{\mathbf{u}}(\mathbf{x}) = f_{\mathbf{u}}(\mathbf{x}_{\mathbf{u}}) := \sum_{\substack{\mathbf{k} \in \mathbb{Z}^d \\ \text{supp } \mathbf{k} = \mathbf{u}}} c_{\mathbf{k}}(f) \phi_{\mathbf{k}}^{\text{trafo}}(\mathbf{x}).$$

The function can then be uniquely decomposed as

$$f(\mathbf{x}) = \sum_{\mathbf{u} \subseteq [d]} f_{\mathbf{u}}(\mathbf{x})$$

into $|\mathcal{P}([d])| = 2^d$ ANOVA terms where $\mathcal{P}([d])$ is the potency set of $[d]$. Here, the exponentially growing number of terms shows an expression of the curse of dimensionality in the decomposition.

It is our goal to obtain information on how important the ANOVA terms $f_{\mathbf{u}}$ are with respect to the function f . In order to measure this, we define the variance of a function f as

$$\sigma^2(f) := \|f\|_{L_2(\mathbb{R}^d, \omega)}^2 - |c_{\mathbf{0}}(f)|^2 = \sum_{\mathbf{k} \in \mathbb{Z}^d \setminus \{\mathbf{0}\}} |c_{\mathbf{k}}(f)|^2.$$

Note that we have the special case $\sigma^2(f_{\mathbf{u}}) = \|f_{\mathbf{u}}\|_{L_2(\mathbb{R}^d, \omega)}^2$, $\mathbf{u} \subseteq [d]$. The relative importance with respect to f is then measured via global sensitivity indices (GSI) or Sobol indices, see [25, 26, 14], defined as

$$\varrho(\mathbf{u}, f) := \frac{\sigma^2(f_{\mathbf{u}})}{\sigma^2(f)}. \quad (8)$$

From the GSI we get a motivation for the concept of effective dimensions, specifically the superposition dimension as one notion of effective dimension. For a given $\alpha \in [0, 1]$ it is defined as

$$d^{(\text{sp})} := \min \left\{ s \in [d] : \frac{1}{\sigma^2(f)} \sum_{\substack{\mathbf{u} \subseteq [d] \\ |\mathbf{u}| \leq s}} \|f_{\mathbf{u}}\|_{L_2(\mathbb{R}^d, \omega)}^2 \geq \alpha \right\}. \quad (9)$$

The superposition dimension $d^{(\text{sp})}$ tells us that we can explain the α -part of the variance of f by terms $f_{\mathbf{u}}$ with $\mathbf{u} \leq d_s$.

Using subsets of ANOVA terms $U \subseteq \mathcal{P}([d])$, it is our goal to find a way to circumvent the curse of dimensionality for efficient approximation. In order to achieve this, we aim to truncate the ANOVA decomposition by taking only the ANOVA terms in U into account. The truncated ANOVA decomposition is then defined as

$$\mathbb{T}_U f(\mathbf{x}) = \sum_{\mathbf{u} \in U} f_{\mathbf{u}}(\mathbf{x}).$$

A specific idea for the truncation comes from the superposition dimension $d^{(\text{sp})}$ in (9). The idea is to take only variable interactions into account that contain d_s or less variables, i.e., the subset of ANOVA terms is

$$U(d, d_s) := \{\mathbf{u} \subseteq [d] : |\mathbf{u}| \leq d_s\}. \quad (10)$$

Since d_s does not necessarily have to coincide to the superposition dimension $d^{(\text{sp})}$, we call it *superposition threshold*. A well-known fact from learning theory is that the number of terms in $U(d, d_s)$ grows only polynomially in d for fixed $d_s < d$, i.e.,

$$|U(d, d_s)| \leq \left(\frac{ed}{d_s}\right)^{d_s}$$

which has reduced the curse of dimensionality.

In the following, we argue why the truncation by a superposition threshold d_s works well in relevant cases. For the approximation of functions that belong to a space $H^s(\mathbb{R}^d, \omega) \subseteq L_2(\mathbb{R}^d, \omega)$ that characterizes the smoothness $s > 0$ by the decay of the basis coefficients $c_{\mathbf{k}}(f)$, we can show upper bounds on the superposition dimension $d^{(\text{sp})}$ for $\alpha \in [0, 1]$, see e.g. [21]. In fact, there are types of smoothness that are proven to yield a low upper bound for the superposition dimension specifically dominating-mixed smoothness with POD (product and order-dependent) weights, cf. [12, 5, 11, 6, 21].

In terms of real data from applications, the situation is much different. Here, we cannot make the assumption that in complete generality we have a low superposition dimension. However, there are many application scenarios where numerical experiments successfully showed that this is indeed the case, see e.g. [3]. Since we generally do not have a-priori information, we work with low superposition thresholds d_s for truncation and validate on our test data.

3.1 Approximation Procedure

In this section, we briefly discuss how the approximation is numerically obtained and how we can interpret the results. In this section, we assume a given subset of ANOVA terms $U \subseteq \mathcal{P}([d])$. This set may be equal to or a subset of $U(d, d_s)$. We have given scattered data in the form of a set $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M\} \subseteq \mathbb{R}^d$ of standard normal distributed nodes and values $\mathbf{y} \in \mathbb{R}^M$, $M \in \mathbb{N}$. Moreover, we assume that there is an $L_2(\mathbb{R}^d, \omega)$ function f of form (7) with $f(\mathbf{x}_i) \approx y_i$ which we want to approximate.

First, we truncate f to the set U such that $f \approx \mathbb{T}_U f$. However, there are still infinitely many coefficients and therefore we perform a truncation to partial sums on finite support index sets

$$\mathcal{I}_{\emptyset} = \{0\}, \text{ and } \mathcal{I}_{\mathbf{u}} = \{1, 2, \dots, N_{|\mathbf{u}|} - 1\}^{|\mathbf{u}|} \quad (11)$$

with order-dependent parameters $N_{|\mathbf{u}|} \in \mathbb{N}$, $|\mathbf{u}| = 1, 2, \dots, d_s$, for every ANOVA term $f_{\mathbf{u}}$, $\mathbf{u} \in U$. Using the projections $P_{\mathbf{u}} \mathcal{I}_{\mathbf{u}} = \{\mathbf{k} \in \mathbb{N}_0^d : \mathbf{k}_{\mathbf{u}} \in \mathcal{I}_{\mathbf{u}}, \mathbf{k}_{\mathbf{u}^c} = \mathbf{0}\}$, we obtain

$$f_{\mathbf{u}}(\mathbf{x}) \approx \sum_{\mathbf{k} \in P_{\mathbf{u}} \mathcal{I}_{\mathbf{u}}} c_{\mathbf{k}}(f) \phi_{\mathbf{k}}^{\text{trafo}}(\mathbf{x}).$$

Now, we taking the union $\mathcal{I}(U) = \bigcup_{\mathbf{u} \in U} P_{\mathbf{u}} \mathcal{I}_{\mathbf{u}}$ yields

$$f(\mathbf{x}) \approx \sum_{\mathbf{k} \in \mathcal{I}(U)} c_{\mathbf{k}}(f) \phi_{\mathbf{k}}^{\text{trafo}}(\mathbf{x}).$$

The unknown coefficients $c_k(f)$ are now to be determined.

We aim to achieve this by solving the regularized least-squares problem

$$\hat{\mathbf{f}} = (\hat{f}_k)_{k \in \mathcal{I}(U)} = \arg \min_{\hat{\mathbf{g}} \in \mathbb{R}^{|\mathcal{I}(U)|}} \|\mathbf{y} - \mathbf{F}(\mathcal{X}, \mathcal{I}(U))\hat{\mathbf{g}}\|_2^2 + \lambda \|\hat{\mathbf{g}}\|_2^2, \quad (12)$$

cf. [21, 23, 2], with the basis matrix $\mathbf{F}(\mathcal{X}, \mathcal{I}(U)) = (\phi_k^{\text{trafo}}(\mathbf{x}))_{\mathbf{x} \in \mathcal{X}, k \in \mathcal{I}(U)}$. We solve the problem using the iterative LSQR solver [19]. In order to apply LSQR, we rewrite (12) by observing the equality

$$\|\mathbf{y} - \mathbf{F}(\mathcal{X}, \mathcal{I}(U))\hat{\mathbf{g}}\|_2^2 + \lambda \|\hat{\mathbf{g}}\|_2^2 = \left\| \begin{pmatrix} \mathbf{y} \\ \mathbf{0} \end{pmatrix} - \begin{pmatrix} \mathbf{F}(\mathcal{X}, \mathcal{I}(U)) \\ \sqrt{\lambda} \mathbf{I} \end{pmatrix} \hat{\mathbf{g}} \right\|_2^2 \quad (13)$$

with $\mathbf{0}$ the zero vector in $\mathbb{R}^{|\mathcal{I}(U)|}$ and $\mathbf{I} \in \mathbb{R}^{|\mathcal{I}(U)|, |\mathcal{I}(U)|}$ the identity matrix. Note that we always have a unique solution in this case since the matrix

$$\begin{pmatrix} \mathbf{F}(\mathcal{X}, \mathcal{I}(U)) \\ \sqrt{\lambda} \mathbf{I} \end{pmatrix}$$

has full column rank. However, the solution depends on the regularization parameter λ .

We apply the matrix-free variant of LSQR, i.e., we never explicitly construct the matrix $\mathbf{F}(\mathcal{X}, \mathcal{I}(U))$. The grouped transformations introduced in [2] provide *oracle* functions for the multiplications of $\mathbf{F}(\mathcal{X}, \mathcal{I}(U))$ and its transposed $\mathbf{F}^\top(\mathcal{X}, \mathcal{I}(U))$ with vectors. For our specific basis functions ϕ_k^{trafo} the grouped transformations are based on the non-equispaced fast cosine transform or NFCT, see [10, 20]. The transformation uses parallelization to separate our multiplication into smaller, up to d_s -dimensional NFCTs which results in an efficient algorithm. For more details we refer to [2].

In order to solve the minimization we employ the iterative LSQR solver [19] which needs a method for efficient multiplication with $\mathbf{F}(\mathcal{X}, \mathcal{I}(U))$ and its adjoint $\mathbf{F}^*(\mathcal{X}, \mathcal{I}(U))$ in the periodic case, otherwise its transposed matrix. This is realized by the Grouped Transformation idea in [2] based on the NFFT or the NFCT, see [10, 20].

One key fact is that the nodes \mathcal{X} have to be distributed according to the probability density ω of the space such that the Moore–Penrose inverse $\mathbf{F}^\dagger(\mathcal{X}, \mathcal{I}(U))$ is well-conditioned. In our case, ω is the density of the standard normal distribution, i.e., the nodes \mathcal{X} have to be distributed accordingly. For a detailed discussion on the properties of those matrices we refer to [9, 15] where our basis is a special case.

We use the global sensitivity indices $\varrho(\mathbf{u}, S(\mathcal{X}, \mathcal{I}(U))f)$, $\mathbf{u} \in U$, from the approximation $S(\mathcal{X}, \mathcal{I}(U))f(\mathbf{x})$ to compute approximations for the global sensitivity indices $\varrho(\mathbf{u}, f)$ of the function f . Here, we do not consider the index to be a good approximation if the values are close together, but rather if their order is identical, i.e., we have

$$\varrho(\mathbf{u}_1, f) \leq \varrho(\mathbf{u}_2, f) \implies \varrho(\mathbf{u}_1, S(\mathcal{X}, \mathcal{I}(U))f) \leq \varrho(\mathbf{u}_2, S(\mathcal{X}, \mathcal{I}(U))f)$$

for any pair $\mathbf{u}_1, \mathbf{u}_2 \in U$. We assume that this is the case for our choices of index sets $\mathcal{I}(U)$. In particular, the quality of the approximation corresponds to the accuracy of this assumption.

In order to rank the influence of the variables x_1, x_2, \dots, x_d we use the ranking score

$$r(i) = \frac{\sum_{\mathbf{u} \in \{\mathbf{v} \in U: i \in \mathbf{v}\}} |\{\mathbf{v} \in U: |\mathbf{u}| = |\mathbf{v}|, i \in \mathbf{v}\}|^{-1} \varrho(\mathbf{u}, S(\mathcal{X}, \mathcal{I}(U))f)}{\sum_{\mathbf{u} \in U} \left(\sum_{i \in \mathbf{u}} |\{\mathbf{v} \in U: |\mathbf{u}| = |\mathbf{v}|, i \in \mathbf{v}\}|^{-1} \right) \varrho(\mathbf{u}, S(\mathcal{X}, \mathcal{I}(U))f)} \quad (14)$$

for $i = 1, 2, \dots, d$ which was introduced in [22]. Note that this score has order-dependent weight and is normalized such that $\sum_{i \in [d]} r(i) = 1$. Computing every score $r(i)$, $i \in [d]$ provides an attribute ranking with respect to U showing the percentage that every variable adds to the variance of the approximation. We then conclude that if we have a good approximation $S(\mathcal{X}, \mathcal{I}(U))f$, the corresponding attribute ranking will be close to the attribute ranking of the function f .

3.2 Active Set

In this section we describe how to obtain a set of ANOVA terms U for approximation. We are still working with the scattered data $\mathcal{X} \subseteq \mathbb{R}^d$ and $\mathbf{y} \in \mathbb{R}^M$, $M \in \mathbb{N}$. The values \mathbf{y} may also contain noise. Our first step is to limit the variable interactions by a superposition threshold $d_s \in [d]$ which may have been estimated by known smoothness properties (or different a-priori information) or set to a sensible value if nothing is known. It is also possible to determine an optimal value through cross-validation. We choose the order-dependent parameters $N_{|\mathbf{u}|}$, $|\mathbf{u}| = 1, 2, \dots, d_s$, cf. (11), to obtain $\mathcal{I}(U(d, d_s))$ and with the procedure described in Section 3.1, the approximation $S(\mathcal{X}, \mathcal{I}(U(d, d_s)))f$.

From the approximation $S(\mathcal{X}, \mathcal{I}(U(d, d_s)))f$ we can then calculate the global sensitivity indices $\varrho(\mathbf{u}, S(\mathcal{X}, \mathcal{I}(U(d, d_s)))f)$, $\mathbf{u} \in U(d, d_s)$, and an attribute ranking $r(i)$, $i \in [d]$, see (14). Then we are able to apply the strategies proposed in [22] to truncate terms from the set $U(d, d_s)$.

One obvious method is the truncation of an entire variable x_i , $i \in [d]$, if the attribute ranking $r(i)$ shows that its influence is insignificant. Specifically, that would translate to an active set $U^* = \{\mathbf{u} \in U(d, d_s) : i \notin \mathbf{u}\}$. This leads to a reduction in dimensionality of the problem and greatly simplifies our model.

A different method is *active set thresholding* where we choose a threshold vector $\varepsilon \in (0, 1)^{d_s}$ and reduce the ANOVA terms to the set

$$U^*(\varepsilon) := \{\mathbf{u} \in U(d, d_s) : \varrho(\mathbf{u}, S(\mathcal{X}, \mathcal{I}(U(d, d_s)))f) > \varepsilon_{|\mathbf{u}|}\}.$$

Here, $\varepsilon_{|\mathbf{u}|}$ denotes the $|\mathbf{u}|$ th entry of the vector ε . The parameter vector ε allows control over how much of the variance may be sacrificed in order to simplify the model function.

In summary, it is necessary to interpret the information from the approximation $S(\mathcal{X}, \mathcal{I}(U(d, d_s)))f$ and decide on strategies for truncating the set of ANOVA terms. One may also use different strategies to obtain an active set or any combination of the multiple approaches, see e.g. [23, 22]. Of course, it is also possible to repeat the procedure multiple times, i.e., through cross-validation.

4 Forest Fire Prevention

We now apply the previously described method to the data set [1] from the UC Irvine machine learning repository. The dataset contains information about forest fires in the Montesinho national park in the Trás-os-Montes northeast region of Portugal. The data was collected from 2002 to 2003. Specifically, we have $d = 12$ attributes about the fires and the target variable is the area of the forest that was destroyed by it. If we obtain an efficient model, it can be possible to predict the risk for a future forest fire using parameters that can be easily measured. This information can then be used to prepare appropriate countermeasures. The data set has been thoroughly considered in [4] and we compare to the results they obtained.

group	name	description
spatial (S)	X	x-coordinate (1 to 9)
	Y	y-coordinate (1 to 9)
temporal (T)	month	month of the year (1 to 12)
	day	day of the week (1 to 7)
FWI	FFMC	FFMC code
	DMC	DMC code
	DC	DC code
	ISI	ISI index
meteorological (M)	temp	outside temperature in °C
	RH	outside relative humidity in %
	wind	outside wind speed in km/h
	rain	outside rain in mm/m ²

Table 1: Attributes and their corresponding groups

We group the 12 attributes into 4 categories as in [4], i.e., spatial, temporal, FWI system, and meteorological data, see Table 1. The spatial attributes describe the spatial location of the fire in a 9 by 9 grid of our considered region. The temporal attributes are the month of the year and the day of the week when the fire occurred. The forest fire weather index (FWI), cf. [27], is the Canadian system for rating fire danger and the datasets collects several components of it. Moreover, four meteorological attributes which are used by the FWI index were selected. The target variable describes the area that was burned by the fire.

In terms of pre-processing, we apply a Z-score transformation to the variables and the logarithmic transformation $\log(1 + \cdot)$ to the burned area. The Z-score transformation achieves that our data has zero mean and unit variance. The logarithmic transformation on the target is necessary since it shows a positive skew with a large number of fires that have a small size. We denote the data $(\mathcal{X}, \mathbf{y})$ with $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M\} \subseteq \mathbb{R}^{12}$, $M = 517$, and $\mathbf{y} \in \mathbb{R}^M$. In the following subsections, we do not use all of the variables, but build models based only on some groups as denoted in Table 1, e.g., **STM** says that we use spatial, temporal and meteorological attributes without the FWI.

Table 2 shows the overall results of our experiment (ANOVA) combined with the benchmark data from [4]. Each value, our ANOVA results as well as the others, were obtained by averaging over executing a 10-fold cross-validation 30 times. This results in a total of 300 experiments. We used a superposition threshold of $d_s = 2$, cf. (10), and therefore needed to detect optimal choices for the parameters N_1 and N_2 from (11), see Table 3. Every experiment utilized 90% of the data as training set $(\mathcal{X}_{\text{train}}, \mathbf{y}_{\text{train}})$ and 10% of the data as test set $(\mathcal{X}_{\text{test}}, \mathbf{y}_{\text{test}})$. The best performing model was selected based on the mean absolute deviation

$$\text{MAD} = \frac{1}{|\mathcal{X}_{\text{test}}|} \sum_{i=1}^{|\mathcal{X}_{\text{test}}|} |(\tilde{\mathbf{y}})_i - (\mathbf{y}_{\text{test}})_i| \quad (15)$$

with $\tilde{\mathbf{y}}$ the predictions of our model for the data points in the test set $\mathcal{X}_{\text{test}}$. As a second error measure, we use the root mean square error

$$\text{RMSE} = \frac{1}{\sqrt{|\mathcal{X}_{\text{test}}|}} \sqrt{\sum_{i=1}^{|\mathcal{X}_{\text{test}}|} |(\tilde{\mathbf{y}})_i - (\mathbf{y}_{\text{test}})_i|^2}. \quad (16)$$

We are able to outperform the previously applied method for every subset of attributes in both MAD and RMSE error. Notably, the difference in the RMSE that penalizes larger deviations in the burned area stronger than the MAD is much more significant.

model	attribute selection			
	S T FWI	S T M	FWI	M
Naive	18.61 (63.7)	18.61 (63.7)	18.61 (63.7)	18.61 (63.7)
MR	13.07 (64.5)	13.04 (64.4)	13.00 (64.5)	13.01 (64.5)
DT	13.46 (64.4)	13.43 (64.6)	13.24 (64.4)	13.18 (64.5)
RF	13.31 (64.3)	13.04 (64.5)	13.38 (64.0)	12.93 (64.4)
NN	13.09 (64.5)	13.92 (68.9)	13.08 (64.6)	13.71 (66.9)
SVM	13.07 (64.7)	13.13 (64.7)	12.86 (64.7)	12.71 (64.7)
ANOVA	12.75 (45.77)	12.81 (46.7)	12.76 (46.09)	12.65 (45.69)

Table 2: MAD and RMSE (in brackets) for the best performing model in the corresponding attribute subset (underline - overall best result, **bold** - best result for this selection).

attribute selection	N_1	N_2	$ I $	λ
S T FWI	2	6	149	e^9
S T M	2	10	261	e^{10}
FWI	2	4	23	e^8
M	2	8	47	e^7

Table 3: Optimal parameter choices for the experiments from Table 2.

While we replicated the setting of [4] for benchmark purposes, it remains our goal to identify the most important attributes for the detection of forest fires. Therefore, we now use all 12 attributes of the dataset in obtaining our approximation and subsequently interpret the results. Figure 3 shows the attribute ranking $r(i)$, $i = 1, 2, \dots, 12$, and the global sensitivity indices $\rho(\mathbf{u}, S(\mathcal{X}_{\text{train}}, \mathcal{I}(U(12, 2)))f)$, $\mathbf{u} \in U(12, 2)$, after computing an approximation with $N_1 = N_2 = 2$ and $\lambda = 1.0$.

The attributes 3, 7, and 9 are clearly the most important. They represent the month of the year (3), the DC code of the FWI (7) and the outside temperature (9). Using only these three attributes and superposition threshold $d_s = 2$, we computed an approximation with $N_1 = 2$, $N_2 = 10$, and $\lambda = e^8$. The resulting model yielded a MAD of 12.64 and a RMSE of 45.57 with 30 times of 10-fold cross validation as before. In summary, we know that the most important information of our problem is contained in only three attributes and we also obtained a better performing model using only those three attributes.

Acknowledgments

The authors thank their colleagues in the research group SAIE for valuable discussions on the contents of this paper. Daniel Potts acknowledges funding by Deutsche Forschungsgemeinschaft (German Research Foundation) – Project-ID




Figure 3: Explainable approximation results with all twelve attributes using $N_1 = N_2 = 2$ and $\lambda = e^8$.

416228727 – SFB 1410. Michael Schmischke is supported by the German Federal Ministry of Education and Research grant 01|S20053A.

References

- [1] Forest Fires. UCI Machine Learning Repository, 2008.
- [2] F. Bartel, D. Potts, and M. Schmischke. Grouped transformations in high-dimensional explainable ANOVA approximation. *ArXiv e-prints 2010.10199*, 2020.
- [3] R. Caflisch, W. Morokoff, and A. Owen. Valuation of mortgage-backed securities using Brownian bridges to reduce effective dimension. *J. Comput. Finance*, 1(1):27–46, 1997.
- [4] P. Cortez and A. Morais. A Data Mining Approach to Predict Forest Fires using Meteorological Data. In J. N. et al., editor, *New Trends in Artificial Intelligence, 13th EPIA 2007 - Portuguese Conference on Artificial Intelligence*, pages 512–523, Guimarães, Portugal, Dec. 2007. APPIA.
- [5] I. G. Graham, F. Y. Kuo, J. A. Nichols, R. Scheichl, C. Schwab, and I. H. Sloan. Quasi-Monte Carlo finite element methods for elliptic PDEs with lognormal random coefficients. *Numer. Math.*, 131(2):329–368, 2014.
- [6] I. G. Graham, F. Y. Kuo, D. Nuyens, R. Scheichl, and I. H. Sloan. Circulant embedding with QMC: analysis for elliptic PDE with lognormal coefficients. *Numer. Math.*, 140(2):479–511, 2018.
- [7] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning - Data Mining, Inference, and Prediction*. Springer Science & Business Media, Berlin Heidelberg, 2013.
- [8] M. Holtz. *Sparse grid quadrature in high dimensions with applications in finance and insurance*, volume 77 of *Lecture Notes in Computational Science and Engineering*. Springer-Verlag, Berlin, 2011.
- [9] L. Kämmerer, T. Ullrich, and T. Volkmer. Worst case recovery guarantees for least squares approximation using random samples. *Constr. Approx.*, 54:295–352, 2021.
- [10] 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.
- [11] F. Y. Kuo and D. Nuyens. Application of Quasi-Monte Carlo Methods to Elliptic PDEs with Random Diffusion Coefficients: A Survey of Analysis and Implementation. *Found. Comput. Math.*, 16(6):1631–1696, 2016.
- [12] F. Y. Kuo, C. Schwab, and I. H. Sloan. Quasi-Monte Carlo finite element methods for a class of elliptic partial differential equations with random coefficients. *SIAM J. Numer. Anal.*, 50:3351 – 3374, 2012.
- [13] F. Y. Kuo, I. H. Sloan, G. W. Wasilkowski, and H. Woźniakowski. On decompositions of multivariate functions. *Math. Comp.*, 79(270):953–966, 2009.
- [14] R. Liu and A. B. Owen. Estimating mean dimensionality of analysis of variance decompositions. *J. Amer. Statist. Assoc.*, 101(474):712–721, 2006.
- [15] M. Moeller and T. Ullrich. L_2 -norm sampling discretization and recovery of functions from RKHS with finite trace. *Sampl. Theory Signal Process. Data Anal.*, 19(2), 2021.

-
- [16] R. Nasdala and D. Potts. Transformed rank-1 lattices for high-dimensional approximation. *Electron. Trans. Numer. Anal.*, 53:239–282, 2020.
- [17] J. A. Nichols and F. Y. Kuo. Fast CBC construction of randomly shifted lattice rules achieving $\mathcal{O}(n^{-1+d})$ convergence for unbounded integrands over \mathbb{R}^s in weighted spaces with POD weights. *J. Complexity*, 30(4):444–468, Aug. 2014.
- [18] A. B. Owen. *Monte Carlo theory, methods and examples*. 2013.
- [19] C. C. Paige and M. A. Saunders. LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Software*, 8:43–71, 1982.
- [20] G. Plonka, D. Potts, G. Steidl, and M. Tasche. *Numerical Fourier Analysis*. Applied and Numerical Harmonic Analysis. Birkhäuser, 2018.
- [21] D. Potts and M. Schmischke. Approximation of high-dimensional periodic functions with fourier-based methods. *SIAM J. Numer. Anal.*, 59(5):2393–2429, 2021.
- [22] D. Potts and M. Schmischke. Interpretable approximation of high-dimensional data. *SIAM J. Math. Data Sci.* (accepted), 2021.
- [23] D. Potts and M. Schmischke. Learning multivariate functions with low-dimensional structures using polynomial bases. *J. Comput. Appl. Math.*, 403:113821, 2022.
- [24] H. Rabitz and O. F. Alis. General foundations of high dimensional model representations. *J. Math. Chem.*, 25:197–233, 1999.
- [25] I. M. Sobol. On sensitivity estimation for nonlinear mathematical models. *Keldysh Applied Mathematics Institute*, 1:112–118, 1990.
- [26] I. M. Sobol. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Math. Comput. Simulation*, 55(1-3):271–280, 2001.
- [27] S. W. Taylor and M. E. Alexander. Science, technology, and human factors in fire danger rating: the canadian experience. *Int. J. Wildland Fire*, 15(1):121, 2006.