

Factorization Machines with Regularization for Sparse Feature Interactions

Kyohei Atarashi

*Graduate School of Information Science and Technology
Hokkaido University
Kita 14, Nishi 9, Kita-ku, Sapporo, Hokkaido, 060-0814, Japan*

KATARASHI@FRONTIER.HOKUDAI.AC.JP

Satoshi Oyama

Masahito Kurihara

*Faculty of Information Science and Technology
Hokkaido University
Kita 14, Nishi 9, Kita-ku, Sapporo, Hokkaido, 060-0814, Japan*

OYAMA@IST.HOKUDAI.AC.JP

KURIHARA@IST.HOKUDAI.AC.JP

Editor: Koji Tsuda

Abstract

Factorization machines (FMs) are machine learning predictive models based on second-order feature interactions and FMs with sparse regularization are called sparse FMs. Such regularizations enable *feature selection*, which selects the most relevant features for accurate prediction, and therefore they can contribute to the improvement of the model accuracy and interpretability. However, because FMs use second-order feature interactions, the selection of features often causes the loss of many relevant feature interactions in the resultant models. In such cases, FMs with regularization specially designed for *feature interaction selection* trying to achieve interaction-level sparsity may be preferred instead of those just for feature selection trying to achieve feature-level sparsity. In this paper, we present a new regularization scheme for feature interaction selection in FMs. For feature interaction selection, our proposed regularizer makes the feature interaction matrix sparse without a restriction on sparsity patterns imposed by the existing methods. We also describe efficient proximal algorithms for the proposed FMs and how our ideas can be applied or extended to feature selection and other related models such as higher-order FMs and the all-subsets model. The analysis and experimental results on synthetic and real-world datasets show the effectiveness of the proposed methods.

Keywords: factorization machines, sparse regularization, feature interaction selection, feature selection, proximal algorithm

1. Introduction

Factorization machines (FMs) (Rendle, 2010, 2012) are machine learning predictive models based on second-order feature interactions, i.e., the products $x_i x_j$ of two feature values. One advantage of FMs compared with linear models with a polynomial term (namely, quadratic regression (QR)) and kernel methods is efficiency. The computational cost of evaluating FMs is linear with respect to the dimension of the feature vector d and independent of the number of training examples N . Another advantage is that FMs can learn well even on a sparse dataset because they can estimate the weights for feature interactions that are not observed from the training dataset. These advantages are due to the low-rank matrix factorization modeling. In FMs, the weight matrix for feature interactions $\mathbf{W} \in \mathbb{R}^{d \times d}$ is factorized as $\mathbf{W} = \mathbf{P}\mathbf{P}^\top$, $\mathbf{P} \in \mathbb{R}^{d \times k}$, where $k \in \mathbb{N}_{>0}$, called the rank hyperparameter, is usually far smaller than d .

Feature selection methods based on sparse regularization (Tibshirani, 1996; Friedman et al., 2010) have been developed to improve the performance and interpretability of FMs (Pan et al., 2016; Xu

Table 1: A summary of sparse regularizers for FMs. $\ell_{2,1}$ (Xu et al., 2016; Zhao et al., 2017) and ℓ_1 (Pan et al., 2016) are those proposed in the existing methods, and $\tilde{\ell}_{1,2}^2$ (TI) and $\ell_{2,1}^2$ (CS) are those proposed in this paper. The TI regularizer is for feature interaction selection and the CS regularizer is for accurate feature selection. TI is an abbreviation of triangular inequality, and CS is an abbreviation of Cauchy-Schwarz.

Regularizer	Formula	Feature interaction selection	Feature selection
$\ell_{2,1}$	$\ \mathbf{P}\ _{2,1} = \sum_{j=1}^d \ \mathbf{p}_j\ _2$	No	Yes
ℓ_1	$\ \mathbf{P}\ _1 = \sum_{j=1}^d \sum_{s=1}^k p_{j,s} $	No	Yes
$\tilde{\ell}_{1,2}^2$ (TI)	$\ \mathbf{P}^\top\ _{1,2}^2 = \sum_{s=1}^k \ \mathbf{p}_{:,s}\ _1^2$	Yes	Yes
$\ell_{2,1}^2$ (CS)	$\ \mathbf{P}\ _{2,1}^2 = \left(\sum_{j=1}^d \ \mathbf{p}_j\ _2\right)^2$	No	Yes

et al., 2016; Zhao et al., 2017). When one uses feature selection methods, it is assumed that there are irrelevant features. However, because FMs use second-order feature interactions, not only feature-level but also interaction-level relevance should be considered. Actually, FMs with feature selection can work well only if all relevant interactions are those among a subset of features, but this is not necessarily the case in practice. In such cases, FMs with regularization specially designed for *feature interaction selection* trying to achieve interaction-level sparsity may be preferred instead of those just for feature selection trying to achieve feature-level sparsity. Technically speaking, the existing methods select features by inducing row-wise sparsity in \mathbf{P} , often leading to undesirable row/column-wise sparsity in \mathbf{W} as a result. We would like to remove such row/column-wise sparsity-pattern restrictions in the interaction-level sparsity modeling.

In this paper, we present a new regularization scheme for feature interaction selection in FMs. The proposed regularizer is intended to make \mathbf{P} sparse and \mathbf{W} sparse without row/column-wise sparsity-pattern restrictions, which means that our proposed regularizer induces interaction-level sparsity in \mathbf{W} . Our basic objectives are to design regularizers of \mathbf{P} imposing a penalty on the density (in some sense) of \mathbf{W} and to develop efficient algorithms for the associated optimization problems. We will see that our regularizer comes from mathematical analyses of norms (of matrices) and their squares in association with upper bounds of the ℓ_1 norm for \mathbf{W} . In addition, we will discuss how our ideas can be applied or extended to better feature selection (in terms of the prediction performance) and other related models such as higher-order FMs and the all-subsets model. The experiments on synthetic datasets demonstrated that one of the proposed methods (but no other existing ones) succeeded in feature interaction selection in FMs and all the proposed methods performed feature selection more accurately in the cases where all relevant interactions are those among a subset of features. Moreover, the experiments on real-world datasets demonstrated that the proposed methods tend to be easier to use and select more relevant interactions and features for predictions than the existing methods.

This paper is organized as follows. In Section 2, we review FMs and sparse FMs. Section 3 presents our basic idea and analyses for constructing regularizers that make \mathbf{W} sparse. In accordance with them, we propose a new regularizer for feature interaction selection in Section 4 and then another regularizer for better feature selection (in terms of the prediction performance) in Section 5, as well as efficient proximal optimization methods for these proposed regularizers. We extend the proposed regularizers to related models: higher-order FMs and the all-subsets model (Blondel et al., 2016a) in Section 6. In Section 7, we discuss related work. In Section 8, we provide the experimental results on two synthetic and four real-world datasets before concluding in Section 9.

Table 1 summarizes existing and proposed methods (regularizers).

Notation. We denote $\{1, 2, \dots, n\}$ as $[n]$. We use \circ for the element-wise product (a.k.a Hadamard product) of the vector and matrix. We denote the ℓ_p norm for vector and matrix as $\|\cdot\|_p$. Given a

matrix \mathbf{X} , we use \mathbf{x}_i for the i -th row vector and $\mathbf{x}_{:,i}$ for the i -th column vector. Given a matrix $\mathbf{X} \in \mathbb{R}^{n \times m}$, we denote the ℓ_q norm of the vector $(\|\mathbf{x}_1\|_p, \dots, \|\mathbf{x}_n\|_p)^\top$ by $\|\mathbf{X}\|_{p,q} := \left(\sum_{i=1}^n \|\mathbf{x}_i\|_p^q\right)^{1/q}$ and call it $\ell_{p,q}$ norm. We use the terms $\tilde{\ell}_p$ and $\tilde{\ell}_{p,q}$ norm for ℓ_p and $\ell_{p,q}$ norm for the transpose matrix, i.e., $\|\mathbf{P}^\top\|_p$ and $\|\mathbf{P}^\top\|_{p,q}$, respectively. For the number of non-zero elements in vectors $(|\{i : x_i \neq 0\}|)$ and matrices $(|\{(i,j) : x_{i,j} \neq 0\}|)$, we use $\text{nnz}(\cdot)$. We define $\text{supp}(\mathbf{x})$, called the support for \mathbf{x} , as the indices of non-zero elements in $\mathbf{x} \in \mathbb{R}^d$: $\{i \in [d] : x_i \neq 0\}$. We define $\text{abs}(\mathbf{x}) : \mathbf{x} \in \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}^n$ as $\text{abs}(\mathbf{x}) = (|x_1|, \dots, |x_n|)^\top$.

2. Factorization Machines and Sparse Factorization Machines

In this section, we briefly review FMs (Rendle, 2010, 2012) and sparse FMs (Pan et al., 2016; Xu et al., 2016; Zhao et al., 2017).

2.1 Factorization Machines

FMs (Rendle, 2010, 2012) are models for supervised learning based on second-order feature interactions. For a given feature vector $\mathbf{x} \in \mathbb{R}^d$, FMs predict the target of \mathbf{x} as

$$f_{\text{FM}}(\mathbf{x}; \mathbf{w}, \mathbf{P}) := \langle \mathbf{w}, \mathbf{x} \rangle + \sum_{j_2 > j_1} \langle \mathbf{p}_{j_1}, \mathbf{p}_{j_2} \rangle x_{j_1} x_{j_2} = \langle \mathbf{w}, \mathbf{x} \rangle + \frac{1}{2} \sum_{j_1=1}^d \sum_{j_2 \in [d] \setminus \{j_1\}} \langle \mathbf{p}_{j_1}, \mathbf{p}_{j_2} \rangle x_{j_1} x_{j_2}, \quad (1)$$

where $\mathbf{w} \in \mathbb{R}^d$ and $\mathbf{P} \in \mathbb{R}^{d \times k}$ are learnable parameters, and $k \in \mathbb{N}_{>0}$ is the rank hyperparameter. The first term in (1) represents the linear relationship, and the second term represents the second-order polynomial relationship between the input and target. For a given training dataset $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$, the objective function of the FM is

$$L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) := \frac{1}{N} \sum_{n=1}^N \ell(f_{\text{FM}}(\mathbf{x}_n), y_n) + \lambda_w \|\mathbf{w}\|_2^2 + \lambda_p \|\mathbf{P}\|_2^2, \quad (2)$$

where $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is the μ -smooth (i.e., its derivative is a μ -Lipschitz) convex loss function, and $\lambda_w, \lambda_p \geq 0$ are the regularization-strength hyperparameters.

The inner product of the j_1 -th and j_2 -th row vectors in \mathbf{P} , $\langle \mathbf{p}_{j_1}, \mathbf{p}_{j_2} \rangle$, corresponds to the weight for the interaction between the j_1 -th and j_2 -th features in the FM. Thus, FMs are equivalent to the following linear model with a second-order polynomial term (we call it (distinct) quadratic regression (QR) in this paper) with factorization of the feature interaction matrix $\mathbf{W} = \mathbf{P}\mathbf{P}^\top$:

$$f_{\text{QR}}(\mathbf{x}; \mathbf{w}, \mathbf{W}) = \langle \mathbf{w}, \mathbf{x} \rangle + \sum_{j_2 > j_1} w_{j_1, j_2} x_{j_1} x_{j_2} = \langle \mathbf{w}, \mathbf{x} \rangle + \frac{1}{2} \sum_{j_1=1}^d \sum_{j_2 \in [d] \setminus \{j_1\}} w_{j_1, j_2} x_{j_1} x_{j_2}, \quad (3)$$

where $\mathbf{W} \in \mathbb{R}^{d \times d}$ is the feature interaction matrix. The computational cost for evaluating FMs is $O(\text{nnz}(\mathbf{x})k)$, i.e., it is linear w.r.t the dimension d of feature vectors, because the second term in Equation (1) can be rewritten as

$$\sum_{j_2 > j_1} \langle \mathbf{p}_{j_1}, \mathbf{p}_{j_2} \rangle x_{j_1} x_{j_2} = \sum_{s=1}^k (\langle \mathbf{p}_{:,s}, \mathbf{x} \rangle^2 - \langle \mathbf{p}_{:,s}, \mathbf{x} \rangle \langle \mathbf{p}_{:,s}, \mathbf{x} \rangle) / 2. \quad (4)$$

On the other hand, the QR clearly requires $O(\text{nnz}(\mathbf{x})^2)$ time and $O(d^2)$ space for storing \mathbf{W} , which is prohibitive for a high-dimensional case. Moreover, this factorized representation enables FMs to

learn the weights for unobserved feature interactions but the QR does not learn such weights (Rendle, 2010).

The objective function in Equation (2) is differentiable, so Rendle (2010) developed a stochastic gradient descent (SGD) algorithm for minimizing (2). Although the objective function is non-convex w.r.t \mathbf{P} , it is multi-convex w.r.t \mathbf{p}_j for all $j \in [d]$. It can thus be efficiently minimized by using a coordinate descent (CD) (a.k.a alternating least squares) algorithm (Rendle et al., 2011; Blondel et al., 2016b). Both the SGD and CD algorithms require $O(\text{nnz}(\mathbf{X})k)$ time per epoch (using all instances at one time in the SGD algorithm and updating all parameters at one time in the CD algorithm), where $\mathbf{X} \in \mathbb{R}^{N \times d}$ is the design matrix. It is linear w.r.t both the number of training examples N and the dimension of feature vector d .

2.2 Sparse Factorization Machines

Feature selection methods based on sparse regularization (Tibshirani, 1996; Yuan and Lin, 2006; Bach et al., 2012) have been developed to improve the performance and interpretability of FMs (Pan et al., 2016; Xu et al., 2016; Zhao et al., 2017). Selecting features necessarily means making the weight matrix \mathbf{W} row/column-wise sparse.

Xu et al. (2016) and Zhao et al. (2017) proposed using $\|\cdot\|_{2,1}$ regularization, it is well-known as group-lasso regularization (Friedman et al., 2010; Yuan and Lin, 2006). We call FMs with this regularization $\ell_{2,1}$ -sparse FMs. The objective function of this FM is $L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \tilde{\lambda}_p \|\mathbf{P}\|_{2,1}$, where $\tilde{\lambda}_p \geq 0$ is the regularization hyperparameter.¹ Xu et al. (2016) and Zhao et al. (2017) proposed the proximal block coordinate descent (PBCD) and proximal gradient descent (PGD) algorithms respectively, for minimizing this objective function. In our setting, at each iteration, the PBCD algorithm updates the j -th row vector \mathbf{p}_j by

$$\mathbf{p}_j \leftarrow \text{prox}_{\eta\tilde{\lambda}_p\|\cdot\|_2}(\mathbf{p}_j - \eta\nabla_{\mathbf{p}_j}L) = \arg \min_{\mathbf{q}} \frac{1}{2} \|\mathbf{q} - (\mathbf{p}_j - \eta\nabla_{\mathbf{p}_j}L)\|_2^2 + \eta\tilde{\lambda}_p \|\mathbf{q}\|_2 \quad (5)$$

$$= \max(1 - \eta\tilde{\lambda}_p / \|\mathbf{p}'_j\|_2, 0) \cdot \mathbf{p}'_j, \quad (6)$$

where $\eta > 0$ is the step size parameter and $\mathbf{p}'_j := \mathbf{p}_j - \eta\nabla_{\mathbf{p}_j}L_{\text{FM}}(\mathbf{P})$. This proximal algorithm produces row-wise sparse parameter matrix \mathbf{P} . When $\mathbf{p}_j = \mathbf{0}$, $\langle \mathbf{p}_j, \mathbf{p}_i \rangle$ clearly equals zero for all $i \in [d]$. This means that the feature interaction matrix $\mathbf{W} = \mathbf{P}\mathbf{P}^\top$ is row/column-wise sparse, so the FM ignores all feature interactions that involve j -th feature, i.e., $\ell_{2,1}$ regularizer enables feature selection in FMs.

Pan et al. (2016) proposed using ℓ_1 ($=\ell_{1,1}$) regularization for \mathbf{P} . We call FMs with ℓ_1 regularized objective function ℓ_1 -sparse FMs². The PGD update for $p_{j,s}$ in ℓ_1 -sparse FMs is given by

$$p_{j,s} \leftarrow \text{prox}_{\eta\tilde{\lambda}_p|\cdot|}(p'_{j,s}) = \arg \min_q \frac{1}{2}(q - p'_{j,s})^2 + \eta\tilde{\lambda}_p |q| \quad (7)$$

$$= \text{sign}(p'_{j,s}) \cdot \max(|p'_{j,s}| - \eta\tilde{\lambda}_p, 0). \quad (8)$$

ℓ_1 -sparse FMs are intended to make not feature interaction matrix $\mathbf{W} = \mathbf{P}\mathbf{P}^\top$ row/column-wise sparse but \mathbf{P} sparse. However, they practically work well for feature selection in FMs (Pan et al., 2016).

1. There are several differences between our formulations of $\ell_{2,1}$ -sparse FMs and the original ones. First, the original formulations (Xu et al., 2016; Zhao et al., 2017) did not introduce the standard ℓ_2^2 regularization $\lambda_p \|\mathbf{P}\|_2^2$ while ours do because setting λ_p to zero reproduces the original formulations. Moreover, they modified the models or assumed some additional information according to some domain specific knowledge. We do not modify the models and do not assume such information since we do not specify any application.

2. Strictly speaking, Pan et al. (2016) considered probabilistic FMs and proposed the use of a Laplace prior. It corresponds to a ℓ_1 regularization in our non-probabilistic formulation.




Figure 1: Comparison of proximal operators associated with **L1**, and **L21** regularizer. We evaluated the proximal operators at a randomly sampled \mathbf{P} with various λ . Left graph shows the number of used feature interactions and right graph shows the number of used features in $\mathbf{Q}^*(\mathbf{Q}^*)^\top$, where \mathbf{Q}^* is the output of the proximal operator: $\mathbf{Q}^* = \text{prox}_{\lambda\Omega(\cdot)}(\mathbf{P})$ and Ω is ℓ_1 (**L1** (Pan et al., 2016)) or $\ell_{2,1}$ (**L21** (Xu et al., 2016; Zhao et al., 2017)).

3. Proposed Scheme for Feature Interaction Selection

Feature Interaction Selection: its Motivation. The existing sparse regularizers (Pan et al., 2016; Zhao et al., 2017; Xu et al., 2016) can improve the performance and interpretability of FMs by selecting only relevant features. However, because FMs use second-order feature interactions, not feature-level but interaction-level relevance should be considered, in other words, feature interaction selection is preferable to feature selection. Assume that a subset S of $[d]$ is selected as a set of relevant features. Then, for all $j \in [d] \setminus S$, all feature interactions with j -th feature are lost but they can contain important feature interactions. Moreover, FMs use all feature interactions from S and they can contain irrelevant feature interactions. Therefore, the existing methods tend to produce all-zeros \mathbf{W} to remove all irrelevant interactions or all-non-zeros (dense) \mathbf{W} to select all relevant interactions. Many relevant interactions are lost in the former case and many irrelevant interactions are used in the latter case. Actually, FMs with feature selection can work well only if all relevant interactions are those among a subset of features, but this is not necessarily the case in practice.

In this section, firstly, we briefly verify whether the existing regularizers based on sparsity-inducing norms can select feature interactions or not, experimentally. We next introduce a preferable but hard to optimize regularizer Ω_* for feature interaction selection in FMs. Then, we present a relationship between norms and Ω_* . We next present a relationship between *squares* of norms and Ω_* , and finalize our scheme: using the square of a *sparsity-inducing* (quasi-)norm.

3.1 Can ℓ_1 and $\ell_{2,1}$ Regularizers Select Feature Interactions?

We verify whether the existing regularizers can select feature interactions or not. Because the objective functions with the existing regularizers are typically optimized by the PGD algorithm, we compared the output of their proximal operators for verification. We sampled $\mathbf{P} \in \mathbb{R}^{200 \times 30}$ with $p_{j,s} \sim \mathcal{N}(0, 1.0)$ for all $j \in [200]$, $s \in [30]$, and evaluated proximal operators $\mathbf{Q}^* = \text{prox}_{\lambda\Omega(\cdot)}(\mathbf{P})$ with various Ω : ℓ_1 (**L1** (Pan et al., 2016)) and $\ell_{2,1}$ (**L21** (Xu et al., 2016; Zhao et al., 2017)). Their corresponding proximal operators are (8) and (6) respectively. Regarding \mathbf{P} as the parameter of FMs, we computed the number of used interactions (i.e., $|\{(j_1, j_2) : \langle \mathbf{q}_{j_1}^*, \mathbf{q}_{j_2}^* \rangle \neq 0, j_2 > j_1\}|$) and the number of used features (i.e., the number of non-zero rows in \mathbf{Q}^*). We set λ to be $2^{-7}, 2^{-6}, \dots$, and 2^7 .

Results are shown in Fig. 1. Both **L1** and **L21** tended to produce completely dense (all-non-zeros) or all-zeros feature interaction matrices. This result indicates that it is difficult for the existing methods to select feature interactions. We will see later in Section 8.1, the proximal operators of the proposed regularizers can produce moderately sparse feature interaction matrices and therefore more useful for feature interaction selection and feature selection in FMs. Moreover, we will show that one of the proposed regularizers can select relevant feature interactions and all of the proposed regularizers can select relevant features in Section 8.2.

3.2 ℓ_1 Norm for Feature Interaction Weight Matrix

We here introduce a preferable but hard to optimize regularizer for feature interaction selection in FMs. Selecting feature interactions necessarily means making the feature interaction weight matrix \mathbf{W} sparse, so our goal is learning sparse $\mathbf{W} = \mathbf{P}\mathbf{P}^\top$ in FMs. Although the existing regularizers are intended to make \mathbf{P} sparse, the sparsity of \mathbf{P} does not necessarily imply the sparsity of \mathbf{W} . Thus, our basic idea is to use a regularization inducing sparsity \mathbf{W} rather than \mathbf{P} . Especially, we propose using ℓ_1 regularization (Tibshirani, 1996) for the strictly upper triangular elements (or equivalently, non-diagonal elements) in \mathbf{W} , i.e.,

$$\sum_{j_2 > j_1} |w_{j_1, j_2}| = \sum_{j_2 > j_1} |\langle \mathbf{p}_{j_1}, \mathbf{p}_{j_2} \rangle| =: \Omega_*(\mathbf{P}), \quad (9)$$

because ℓ_1 regularization is the well-known and one of the promising regularization for inducing sparsity. The corresponding objective function is

$$L_{\text{FM}}^*(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p, \tilde{\lambda}_p) := L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \tilde{\lambda}_p \Omega_*(\mathbf{P}). \quad (10)$$

Unfortunately, this objective function is hard to optimize w.r.t \mathbf{P} . In the following, we introduce three well-known algorithms for minimizing a sum of a differentiable loss and a non-smooth regularization like (10) and show that the use of them is unrealistic.

Subgradient Descent Algorithm. Consider the use of the subgradient descent (SubGD) algorithm for minimizing (10). Ω_* is non-convex and thus its subdifferential cannot be defined. Fortunately, $\|\mathbf{P}\mathbf{P}^\top\|_1/2 = \Omega_*(\mathbf{P}) + \|\mathbf{P}\|_2^2/2$ is convex, so its subdifferential can be defined. Therefore, consider the minimization of $L_{\text{FM}}^*(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p + \tilde{\lambda}_p/2, \tilde{\lambda}_p)$, i.e.,

$$L_{\text{FM}}^*(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p + \tilde{\lambda}_p/2, \tilde{\lambda}_p) = L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \frac{\tilde{\lambda}_p}{2} \|\mathbf{P}\mathbf{P}^\top\|_1. \quad (11)$$

At each iteration, the SubGD algorithm for minimizing (11) picks a subgradient $\mathbf{G} \in \partial_{\mathbf{P}} \|\mathbf{P}\mathbf{P}^\top\|_1$ and updates the parameter \mathbf{P} as

$$\mathbf{P} \leftarrow \mathbf{P} - \eta \left(\nabla_{\mathbf{P}} L + \frac{\tilde{\lambda}_p}{2} \mathbf{G} \right). \quad (12)$$

The subdifferential of $\|\mathbf{P}\mathbf{P}^\top\|_1$ is defined as (Li et al., 2020)

$$\partial \|\mathbf{P}\mathbf{P}^\top\|_1 = \{2\mathbf{Z}\mathbf{P} : \mathbf{Z} \in \partial_{\mathbf{S}} \|\mathbf{S}\|_1, \mathbf{S} = \mathbf{P}\mathbf{P}^\top\}. \quad (13)$$

Therefore, picking a subgradient $\mathbf{G} \in \partial_{\mathbf{P}} \|\mathbf{P}\mathbf{P}^\top\|_1$ requires $O(d^2k)$ computational cost (for computing $\mathbf{P}\mathbf{P}^\top$), so it might be prohibitive to use the SubGD algorithm for a high-dimensional case. To be more precise, the computational cost of the SubGD algorithm at each iteration is $O(T(\text{mz}(\mathbf{X})k + d^2k))$, where T is the number of line search iterations. Moreover, in general, the SubGD algorithm cannot produce a sparse solution and therefore it is not suitable for feature interaction selection (Bach et al., 2012).

Inexact PGD Algorithm. To obtain a sparse \mathbf{W} , we consider the use of a PGD algorithm for (11). At each iteration, the PGD algorithm for minimizing (11) requires the evaluation of the following proximal operator

$$\text{prox}_{\lambda\|\cdot\|_1}(\mathbf{P}) := \arg \min_{\mathbf{Q} \in \mathbb{R}^{d \times k}} \frac{1}{2} \|\mathbf{P} - \mathbf{Q}\|_2^2 + \lambda \|\mathbf{Q}\mathbf{Q}^\top\|_1 \quad (14)$$

with some $\lambda > 0$. This proximal problem (14) is convex but unfortunately cannot be evaluated analytically. The PGD algorithm with *inexact* evaluation of proximal operator is called inexact PGD algorithm and we here consider the use of the inexact PGD algorithm (Gu et al., 2018; Yao et al., 2017; Schmidt et al., 2011). Because (14) is a $(d \times k)$ -dimensional non-smooth convex optimization problem, for example, it can be optimized using the SubGD algorithm (Gu et al., 2018). Unfortunately, the SubGD algorithm for (14) requires $O(d^2k)$ computational cost at each iteration, which is free from \mathbf{X} but depends on d quadratically. Thus, one iteration of this inexact PGD algorithm for (11) using the SubGD algorithm for inexact evaluation of (14) takes $O(T(\text{nnz}(\mathbf{X})k + T'd^2k))$ time, where T is the number of line search iteration and T' is the number of iterations of the SubGD algorithm for (14). Moreover, the precision of the inexact proximal operator should be high in practice and must be controlled carefully for convergence. Furthermore, the convergence rate of the SubGD algorithm for a convex optimization problem is $O(1/\epsilon^2)$ (Nesterov et al., 2018). Thus, we must set T' to be a large value for a good solution, and then it is also not practical for (11) to use the inexact PGD algorithm for a high-dimensional case.

Inexact PCD Algorithm. If we assume the use of proximal CD (PCD) algorithm for the regularized objective (10), the proximal operator evaluated at $p'_{j,s} := p_{j,s} - \eta \partial L_{\text{FM}}(\mathbf{P}) / \partial p_{j,s}$ is

$$\arg \min_q \left\{ \frac{1}{2} (q - p'_{j,s})^2 + \eta \tilde{\lambda}_p \sum_{i \in [d] \setminus \{j\}} |qp_{i,s} + r_{j,i,s}| \right\}, \quad (15)$$

where $r_{j,i,s} = \sum_{s' \in [k] \setminus \{s\}} p_{j,s'} p_{i,s'}$. The second term in Equation (15) takes the form of the sum of the absolute deviations and can be rewritten as a d -dimensional linear programming problem with inequality constraints (Boyd and Vandenberghe, 2004). Thus, the optimization problem in this proximal operator is a typical quadratic programming problem and can be solved by some well-known methods (e.g., an interior-point method). Alternatively, one can also solve (15) using the SubGD algorithm or the alternating direction of direction method of multipliers (ADMM) algorithm (Boyd et al., 2011). However, in any case, $r_{i,j,s}$ must be computed for all $i \in [d] \setminus \{j\}$ and it requires $O(dk)$ computational cost. Thus, the inexact PCD algorithm for (10) requires $\Omega(dk(dk)) = \Omega(d^2k^2)$ computational cost only for evaluating (15) per epoch. It might be prohibitive for a high-dimensional case.

3.3 Upper Bound Regularizers of Ω_*

As described above, Ω_* regularizer seems appropriate for feature interaction selection but unfortunately it is hard to optimize. Thus, we consider the use of an upper bound regularizer being easy to optimize and ensuring sparsity to \mathbf{W} . The use of an easy-to-optimize upper bound is a common approach for minimizing a hard-to-optimize objective function (Zhou et al., 2012; Liu, 2011; Chen et al., 2019a).

3.3.1 NON-EQUIVALENCE OF NORMS AND Ω_*

Firstly, are existing $\ell_{2,1}$ and ℓ_1 regularizers upper bounds of Ω_* ? Unfortunately, not only them but also any norm on $\mathbb{R}^{d \times k}$ can be neither an upper bound nor a lower bound; i.e., all norms are not equivalent to Ω_* .

Theorem 1 Let $\|\cdot\|$ be a norm on $\mathbb{R}^{d \times k}$. Then, for any $C > 0$, there exist $\mathbf{P}, \mathbf{Q} \in \mathbb{R}^{d \times k}$ such that

$$\Omega_*(\mathbf{P}) < C \|\mathbf{P}\|, \Omega_*(\mathbf{Q}) > C \|\mathbf{Q}\|. \quad (16)$$

Proof Since $\|\cdot\|$ is a norm on $\mathbb{R}^{d \times k}$, it is absolutely homogeneous $\|a\mathbf{P}\| = |a| \|\mathbf{P}\|$ for all $a \in \mathbb{R}$ and $\|\mathbf{P}\| = 0 \iff \mathbf{P} = \mathbf{0}$. On the other hand, Ω_* is 2-homogeneous:

$$\Omega_*(a\mathbf{P}) = \sum_{j_1=1}^d \sum_{j_2 > j_1}^k |\langle a\mathbf{p}_{j_1}, a\mathbf{p}_{j_2} \rangle| = a^2 \Omega_*(\mathbf{P}), \quad (17)$$

and $\Omega_*(\mathbf{P}) = 0 \iff \langle \mathbf{p}_{j_1}, \mathbf{p}_{j_2} \rangle = 0$ for all $j_1 \neq j_2 \in [d]$. Thus, we can take $\mathbf{P}' \in \mathbb{R}^{d \times k}$ such that $\|\mathbf{P}'\| \neq 0$ and $\Omega_*(\mathbf{P}') \neq 0$. Given $C > 0$, we take a positive number a such that $0 < a < C \|\mathbf{P}'\| / \Omega_*(\mathbf{P}')$. Then, $\Omega_*(a\mathbf{P}') = a^2 \Omega_*(\mathbf{P}') < a [C \|\mathbf{P}'\| / \Omega_*(\mathbf{P}')] \Omega_*(\mathbf{P}') = C \|a\mathbf{P}'\|$, which is surely the first inequality in (16) ($\mathbf{P} = a\mathbf{P}'$). Similarly, if we take $a > C \|\mathbf{P}'\| / \Omega_*(\mathbf{P}')$, we can derive the second inequality in (16). \blacksquare

In some cases, the fact that any norm cannot be an upper bound of Ω_* is crucial. Suppose that one wants FMs with \mathbf{P} such that $\Omega_*(\mathbf{P}) \leq \lambda$; i.e., one solves the constrained minimization problem. Since this problem is also hard to optimize, one can replace Ω_* with $\|\cdot\|$, and the revised problem may be easier to optimize. However, it is not guaranteed that the solution \mathbf{P} satisfies $\Omega_*(\mathbf{P}) \leq \lambda$ because $\|\cdot\|$ cannot be an upper bound of Ω_* .

The existing methods using sparsity-inducing norms produce completely dense (all-non-zeros) or all-zeros feature interaction matrices as shown in Fig. 1. This phenomenon can be explained by Theorem 1. From the proof of Theorem 1, we have $\|\mathbf{P}\| \gg \Omega_*(\mathbf{P})$, i.e, the regularization strength of norm is much greater than that of Ω_* , if the absolute value of each element in \mathbf{P} is sufficiently small. Thus, when λ is large, the existing methods using norm regularizers can produce all-zeros matrices. Similarly, we have $\|\mathbf{P}\| \ll \Omega_*(\mathbf{P})$ if the absolute value of each element in \mathbf{P} is sufficiently large. Thus, when λ is small, the existing methods using norm regularizers can produce completely dense matrices.

3.3.2 UPPER BOUND REGULARIZERS OF Ω_* BY SQUARES OF (QUASI-)NORMS

In this section, we present how to construct an upper bound of Ω_* . We first define m -homogeneous quasi-norms.

Definition 2 We say a function $\Omega : \mathbb{R}^{d \times k} \rightarrow \mathbb{R}_{\geq 0}$ is an m -homogeneous quasi-norm if, for all $\mathbf{P}, \mathbf{Q} \in \mathbb{R}^{d \times k}$, (i) $\Omega(\mathbf{P}) \geq 0$, $\Omega(\mathbf{P}) = 0 \iff \mathbf{P} = \mathbf{0}$, (ii) there exists $m \in \mathbb{N}_{>0}$ for all $a \in \mathbb{R}$ such that $\Omega(a\mathbf{P}) = |a|^m \Omega(\mathbf{P})$, and (iii) there exists $K > 0$ ($K \geq 2^{m-1}$) such that $\Omega(\mathbf{P} + \mathbf{Q}) \leq K(\Omega(\mathbf{P}) + \Omega(\mathbf{Q}))$. Note that $m = 1$ implies that Ω is a quasi-norm.

There is an important relationship between Ω_* and 2-homogeneous quasi-norms: unlike norms, any 2-homogeneous quasi-norm can be an upper bound of Ω_* .

Theorem 3 For any 2-homogeneous quasi-norm Ω , there exists $C > 0$ such that $\Omega_*(\mathbf{P}) \leq C\Omega(\mathbf{P})$ for all $\mathbf{P} \in \mathbb{R}^{d \times k}$.

Moreover, one can construct an m -homogeneous quasi-norm by m -th power of a (quasi-)norm.

Theorem 4 $\Omega : \mathbb{R}^{d \times k} \rightarrow \mathbb{R}_{\geq 0}$ is an m -homogeneous quasi-norm if and only if there exists a quasi-norm $\|\cdot\|'$ such that $\Omega(\cdot) = (\|\cdot\|')^m$.

Thus, one can construct an upper bound of Ω_* by the square of a (quasi-)norm. The regularizer in canonical FMs, ℓ_2^2 , is clearly the square of the norm but it does not produce sparse feature interaction matrix $\mathbf{W} = \mathbf{P}\mathbf{P}^\top$ (as indicated in our experimental results in Section 8). It means that using an

upper bound of Ω_* is not sufficient for feature interaction selection. We thus propose especially using the square of a *sparsity-inducing* (quasi-)norm, which can make $\mathbf{W} = \mathbf{P}\mathbf{P}^\top$ sparse by making \mathbf{P} sparse since $\text{supp}(\mathbf{p}_{j_1}) \cap \text{supp}(\mathbf{p}_{j_2}) = \emptyset$ implies $w_{j_1, j_2} = 0$. As described above, the sparsity of \mathbf{P} does not necessarily imply the sparsity of \mathbf{W} . However, the square of a sparsity-inducing (quasi-)norm can be more useful for feature interaction selection than a sparsity-inducing norm because the square of a norm can be an upper bound of Ω_* . In the following sections, we present such regularizers.

3.4 Comparison of Norm and Squared Norm

Before presenting the proposed regularizers, we discuss relationships between regularizations based on squared norms and those based on norms. Consider the optimization problem with the regularization based on the squared norm

$$\min_{\mathbf{w}, \mathbf{P}} L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \lambda \|\mathbf{P}\|^2 \quad (18)$$

and one of its stationary points $\{\hat{\mathbf{w}}, \hat{\mathbf{P}}\}$. Then, $\{\hat{\mathbf{w}}, \hat{\mathbf{P}}\}$ is also one of the stationary points of the optimization problem with the regularization based on the (non-squared) norm

$$\min_{\mathbf{w}, \mathbf{P}} L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \left(2\lambda \|\hat{\mathbf{P}}\|\right) \|\mathbf{P}\| \quad (19)$$

since $\partial \|\mathbf{P}\|^2 = 2\|\mathbf{P}\| \partial \|\mathbf{P}\|$ (from Theorem 12 in Appendix A.1). Therefore, one might consider that the use of the squared norm $\|\cdot\|^2$ is essentially equivalent to the use of the (non-squared) norm $\|\cdot\|$. However, the optimal solutions of (18) and (19) are not necessarily equivalent to each other because $L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p)$ is non-convex w.r.t \mathbf{P} . To the best of our knowledge, there are no known relationships between the optimal solutions of the existing methods (Pan et al., 2016; Xu et al., 2016; Zhao et al., 2017) and those of the QR with the ℓ_1 regularization for \mathbf{W} . On the other hand, under some conditions, the optimal solutions of one of the proposed methods are equivalent to those of the QR with the ℓ_1 regularization for \mathbf{W} , which is shown in Section 4.

As described above, the squared norm regularization is not necessarily equivalent to the norm regularization. On the other hand, the squared norm regularization can be interpreted as the norm regularization with an adaptive regularization-strength since $\lambda \|\mathbf{P}\|^2 = (\lambda \|\mathbf{P}\|) \|\mathbf{P}\|$. Indeed, as we will see later in Section 4 and Section 5, the proposed TI (CS) regularizer based on squared norms shrink elements (row vectors) of relatively small absolute value (ℓ_2 norm) to 0 ($\mathbf{0}$). On the other hand, the existing methods shrink absolutely small elements to 0 (i.e., elements that are smaller than a fixed threshold are shrunk to 0). Therefore, regularizers based on squared norms (the proposed methods) can be less sensitive to the choice of the regularization-strength hyperparameter than those based on (non-squared) norms (the existing methods).

4. TI Upper Bound Regularizer

We first propose using $\tilde{\ell}_{1,2}^2$ as an upper bound regularizer of Ω_* . We call this regularizer $\tilde{\ell}_{1,2}^2$ regularizer or triangular inequality (TI) regularizer and call FMs with this regularization $\tilde{\ell}_{1,2}^2$ -sparse FMs or TI-sparse FMs since this regularization can also be derived using the triangle inequality:

$$\frac{1}{2} \|\mathbf{P}^\top\|_{1,2}^2 = \frac{1}{2} \sum_{s=1}^k \|\mathbf{p}_{:,s}\|_1^2 = \sum_{j_2 > j_1} \sum_{s=1}^k |p_{j_1,s} p_{j_2,s}| + \frac{1}{2} \|\mathbf{P}\|_2^2 \quad (20)$$

$$=: \Omega_{\text{TI}}(\mathbf{P}) + \frac{1}{2} \|\mathbf{P}\|_2^2 \geq \Omega_*(\mathbf{P}) + \frac{1}{2} \|\mathbf{P}\|_2^2. \quad (21)$$

Because $\|\mathbf{P}^\top\|_{1,2}^2 = 2\Omega_{\text{TI}}(\mathbf{P}) + \|\mathbf{P}\|_2^2$ and $\|\mathbf{P}\|_2^2$ can be taken into L_{FM} , we will sometimes discuss not TI-sparse FMs but rather Ω_{TI} -sparse FMs (FMs with Ω_{TI} regularization).

We here discuss the relationship between Ω_{TI} -sparse FMs and the QR (3) with ℓ_2^2 regularization for \mathbf{w} and ℓ_1 regularization for \mathbf{W} . Theorem 5 states that the optimal Ω_{TI} -sparse FMs are equivalent (or better in the sense of the objective value) to the optimal QR with such regularizations when the rank hyperparameter k is sufficiently large. We also obtain a similar relationship between Ω_{TI} -sparse FMs and Ω_* -sparse FMs. It is shown in Appendix A.5. Note that Ω_{TI} -sparse FMs can be regarded as TI-sparse FMs ($\tilde{\ell}_{1,2}^2$ -sparse FMs) when $\lambda_p \geq \tilde{\lambda}_p/2$.

Theorem 5 *Let $L_{\text{QR}}(\mathbf{w}, \mathbf{W}; \lambda_w)$ be the objective function of the QR with ℓ_2^2 regularization for \mathbf{w} :*

$$L_{\text{QR}}(\mathbf{w}, \mathbf{W}; \lambda_w) := \sum_{n=1}^N \ell(f_{\text{QR}}(\mathbf{x}_n), y_n)/N + \lambda_w \|\mathbf{w}\|_2^2. \quad (22)$$

Then, for any $\lambda_w, \lambda_p, \tilde{\lambda}_p \geq 0$, there exists $k' \leq d(d-1)/2$ such that for all $k \geq k'$,

$$\begin{aligned} & \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{P} \in \mathbb{R}^{d \times k}} L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \tilde{\lambda}_p \Omega_{\text{TI}}(\mathbf{P}) \\ & \leq \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{W} \in \mathbb{R}^{d \times d}} L_{\text{QR}}(\mathbf{w}, \mathbf{W}; \lambda_w) + (\tilde{\lambda}_p + 2\lambda_p) \|\mathbf{W}\|_1. \end{aligned} \quad (23)$$

Moreover, if $\lambda_p = 0$, the equality holds, and $f_{\text{FM}}(\mathbf{x}; \mathbf{w}_{\text{TI}}^, \mathbf{P}_{\text{TI}}^*) = f_{\text{QR}}(\mathbf{x}; \mathbf{w}_{\text{QR}}^*, \mathbf{W}_{\text{QR}}^*)$ for all $\mathbf{x} \in \mathbb{R}^d$, where $\{\mathbf{w}_{\text{TI}}^*, \mathbf{P}_{\text{TI}}^*\}$ and $\{\mathbf{w}_{\text{QR}}^*, \mathbf{W}_{\text{QR}}^*\}$ are the solutions of the left- and right-hand sides, respectively, of Equation (23).*

We next consider TI ($\tilde{\ell}_{1,2}^2$) regularized objective function, i.e., $L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \tilde{\lambda}_p \|\mathbf{P}^\top\|_{1,2}^2$. Since $\|\mathbf{P}^\top\|_{1,2}^2 = \|\mathbf{P}\|_2^2 + 2\Omega_{\text{TI}}(\mathbf{P})$, it can be written as the Ω_{TI} regularized objective function: $L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p + \tilde{\lambda}_p) + 2\tilde{\lambda}_p \Omega_{\text{TI}}(\mathbf{P})$. The optimization problem of the TI regularized objective function can be written as

$$\min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{P} \in \mathbb{R}^{d \times k}} L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \tilde{\lambda}_p \|\mathbf{P}^\top\|_{1,2}^2 \quad (24)$$

$$= \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{P} \in \mathbb{R}^{d \times k}} L_{\text{QR}}(\mathbf{w}, \mathbf{P}\mathbf{P}^\top; \lambda_w) + \sum_{s=1}^k \left\{ \lambda_p \|\mathbf{p}_{:,s}\|_2^2 + \tilde{\lambda}_p \|\mathbf{p}_{:,s}\|_1^2 \right\}, \quad (25)$$

$$= \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{W} \in \mathbb{R}^{d \times d}} L_{\text{QR}}(\mathbf{w}, \mathbf{W}; \lambda_w) + \Omega_{\text{MF}}(\mathbf{W}; k, \lambda_p, \tilde{\lambda}_p), \quad (26)$$

where

$$\Omega_{\text{MF}}(\mathbf{W}; k, \lambda_p, \tilde{\lambda}_p) := \inf_{r \in [k]} \inf_{\substack{\mathbf{U}, \mathbf{V} \in \mathbb{R}^{d \times r} \\ \mathbf{U}\mathbf{V}^\top = \mathbf{W}}} \sum_{s=1}^r \left\{ \lambda_p \|\mathbf{u}_{:,s}\|_2^2 + \tilde{\lambda}_p \|\mathbf{u}_{:,s}\|_1^2 + I_{\{\mathbf{u}_{:,s}\}}(\mathbf{v}) \right\}, \quad (27)$$

and $I_C : \mathbb{R}^d \rightarrow \{0, \infty\}$ is the indicator function on $C \subseteq \mathbb{R}^d$. When the rank hyperparameter k is sufficiently large, $\Omega_{\text{MF}}(\mathbf{W}; k, \lambda_p, \tilde{\lambda}_p)$ satisfies the condition as matrix factorization regularizer (Haeffele and Vidal, 2019). Thus, from Proposition 2, Theorem 1, and Theorem 2 in (Haeffele and Vidal, 2019), we conclude that all local minima of the TI regularized objective function are global minima when k is sufficiently large.

Theorem 6 (Haeffele and Vidal (2019)) *When $k \geq d^2$, all local minima of (24) are global minima.*

Note that for any $m \in \mathbb{N}_{>0}$ all local minima of $\tilde{\ell}_{m,2}^2$ regularized optimization problem are global minima under the same condition.

4.1 PGD/PSGD-based Algorithm for TI Regularizer

Because the TI regularizer is continuous but non-differentiable, we consider the use of the PGD algorithm for optimizing TI-sparse FMs similarly to ℓ_1 -sparse FMs and $\ell_{2,1}$ -sparse FMs. The PGD algorithm for TI-sparse FMs solves the following optimization problem at t -th iteration:

$$\min_{\mathbf{Q} \in \mathbb{R}^{d \times k}} \left\langle \nabla \ell_{\mathcal{D}}^{(t-1)}, \mathbf{Q} - \mathbf{P}^{(t-1)} \right\rangle + \frac{1}{2\eta} \left\| \mathbf{Q} - \mathbf{P}^{(t-1)} \right\|_2^2 + \lambda_p \|\mathbf{Q}\|_2^2 + \tilde{\lambda}_p \|\mathbf{Q}^\top\|_{1,2}^2, \quad (28)$$

where $\nabla \ell_{\mathcal{D}}^{(t-1)} := \sum_{n=1}^N \nabla \ell(f_{\text{FM}}(\mathbf{x}_n; \mathbf{P}^{(t-1)}), y_n) / N$ and $\mathbf{P}^{(t-1)}$ is the parameter matrix \mathbf{P} after $(t-1)$ -th iteration (we omit $\mathbf{w}^{(t-1)}$ for simplicity). Fortunately, the TI regularizer is separable w.r.t each column vector: $\|\mathbf{P}^\top\|_{1,2}^2 = \sum_{s=1}^k \|\mathbf{p}_{:,s}\|_1^2$. We thus can separate the optimization problem (28) w.r.t each column as

$$\min_{\mathbf{q}_{:,s} \in \mathbb{R}^d} \left\langle \nabla_{\mathbf{p}_{:,s}} \ell_{\mathcal{D}}^{(t-1)}, \mathbf{q}_{:,s} - \mathbf{p}_{:,s}^{(t-1)} \right\rangle + \frac{1}{2\eta} \left\| \mathbf{q}_{:,s} - \mathbf{p}_{:,s}^{(t-1)} \right\|_2^2 + \lambda_p \|\mathbf{q}_{:,s}\|_2^2 + \tilde{\lambda}_p \|\mathbf{q}_{:,s}\|_{1,2}^2, \quad (29)$$

where $\nabla_{\mathbf{p}_{:,s}} \ell_{\mathcal{D}}^{(t-1)} := \sum_{n=1}^N \nabla_{\mathbf{p}_{:,s}} \ell(f_{\text{FM}}(\mathbf{x}_n; \mathbf{P}^{(t-1)}), y_n) / N$ and solve this separated problem for all $s \in [k]$ at t -th iteration. Therefore, the PGD algorithm for TI-sparse FMs surely solves the following type of optimization problem for all $s \in [k]$:

$$\text{prox}_{\lambda \|\cdot\|_1^2}(\mathbf{p}) = \arg \min_{\mathbf{q} \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{q} - \mathbf{p}\|_2^2 + \lambda \|\mathbf{q}\|_1^2 =: \mathbf{q}^*, \quad (30)$$

where $\lambda \geq 0$. The following theorem gives us the insight needed for constructing an algorithm for computing this proximal operator (Martins et al., 2011).

Theorem 7 (Martins et al. (2011)) *Assume that $\mathbf{p} \in \mathbb{R}^d$ is sorted in descending order by absolute value: $|p_1| \geq |p_2| \geq \dots \geq |p_d|$. Then, the solution to the proximal problem (30) $\mathbf{q}^* \in \mathbb{R}^d$ is*

$$q_j^* = \text{sign}(p_j) \max\{|p_j| - 2\lambda S_\theta, 0\} \quad \forall j \in [d], \quad (31)$$

where $\theta = \max\{j : |p_j| - 2\lambda S_j \geq 0\}$, and $S_j = \sum_{i=1}^j |p_i| / (1 + 2\lambda j)$.

This theorem states that, for arbitrary $\mathbf{p} \in \mathbb{R}^d$, the proximal operator (30) can be computed in $O(d \log d)$ time by first sorting \mathbf{p} by absolute value and then computing S_j for all $j \in [d]$ and θ . In fact, this proximal operator can be evaluated in $O(d)$ time in expectation by randomized-median-finding-like method (Duchi et al., 2008). For more detail, please see our appendix (Algorithm 2 and Algorithm 3 in Appendix B). The proximal operator of ℓ_1 regularizer (8) shrinks each element in $\mathbf{p}_{:,s}$: $q_{j,s}^* = 0$ if $|p_{j,s}| \leq \lambda$ for all $j \in [d]$ and the threshold λ does not depend on \mathbf{P} . Similarly, the proximal operator of $\tilde{\ell}_{1,2}^2$ (30) also shrinks each element in $\mathbf{p}_{:,s}$. However, its threshold depends on $\mathbf{p}_{:,s}$: $q_{j,s}^* = 0$ if $|p_{j,s}| \leq 2\lambda S_\theta$ and S_θ depends on $\mathbf{p}_{:,s}$. That is, intuitively, $\tilde{\ell}_{1,2}^2$ regularizer shrinks elements of relatively small absolute value among $p_{1,s}, \dots, p_{d,s}$ to 0.

Clearly, one can construct a proximal SGD (PSGD) algorithm by replacing $\nabla \ell_{\mathcal{D}}^{(t-1)}$ in (28) with a stochastic gradient. The PSGD-based algorithms are typically more useful than the PGD-based (i.e., batch) algorithms when the number of instances N is large (Bottou, 2012; Bottou et al., 2018; Allen-Zhu, 2018). For the (not necessarily convex) smooth optimization problem $\min_{\mathbf{z} \in \mathbb{R}^d} \sum_{n=1}^N f_i(\mathbf{z}) / N$, $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ for all $i \in [N]$, the GD and the SGD require respectively $O(1/\varepsilon^2)$ and $O(V/\varepsilon^4)$ iterations to get an ε -approximate stationary point (Allen-Zhu, 2018), where V is the upper bound of the variance of the stochastic gradients. In most cases, the GD requires to compute N gradients while the SGD requires one gradient at each iteration. Thus, when N is large and ε is moderate, the SGD is superior to the GD. Unfortunately, the PSGD algorithm and

its variants with Algorithm 3 cannot leverage the sparsity of data: they require $O(dk)$ time for each iteration and $O(Ndk)$ time for each epoch even if a dataset \mathbf{X} is sparse. The $O(dk)$ computational cost is due to the evaluation of the proximal operator (Algorithm 3). A workaround for this issue is the use of mini-batches. The use of mini-batch reduces the variance of the stochastic gradient and it hence reduces the number of iteration for convergence, but in general it increases the computational cost for each iteration. In our setting, if one chooses a mini-batch such that its number of non-zero elements is $O(d)$, the mini-batch PSGD also runs in $O(dk)$ for one iteration, that is, the use of appropriate-size-mini-batches can reduce the variance of the stochastic gradients without changing the computational complexity for one iteration. However, it does not solve the issue completely: the cost for one iteration, $O(dk)$, is not so improved compared to the that of the PGD algorithm $O(\text{nnz}(\mathbf{X})k)$ when \mathbf{X} is very sparse. Thus, the (mini-batch) PSGD algorithm should be used only when $\text{nnz}(\mathbf{X})/d$ is large.

4.2 Efficient PCD Algorithm for TI Regularizer

Here, we present an efficient PCD algorithm for TI-sparse FMs, which is often used for minimizing the objective function with non-smooth regularization and has several advantages compared to the PGD algorithm introduced in Section 4.1. Firstly, it does not require tuning nor using line search technique for the step size, and this is its most important advantage compared with the PGD/PSGD-based algorithms. Secondly, it can leverage sparsity of data: it runs in $O(\text{nnz}(\mathbf{X})k)$ for one epoch. Thirdly, it is easy to implement: its implementation is simple and almost the same as the CD algorithm for canonical FMs and the PCD algorithm for ℓ_1 -sparse FMs. Fourthly, it can be easily extended to other related models as shown in Section 6. Strictly speaking, $\tilde{\ell}_{1,2}^2$ is not separable w.r.t $p_{1,1}, \dots, p_{d,k}$ and thus the convergence of the PCD algorithm is not guaranteed. However, actually, it doesn't matter much that the convergence of the PCD algorithm is not guaranteed. Because the objective function of FMs is non-convex, a global minimum solution cannot be obtained even if the PGD algorithm is used.

Because $\|\mathbf{P}^\top\|_{1,2}^2 = 2\Omega_{\text{TI}}(\mathbf{P}) + \|\mathbf{P}\|_2^2$ and $\|\mathbf{P}\|_2^2$ can be taken into L_{FM} , for simplify we consider Ω_{TI} regularized objective function $L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \tilde{\lambda}_p \Omega_{\text{TI}}(\mathbf{P})$, focusing on one parameter $p_{j,s}$ as the optimized parameter. Then, the Ω_{TI} regularizer can be regarded as a ℓ_1 regularizer such that its regularization strength is $\sum_{i \in [d] \setminus \{j\}} |p_{i,s}|$:

$$\Omega_{\text{TI}}(\mathbf{P}) = \left(\sum_{i \in [d] \setminus \{j\}} |p_{i,s}| \right) |p_{j,s}| + \text{const.} \quad (32)$$

Therefore, given this regularization strength, $\sum_{i \in [d] \setminus \{j\}} |p_{i,s}|$, the procedure of the PCD algorithm for the Ω_{TI} (and of course TI ($\tilde{\ell}_{1,2}^2$)) regularized objective function is the same as that for the ℓ_1 regularized objective function. Fortunately, by caching $c_s := \sum_{j=1}^d |p_{j,s}|$ before updating $p_{1,s} \dots p_{d,s}$, the algorithm can compute $\sum_{i \in [d] \setminus \{j\}} |p_{i,s}|$ in $O(1)$ time at each iteration. Given c_s , one can compute $\sum_{i \in [d] \setminus \{j\}} |p_{i,s}|$ by using $c_s - |p_{j,s}|$, and c_s for the next iteration (i.e., for updating $p_{j+1,s}$) can be computed by using $c_s - |p_{j,s}^{\text{old}}| + |p_{j,s}^{\text{new}}|$. Algorithm 1 shows the procedure for the PCD algorithm for the TI regularized objective function. The computational cost of Algorithm 1 for one epoch is $O(\text{nnz}(\mathbf{X})k)$, which is the same as those of the CD and SGD algorithms for canonical FMs. For more detail about the CD algorithm for canonical FMs, please see (Blondel et al., 2016a; Rendle et al., 2011; Rendle, 2012) or Appendix B.

5. CS Upper Bound Regularizer

We next propose using $\ell_{2,1}^2$ as an upper bound regularizer of Ω_* . Because $\ell_{2,1}^2$ is an upper bound of Ω_* and its corresponding proximal operator outputs row-wise sparse \mathbf{P} (we will see it in Section 5.1),

Algorithm 1 PCD algorithm for TI-sparse FMs

Input: $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$, $k \in \mathbb{N}_{>0}$, $\lambda_w, \lambda_p, \tilde{\lambda}_p \geq 0$
 1: Initialize $\mathbf{P} \in \mathbb{R}^{d \times k}$, $\mathbf{w} \in \mathbb{R}^d$;
 2: Compute caches as in canonical FM;
 3: **while** not convergence **do**
 4: Optimize \mathbf{w} and update caches as in canonical FM;
 5: **for** $s = 1, \dots, k$ **do**
 6: $c_s \leftarrow \sum_{j=1}^d |p_{j,s}|$; ▷ Cache for Ω_{TI}
 7: **for** $j = 1, \dots, d$ **do**
 8: $c_s \leftarrow c_s - |p_{j,s}|$; ▷ $c_s = \sum_{i \in [d] \setminus \{j\}} |p_{i,s}|$
 9: $\eta \leftarrow (\mu \sum_{n \in \text{supp}(\mathbf{x}_{:,j})} (\partial f_{\text{FM}}(\mathbf{x}_n) / \partial p_{j,s})^2 / N + 2\lambda_p)^{-1}$;
 10: Update $p_{j,s}$ as in canonical FM;
 11: $p_{j,s} \leftarrow \text{prox}_{\eta \tilde{\lambda}_p c_s |\cdot|}(p_{j,s}) = \text{sign}(p_{j,s}) \max(|p_{j,s}| - \eta \tilde{\lambda}_p c_s, 0)$;
 12: Update caches as in canonical FM;
 13: $c_s \leftarrow c_s + |p_{j,s}|$; ▷ Update cache for $p_{j+1,s}$
 14: **end for**
 15: **end for**
 16: **end while**
Output: Learned \mathbf{P} and \mathbf{w}

it can be better for feature selection in FMs (i.e., it can select better features in terms of the prediction performance) than the existing regularizers and TI regularizer. We call this regularizer $\ell_{2,1}^2$ regularizer or Cauchy–Schwarz (CS) regularizer and call FMs with this regularization $\ell_{2,1}^2$ -sparse FMs or CS-sparse FMs since this regularization is derived using the Cauchy–Schwarz inequality:

$$\frac{1}{2} \|\mathbf{P}\|_{2,1}^2 = \sum_{j_2 > j_1} \|\mathbf{p}_{j_2}\|_2 \|\mathbf{p}_{j_1}\|_2 + \frac{1}{2} \|\mathbf{P}\|_2^2 \quad (33)$$

$$=: \Omega_{\text{CS}}(\mathbf{P}) + \frac{1}{2} \|\mathbf{P}\|_2^2 \geq \Omega_*(\mathbf{P}) + \frac{1}{2} \|\mathbf{P}\|_2^2. \quad (34)$$

5.1 PGD/PSGD-based Algorithm for CS Regularizer

Similarly to TI regularizer, we consider the use of the PGD algorithm for optimizing CS-sparse FMs. The PGD algorithm for CS-sparse FMs requires to compute the following proximal operator for a $d \times k$ matrix:

$$\text{prox}_{\lambda \|\cdot\|_{2,1}^2}(\mathbf{P}) = \arg \min_{\mathbf{Q} \in \mathbb{R}^{d \times k}} \frac{1}{2} \|\mathbf{Q} - \mathbf{P}\|_2^2 + \lambda \|\mathbf{Q}\|_{2,1}^2. \quad (35)$$

The following theorem states that the proximal operator (35) can be computed in $O(dk + d \log d)$ by Algorithm 2 or $O(dk)$ time in expectation by Algorithm 3.

Theorem 8 *The solution to the proximal problem (35) $\mathbf{Q}^* \in \mathbb{R}^{d \times k}$ is*

$$\mathbf{q}_j^* = \begin{cases} \frac{c_j^*}{\|\mathbf{p}_j\|_2} \mathbf{p}_j & \|\mathbf{p}_j\|_2 \neq 0, \\ \mathbf{0} & \|\mathbf{p}_j\|_2 = 0, \end{cases} \quad (36)$$

where $\mathbf{c}^* = \text{prox}_{\lambda \|\cdot\|_1^2}((\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2)^\top)$.

The proximal operator of $\ell_{2,1}$ (6) shrinks row vectors in \mathbf{P} : $\mathbf{q}_j^* = \mathbf{0}$ if $\|\mathbf{p}_j\|_2 \leq \lambda$ and the threshold λ does not depend on \mathbf{P} , so it sets row vectors of absolutely small ℓ_2 norm to be $\mathbf{0}$. Since $\mathbf{c}^* =$

$\text{prox}_{\lambda\|\cdot\|_1^2} \left((\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2)^\top \right)$ can be sparse, $\ell_{2,1}^2$ (35) also shrinks row vectors in \mathbf{P} : if $\|\mathbf{p}_j\|_2$ is smaller than the threshold, then $\mathbf{q}_j^* = \mathbf{0}$. However, the threshold depends on \mathbf{P} : $\mathbf{q}_j^* = \mathbf{0}$ if $\|\mathbf{p}_j\|_2 \leq 2\lambda S_\theta$ and S_θ depends on $\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2$ as described in Theorem 7. That is, $\ell_{2,1}^2$ regularizer shrinks row vectors of relatively small ℓ_2 norm among $\mathbf{p}_1, \dots, \mathbf{p}_d$ to $\mathbf{0}$. Clearly, CS-sparse FMs also cannot achieve feature interaction selection like $\ell_{2,1}$ -sparse FMs (Xu et al., 2016): they produce a row-wise sparse \mathbf{P} . However, the CS regularizer is more useful than the $\ell_{2,1}$ one for feature selection in FMs since it is also an upper bound of the $\ell_{2,1}$ norm for $\mathbf{P}\mathbf{P}^\top$ without diagonal elements, which seems appropriate for feature selection in FMs.

Unfortunately, PSGD-based algorithms for the CS regularizer using Algorithm 3 cannot leverage the sparsity of dataset \mathbf{X} . Thus, PSGD-based algorithms should be used only when the number of instances is large and dataset \mathbf{X} is dense like those for the TI regularized objective function as described in Section 4.

5.2 PBCD Algorithm for CS Regularizer

We here propose an efficient PBCD algorithm for CS-sparse FMs that optimizes each row vector in \mathbf{P} at each iteration. Like the PCD algorithm for TI-sparse FMs, strictly speaking, $\ell_{2,1}^2$ is not separable w.r.t $\mathbf{p}_1, \dots, \mathbf{p}_d$ and thus it should not be used but it has several advantages: the PBCD algorithm does not require tuning nor using line search technique for the step size although its convergence is not guaranteed due to non-separability.

We consider Ω_{CS} regularized objective function $L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \tilde{\lambda}_p \Omega_{\text{CS}}(\mathbf{P})$ because $\|\mathbf{P}\|_2^2$ can be taken into L_{FM} as in Section 4.1, focusing on the j -th row vector in \mathbf{P} as the optimized parameter. Then, the Ω_{CS} regularizer can be regarded as a ℓ_2 regularizer such that its regularization strength is $\sum_{i \in [d] \setminus \{j\}} \|\mathbf{p}_i\|_2$:

$$\Omega_{\text{CS}}(\mathbf{P}) = \left(\sum_{i \in [d] \setminus \{j\}} \|\mathbf{p}_i\|_2 \right) \|\mathbf{p}_j\|_2 + \text{const.} \quad (37)$$

Therefore, as the PCD algorithm for TI-sparse FMs is almost the same as that for ℓ_1 -sparse FMs, the PBCD algorithm for CS-sparse FMs is almost the same as that for $\ell_{2,1}$ -sparse FMs (Xu et al., 2016). Our remaining task is to design an algorithm for computing regularization strength $\sum_{i \in [d] \setminus \{j\}} \|\mathbf{p}_i\|_2$ in $O(k)$ time. Fortunately, given $c := \sum_{i=1}^d \|\mathbf{p}_i\|_2$, we can compute $c_j := \sum_{i \in [d] \setminus \{j\}} \|\mathbf{p}_i\|_2$ in $O(k)$ time as $c_j = c - \|\mathbf{p}_j\|_2$. The gradient of $\sum_{n=1}^N \ell(f(\mathbf{x}_n), y_n)/N$ w.r.t \mathbf{p}_j is Lipschitz continuous with constant $\mu \sum_{n \in \text{supp}(\mathbf{x}_{:,j})} \|\nabla_{\mathbf{p}_j} f_{\text{FM}}(\mathbf{x}_n)\|_2^2 / N$ since FMs are linear w.r.t \mathbf{p}_j . Thus, for determining step size η , we do not have to use the line search technique (Tseng and Yun, 2009), which might require a high computational cost. For more detail, please see Appendix B.

6. Extensions to Related Models

In this section, we extend $\Omega_{\text{TI}}, \tilde{\ell}_{1,2}^2, \Omega_{\text{CS}}$ and $\ell_{2,1}^2$ to other related models.

6.1 Higher-order FMs

Blondel et al. (2016a) proposed higher-order FMs (HOFMs), which use not only second-order feature interactions but also higher-order feature interactions. M -order HOFMs predict the target of \mathbf{x} as

$$f_{\text{HOFM}}^M(\mathbf{x}; \mathbf{w}, \mathbf{P}^{(2)}, \dots, \mathbf{P}^{(M)}) := \langle \mathbf{x}, \mathbf{w} \rangle + \sum_{m=2}^M \sum_{s=1}^k K_A^m(\mathbf{x}, \mathbf{p}_{:,s}^{(m)}), \quad (38)$$

where $\mathbf{P}^{(2)}, \dots, \mathbf{P}^{(M)} \in \mathbb{R}^{d \times k}$ are learnable parameters for $2, \dots, M$ -order feature interactions, respectively, and $K_A^m : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is the m -order ANOVA kernel:

$$K_A^m(\mathbf{x}, \mathbf{p}) := \sum_{j_m > j_{m-1} > \dots > j_1} x_{j_1} p_{j_1} \cdots x_{j_m} p_{j_m}. \quad (39)$$

M -order HOFMs clearly use from second to M -order feature interactions. Although the evaluation of HOFMs (38) seems to take $O(d^m k)$ time at first glance, it can be completed in $O(dkM^2)$ (strictly speaking, $O(\text{nnz}(\mathbf{x})kM^2)$) time since m -order ANOVA kernels can be evaluated in $O(dm)$ (strictly speaking, $O(\text{nnz}(\mathbf{x})m)$) time by using dynamic programming (Blondel et al., 2016a; Shawe-Taylor and Cristianini, 2004). Blondel et al. (2016a) also proposed efficient CD and SGD-based algorithms.

6.1.1 EXTENSION OF Ω_{TI} AND $\tilde{\ell}_{1,2}^2$ FOR HOFMS

The output of M -order HOFMs (38) can be rearranged as

$$f_{\text{HOFM}}^M(\mathbf{x}) := \langle \mathbf{x}, \mathbf{w} \rangle + \sum_{m=2}^M \sum_{j_m > \dots > j_1} w_{j_1, \dots, j_m} x_{j_1} \cdots x_{j_m}, \quad \text{where } w_{j_1, \dots, j_m} = \sum_{s=1}^k p_{j_1, s}^{(m)} \cdots p_{j_m, s}^{(m)}. \quad (40)$$

Thus, we extend Ω_* and Ω_{TI} for higher-order feature interactions as

$$\Omega_*^m(\mathbf{P}) := \sum_{j_m > \dots > j_1} \left| \sum_{s=1}^k p_{j_1, s} \cdots p_{j_m, s} \right| \leq \sum_{j_m > \dots > j_1} \sum_{s=1}^k |p_{j_1, s} \cdots p_{j_m, s}| =: \Omega_{\text{TI}}^m(\mathbf{P}). \quad (41)$$

Obviously, $\Omega_{\text{TI}} = \Omega_{\text{TI}}^2$ holds. We hence propose using the following regularization for higher-order feature interactions in M -order HOFMs:

$$\Omega_{\text{TI}}^2(\mathbf{P}^{(2)}) + \cdots + \Omega_{\text{TI}}^M(\mathbf{P}^{(M)}). \quad (42)$$

Ω_{TI}^m can be represented by the ANOVA kernel:

$$\Omega_{\text{TI}}^m(\mathbf{P}) = \sum_{s=1}^k K_A^m(\text{abs}(\mathbf{p}_{:,s}), \mathbf{1}). \quad (43)$$

By using multi-linearity of the ANOVA kernel, we can rewrite $\Omega_{\text{TI}}^m(\mathbf{P})$ as

$$\Omega_{\text{TI}}^m(\mathbf{P}) = |p_{j,s}| \frac{\partial}{\partial |p_{j,s}|} K_A^m(\text{abs}(\mathbf{p}_{:,s}), \mathbf{1}) + \text{const}. \quad (44)$$

Hence, $\Omega_{\text{TI}}^m(\mathbf{P})$ is also regarded as ℓ_1 regularization for one parameter $p_{j,s}$, like Ω_{TI} , and the PCD algorithm for the Ω_{TI}^m regularized objective function is almost the same as that for TI-sparse FMs: at each iteration, the algorithm first updates the parameter as in the canonical HOFMs and next applies $\text{prox}_{\eta \bar{\lambda}_p \cdot c|\cdot|}$ to updated $p_{j,s}^{(m)}$, where $c = \partial K_A^m(\text{abs}(\mathbf{p}_{:,s}^{(m)}), \mathbf{1}) / \partial |p_{j,s}^{(m)}|$. Given $K_A^m(\text{abs}(\mathbf{p}_{:,s}^{(m)}), \mathbf{1})$ and additional caches, we can compute $\partial K_A^m(\text{abs}(\mathbf{p}_{:,s}^{(m)}), \mathbf{1}) / \partial |p_{j,s}^{(m)}|$ in $O(m)$ time (Atarashi et al., 2020). Therefore, we can update $\mathbf{p}_{j,s}^{(m)}$ in $O(m \text{nnz}(\mathbf{x}_{:,j})k)$ time, which is the same as that for the PCD algorithm for canonical HOFMs.

Next, we consider the extension of $\tilde{\ell}_{1,2}^2$ for higher-order feature interactions. We first present a generalization of Theorem 3.

Theorem 9 *For any m -homogeneous quasi-norm Ω^m , there exists $C > 0$ such that $\Omega_*^m(\mathbf{P}) \leq C\Omega^m(\mathbf{P})$ for all $\mathbf{P} \in \mathbb{R}^{d \times k}$.*

Thus, we simply propose the use of $\tilde{\ell}_{1,m}^m$ as an extension of $\tilde{\ell}_{1,2}^2$ for $\mathbf{P}^{(m)}$, $m \in \{2, \dots, M\}$. