

Scalable Higher-Order Tensor Product Spline Models

David Rügamer

Department of Statistics, LMU Munich
Munich Center for Machine Learning

Abstract

In the current era of vast data and transparent machine learning, it is essential for techniques to operate at a large scale while providing a clear mathematical comprehension of the internal workings of the method. Although there already exist interpretable semi-parametric regression methods for large-scale applications that take into account non-linearity in the data, the complexity of the models is still often limited. One of the main challenges is the absence of interactions in these models, which are left out for the sake of better interpretability but also due to impractical computational costs. To overcome this limitation, we propose a new approach using a factorization method to derive a highly scalable higher-order tensor product spline model. Our method allows for the incorporation of all (higher-order) interactions of non-linear feature effects while having computational costs proportional to a model without interactions. We further develop a meaningful penalization scheme and examine the induced optimization problem. We conclude by evaluating the predictive and estimation performance of our method.

1 INTRODUCTION

Two of the core principles of statistical regression models are additivity and linearity of the predictors. These properties allow estimated feature effects to be easily interpreted, which also led to (revived) interest in such models in the machine learning and information retrieval community (Chen et al., 2017; Wang et al.,




Figure 1: Comparison of memory consumption (first row) and time consumption (second row) between the state-of-the-art big additive model (BAM) implementation (in red) and our proposal (in gray) when fitting a model for all $\binom{p}{2}$ tensor product splines using different numbers of features p (x-axis) and observations (columns).

2020; Zhuang et al., 2021; Chang et al., 2021). A frequently used and cited example of an interpretable yet flexible statistical regression model is the generalized additive model (GAM; Hastie and Tibshirani, 2017; Wood, 2017). Using basis functions to approximate non-linear functions, these models can represent non-linear feature effects in one or a moderate number of dimensions. Applying this principle in settings with many features and higher-order interactions, however, comes with considerable downsides. For univariate non-linear effects, the number of basis functions M for each feature typically lies in the range of 10 to 20 and needs to be evaluated prior to model fitting. Representing and fitting all available features using basis functions will not only result in a notable increase in training time, but also requires a considerable amount of additional memory. In higher dimensions D , these problems carry even more weight as D -variate non-linear representations are typically constructed using Kronecker or tensor product splines (TPS), i.e., a (row-wise) Kronecker product of all involved bases. This results in computational costs of $\mathcal{O}((pM)^D)$ for TPS models with p features. Computational feasibility is thus one of the main reasons statistical applications are often restricted to only uni- and bivariate (tensor product) splines. While several approaches to tackle this problem have been proposed (e.g., Wood et al.,

2017), existing solutions still either suffer from extensive memory or runtime costs.

Our Contribution: In order to efficiently scale additive models in higher dimensions (cf. also Figure 1), we propose an approach for modeling higher-order TPS that only has linear complexity in the number of features p based on the idea of factorization machines (FMs; Rendle, 2010). While this effectively addresses existing scaling problems of GAMs, our approach also extends (higher-order) factorization machines by allowing for non-linear relationships. In addition to deriving the resulting computational complexity, we also propose an efficient way of computing the model, suggest a suitable penalization scheme and provide an optimization routine for our approach. Our experimental section demonstrates that models with higher-order TPS work well in practice, show desirable approximation behavior, and yield competitive results on real-world data.

2 RELATED LITERATURE

GAMs and TPS GAMs extend generalized linear models (Nelder and Wedderburn, 1972) by allowing feature effects to be non-linear, typically achieved by using a spline bases representation. Next to basic principles (see, e.g., Hastie and Tibshirani, 2017; Wood, 2017), many extensions of GAMs have been discussed in the literature. In order to include non-linear functions of more than one variable, various options exist, e.g., by using spline bases in multiple dimensions. Wood (2006) proposed a flexible way of constructing multivariate non-linear functions in GAMs using TPS as an alternative option, which we will outline in more detail in Section 3.1.2. Although GAM software is usually optimized in terms of efficiency, computational costs can be a bottleneck for large data sets or complex model formulations. While there exist approaches work that allows GAM estimation for data with many observations (Wood et al., 2017), GAMs still scale unfavorably with many features or feature interactions. Recent approaches (e.g., Rügamer et al., 2023c; Rügamer, 2023) suggest fitting structured regression models as (part of) a neural network. This can result in a better space complexity in situations with many data points and allows for more flexibility in the additive predictors of models.

Factorization Approaches Similar to GAMs, factorization approaches have been studied extensively. Popularized for recommender systems, different (matrix) factorization approaches have been proposed in the early 2000s (see, e.g., Srebro et al., 2004; Adomavicius and Tuzhilin, 2005; Koren et al., 2009) and are still considered state-of-the-art in terms of performance

and efficiency (Rendle et al., 2020; Jin et al., 2021). Factorization approaches are also used as part of additive models in order to improve scalability (Rügamer et al., 2023a). Closely related to matrix factorization are factorization machines (FMs; Rendle, 2010). FMs are based on a linear model formulation with pairwise interactions between all features and use a factorization trick to overcome unfavorable scaling when the number of features is large. Various extensions have been developed over the past years, including convex FMs (Blondel et al., 2015) and an efficient calculation of higher-order FMs (HOFMs; Blondel et al., 2016). Other extensions include boosted FMs (Yuan et al., 2017), FMs with special personalized feature treatment (Chen et al., 2019) or interaction-aware FMs (Hong et al., 2019). Similar to our proposal, Lan and Geng (2019) use a non-parametric subspace feature mapping to encode interactions and account for non-linearity, but rely on binning the features.

Boosting Apart from FMs and GAMs, various other approaches exist that model non-linearity and/or interactions while preserving an additivity structure of the model. One of the most prominent approaches in machine learning that combines additivity and predictive performance is boosting. Already in the seminal work of Friedman (Friedman, 2001), (functional gradient) boosting was proposed to optimize additive models (potentially with interactions). This idea is used, e.g., for GAMboost (Hothorn et al., 2010) and lays the foundation for other interpretable boosting frameworks such as GA²M (Lou et al., 2013), explainable boosting machines (Nori et al., 2019), or functional regression boosting (Brockhaus et al., 2020). While extensions for scalability have been suggested (see, e.g. Schalk et al., 2023), boosting GAMs with all possible interaction terms remains challenging.

3 BACKGROUND

We first introduce our notation and then give a short introduction into GAMs in Section 3.1. For more details, see, e.g., Wood (2017).

In the following, we write scalar values in small or capital letters without formatting, vectors in small bold letters, matrices in capital bold letters, and tensors using fraktur typeface, e.g., \mathfrak{X} . Calligraphic letters will have different meaning depending on the context, while \mathcal{O} is reserved to describe the complexity of calculations in terms of computing time or memory. For better readability, we will denote the sequence from $1, \dots, x$ with $[x]$. We further use \otimes for the Kronecker product.

3.1 Generalized Additive Models

Given the response random variable Y and p features $\mathbf{x} = (x_1, \dots, x_p)$, an additive model with linear and non-linear effects for all features assumes the following relationship:

$$Y = \eta(\mathbf{x}) + \varepsilon = \alpha_0 + \sum_{j=1}^p x_j \alpha_j + \sum_{j=1}^p f_j(x_j) + \varepsilon, \quad (1)$$

where $\alpha_0, \alpha_1, \dots, \alpha_p$ are linear regression coefficients, f_1, \dots, f_p univariate non-linear functions, $\varepsilon \sim \mathcal{N}(0, \sigma^2)$ is a zero-mean Gaussian random variable with variance $\sigma^2 > 0$ and η the model predictor. GAMs, the generalization of additive models, replace the distribution assumption in (1) using a more general distribution by assuming that $Y|\mathbf{x}$ has some exponential family distribution and $\mathbb{E}(Y|\mathbf{x}) = h(\eta(\mathbf{x}))$ for some monotonic (response) function h . A prediction $\hat{y} = h(\hat{\eta}(\mathbf{x}))$ for the observed value y in GAMs is formed by estimating the regression coefficients and functions f_j . GAMs can be optimized using (different types of) maximum likelihood estimation. Alternatively, using the negative log-likelihood as (convex) loss function ℓ , their optimization can also be framed as an empirical risk minimization problem.

