

Factorized Structured Regression for Large-Scale Varying Coefficient Models

David Rügamer[✉]^{1,2}[0000-0002-8772-9202], Andreas
Bender¹[0000-0001-5628-8611], Simon Wiegrebe¹[0000-0003-3385-6879], Daniel
Racek¹[0000-0002-9564-7880], Bernd Bischl¹[0000-0001-6002-6980], Christian L.
Müller^{1,3,4}[0000-0002-3821-7083], and Clemens Stachl⁵[0000-0002-4498-3067]

¹ Department of Statistics, LMU Munich

² Institute of Statistics, RWTH Aachen

³ ICB, Helmholtz Zentrum Munich

⁴ CCM, Flatiron Institute

⁵ Institute of Behavioral Science and Technology, University of St.Gallen

Abstract. Recommender Systems (RS) pervade many aspects of our everyday digital life. Proposed to work at scale, state-of-the-art RS allow the modeling of thousands of interactions and facilitate highly individualized recommendations. Conceptually, many RS can be viewed as instances of statistical regression models that incorporate complex feature effects and potentially non-Gaussian outcomes. Such structured regression models, including time-aware varying coefficients models, are, however, limited in their applicability to categorical effects and inclusion of a large number of interactions. Here, we propose Factorized Structured Regression (FaStR) for scalable varying coefficient models. FaStR overcomes limitations of general regression models for large-scale data by combining structured additive regression and factorization approaches in a neural network-based model implementation. This fusion provides a scalable framework for the estimation of statistical models in previously infeasible data settings. Empirical results confirm that the estimation of varying coefficients of our approach is on par with state-of-the-art regression techniques, while scaling notably better and also being competitive with other time-aware RS in terms of prediction performance. We illustrate FaStR’s performance and interpretability on a large-scale behavioral study with smartphone user data.

Keywords: Recommender systems · Neural networks · Tensor regression · Time-varying effects · Generalized additive models

1 Introduction

From buying products online to selecting a movie to watch, recommender systems (RS) are part of our everyday life. RS are used to suggest those items that are most appealing to a given user based on the user’s past preference data or the similarity of a user to other users. One big advantage of RS is their scalability, as they allow for modeling thousands of interactions, e.g., between users

and items, and thereby facilitate individual recommendations in many dimensions (see, e.g., [33] for a recent implementation framework). Many RS can be represented as a regression model with the user and item as covariates. This makes it straightforward to include further features into the model and extend the method by other structural components.

At the same time, the increasing amount of available data and the possibility to model increasingly complex data generating processes calls for efficient methods to fit flexible regression models on large-scale data sets with many observations and features. In the past, several advanced statistical regression models have been proposed to incorporate complex feature effects. One of the most common approaches are generalized additive models (GAMs), widely considered to be state-of-the-art (SotA) for statistical modeling [26]. These models allow the incorporation of time-varying feature effects and spatial effects, among others, and can also deal with non-Gaussian outcomes (see [26] for more details). While well-working adaptations of GAMs for large data scenarios exist [28], both methodology and software reach their limits when modelling categorical effects or categorical interactions of several variables where features comprise hundreds or thousands of categories. An amalgamation of methods from RS and statistical regression can overcome the limitations of statistical regression models on large-scale data sets with many categorical effects and interactions. In this work, we combine smoothing approaches with factorization terms to overcome the limitations of varying coefficient models for categorical features with many factor levels. Our idea arises from the statistical analysis of a large-scale behavioral dataset (Section 7). In this dataset, smartphone usage behavior of participants was tracked for several weeks. Domain experts are interested in various structured regression effects, such as the continuous activity levels over time. While standard regression software allows to fit some of these effects, several hundred activities and users make it infeasible to fit a model that learns interaction effects of users and activities or smoothly varying time effects for one or both of these categorical variables.

Our contribution We propose Factorized Structured Regression (FaStR) for scalable varying coefficient models. This combined approach has the flexibility that has proven successful in additive regression models, while also being able to deal with high-dimensional categorical effects and interactions. More specifically, we 1) derive a general model formulation in (4) to combine GAMs and factorization approaches, 2) derive a varying factorization interaction in Section 4.1 that reduces the number of parameters and therefore computations by a factor of $(I + U)/(IU)$ for given numbers of category levels I and U , and 3) propose an efficient implementation of this fusion approach that a) can reduce the storage cost by a factor of $1/(I + U)$ and b) circumvents computations quadratic in I and U by using stochastic optimization, an array reformulation, as well as dynamic feature encoding. In numerical experiments, we moreover show that our approach 4) leads to an estimation performance comparable with a SotA implementation and 5) has the desired computational complexity. Finally, we 6) demonstrate its interpretability and applicability to large-scale data sets.

2 Related Literature

Multiple different RS have been proposed over the last years, many of them based on matrix factorization (MF) [16] or collaborative filtering [24]. While recent methods increasingly rely on neural network-based factorization or recommendation, e.g., [29], it remains debatable whether they yield superior results, e.g., with respect to performance and efficiency [20]. Factorization Machines (FM) represent another line of research which is closely related to MF. Initially proposed by [19], FM are based on a linear model problem formulation with pairwise (or order $d \geq 3$) interactions between all features. In particular, the formulation as regression model is the basis for extensions to other (prediction) tasks, with many different FM-type models having been developed in recent years (e.g., [4]). An important influence (context variable) in RS is time. Various methods for controlling for short- and long-term temporal dynamics, cyclic patterns, drift or time decay exist [2]. While short-term approaches either divide time into smaller periods or integrate time features into the factorization of the neighbourhood, long-term effects are accounted for by some form of distance calculations between the current and other designated time points. Some approaches also combine the factorization with the time context, and for instance assume smoothness in the factorization, e.g., for video pixel completion [13]. Specific time-aware methods include collaborative filtering with temporal dynamics [15], dynamic MF [7], temporal regularized MF [30], or sequence-aware FM [6]. The common ground of these methods is to account for the time context in the factorization.

Statistical approaches and interpretability Several combinations of statistical approaches and RS have been proposed in past years. Already in 1999, [8] proposed a (Bayesian) generalized mixed-effects model as RS Likelihood approximation approach. [31] proposed the GLMix model that combines the idea of generalized linear models (GLMs; [17]) with RS for large-scale response prediction. Our work is most similar to RS approaches which facilitate interpretability by making connections to generalized additive models (GAMs) [10]. In contrast to our approach, however, past work does not include smoothing splines directly into the models, nor does it address varying coefficient models. An exception is the time-varying tensor decomposition by [32] which is inspired by varying coefficient models. While similar in motivation, their work does not focus on scaling aspects and only considers approximate varying coefficients with separately learned basis coefficients. Our approach implements the full varying coefficient model with exact single-varying coefficients as well as a doubly-varying coefficient with jointly-learned latent basis coefficients.

3 Background

We will first describe the necessary background on factorization approaches in RS, structured additive regression models, and introduce our notation.

3.1 General Notation

In the following, we use $Y \sim \mathcal{F}$ to denote a random outcome value (e.g., a rating) from distribution \mathcal{F} and its observation $y \in \mathcal{Y} \subseteq \mathbb{R}$ for which the model generates a prediction \hat{y} . We reserve the indices $i \in \mathcal{I}$ and $u \in \mathcal{U}$ for two categorical features (exemplarily referred to as *item* and *user*) and $t \in \mathcal{T}$ for the context variable *time* on a given time interval \mathcal{T} . The features associated with i and u are assumed to be binary indicator variables and are only implicitly referenced using their index. In Section 3.3, we however use an integer representation to introduce a memory-efficient storage representation. Other context features are summarized by $\mathbf{x} \in \mathbb{R}^p$. We use b, w and v to denote weights in the model that relate to items and users, and make their dependence explicit by indexing these weights correspondingly with i and u . To distinguish between dependencies of categorical features and the (continuous) feature t , we highlight time-dependency by writing objects as functions of the time t . We assume that we are given a dataset $\mathcal{D} = \{(y_{iu}(t), \mathbf{x}_{iu}(t))\}_{i \in \mathcal{I}, u \in \mathcal{U}, t \in \mathcal{T}}$ of total size N and allow observations to be sparse, i.e., for \mathcal{D} to be a true subset of $\mathcal{I} \times \mathcal{U} \times \mathcal{T}$. For matrix computations in later sections, let $\mathbf{A} \odot \mathbf{B} \in \mathbb{R}^{N \times a \times b}$ define the row-wise tensor product (RWTP) of matrices $\mathbf{A} \in \mathbb{R}^{N \times a}$ and $\mathbf{B} \in \mathbb{R}^{N \times b}$, i.e., a Kronecker product applied to every pair of rows of both matrices $[\mathbf{A}_{[1,]} \otimes \mathbf{B}_{[1,]} \dots \mathbf{A}_{[N,]} \otimes \mathbf{B}_{[N,]}]$. Further, for $a = b$, let $\mathbf{A} \bullet \mathbf{B} := (\mathbf{A} * \mathbf{B}) \mathbb{1}_b$, where $*$ is the Hadamard product and $\mathbb{1}_b \in \mathbb{R}^b$ a vector of ones. The operation defined by \bullet can be exploited in models with Kronecker product structures such as array models for fast computation (see Section 4.2).

3.2 Model-based Recommender Systems

The basic MF model generates its predictions as

$$\hat{y}_{iu} = \langle \mathbf{v}_{1,i}, \mathbf{v}_{2,u} \rangle \quad (1)$$

using a dot product $\langle \mathbf{v}_{1,i}, \mathbf{v}_{2,u} \rangle = \mathbf{v}_{1,i}^\top \mathbf{v}_{2,u}$ of two latent factors $\mathbf{v}_{1,i}, \mathbf{v}_{2,u} \in \mathbb{R}^D$ from a D -dimensional joint latent factor space. After learning the mapping from each item i and user u to the respective latent factor vector, the dot product describes the interplay between user and item and is used to estimate the outcome y (ratings). If the combination $\mathcal{I} \times \mathcal{U}$ is observed completely, common matrix decomposition approaches such as a singular value decomposition can be applied. If the matrix containing the ratings for all user-item combinations is sparse, missing values can be imputed. This, however, can be inaccurate and computationally expensive. The common alternative is to use (1) to only model $(i, u) \in \mathcal{K} \subset (\mathcal{I} \times \mathcal{U})$ in the set \mathcal{K} of observed combinations. The solutions \hat{y} can be found by least squares estimation where an additional L_2 -penalty for \mathbf{v}_i and \mathbf{v}_u is typically added to the objective function [16]. In order to account for systematic user- and item-level trends, biases are further added to (1), yielding

$$\hat{y}_{iu} = \mu + b_i + b_u + \langle \mathbf{v}_{1,i}, \mathbf{v}_{2,u} \rangle, \quad (2)$$

where μ is a global intercept representing the average rating, and b_i, b_u represent the item and user tendencies. The latter two bias terms are again penalized using

a ridge penalty. Together with this penalization, b_i and b_u can also be interpreted as a random effect for the item and user (see, e.g., [31]).

Time-aware recommender systems Contexts such as the location or time in which data has been observed can make a crucial difference (see, e.g., [3]). RS therefore often include a context dependency. One of the most common context-aware RS are time-aware RS [5]. Time-aware model-based approaches assume the following relationship

$$\hat{y}_{iu}(t) = \mu + b_i(t) + b_u(t) + \mathbf{v}_{2,u}^\top(t)\mathbf{v}_{1,i}, \quad (3)$$

where both biases and the dot product are time-dependent. The rationale behind a time-varying latent user effect is that users change their behavior over time, whereas influences of items should be time-independent [16, 5]. While time is often assumed to be continuous, categorical time-aware models are used if time information is represented as discrete contextual values.

The time-varying latent user effect in (3) has a similar role as varying-coefficients in structured additive regression discussed in the following section.

3.3 Structured Additive Regression and Varying Coefficient Model

In statistical modeling, structured additive regression is a technique for estimating the relationships between a dependent variable (outcome value) and one or more independent variables (predictors, features). While the most common form of regression follows an additive structure as introduced in (2) and (3), in particular including linear effects $\mathbf{x}^\top \boldsymbol{\beta}$ of features \mathbf{x} or pairwise interactions, factorization terms are usually not present in classical regression models. Instead, to adapt models for complex relationships between features and outcome value, smooth, non-linear, additive terms $f(\cdot)$ of one or more features are included into the model. These terms are represented by a linear combination of L appropriate basis functions. A univariate non-linear effect of feature z is, e.g., approximated by $f(z) \approx \sum_{l=1}^L B_l(z)w_l$, where $B_l(z)$ is the l -th basis function (such as regression splines, polynomial bases or B-splines) evaluated at z and w_l is the corresponding basis coefficient. Similarly, tensor product basis representations allow for two- or moderate-dimensional non-linear interactions.

One important part of additive regression is the so-called *varying coefficient* model [12]. The rationale for these models is the same as for time-varying RS: effects of features in the model naturally vary over time. Therefore these models include effects $xf(t)$, such that the effect (coefficient) of x is given by f evaluated at time t , and f is estimated from the data. A special case is a varying coefficient $f_i(t), i = 1, \dots, I$ where a separate function f_i is estimated for all I levels of a categorical feature. Existing software to model varying coefficients with smooth time-effects is, however, not scalable to features with many categories. The bottleneck is an RWTP of the matrix of evaluated basis functions $\mathbf{B} := (B_1(t) \dots B_L(t))$ and a (one hot-)encoded matrix for a categorical variable (e.g., item with I levels).

Computational complexity Assuming equal number of basis functions L for every smooth term f_i in a varying coefficient model with $i = 1, \dots, I$ levels and N observations, the storage required for the model matrix is $\mathcal{O}(NLI)$ and the computations $\mathcal{O}(N(LI)^2)$ (cf. [27]). Similar, for a model with an interaction effect of, e.g., item and user (with U levels), the storage is $\mathcal{O}(NLIU)$ and $\mathcal{O}(N(LIU)^2)$.

4 Factorized Structured Regression

In order to address the computational limitations of statistical regression techniques, we will first introduce the general idea to obtain predictions from a FaStR model and then go into more specific details and merits. We use the RS notation to define the model by means of a classical recommendation task with items, users, time context and an outcome y such as a rating. As in a typical regression setting, further features \mathbf{x} might exist that the modeler is interested in. We assume that the outcome $Y_{iu}(t)$ and all features $\mathbf{x}_{iu}(t)$ are observed on a grid for item $i = 1, \dots, I$, user $u = 1, \dots, U$, time $t = 1, \dots, T$. While our approach also works for sparsely and irregularly observed data, we assume a grid of observations to simplify the notation. Conditional on the item, user, time and further features, $Y_{iu}(t)$ is assumed to follow a parametric distribution \mathcal{F} . We model the expectation of $Y_{iu}(t)$ as

$$\begin{aligned} \mathbb{E}(Y_{iu}(t)|\mathbf{x}_{iu}(t)) &= h(\eta_{iu}(t)), \\ \eta_{iu}(t) &= \mu + b_i + b_u + b_{iu} + f^{[0]}(t) + f_i^{[1]}(t) + f_u^{[2]}(t) + f_{iu}^{[3]}(t) \\ &\quad + \sum_{o=1}^O g_o(\mathbf{x}_{iu}(t)). \end{aligned} \quad (4)$$

Here, h is an activation or response function mapping the additive predictor $\eta_{iu}(t)$ onto the correct domain (e.g., \mathbb{R}^+ for a positive outcome variable). All terms indicated with b are (regularized) bias terms. μ is a global bias term, b_i an item-specific bias, b_u a user-specific bias and b_{iu} an item-user-specific one. Terms denoted by f are smooth non-linear functions of time t represented by (penalized) basis functions. These include a global trend $f^{[0]}$, a subject and activity trend, $f^{[1]}$ and $f^{[2]}$, respectively, and a joint trend $f^{[3]}$. Additional covariates \mathbf{x} can be modeled using other (smooth) functions g_o . In the Supplementary Material A we provide further details on smoothness penalties and model optimization.