Since linear effects α_j can be incorporated in the functions f_j , we will drop the linear model part in the following. The non-linear functions f_j in GAMs are usually approximated using a (spline) basis representation, i.e.,

$$f_j(x_j) \approx \sum_{m=1}^{M_j} B_{m,j}(x_j) \beta_{m,j} = \mathbf{B}_j^\top \boldsymbol{\beta}_j, \quad (2)$$

where $B_{m,j}$ are pre-defined basis functions (e.g., truncated polynomials or B-splines) and $\beta_{m,j}$ the corresponding basis coefficients. We summarize all basis functions and coefficients using $\mathbf{B}_j = (B_{1,j}, \dots, B_{M_j,j}) \in \mathbb{R}^{M_j}$ and $\boldsymbol{\beta}_j = (\beta_{1,j}, \dots, \beta_{M_j,j})^\top \in \mathbb{R}^{M_j}$, respectively. To enforce smoothness of the functions f_j , the $\boldsymbol{\beta}_j$ coefficients are typically estimated using a smoothness penalty.

3.1.1 Smoothness Penalties

One of the most common approaches to estimate smooth functions f_j is to employ a difference penalty for successive basis coefficients $\beta_{m,j}, \beta_{m+1,j}$ of basis functions $B_{m,j}(x), B_{m+1,j}(x)$, which penalizes deviating behavior in neighboring basis functions. The penalty term for the penalized loss function is then given by $\mathcal{P} = \sum_{j=1}^p \lambda_j \int (f_j''(x))^2 dx$, which is a trade-off between goodness-of-fit and roughness of the func-

tions f_j . The penalized loss can be written as

$$\ell(y, \hat{y}) + \sum_{j=1}^p \lambda_j \boldsymbol{\beta}_j^\top \mathbf{P}_j \boldsymbol{\beta}_j, \quad (3)$$

where $\mathbf{P}_j \in \mathbb{R}^{M_j \times M_j}$ is a squared penalty matrix depending on the evaluated basis \mathbf{B}_j for the j th feature and usually penalizes first or second differences in the coefficients $\boldsymbol{\beta}_j$.

GAMs also allow for higher dimensional non-linear functions, e.g., bivariate smooth terms $f_{k,l}(x_k, x_l)$. A common approach for their construction are tensor product splines.

3.1.2 Tensor Product Splines

While there are various approaches to model smooth functions of several features, TPS constructed from marginal univariate bases constitute an attractive option. The resulting smooth terms are very flexible, scale-invariant, relatively low rank as well as easy to construct and interpret (see Wood, 2006). For a model with all $\binom{p}{2}$ possible bivariate effects, the TPS part is given by

$$\begin{aligned} \sum_{k=1}^p \sum_{l=k+1}^p f_{k,l}(x_k, x_l) &\approx \sum_{k=1}^p \sum_{l=k+1}^p (\mathbf{B}_k \otimes \mathbf{B}_l) \boldsymbol{\beta}_{k,l} \\ &= \sum_{k=1}^p \sum_{l=k+1}^p \sum_{m=1}^{M_k} \sum_{o=1}^{O_l} B_{m,k}(x_k) B_{o,l}(x_l) \beta_{m,k,o,l} \end{aligned} \quad (4)$$

with univariate spline basis functions $B_{m,k}, B_{o,l}$ and basis coefficients $\beta_{m,k,o,l}$, summarized in $\boldsymbol{\beta}_{k,l} \in \mathbb{R}^{M_k O_l}$. Bivariate TPS are penalized using

$$\mathcal{P}(f_{k,l}) = \int_{x_k, x_l} \lambda_k (\partial^2 f / \partial x_k^2)^2 + \lambda_l (\partial^2 f / \partial x_l^2)^2 dx_k dx_l,$$

which can be written as

$$\mathcal{P}(f_{k,l}) = \boldsymbol{\beta}_{k,l}^\top (\lambda_k \mathbf{P}_k \oplus \lambda_l \mathbf{P}_l) \boldsymbol{\beta}_{k,l}, \quad (5)$$

where \oplus is the Kronecker sum operator.

This principle can be generalized to D -variate smooths for variables $\mathcal{J} := \{j_1, \dots, j_D\}$, which are approximated by

$$f_{j_1, \dots, j_D}(x_{j_1}, \dots, x_{j_D}) \approx (\otimes_{j \in \mathcal{J}} \mathbf{B}_j) \boldsymbol{\beta}_{\mathcal{J}}, \quad (6)$$

where $\boldsymbol{\beta}_{\mathcal{J}} \in \mathbb{R}^{\prod_{t=1}^D M_{j_t}}$ contains the coefficients for all combinations of the D basis functions. The corresponding penalty term for (6) is constructed analogously to (5) by $\boldsymbol{\beta}_{\mathcal{J}}^\top (\oplus_{j \in \mathcal{J}} \lambda_j \mathbf{P}_j) \boldsymbol{\beta}_{\mathcal{J}}$.

In the following, we assume that the number of spline basis functions is roughly equal across different features and use $M := \max_{k,l} \{M_k, O_l\}$ to denote the spline basis in multivariate splines that uses the most basis functions.

4 SCALABLE HIGHER-ORDER TENSOR PRODUCT SPLINE MODELS

As can be directly inferred from (4), the cost of fitting a bivariate TPS is $\mathcal{O}(p^2M^2)$. While M is usually kept fixed and of moderate size (e.g., $M = 10$), this implies that models will be increasingly expensive for both a growing number of basis evaluations and number of features p . For models with (up to) D -variate TPS, the computational cost increases to $\mathcal{O}(p^D M^D)$. This makes GAMs infeasible both in terms of computing time and also in terms of memory storage.

4.1 Additive Factorization Machines

To overcome the unfavorable scaling of GAMs with many (or higher-order) TPS, we introduce additive factorization machines (AFMs). Based on the idea of factorization machines, we approximate $f_{k,l}(x_k, x_l)$ in (4) by $\phi_{k,l}(x_k, x_l)$ defined as

$$\phi_{k,l}(x_k, x_l) = \sum_{m=1}^{M_k} \sum_{o=1}^{O_l} B_{m,k}(x_k) B_{o,l}(x_l) \sum_{f=1}^F \gamma_{m,k,f} \gamma_{o,l,f}, \quad (7)$$

where $\gamma_{m,k,f} \in \mathbb{R}$ are latent factors approximating the joint effect $\beta_{m,k,o,l}$. When approximating every bivariate interaction term in (4) with the term defined in (7), we can derive the following representation.

Lemma 1 (AFM Representation). *The approximation of (7) using (4) can be written as*

$$\sum_{k=1}^p \sum_{l=k+1}^p f_{k,l}(x_k, x_l) \approx \frac{1}{2} \sum_{f=1}^F \left\{ \left[\sum_{k=1}^p \varphi_{k,f} \right]^2 - \sum_{k=1}^p \varphi_{k,f}^2 \right\}, \quad (8)$$

with $\varphi_{k,f} = \sum_{m=1}^{M_k} B_{m,k}(x_k) \gamma_{m,k,f}$.

As a direct result of Lemma 1, we obtain the scaling of computing AFMs.

Proposition 1 (Linear Scaling of AFMs). *Computations for AFMs scale with $\mathcal{O}(pMF)$.*

Lemma 1 and Proposition 1 are natural extensions of linearity results from FMs. A proof of the corollary is provided in the Supplementary Material. Roughly speaking, the factorization trick from FMs also works for AFMs in a similar manner as the additional basis function dimension only depends on the respective feature dimension. In particular, this means that AFMs scale linearly both in the number of features p and the spline basis dimension M . Another direct result of this representation and noteworthy property unique to

AFMs is given in the following proposition for a data set of n observations.

Proposition 2 (Basis Evaluations in AFMs). *If every feature in AFMs is represented by only one basis, it suffices to evaluate all univariate basis functions once for each feature and the memory costs for storing all features are $\mathcal{O}(npM)$.*

While this seems inconspicuous at first glance, a naïve approach for bivariate models requires storing $\mathcal{O}(np^2M^2)$ entries, which is infeasible if n or p is large. In contrast, AFMs only require the same amount of storage as for a univariate spline model. Note that the number of parameters also (linearly) increases with F and needs to be taken into account for the total required storage. However, Proposition 2 specifically looks at the costs of storing basis evaluated features in memory as this can be a storage bottleneck during the pre-processing of GAMs.

4.2 Additive Higher-Order Factorization Machines

We now extend previous results to the general case of a D -variate interaction GAM. By analogy to higher-order FMs (Blondel et al., 2015), we refer to our approximation as *additive higher-order factorization machines* (AHOFMs). The goal of AHOFMs is to provide a scalable version of a TPS model which includes multivariate splines up to D -variate smooths, i.e.,

$$\eta(\mathbf{x}) = \alpha_0 + \sum_{j=1}^p f_j(x_j) + \sum_{j'>j} f_{j',j}(x_{j'}, x_j) + \dots + \sum_{j_D > \dots > j_1} f_{j_1, \dots, j_D}(x_{j_1}, \dots, x_{j_D}). \quad (9)$$

Assume a TPS representation for all smooth terms in (9) with a maximum number of M basis functions. Then, the cost of computing only the last term in (9) is already $\mathcal{O}(p^D M^D)$ (and analogous for memory costs).

Inspired by Vieta’s formula and the ANOVA kernel, we can derive a similar result as given in Blondel et al. (2016) to reduce the cost of computing the d th degree term in AHOFMs. We will make the degree d explicit for γ and φ using the superscript (d) .

Definition 1 (Additive Higher-order Term (AHOT)). The f th additive higher-order term (AHOT) of degree $2 \leq d \leq D$ in AHOFMs is given by

$$\Phi_f^{(d)} = \sum_{j_d > \dots > j_1} \prod_{t=1}^d \sum_{m=1}^{M_{j_t}} B_{m,j_t}(x_{j_t}) \gamma_{m,j_t}^{(d)}. \quad (10)$$

We use F_d AHOTs to approximate a d -variate smooth:

$$\sum_{f=1}^{F_d} \Phi_f^{(d)} \approx \sum_{j_d > \dots > j_1} f_{j_1, \dots, j_d}(x_{j_1}, \dots, x_{j_d}). \quad (11)$$

and estimate the D -variate TPS model (9) with an AHOFM of degree D , defined as follows.

Definition 2 (AHOFM of Degree D). The predictor $\eta(\mathbf{x})$ of an AHOFM of degree D is defined by

$$\alpha_0 + \sum_{j=1}^p B_{m,j}(x_j) \beta_{m,j} + \sum_{d=2}^D \sum_{f=1}^{F_d} \Phi_f^{(d)}. \quad (12)$$

The following corollary defines how to recursively describe all AHOTs for $d \geq 2$ using univariate spline representations $\varphi_{j,f}$ as defined in (8).

Lemma 2 (Representation AHOT of Degree d). Let $\Phi_f^{(0)} \equiv 1$ as well as $\Phi_f^{(1)} = \sum_{j=1}^p \varphi_{j,f}^{(d)}$. The degree $d \geq 2$ AHOT can be recursively defined by

$$\Phi_f^{(d)} = \frac{1}{d} \sum_{t=1}^d (-1)^{t+1} \Phi_f^{(d-t)} \left\{ \sum_{j=1}^p [\varphi_{j,f}^{(d)}]^t \right\}. \quad (13)$$

The recursive representation (13) allows us to efficiently calculate AHOTs of higher order. The corresponding proof can be found in the Supplementary Material. As another consequence of this representation, we have the following scaling properties.

Proposition 3 (Linear Scaling of AHOFMs). Computations for AHOFMs scale with $\mathcal{O}(pMFD + FD^2)$ where $\mathcal{F} = \sum_{d=1}^D F_d$.

Since D is usually small, computations again roughly scale linearly with the number of features, the basis, and the latent factor dimension. We also recognize that despite the increased dimension D , every feature basis has to be evaluated only once.

Proposition 4 (Basis Evaluations in AHOFMs). If every feature in AHOFMs is represented by only one basis, it suffices to evaluate all univariate basis functions once for each feature and the memory costs for storing all features are $\mathcal{O}(npM)$.

While the memory consumption also increases with \mathcal{F} when considering the storage of all γ parameters, Proposition 4 again focuses on the storage of all features after applying the basis evaluations. Similar to the kernel trick, higher-order features are not actually calculated and stored in memory, as AHOFMs only work on the (basis evaluated) univariate features independent of D .

In contrast to FMs and HOFMs, AHOFMs require additional considerations to enforce appropriate smoothness of all non-linear functions in (9) without impairing favorable scaling properties.

4.3 Penalization and Optimization

Learning the smoothness of many (higher-order) TPS is an additional challenge discussed below.

4.3.1 Penalization

Following the penalization scheme of TPS described in Section 3.1.2, we propose a smoothness penalization for AHOFMs based on the penalties of involved marginal bases $B_{m,j}$ with corresponding difference penalty matrices \mathbf{P}_j . Let $\mathfrak{G}^{(d)}$ be the array of all coefficients $\gamma_{m,j,f}^{(d)}$ for all $m \in [M], j \in [p], f \in [F_d]$. For simplicity¹, we assume $M_j \equiv M$, so that $\mathfrak{G}^{(d)} \in \mathbb{R}^{M \times p \times F_d}$. Further, let $\mathfrak{G}_{[D]} = \mathfrak{G}^{(1)}, \dots, \mathfrak{G}^{(D)}$ and $\boldsymbol{\gamma}_{j,f}^{(d)} = (\gamma_{1,j,f}^{(d)}, \dots, \gamma_{M,j,f}^{(d)})^\top$, i.e., the Mode-1 (column) fibers of $\mathfrak{G}^{(d)}$, and $\boldsymbol{\Theta} = (\alpha_0, \beta_1, \dots, \beta_p)$.

We define the penalty of AHOFMs as follows.

Definition 3 (AHOFM Penalty). The smoothing penalty of AHOFMs is defined as

$$\begin{aligned} \mathcal{P}(\mathfrak{G}_{[D]}, \boldsymbol{\Theta}) &= \sum_{j=1}^p \lambda_j \beta_j^\top \mathbf{P}_j \beta_j + \\ &\sum_{d=2}^D \sum_{f=1}^{F_d} \sum_{j=1}^p \lambda_{j,f}^{(d)} \boldsymbol{\gamma}_{j,f}^{(d)\top} \mathbf{P}_j \boldsymbol{\gamma}_{j,f}^{(d)}. \end{aligned} \quad (14)$$

Due to the independence assumption of all latent factors involved in every factorization, it is natural to only penalize univariate directions as expressed in Definition 3. A regularization that involves multiple dimensions would further result in a non-decomposable penalty w.r.t. the $\boldsymbol{\gamma}_{j,f}^{(d)}$ and make the optimization of AHOFMs more challenging (see Section 4.3.3 for details). The penalized optimization problem for n i.i.d. data points $(y_i, \mathbf{x}_i)_{i \in [n]}$ with $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,p})$ and loss function ℓ is then given by

$$\arg \min_{\mathfrak{G}_{[D]}, \boldsymbol{\Theta}} \sum_{i=1}^n \ell(y_i, \eta_i(\mathbf{x})) + \frac{1}{2} \mathcal{P}(\mathfrak{G}_{[D]}, \boldsymbol{\Theta}). \quad (15)$$

In (15), the smoothing parameters are considered to be tuning parameters. While it is possible in univariate GAMs to estimate the smoothing parameters $\lambda_1, \dots, \lambda_p$ directly, this becomes computational challenging for higher-order TPS due to the exponentially increasing amount of parameters. HOFMs avoid this combinatorial explosion of hyperparameters by setting all the parameters to the same value (Blondel et al., 2016). This, however, is not a meaningful approach for smoothing parameters as every smooth term can

¹Next to a ragged tensor definition that allows for a varying first dimension, padding the tensor can also be an option to always have M dimensions for every feature j .

potentially live on a completely different domain (e.g., $\lambda = 1$ can imply no penalization for one smooth, but maximum penalization for another term).

4.3.2 Scalable Smoothing

To derive a meaningful penalization in AHOFMs, we exploit the definition of *degrees-of-freedom* for penalized linear smoothers (Buja et al., 1989). Given a matrix of basis evaluations $\mathbf{B}_j \in \mathbb{R}^{n \times M_j}$ with entries $B_{m,j}(x_{i,j})$ and a squared penalty matrix \mathbf{P}_j for the penalization of differences in neighboring basis coefficients, there exists a one-to-one map between $\lambda_{j,f}^{(d)}$ and the respective degrees-of-freedom

$$df_{j,f}^{(d)}(\lambda_{j,f}^{(d)}) = \text{tr}(2\mathbf{H}_j(\lambda_{j,f}^{(d)}) - \mathbf{H}_j(\lambda_{j,f}^{(d)})^\top \mathbf{H}_j(\lambda_{j,f}^{(d)})) \quad (16)$$

with $\mathbf{H}_j(\lambda_{j,f}^{(d)}) = \mathbf{B}_j(\mathbf{B}_j^\top \mathbf{B}_j + \lambda_{j,f}^{(d)} \mathbf{P}_j)^{-1} \mathbf{B}_j^\top$. While the exact degrees-of-freedom only hold for a linear model with a single smooth term, this approach allows to define a meaningful a priori amount of penalization for all smooth terms by restricting their degrees-of-freedom to the same global $df^{(d)}$ value as follows.

Proposition 5 (Homogeneous AHOFM Smoothing). *Given a global $df^{(d)}$ value, an equal amount of penalization for all $\binom{p}{d}$ AHOTs in $\Phi_f^{(d)}$ is achieved by choosing $\lambda_{j,f}^{(d)}$ such that $df_{j,f}^{(d)}(\lambda_{j,f}^{(d)}) \equiv df^{(d)} \forall j \in [p], f \in [F_d]$.*

The Demmler-Reinsch orthogonalization (Ruppert et al., 2003) can be used to efficiently solve (16) for $\lambda_{j,f}^{(d)}$, i.e., calculate $\lambda_{j,f}^{(d)}$ based on a given value $df_{j,f}^{(d)}$. The orthogonalization involves the calculation of singular values \mathbf{s}_j of a squared $M_j \times M_j$ matrix. Once \mathbf{s}_j are computed, (16) can also be solved multiple times for different df values without additional costs (see Supplement C.1 for details). Moreover, as the factorization only requires univariate smooth terms, $\lambda_{j,f}^{(d)}$ can be calculated for every feature separately at the cost of $\mathcal{O}(M_j^3)$ due to our factorization approach. This cost is comparatively small compared to a computation for all features in a D -variate interaction term with $\mathcal{O}(M^{3D})$. Also note that in Proposition 5, $df_{j,f}^{(d)}(\lambda_{j,f}^{(d)})$ only needs to be calculated once for every j as all involved matrices in (16) are independent of f , and can be done prior to the optimization with no additional costs during training. Algorithm 4 summarizes the routine. Homogeneous AHOFM smoothing amounts to equally flexible non-linear interactions for every order- d AHOT and hence implies isotropic smoothing for all TPS. Given no a priori information on the non-linear interactions of all features, this is a natural choice. In contrast, if we choose different values for one or more features, i.e., $\exists j : df_{j,f} \neq df$, all TPS involving the j th feature are subject to anisotropic smoothing.




Figure 2: Estimated and true surfaces for 7 different TP splines (columns) and methods (rows) visualized by contour plots. Colors represent the partial effect values.

4.3.3 Optimization

In order to scale also for large numbers of observations, we propose a stochastic mini-batch gradient descent routine for the optimization of A(HO)FMs. We discuss the optimization problem of A(HO)FMs in Supplementary Material B and suggest a block-coordinate descent as an alternative optimization routine by showing that the problem in (15) is coordinate-wise convex in $\gamma_{j,f}^{(d)}$ (Lemma 3). In practice, however, different block-coordinate descent variants showed slow convergence and finding a good choice for hyperparameters such as the learning rate proved to be challenging. In contrast, sophisticated stochastic gradient descent routines such as Adam (Kingma and Ba, 2014) showed similar or even better results.

5 NUMERICAL EXPERIMENTS

We investigate the following research questions to empirically analyze theoretical findings and better understand the practical aspects of A(HO)FMs:

1. Do they provide a reasonable approximation of TPS surfaces (for increasing F)?
2. Do they provide a reasonable prediction performance (for increasing F) within the class of GAMs?
3. Can the theoretical model complexity be verified in practice?
4. How well do they perform when applied to real-world data? Which are influencing factors for pre-

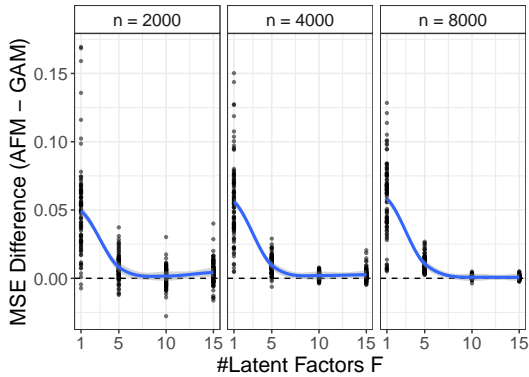


Figure 3: Estimation quality measures by the MSE difference between a GAM estimation (gold standard) and our proposal with different numbers of latent dimensions F (x-axis) and different numbers of observations (columns). Points correspond to different simulation replications and surfaces. A blue smoother function visualizes the trend in F . For smaller data sets, however, smaller F can be beneficial by inducing additional regularization.

dictive accuracy and can we identify limitations when compared to alternative fitting methods?

We will answer these questions in consecutive order in the following subsections numbered accordingly.

5.1 Estimation Performance

We first compare the estimation performance of our proposal with the SotA for fitting GAMs with TPS. More specifically, we simulate features and generate bivariate non-linear effects for every possible feature pair. The response is generated by adding random noise with a signal-to-noise ratio of 0.5 to the sum of all bivariate effects. We then compare the estimation performance of all feature effects qualitatively by inspecting the estimated non-linear effects visually (cf. Figure 2 for $n = 2000$), and quantitatively by computing the mean squared error (MSE) between the estimated and true surfaces (Figure 3). For $n \in \{2000, 4000, 8000\}$ and $p = 5$ (resulting in 10 bivariate effects), we run AFMs with $F \in \{1, 5, 15\}$. We repeat every setting 10 times with different random seeds. The GAM estimation can be thought of as a gold standard which is only subject to an estimation error, but no approximation error. In contrast, AFMs are also subject to an approximation error. Our quantitative analysis of results confirms this hypothesis. Figure 3 depicts the MSE differences for all analyzed settings to compare the estimation performance of AFMs and GAMs when calculating the average point-wise differences between the estimated bivariate surface and the true surface. Results suggest that for increasing F our approach will approach the estimation performance of GAMs. With more data (larger n), this effect becomes even more clear.

5.2 Prediction Performance

We follow the setup from the previous section and compare the prediction performance of the exact GAM and our approach to quantify the approximation error made by the factorization. Figure 4 depicts the results, confirming that the approximation error will tend to zero when increasing the number of observations n or latent factors F .




Figure 4: Prediction error of the GAM (lower bound) and our proposal with different numbers of latent dimensions F (colors) for different numbers of observations (x-axis).

5.3 Scalability

Our next experiment investigates the scaling behavior of our approach and compares it to the SotA implementation for big additive models (BAMs Wood et al., 2017). We simulate $p \in \{3, 6, 9, 12\}$ standard normal distributed features for $n \in \{6000, 12000, 18000\}$ observations and fit both BAM and an AFM to learn a GAM with TP splines for all possible combinations of the p features. Figure 1 (first page) summarizes the results by comparing the memory consumption (in megabytes) and the computing time (in seconds). The results reflect our initial motivation to propose AFMs. While both time and memory consumption grow exponentially in the number of features for BAMs, we observe a linear scaling for AFMs both for memory consumption and computation time.

5.4 Benchmark Studies

To assess prediction performance and the limitations of our approach when applied to real-world data, we compare A(HO)FMs against a variety of alternative approaches with similar properties. In particular, all methods use linear or spline feature effects. The similarity of comparison methods allows us to objectively reason about our methods' pros and cons. We compare all methods on commonly used benchmark data sets using 10 train-test splits and report average MSE values as well as their standard deviation. Further details on hyperparameters and benchmark data sets can be found in Section E.2 in the Supplementary Material.

Table 1: Average test performance (MSE \downarrow) with standard deviation in brackets for FM methods (columns) and data sets (rows) over 10 different train-test splits. The best method per data set is highlighted in bold.

	GAM	FM	HOFM($D = 3$)	AFM	AHOFM($D = 3$)
Airfoil	119.3 (2.490)	90.70 (45.55)	79.03 (51.29)	4.314 (0.576)	4.181 (0.699)
Concrete	9.768 (8.408)	10.55 (0.970)	10.56 (0.955)	6.100 (0.648)	6.127 (0.491)
Diabetes	142.1 (9.626)	143.7 (10.04)	143.3 (9.312)	57.12 (6.243)	63.65 (9.672)
Energy	3.409 (0.363)	7.487 (0.620)	7.480 (0.632)	3.137 (0.332)	3.183 (0.359)
Forest Fire	1.465 (0.163)	1.403 (0.105)	1.398 (0.107)	1.573 (0.325)	1.773 (0.307)
Naval	0.002 (0.000)	0.009 (0.001)	0.009 (0.001)	0.004 (0.002)	0.003 (0.001)
Yacht	3.027 (0.517)	8.883 (0.990)	8.898 (0.985)	2.401 (0.618)	1.743 (0.624)

5.4.1 Ablation Study

We first compare A(HO)FMs with their linear counterpart (HO)FMs and univariate GAMs. This study will provide statements to what extent A(HO)FMs’ performance is driven by non-linearity and interactions. For this, we use three different numbers of latent dimensions ($F \in \{1, 5, 15\}$) and fix the degrees-of-freedom $df = 15$. We implement all methods in a unifying framework in TensorFlow (Abadi et al., 2016), use the Adam optimizer with default settings for optimization and find the best epoch via early stopping based on a 15%-validation split. Table 1 summarizes the results when choosing the best-performing hyperparameter set per method and data set. Our results show that A(HO)FMs perform well and indicate that the extension to non-linearities is in some cases beneficial (Energy, Naval, Yacht), and the combination of non-linearity and interaction can improve performance notably (for Airfoil, Concrete, Diabetes, and Yacht).

5.4.2 Comparison with Sparse Approaches

As A(HO)FMs model all splines and possible interactions thereof, the natural question arises if this can also be a limitation. To examine this question, we compare our approach against methods that still represent an additive smooth model but allow for variable selection. The results of this study thus additionally help in better understanding A(HO)FMs and can pave the way for future research by highlighting potential improvements to the presented approach. More specifically, we compare our method against boosting with univariate splines (GAMBoost; Hothorn et al., 2010) and boosting with univariate splines and bivariate tensor products (GA²MBoost). These methods optimize GAMs using a greedy coordinate-wise algorithm that inherently provides variable selection when used in combination with early stopping. We use the same settings as in the previous study for A(HO)FMs and select the best stopping iteration for boosting using also the same train/validation-split as for the FM approaches. Table 2 summarizes the results. As before, A(HO)FMs perform well, in particular for data sets where inter-

 Table 2: Average test performance (MSE \downarrow) with standard deviation in brackets for different methods (columns) and data sets (rows) over 10 different train-test splits. The best method is highlighted in bold.

	GAMBoost	GA ² MBoost	A(HO)FM
Airfoil	79.03 (51.29)	4.628 (0.213)	4.181 (0.699)
Concrete	6.280 (0.486)	6.292 (0.586)	6.100 (0.648)
Diabetes	55.18 (5.020)	55.16 (5.618)	57.12 (6.243)
Energy	3.780 (0.427)	3.759 (0.368)	3.137 (0.332)
ForestF	1.388 (0.113)	1.397 (0.098)	1.573 (0.325)
Naval	0.012 (0.000)	0.003 (0.001)	0.003 (0.001)
Yacht	1.532 (0.546)	1.504 (0.504)	1.743 (0.624)

actions have been found beneficial in the previous ablation study (e.g., for Airfoil). Further, we see that boosting shows slightly better performance for three out of the 7 data sets, suggesting that feature selection can potentially improve prediction performance. As model-based boosting methods usually do not scale well with the number of features, a sparse A(HO)FM approach could thus be a promising future research direction. We note, however, that due to the nature of approximation used in A(HO)FMs (i.e., the shared latent factors across all interactions), incorporating a sparsity-inducing penalty is not straightforward.

6 SUMMARY AND OUTLOOK

We presented an additive model extension of HOFMs for scalable higher-order smooth function estimation based on TPS. The proposed approach allows fitting GAMs with D -variate smooth terms at costs similar to a univariate GAM. Our simulation studies showed that these machines approximate TPS surfaces very well and can match the exact GAM both in estimation and prediction performance. A(HO)FMs thereby not only allow to fit higher-order GAMs but also make additive models more competitive in their predictive performance shown in our benchmark study. A promising future research direction therefore is the sparsification of our proposed approach. This would not only allow for improving prediction performance as suggested

by our benchmark study but is also closely related to the interpretability of A(HO)FMs, which we discuss in Supplement D in more detail. A sparsification approach would require being compatible with the first-order optimization techniques used in this work. Hence, a possible interesting direction could be an optimization transfer (Kolb et al., 2023).

References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, pages 265–283, 2016.
- Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE transactions on knowledge and data engineering*, 17(6):734–749, 2005.
- Mathieu Blondel, Akinori Fujino, and Naonori Ueda. Convex factorization machines. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 19–35. Springer, 2015.
- Mathieu Blondel, Akinori Fujino, Naonori Ueda, and Masakazu Ishihata. Higher-order factorization machines. In *Advances in Neural Information Processing Systems*, pages 3351–3359, 2016.
- Sarah Brockhaus, David Rügamer, and Sonja Greven. Boosting functional regression models with fdboost. *Journal of Statistical Software*, 94(10):1 – 50, 2020.
- Andreas Buja, Trevor Hastie, and Robert Tibshirani. Linear Smoothers and Additive Models. *The Annals of Statistics*, 17(2):453 – 510, 1989.
- Chun-Hao Chang, Rich Caruana, and Anna Goldenberg. Node-gam: Neural generalized additive model for interpretable deep learning, 2021.
- Hong Chen, Xiaoqian Wang, Cheng Deng, and Heng Huang. Group sparse additive machine. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 197–207, 2017.
- Yifan Chen, Pengjie Ren, Yang Wang, and Maarten de Rijke. Bayesian personalized feature interaction selection for factorization machines. SIGIR’19, page 665–674, New York, NY, USA, 2019. Association for Computing Machinery.
- Andrea Coraddu, Luca Oneto, Alessandro Ghio, Stefano Savio, Davide Anguita, and Massimo Figari. Machine learning approaches for improving condition-based maintenance of naval propulsion plants. *Journal of Engineering for the Maritime Environment*, 2014.
- Paulo Cortez and Aníbal de Jesus Raimundo Morais. A data mining approach to predict forest fires using meteorological data. 2007.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.
- Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- Trevor J Hastie and Robert J Tibshirani. *Generalized additive models*. Routledge, 2017.
- Fuxing Hong, Dongbo Huang, and Ge Chen. Interaction-aware factorization machines for recommender systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3804–3811, 2019.
- Torsten Hothorn, Peter Bühlmann, Thomas Kneib, Matthias Schmid, and Benjamin Hofner. Model-based boosting 2.0. *Journal of Machine Learning Research*, 11(71):2109–2113, 2010.
- Ruoming Jin, Dong Li, Jing Gao, Zhi Liu, Li Chen, and Yang Zhou. Towards a better understanding of linear models for recommendation. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD ’21*, page 776–785, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383325.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Chris Kolb, Christian L. Müller, Bernd Bischl, and David Rügamer. Smoothing the Edges: A General Framework for Smooth Optimization in Sparse Regularization using Hadamard Overparametrization. 2023. doi: <https://doi.org/10.48550/arXiv.2307.03571>.
- Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- Liang Lan and Yu Geng. Accurate and interpretable factorization machines. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4139–4146, 2019.
- Yin Lou, Rich Caruana, Johannes Gehrke, and Giles Hooker. Accurate intelligible models with pairwise interactions. In *Proceedings of the 19th ACM*

- SIGKDD international conference on Knowledge discovery and data mining*, pages 623–631, 2013.
- John Ashworth Nelder and Robert WM Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135(3):370–384, 1972.
- Harsha Nori, Samuel Jenkins, Paul Koch, and Rich Caruana. Interpretml: A unified framework for machine learning interpretability. *arXiv preprint arXiv:1909.09223*, 2019.
- I Ortigosa, R Lopez, and J Garcia. A neural networks approach to residuary resistance of sailing yachts prediction. In *Proceedings of the international conference on marine engineering MARINE*, volume 2007, page 250, 2007.
- Steffen Rendle. Factorization machines. In *2010 IEEE International conference on data mining*, pages 995–1000. IEEE, 2010.
- Steffen Rendle, Walid Krichene, Li Zhang, and John Anderson. Neural collaborative filtering vs. matrix factorization revisited. In *Fourteenth ACM Conference on Recommender Systems*, pages 240–248, 2020.
- David Rügamer. A new PHO-rmula for improved performance of semi-structured networks. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 29291–29305. PMLR, 23–29 Jul 2023.
- David Rügamer, Andreas Bender, Simon Wiegerebe, Daniel Racek, Bernd Bischl, Christian Müller, and Clemens Stachl. Factorized structured regression for large-scale varying coefficient models. In *Machine Learning and Knowledge Discovery in Databases*, pages 20–35, Cham, 2023a. Springer International Publishing.
- David Rügamer, Chris Kolb, Cornelius Fritz, Florian Pfisterer, Philipp Kopper, Bernd Bischl, Ruolin Shen, Christina Bukas, Lisa Barros de Andrade e Sousa, Dominik Thalmeier, Philipp Baumann, Lucas Kook, Nadja Klein, and Christian L. Müller. deepregression: a flexible neural network framework for semi-structured deep distributional regression. *Journal of Statistical Software*, 105(1):1–31, 2023b.
- David Rügamer, Chris Kolb, and Nadja Klein. Semi-Structured Deep Distributional Regression: A Combination of Additive Models and Deep Learning. *The American Statistician*, 2023c.
- D. Ruppert, M.P. Wand, and R.J. Carroll. *Semiparametric Regression*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2003. ISBN 9780521785167.
- Daniel Schalk, Bernd Bischl, and David Rügamer. Accelerated componentwise gradient boosting using efficient data representation and momentum-based optimization. *Journal of Computational and Graphical Statistics*, 32(2):631–641, 2023.
- Nathan Srebro, Jason DM Rennie, and Tommi S Jaakkola. Maximum-margin matrix factorization. In *NIPS*, volume 17, pages 1329–1336. Citeseer, 2004.
- Almond Stöcker and Sonja Greven. Functional additive regression on shape and form manifolds of planar curves. *arXiv preprint arXiv:2109.02624*, 2021.
- Athanasios Tsanas and Angeliki Xifara. Accurate quantitative estimation of energy performance of residential buildings using statistical machine learning tools. *Energy and Buildings*, 49:560–567, 2012.
- Yingjie Wang, Hong Chen, Feng Zheng, Chen Xu, Tieliang Gong, and Yanhong Chen. Multi-task additive models for robust estimation and automatic structure discovery. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11744–11755. Curran Associates, Inc., 2020.
- Simon N. Wood. Low-rank scale-invariant tensor product smooths for generalized additive mixed models. *Biometrics*, 62(4):1025–1036, 2006.
- Simon N Wood. *Generalized additive models: an introduction with R*. Chapman and Hall/CRC, 2017.
- Simon N. Wood, Zheyuan Li, Gavin Shaddick, and Nicole H. Augustin. Generalized additive models for gigadata: Modeling the u.k. black smoke network daily data. *Journal of the American Statistical Association*, 112(519):1199–1210, 2017.
- I-C Yeh. Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete research*, 28(12):1797–1808, 1998.
- Fajie Yuan, Guibing Guo, Joemon M. Jose, Long Chen, Haitao Yu, and Weinan Zhang. Boostfm: Boosted factorization machines for top-n feature-based recommendation. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces, IUI '17*, page 45–54, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450343480.
- Honglei Zhuang, Xuanhui Wang, Michael Bendersky, Alexander Grushetsky, Yonghui Wu, Petr Mitrichev, Ethan Sterling, Nathan Bell, Walker Ravina, and Hai Qian. Interpretable ranking with generalized additive models. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining, WSDM '21*, page 499–507, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450382977.

Scalable Higher-Order Tensor Product Spline Models: Supplementary Materials

A Proofs

A.1 Proof of Lemma 1

$$\begin{aligned}
& \sum_{k=1}^p \sum_{l=k+1}^p \sum_{m=1}^{M_k} \sum_{o=1}^{O_l} B_{m,k}(x_k) B_{o,l}(x_l) \beta_{m,k,o,l} \\
& \approx \sum_{k=1}^p \sum_{l=k+1}^p \sum_{m=1}^{M_k} \sum_{o=1}^{O_l} \sum_{f=1}^F B_{m,k}(x_k) B_{o,l}(x_l) \gamma_{m,k,f} \gamma_{o,l,f} \\
& = \sum_{k=1}^p \sum_{l=k+1}^p \sum_{m=1}^{M_k} \sum_{o=1}^{O_l} \sum_{f=1}^F B_{m,k}(x_k) \gamma_{m,k,f} B_{o,l}(x_l) \gamma_{o,l,f} \\
& = \sum_{f=1}^F \left\{ \sum_{k=1}^p \sum_{m=1}^{M_k} B_{m,k}(x_k) \gamma_{m,k,f} \sum_{l=k+1}^p \sum_{o=1}^{O_l} B_{o,l}(x_l) \gamma_{o,l,f} \right\} \\
& = \sum_{f=1}^F \left\{ \sum_{k=1}^p \sum_{m=1}^{M_k} B_{m,k}(x_k) \gamma_{m,k,f} \sum_{l=k+1}^p \varphi_{l,f} \right\} \\
& = \sum_{f=1}^F \left\{ \sum_{k=1}^p \sum_{m=1}^{M_k} \sum_{l=k+1}^p \underbrace{B_{m,k}(x_k) \gamma_{m,k,f} \varphi_{l,f}}_{c_{m,k,f}} \right\} \tag{17} \\
& = \frac{1}{2} \sum_{f=1}^F \left\{ \sum_{k=1}^p \sum_{m=1}^{M_k} \sum_{l=1}^p c_{m,k,f} \varphi_{l,f} - \sum_{k=1}^p \sum_{m=1}^{M_k} c_{m,k,f} \varphi_{k,f} \right\} \\
& = \frac{1}{2} \sum_{f=1}^F \left\{ \left[\sum_{k=1}^p \sum_{m=1}^{M_k} c_{m,k,f} \right] \left[\sum_{l=1}^p \sum_{o=1}^{O_l} c_{o,l,f} \right] - \sum_{k=1}^p \varphi_{k,f}^2 \right\} \\
& = \frac{1}{2} \sum_{f=1}^F \left\{ \left[\sum_{k=1}^p \sum_{m=1}^{M_k} c_{m,k,f} \right]^2 - \sum_{k=1}^p \varphi_{k,f}^2 \right\} \\
& = \frac{1}{2} \sum_{f=1}^F \left\{ \left[\sum_{k=1}^p \varphi_{k,f} \right]^2 - \sum_{k=1}^p \varphi_{k,f}^2 \right\}.
\end{aligned}$$

A.2 Proof of Proposition 1 and 2

Both propositions directly follow from the fact that (17) only sums over f , k and m once, and every basis function $B_{m,k}$ is therefore also only evaluated once.

A.3 Proof of Lemma 2

An alternative representation of (2) is given by

$$\Phi_f^{(d)} = \sum_{j_d > \dots > j_1} \prod_{t=1}^d \varphi_{j_t, f} \quad (18)$$

by just plugging in the definition for φ . We can consider (18) as an ANOVA kernel of degree d in the new feature space given by all φ s. As a result, the multi-linearity property of the ANOVA kernel holds (see Blondel et al., 2016, Appendix B.1), i.e.,

$$\Phi_f^{(d)} = \Phi_{f, \neg j}^{(d)} + \varphi_{j, f} \Phi_{f, \neg j}^{(d-1)} \quad (19)$$

where $\Phi_{f, \neg j}^{(d)} = \sum_{\{j_d > \dots > j_1\} \setminus j} \prod_{t=1}^d \varphi_{j_t, f}$ and therefore AHOFMs can be represented as in Lemma 2 by using the same arguments as for HOFMs (Blondel et al., 2016).

A.4 Proof of Lemma 3

For simplicity assume that the model predictor only consists of a single AHOT, i.e., $\eta(\mathbf{x}) = \sum_{f=1}^{F_d} \Phi_f^{(d)}$. The generalization of the following statement to several AHOTs follows due to the additivity of the model predictor. Using Equation (19) and constants ξ_f, ζ_f , it follows

$$\begin{aligned} \eta(\mathbf{x}) &= \sum_{f=1}^{F_d} \Phi_f^{(d)} = \sum_{f=1}^{F_d} \left\{ \Phi_{f, \neg j}^{(d)} + \varphi_{j, f} \Phi_{f, \neg j}^{(d-1)} \right\} = \sum_{f=1}^{F_d} \{ \xi_f + \varphi_{j, f} \zeta_f \} \\ &= \sum_{f=1}^{F_d} \{ \xi_f + \mathbf{B}_j^\top \boldsymbol{\gamma}_{j, f} \zeta_f \} = \sum_{f=1}^{F_d} \xi_f + (\mathbf{B}_j \otimes \boldsymbol{\zeta})^\top \boldsymbol{\Gamma}_j^{(d)} = \text{const.} + \langle \tilde{\mathbf{B}}_j, \boldsymbol{\Gamma}_j^{(d)} \rangle, \end{aligned} \quad (20)$$

where $\boldsymbol{\Gamma}_j^{(d)} = \text{vec}(\boldsymbol{\Theta}_{:,j,:}^{(d)}) \in \mathbb{R}^{MF_d}$, $\tilde{\mathbf{B}} = \mathbf{B}_j \otimes \boldsymbol{\zeta} \in \mathbb{R}^{MF_d}$ and $\boldsymbol{\zeta} = (\zeta_1, \dots, \zeta_{F_d}) \in \mathbb{R}^{F_d}$. (20) shows that η is an affine function in $\boldsymbol{\Gamma}_j^{(d)} \forall j \in [p]$. Now let ℓ be a convex loss function of η and note that $\mathcal{P}(\boldsymbol{\Theta}, \boldsymbol{\Theta})$ in Definition 3 is decomposable across the parameters $\boldsymbol{\gamma}_{j, f}^{(d)}$. The composition of ℓ and η is convex and as \mathcal{P} is decomposable in $\boldsymbol{\gamma}_{j, f}^{(d)}$ and hence also in $\boldsymbol{\Gamma}_j^{(d)}$, the penalized objective in (15) is convex w.r.t. every $\boldsymbol{\Gamma}_j^{(d)}$ and thus every $\boldsymbol{\gamma}_{j, f} \forall j \in [p], f \in [F_d]$.

B Optimization

Having defined the objective in (15), we obtain the following result.

Lemma 3. *The optimization problem in (15) is coordinate-wise convex in $\boldsymbol{\gamma}_{j, f}^{(d)}$.*

A corresponding proof is given in Appendix A. Using this finding suggests a block coordinate descent (BCD) solver as an alternative approach to optimize (penalized) AHOFMs with block updates for $\boldsymbol{\gamma}_{j, f}^{(d)}$. In contrast to (HO)FMs, we perform block updates instead of plain coordinate descent as the AHOFM penalty is only

decomposable w.r.t. all $\gamma_{j,f}^{(d)}$, but not w.r.t. $\gamma_{m,j,f}^{(d)}$. For BCD we require several quantities:

$$\begin{aligned}
 \nabla \varphi_{j,f}(x_{i,j}) &:= \partial \varphi_{j,f}(x_{i,j}) / \partial \gamma_{j,f} = \mathbf{B}_j(x_{i,j}), \\
 \nabla \Phi_f^{(d)}(\mathbf{x}_i) &:= \frac{\partial \Phi_f^{(d)}(\mathbf{x}_i)}{\partial \gamma_{j,f}^{(d)}} \\
 &= \frac{1}{d} \sum_{t=1}^d (-1)^{t+1} \left\{ \frac{\partial \Phi_f^{(d-t)}(\mathbf{x}_i)}{\partial \gamma_{j,f}^{(d)}} \left\{ \sum_{j=1}^p [\varphi_{j,f}^{(d)}(x_{i,j})]^t \right\} \right. \\
 &\quad \left. + \Phi_f^{(d-t)}(\mathbf{x}_i)_t [\varphi_{j,f}^{(d)}(x_{i,j})]^{t-1} \nabla \varphi_{j,f}(x_{i,j}) \right\}, \\
 \nu &= \left\{ \sum_{i=1}^n \frac{\partial^2 \ell(\mathbf{x}_i, y_i)}{\partial (\gamma_{j,f}^{(d)})^2} + \lambda_{j,f}^{(d)} \mathbf{P}_j \right\}^{-1} \\
 \frac{\partial \mathcal{L}(\mathbf{x}_i, y_i)}{\partial \gamma_{j,f}^{(d)}} &= \frac{\partial \ell(\mathbf{x}_i, y_i)}{\partial \gamma_{j,f}^{(d)}} \nabla \Phi_f^{(d)}(\mathbf{x}_i) + \lambda_{j,f}^{(d)} \mathbf{P}_j \gamma_{j,f}^{(d)}.
 \end{aligned} \tag{21}$$

Note that the first term is involved in every update step, but independent of f and the iteration. It is therefore possible to cache the result once at the beginning of the training routine as also mentioned in Algorithm 5. For reverse-mode differentiation, note that the second term can be calculated efficiently by caching intermediate results.

A high-level routine is described in Algorithm 1 (for simplicity for the case with only a single D -variate smooth). In Algorithm 1, we require gradients

$$\nabla \mathcal{L}(\gamma_{j,f}^{(d)}) := \sum_{i=1}^n \frac{\partial \mathcal{L}(\mathbf{x}_i, y_i)}{\partial \gamma_{j,f}^{(d)}},$$

where \mathcal{L} is the objective function from (15) and a learning rate ν , which is defined above. Various terms involved in the update step can be pre-computed or cached (see Appendix C for details).

Algorithm 1 BCD AHOFM Optimization

Input: Data $(y_i, \mathbf{x}_i)_{i \in [n]}$; $B_{m,j}, \mathbf{P}_j, \forall j \in [p], m \in [M_j]$; D ; $\text{df}^{(D)}$; F_D ; BCD convergence criterion

Initialization: $\hat{\eta}, \mathfrak{G}^{(D)} = \text{init}(\text{Input})$ (Appendix C.3)

repeat

for $f = 1$ **to** F_D **do**

for $j = 1$ **to** p **do**

 Calculate step-size ν

 Update $\gamma_{j,f}^{(d)} \leftarrow \gamma_{j,f}^{(d)} - \nu \nabla \mathcal{L}(\gamma_{j,f}^{(d)})$

 Synchronize $\hat{\eta}_i, i \in [n]$

end for

end for

until convergence

C Algorithmic Details

C.1 Demmler-Reinsch Orthogonalization

We here describe the Demmler-Reinsch Orthogonalization (DRO; Algorithm 2) and sv2la (Algorithm 3) routine proposed to efficiently compute smoothing penalties. Details can be found in Ruppert et al. (2003), Appendix B.1.1. We use Chol to denote the Cholesky decomposition of a matrix and SVD for the singular value decomposition of a matrix.

Algorithm 2 DRO

Input: Feature matrix $\mathbf{B} \in \mathbb{R}^{n \times M}$, penalty matrix $\mathbf{P} \in \mathbb{R}^{M \times M}$

Compute:

1. $\mathbf{R}^\top \mathbf{R} \leftarrow \text{Chol}(\mathbf{B}^\top \mathbf{B})$
2. $\mathbf{U} \text{diag}(\mathbf{s}) \mathbf{U}^\top \leftarrow \text{SVD}(\mathbf{R}^{-\top} \mathbf{P} \mathbf{R}^{-1})$

Output: singular values \mathbf{s}

Algorithm 3 sv2la

Input: Singular values $\mathbf{s} \in \mathbb{R}^M$, df

Define $\text{dffun}(l) = \sum_{j=1}^M (1 + ls_j)^{-1}$;

Compute: λ for which $\text{dffun}(\lambda) = \text{df}$ using a uniroot search;

Output: λ

C.2 Homogeneous AHOFM Smoothing

Given the previous algorithms, we can implement homogeneous AHOFM smoothing as described in Algorithm 4.

Algorithm 4 Homogeneous AHOFM Smoothing

Input: $\mathbf{B}_j, \mathbf{P}_j \forall j \in [p]$; $\text{df}^{(d)} \forall d \in [D]$

for $j = 1$ **to** p **do**

Compute $\mathbf{s}_j = \text{DRO}(\mathbf{B}_j, \mathbf{P}_j)$ (costs: $\mathcal{O}(M_j^3)$)

for $d = 1$ **to** D **do**

Compute $\lambda_{j,1}^{(d)} = \text{sv2la}(\mathbf{s}_j, \text{df}^{(d)})$ (negligible costs);

Set $\lambda_{j,f}^{(d)} = \lambda_{j,1}^{(d)}$ for $f \in [F_d]$;

end for

end for

Output: $\lambda_{j,f}^{(d)}$ for all $j \in [p], d \in [D], f \in [F_d]$

C.3 AHOFM Initialization

Putting everything together, the initialization of AHOFMs is given in Algorithm 5.

D Interpretability of A(HO)FMs

Due to their additivity assumption, every additive feature effect in GAMs can be interpreted on its own (*ceteris paribus*). Although AHOFMs inherit some of the interpretability properties from GAMs, e.g., their additivity, interpreting higher-order (non-linear) interaction terms remains challenging and cannot be done without considering lower-order effects of the same feature. For larger values of p , the quickly growing number of additive terms further makes it infeasible to grasp the influence of certain features or interactions. While this is a limitation of the current approach, we here propose three ways to check effects for models with small to moderate p . The first approach examines interaction terms by visualizing the single univariate smooth terms $\varphi_{j,f}^d$ for $f = 1, \dots, F^d$ and all involved feature dimensions j . Analyzing the univariate latent dimensions separately is not a new approach and the use of factorization approaches can even be motivated by the need to interpret higher-dimensional interactions in lower dimensions (see, e.g., Stöcker and Greven, 2021). Another approach that focuses only on the interaction effects itself is to visualize the actual approximations ϕ_{j_1, \dots, j_d} as defined in (7). This reduces the number of terms to analyze by the factor F_d , but requires a method for presenting the d -variate effect. As shown in Figure 5, a third approach is to visualize the marginals of these multivariate functions together with their variation across the respective other dimensions.

Algorithm 5 AHOFM Init

Input: Data $(y_i, \mathbf{x}_i), i \in [n]$; order D , bases functions $B_{m,j}, \mathbf{P}_j, j \in [p], m \in [M_j]$; $\text{df}^{(d)}$
 Compute the following quantities:

- $\lambda_{j,f}^{(d)} \forall j \in [p], f \in [F_d]$ using Algorithm 4;
- $\nabla \varphi_{j,f}(x_{i,j}) \forall i \in [n]$ as in (21);

Cache the derivatives $\nabla \varphi_{j,f}(x_{i,j})$ for later update steps;

Randomly initialize $\gamma_{j,f}^{(d)}$ and calculate $\varphi_{j,f}^{(d)}$ for all $j \in [p], f \in [F_d]$;

Compute $\hat{\eta}$

Output: $\hat{\eta}, \mathfrak{G}^{(d)}$

Experiments We simulate a toy example for $D = 3$ with four features to demonstrate the third approach. For better understanding, the features are referred to as *time*, *lat*, *lon* and *rate*. The outcome is assumed to be normally distributed with $\sigma = 0.1$ and the mean given by the sum of all possible smooth 3-way interactions of the features. To simulate non-linear three-dimensional functions, we use a basis evaluation of features with 4 degrees-of-freedom for *time*, 5 degrees-of-freedom for *lat*, 7 degrees-of-freedom for *lon* and 5 degrees-of-freedom for *rate*. Partial effects for each three-dimensional smooth are generated by calculating the TP for these basis and randomly drawing coefficients for the resulting TPS. We generate 10^4 observations, fit a AHOFM($D = 3$) and visualize the resulting effects in two different ways. Figure 5 depicts the marginal univariate effects of all

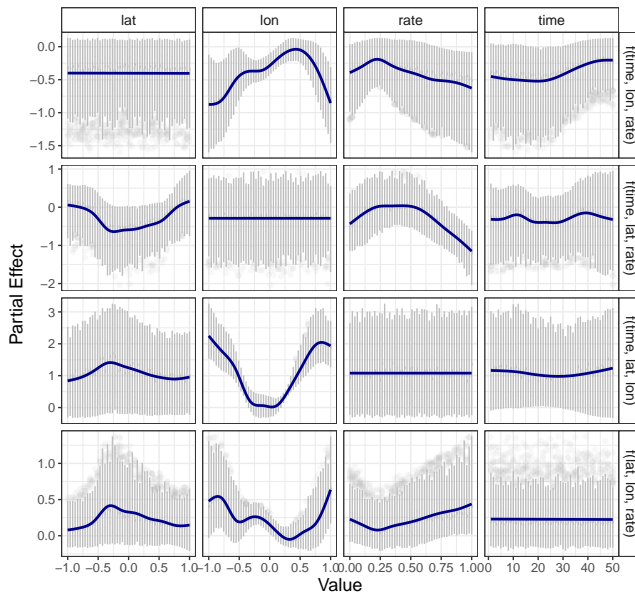


Figure 5: Estimated partial effects of different features (columns) for different three-dimensional functions (rows). Blue lines indicate the marginal average in the respective feature direction while gray vertical lines show the spread across the other two dimensions. More variation indicates larger variation across the other two dimensions. Features not involved in a partial effect (diagonal from top left to bottom right) naturally have a constant effect.

four features for all four 3-way interaction effects. This allows us to see how each feature marginally affects each of the interaction terms. Additionally, it shows how much variation the marginal effects have in the respective other two dimensions and thereby provides information on how much the features interact with the respective two other variables.

Table 3: Data set characteristics, additional pre-processing and references.

Data set	# Obs.	# Feat.	Reference
Airfoil	1503	5	Dua and Graff (2017)
Concrete	1030	8	Yeh (1998)
Diabetes	442	10	Efron et al. (2004)
Energy	768	8	Tsanas and Xifara (2012)
ForestF	517	12	Cortez and Morais (2007)
Naval	11934	16	Coraddu et al. (2014)
Yacht	308	6	Ortigosa et al. (2007); Dua and Graff (2017)

E Numerical Experiments Details

E.1 Implementation

All methods have been implemented in TensorFlow (Abadi et al., 2016) using the package deepregression (Rügamer et al., 2023b), except for GAMBoost, where we used the package mboost (Hothorn et al., 2010) and the SotA GAM, implemented in mgcv (Wood, 2017).

E.2 Benchmark Details

As described in Section 5.4, neither GAMs, (HO)FMs nor AHOFMs have many hyperparameters to tune. We investigate the influence of the latent dimension by testing $F = 1, 5, 10$ for all approaches and, for a fair comparison between GAMs and A(HO)FMs, set the df values for all methods to the same value 15. All methods use early stopping on 15% validation data with a patience of 50.

Table 3 further lists the data characteristics for our benchmark data. Pre-processing is only done for ForestF, using a logp1 transformation for `area` and a numerical representation for `month` and `day`.