The model in (4) can be seen as an alternative notation for a varying coefficient model, or also as a time-aware RS with additional exogenous terms. As \mathcal{F} is not required to be Gaussian, it has, however, a more general applicability (e.g., binary, count or interval data). What further distinguishes (4) from existing approaches is the smoothness assumption of terms denoted with f , combined with the efficient implementation of terms $f^{[1]}$, $f^{[2]}$ and a factorization assumption for $f^{[3]}$. These aspects are explained in more detail in the following.

4.1 Varying Factorized Interactions

For high-dimensional data, such as the mobile phone data in our example, estimating the 2- or 3-way interaction terms is computationally not feasible. We

thus propose to define $\eta_{iu}(t)$ in (4) using latent factorization representations. We therefore decompose the discrete interaction term(s) into an inner product

$$b_{iu} = \langle \mathbf{v}_{1,i}, \mathbf{v}_{2,u} \rangle = \sum_{d=1}^D v_{1,i,d} \cdot v_{2,u,d} \quad (5)$$

with $D \ll \min(I, U)$ latent factors, resulting in the estimation of $D \cdot (I + U)$ instead of $I \cdot U$ parameters. If $I = U = 1000$ and $D = 5$, for example, this reduces the number of parameters by a factor of 100 from 10^6 to 10^4 . While this is the common approach to model interactions in factorization approaches, we here propose to proceed in a similar fashion to model time-dependent interactions and approximate time-varying interactions by a factorization approach:

$$f_{iu}^{[3]}(t) \approx \tilde{f}^{[3]}(t, \mathbf{V}_{1,i}, \mathbf{V}_{2,u}) = \sum_{l=1}^L B_l(t) \sum_{d=1}^D v_{1,i,l,d} \cdot v_{2,u,l,d}, \quad (6)$$

where $\mathbf{V}_{\cdot, \cdot} \in \mathbb{R}^{L \times D}$ are matrices with rows corresponding to the L basis functions for one categorical effect and columns to the D latent factors. In other words, we approximate the interaction of the smooth basis \mathbf{B} of dimension L and the two latent matrices $\mathbf{B}^\top (\mathbf{V}_{1,i} \bullet \mathbf{V}_{2,u})$, which can be computed efficiently for all N rows. The representation via latent factors requires the estimation of $L \cdot D \cdot (I + U)$ instead of $L \cdot D \cdot I \cdot U$ parameters (a multiplicative reduction of $(I + U)/(IU)$ in parameters and computations). This principle is general and can be applied to various types of additive effects, also of two or higher dimensions such as tensor-product splines or Markov random field smooths (see [26]). The dimension D is usually tuned to maximize prediction performance or estimation quality. While an increase in D will result in more parameters, the computational cost at this point only plays a minor role, as $D < 20$ often works well in practice.

Penalization In order to enforce smoothness of the varying coefficients in the time-dimension, a quadratic Kronecker sum penalty J can be added to the loss function [26]. In a similar manner, we can promote smoothness of the latent factors $\mathbf{V}_{iu} = (\mathbf{V}_{1,i} \bullet \mathbf{V}_{2,u}) \in \mathbb{R}^L$ in our adaption using a symmetric difference penalty matrix $\mathbf{P}_V \in \mathbb{R}^{L \times L}$. \mathbf{P}_V penalizes the time-dimension of the factorized varying-coefficients, where the penalized differences (its entries) depend on the chosen basis B . We further allow for an L_2 -regularization of the latent factors in the i - and u -dimension, yielding

$$J = \lambda_t \cdot \mathbf{V}_{iu}^\top \mathbf{P}_V \mathbf{V}_{iu} + \lambda_{iu} \cdot (\|\mathbf{V}_{1,i}\|_F + \|\mathbf{V}_{2,u}\|_F), \quad (7)$$

where λ_t controls the smoothness of the non-linearity in the direction of the time t , $\|\cdot\|_F$ is the Frobenius norm, and λ_{iu} the regularization for items and users.

4.2 Efficient Implementation

While the previous section allows to efficiently model (smooth) interactions of two or more categorical features with many categories, the factorization is not

a solution for coefficients $f^{[1]}$ and $f^{[2]}$, as these only vary with a single category. Many use cases also require to estimate one effect for each of the levels of a categorical (interaction) effect. One bottleneck if I and/or U is large, is the computation of their dummy-encoded design matrices. For example, for item the matrix \mathbf{X}_I of size $N \times I$ contains binary entries indicating which observation (rows) belong to which item category (column), and analogous for a user matrix \mathbf{X}_U . Second, the execution of operations involving such large matrices is another computational bottleneck. Computations become even more challenging if the model includes interactions, resulting in $\mathcal{O}(NIU)$ storage and $\mathcal{O}(N(IU)^2)$ computations (cf. Section 3.3). These interactions are created by calculating the RWTP between both matrices, i.e., $\mathbf{X}_I \odot \mathbf{X}_U \in \mathbb{R}^{N \times I \cdot U}$. To circumvent creating, storing and processing \mathbf{X}_I and \mathbf{X}_U as a whole, we propose two simple yet effective implementation tricks explained in the following.

Stochastic optimization The first bottleneck in computations of varying coefficient models at scale is the number of observations N . We therefore implement FaStR in a neural network and thereby can make use of stochastic gradient descent optimization routines with mini-batches of size $M \ll N$. This reduces the original cost of computations from $\mathcal{O}(N(IU)^2)$ to $\mathcal{O}(EMIU)$ where E is the number of model updates. It also allows us to leverage high-performance computing platforms such as TensorFlow [1] that support GPU computations.

Array computations The second bottleneck is computing the RWTP of two- or higher-dimensional interaction terms. Our proposal is to use an array reformulation that does not require to compute the RWTP design matrix in the first place. More specifically, a two-dimensional interaction effect $(\mathbf{X}_I \odot \mathbf{X}_U)\mathbf{w}$ with weights \mathbf{w} can be equally represented by the array reformulation (cf. Section 3.1)

$$(\mathbf{X}_I \mathbf{W}) \bullet \mathbf{X}_U, \quad (8)$$

where $\mathbf{W} \in \mathbb{R}^{I \times U}$ is a matrix of weights with the i th row and u th column being the weight for the interaction of the i th level in \mathbf{X}_I and the u th level in \mathbf{X}_U . By using (8) instead of a plain linear effect, we circumvent the construction of the large RWTP and the storage cost is reduced from $\mathcal{O}(NIU)$ to $\mathcal{O}(N(I+U))$ without increasing the time complexity (as the \bullet operation for M observations is neglectable with $\mathcal{O}(MU)$ compared to the matrix multiplication with $\mathcal{O}(MIU)$). This array formulation can also be defined for higher-order interactions [9].

Dynamic feature encoding Although array computations can reduce the storage problem notably by not constructing the RWTP in the first place, a third bottleneck is storing the large dummy-encoded matrices \mathbf{X}_I and \mathbf{X}_U themselves. We circumvent this extra space complexity, by evaluating categorical features dynamically during network training and only constructing the one-hot encoding for categorical features on a given mini-batch m . Thereby, only a matrix $\mathbf{X}_I^{(m)}$ of size $M \times I$ needs to be loaded into memory and the full $N \times I$ matrix is never created explicitly. This effectively reduces the storage from $\mathcal{O}(NIU)$ to $\mathcal{O}(N)$

(exactly $2N$ for two integer vectors). While this potentially results in redundant computations as it will create the encoding for a specific observation multiple times (if the number of epochs is greater than 1), deterministic integer encoding is cheap. Hence, the resulting computation overhead is usually neglectable and both forward- and backward pass can make use of the sparse representation.

When evaluating the two matrix operators (matrix product and \bullet) in (8) sequentially, encodings can again be created dynamically and the largest matrix involved only contains $\max(I, U)$ instead of $I \cdot U$ columns.

5 Numerical Experiments

Our numerical experiments investigate 1) whether FaStR can estimate (factorized) varying coefficients as proposed in Section 4.1 with performance comparable to other SotA methods, and 2) whether the model presented in (4) can be initialized and fitted with constant memory scaling w.r.t. the number of factor levels. In addition, in the Supplementary Material B we investigate whether FaStR can recover GAMs in general using our implementation techniques. Details on the data generating processes for each simulation study, method implementation and the computing environment are given in the Supplementary Material C. In the first experiment, we compare the estimated and true model coefficients using the mean squared error (MSE) and the estimated vs. the true functions f using the squared error, integrated and averaged over the domain of the predictor variables (MISE). We repeat the data generating process 10 times to account for variability in the data and model estimation.

5.1 Estimation of Factorized Smooth Interactions

We first investigate how well our models can recover smooth factorized terms from Section 4.1 of two categorical variables j and u , with 4 and 5 levels, respectively, i.e., 20 different smooth effects. We use a relatively small number of levels to be able to fit all possible interactions also in a classic structured regression approach. We define one of the two true factorized smooth effects as a varying coefficient term, i.e., $f_{1,iu}(t)$, and one stemming from an actual factorization, i.e., $f_2(t, \mathbf{V}_{1,i}, \mathbf{V}_{2,u})$. We use the Bernoulli and Gaussian distributions and investigate factorized terms in the distributions' mean for different data sizes $N \in \{500, 1000, 2000\}$. While the true model is generated using three latent factor dimensions, we additionally investigate a model with six dimensions to see how the misspecification of the latent dimension size influences estimation performance. FaStR is trained using a batch size of 250 with early stopping on 10% of the training data and a patience of 50 epochs. Larger batch sizes do not change estimation performance notably, but slow down the convergence speed.

Results All results show that FaStR can estimate varying coefficients equally well compared to a classic GAM estimation routine. Figure 1 shows the resulting M(I)SE values for all data settings. Note that the GAM implicitly assumes as

many latent dimensions as there are factor levels, but can also shrink single smooth functions to zero. In cases where it is feasible to fit smooth effects for every factor level (as is the case here), GAM can thus be seen as gold standard. For the case of a normal distribution, we observe that GAM yields better results, but also that the performance of FaStR converges to the one from GAM with increasing number of observations. For the Bernoulli case, our approach benefits from the optimization in a neural network and even outperforms the classic GAM estimation routine which requires a multiple of observations compared to a Gaussian setting for good estimation results (as, e.g., also found in [22]).

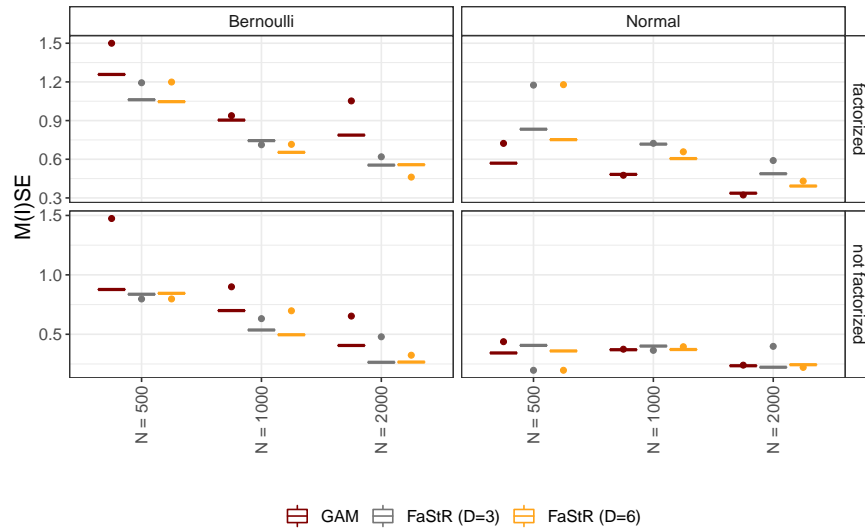


Fig. 1. Comparison of M(I)SE values for the estimated partial effects of different methods (colors) for factorized and non-factorized smooth terms (x-axis), different distributions (columns) and different data sizes (rows). Values > 1.5 are set to the value 1.5 to improve readability.

5.2 Memory Consumption

Finally, we compare the memory consumption of our implementation against the SotA implementation for big additive models (BAM) in the R [18] package `mgcv` [25] for an increasing number of category levels (20, 40, 60, 80) when using a categorical effect or a varying coefficient based on the representation proposed in Section 4.2, $N \in \{1000, 2000, 4000\}$ and optimization as in Section 5.1. While the improvement in memory consumption is expected to be even larger when using factorized terms instead of interaction terms with weights for each category combination, we do not use factorization in this experiment as there is

no equivalent available in software for additive regression. Additionally, we also track the time when running FaStR for 10 epochs to see if there are notable changes in the time consumption for varying data generating settings.

Results Figure 2 visualizes the results for all different settings and compares run times and memory consumption of the two methods for factor variables (single) and varying coefficient effects (varying). Results show that FaStR has both, almost constant time and memory consumption while the SotA method requires exponentially more memory for growing numbers of factor levels (as the whole encoded matrix must be loaded into memory). These results confirm that our implementation works as intended to allow for the estimation of varying coefficient models in large-scale settings.




Fig. 2. Memory and time consumption (y-axis; \log_{10} -scale) comparison between the SotA big additive model (BAM) implementation and our method (in different colours) for an increasing number of categorical levels (x-axis) of a factor effect (single) and varying coefficient term (varying) for different data sizes (columns).

6 Benchmarks

Although the focus of this work is to provide scalable and interpretable regression models, prediction performance of our models is also of (secondary) interest. We aim for a similar performance compared to SotA time-aware RS techniques, yet without the ambition to outperform these methods. We use the *MovieLens 10M movie ratings* benchmark data set [11], which is sparse in terms of user-item combinations, with *items* corresponding to rated movies. In addition, we benchmark models on a subset of the densely observed *PhoneStudy behavior* data set [23], analyzed in more detail in Section 7. In both cases, we use single

train-test splits (90%/10% and 70%/30%, respectively) and evaluate the models predictive performance with the root mean squared error (RMSE). The different characteristics of both data sets are given in Table 1.

Table 1. Descriptive statistics of the two benchmark data sets.

	Movies	Phone Study
# Observations (N)	≈ 9 m	≈ 8.7 m
# Users (U)	69,878	342
# Movies/Activities (I)	10,677	176
# Unique Time Points (T)	≈ 6.5 m	348

Methods As comparison we use Bayesian timeSVD and Bayesian timeSVD++ flipped, two variations of the SVD++ method [14], a latent factor model whose key innovative feature is the incorporation of implicit user information. Both Bayesian timeSVD and Bayesian timeSVD++ flipped have been extended to be time-aware [15] and optimized by Gibbs sampling using a Bayesian reformulation [21]. Bayesian timeSVD++ flipped integrates both implicit user and item information and has been reported to be the best-performing model among multiple SotA methods in a recent benchmark study [21]. The second variation, Bayesian timeSVD, is still a time-aware latent factor model, yet it does not incorporate implicit user or item information. As we are mainly interested in the performance of the proposed time-varying coefficient model, the Bayesian timeSVD provides a much fairer comparison with FaStR as it does not include the aforementioned types of implicit information. We use tuning parameter settings as given in [21] for the two benchmark methods (i.e., we use the already tuned models). We

Table 2. Benchmark results based on the RMSE on the respective test data set for both benchmark data sets (columns) for different methods (rows). Best results are in bold and best results of the respective other method is underlined.

	Movies	PhoneStudy
timeSVD	0.872	0.089
timeSVD++ flipped	0.856	<u>0.087</u>
FaStR ($D = 10$)	0.984	0.076
FaStR ($D = 3$)	0.975	0.080
FaStR ($D = 1$)	<u>0.890</u>	0.087
FaStR (w/o $f^{[3]}$)	1.027	0.093

compare these two models against our method as proposed in (4) and thereby not only test its predictive performance, but also its capability to scale well to high-dimensional data sets. We do not tune FaStR extensively, but perform a

small ablation study by testing different latent dimensions ($D \in 1, 3, 10$) for the factorized varying coefficient term $f^{[3]}$ and by excluding the whole term. All models use early stopping based on a validation data split with the same proportion as the train-test split. Considering further model definitions and tuning the smoothing parameters could slightly improve results, but at the cost of having to fit additional models.

Results Table 2 shows the performance of all methods. Interestingly, FaStR is competitive with the SotA timeSVD approaches, even though the premise of this paper was merely to develop a scalable variant of the varying coefficient model, not to propose a method with SotA performance on RS tasks.

7 User Behavior Phone Study




Fig. 3. One dimension of the latent time-varying coefficients for different activities (left plot with “locked screen” in blue) and users (right plot with exemplary user in green) over 87 study days (x-axis). Vertical lines show Christmas and New Year’s Eve.

We finally turn to the motivating case study. A more detailed description of the study and data set can be found in the Supplementary Material D. To analyze the activity levels (a value between 0 and 1, indicating the intensity of the activity in the given aggregation window) of participants in the study, we model the expected activity levels using user and activity effects, their interaction based on a factorization, an effect of the day of the week (Mon - Sun), the time of the day (6-hour windows), a factorized interaction of users and weekday as well as users and daytime, and a smooth time-dependent study day effect varying by user and/or activity. All factorizations use a three-dimensional latent space.

Results Results are generally plausible and in line with prior expectations. Various model effects are examined in the Supplementary Material. We here briefly

analyze the factorized varying coefficient interaction by analyzing its latent factors. One of the dimensions is depicted in Figure 3. Most activities follow a global pattern (darker concentration of lines), while a few show very specific sequences. The “locked screen” event, e.g., is observed less in the days around New Year’s Eve. At the user level, no general pattern is visible due to the different study starting dates, but we observe still common changes in activity around holidays, e.g., Christmas.

8 Conclusion

In this work we presented an amalgamation of structured statistical regression and RS to allow for large-scale varying coefficient models with many categorical levels. For this, we leveraged factorization approaches combined with an efficient neural network-based implementation. Empirical results confirm the efficacy of our proposal. In order to make the proposed approach as flexible as commonly used statistical regression software, we used a model-based point of view and cast the approach as a generalized additive model.

Acknowledgement

This work has been partially supported by the German Federal Ministry of Education and Research (BMBF) under Grant No. 01IS18036A. We also thank four anonymous reviewers for their helpful suggestions and comments.

References

1. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al.: Tensorflow: A system for large-scale machine learning. In: 12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16). pp. 265–283 (2016)
2. Al-Hadi, I., Sharef, N.M., Sulaiman, M.N., Mustapha, N.: Review of the temporal recommendation system with matrix factorization. *Int. J. Innov. Comput. Inf. Control* **13**(5), 1579–1594 (2017)
3. Baltrunas, L., Ricci, F.: Experimental evaluation of context-dependent collaborative filtering using item splitting. *User Modeling and User-Adapted Interaction* **24**(1-2), 7–34 (2014). <https://doi.org/10.1007/s11257-012-9137-9>
4. Blondel, M., Fujino, A., Ueda, N., Ishihata, M.: Higher-order factorization machines. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 29. Curran Associates, Inc. (2016). <https://doi.org/10.5555/3157382.3157473>
5. Campos, P.G., Díez, F., Cantador, I.: Time-aware recommender systems: a comprehensive survey and analysis of existing evaluation protocols. *User Modeling and User-Adapted Interaction* **24**(1), 67–119 (2014). <https://doi.org/10.1007/s11257-012-9136-x>

6. Chen, T., Yin, H., Nguyen, Q.V.H., Peng, W.C., Li, X., Zhou, X.: Sequence-aware factorization machines for temporal predictive analytics. In: 2020 IEEE 36th International Conference on Data Engineering (ICDE). pp. 1405–1416. IEEE (2020). <https://doi.org/10.1109/ICDE48307.2020.00125>
7. Chua, F.C.T., Oentaryo, R.J., Lim, E.P.: Modeling temporal adoptions using dynamic matrix factorization. In: 2013 IEEE 13th International Conference on Data Mining. pp. 91–100 (2013). <https://doi.org/10.1109/ICDM.2013.25>
8. Condl, M.K., Lewis, D.D., Madigan, D., Posse, C.: Bayesian mixed-effects models for recommender systems. In: ACM SIGIR. vol. 99 (1999)
9. Currie, I.D., Durban, M., Eilers, P.H.: Generalized linear array models with applications to multidimensional smoothing. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **68**(2), 259–280 (2006). <https://doi.org/10.1111/j.1467-9868.2006.00543.x>
10. Guo, Y., Su, Y., Yang, Z., Zhang, A.: Explainable recommendation systems by generalized additive models with manifest and latent interactions (2020)
11. Harper, F.M., Konstan, J.A.: The movielens datasets: History and context **5**(4) (2015). <https://doi.org/10.1145/2827872>
12. Hastie, T., Tibshirani, R.: Varying-coefficient models. *Journal of the Royal Statistical Society: Series B (Methodological)* **55**(4), 757–779 (1993). <https://doi.org/10.1111/j.2517-6161.1993.tb01939.x>
13. Imaizumi, M., Hayashi, K.: Tensor decomposition with smoothness. In: International Conference on Machine Learning. pp. 1597–1606. PMLR (2017). <https://doi.org/10.5555/3305381.3305546>
14. Koren, Y.: Factorization meets the neighborhood: a multifaceted collaborative filtering model. In: Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 426–434 (2008). <https://doi.org/10.1145/1401890.1401944>
15. Koren, Y.: Collaborative filtering with temporal dynamics. In: Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 447–456 (2009). <https://doi.org/10.1145/1721654.1721677>
16. Koren, Y., Bell, R., Volinsky, C.: Matrix factorization techniques for recommender systems. *Computer* **42**(8), 30–37 (2009). <https://doi.org/10.1109/MC.2009.263>
17. Nelder, J.A., Wedderburn, R.W.: Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)* **135**(3), 370–384 (1972). <https://doi.org/10.2307/2344614>
18. R Core Team: R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria (2021)
19. Rendle, S.: Factorization machines. In: 2010 IEEE International conference on data mining. pp. 995–1000. IEEE (2010). <https://doi.org/10.1109/ICDM.2010.127>
20. Rendle, S., Krichene, W., Zhang, L., Anderson, J.: Neural collaborative filtering vs. matrix factorization revisited. In: Fourteenth ACM Conference on Recommender Systems. pp. 240–248 (2020). <https://doi.org/10.1145/3383313.3412488>
21. Rendle, S., Zhang, L., Koren, Y.: On the Difficulty of Evaluating Baselines: A Study on Recommender Systems. *arXiv preprint arXiv:1905.01395* (2019). <https://doi.org/10.48550/ARXIV.1905.01395>
22. Rügamer, D., Kolb, C., Klein, N.: Semi-Structured Deep Distributional Regression: A Combination of Additive Models and Deep Learning. *arXiv preprint arXiv:2002.05777* (2020). <https://doi.org/10.48550/ARXIV.2002.05777>
23. Stachl, C., Au, Q., Schoedel, R., Gosling, S.D., Harari, G.M., Buschek, D., Völkel, S.T., Schuwerk, T., Oldemeier, M., Ullmann, T., Hussmann, H., Bischl, B., Büh-

- ner, M.: Predicting personality from patterns of behavior collected with smart-phones. *Proceedings of the National Academy of Sciences* **117**, 17680–17687 (2020). <https://doi.org/10.1073/pnas.1920484117>
24. Thorat, P.B., Goudar, R., Barve, S.: Survey on collaborative filtering, content-based filtering and hybrid recommendation system. *International Journal of Computer Applications* **110**(4), 31–36 (2015). <https://doi.org/10.5120/19308-0760>
 25. Wood, S.N.: Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society (B)* **73**(1), 3–36 (2011). <https://doi.org/10.1111/j.1467-9868.2010.00749.x>
 26. Wood, S.N.: *Generalized additive models: an introduction with R*. Chapman and Hall/CRC (2017). <https://doi.org/10.1201/9781315370279>
 27. Wood, S.N.: Inference and computation with generalized additive models and their extensions. *Test* **29**(2), 307–339 (2020). <https://doi.org/10.1007/s11749-020-00716-0>
 28. Wood, S.N., Li, Z., Shaddick, G., Augustin, N.H.: 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). <https://doi.org/10.1080/01621459.2016.1195744>
 29. Wu, C., Lian, D., Ge, Y., Zhu, Z., Chen, E., Yuan, S.: Fight fire with fire: Towards robust recommender systems via adversarial poisoning training. In: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. p. 1074–1083. SIGIR '21, Association for Computing Machinery, New York, NY, USA (2021). <https://doi.org/10.1145/3404835.3462914>
 30. Yu, H.F., Rao, N., Dhillon, I.S.: Temporal regularized matrix factorization for high-dimensional time series prediction. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 29. Curran Associates, Inc. (2016)
 31. Zhang, X., Zhou, Y., Ma, Y., Chen, B.C., Zhang, L., Agarwal, D.: Glmix: Generalized linear mixed models for large-scale response prediction. p. 363–372. *KDD '16*, Association for Computing Machinery, New York, NY, USA (2016). <https://doi.org/10.1145/2939672.2939684>
 32. Zhang, Y., Bi, X., Tang, N., Qu, A.: Dynamic tensor recommender systems. *Journal of Machine Learning Research* **22**(65), 1–35 (2021). <https://doi.org/10.11159/icsta19.09>
 33. Zhao, W.X., Mu, S., Hou, Y., Lin, Z., Li, K., Chen, Y., Lu, Y., Wang, H., Tian, C., Pan, X., Min, Y., Feng, Z., Fan, X., Chen, X., Wang, P., Ji, W., Li, Y., Wang, X., Wen, J.R.: Recbole: Towards a unified, comprehensive and efficient framework for recommendation algorithms (2020). <https://doi.org/10.1145/3459637.3482016>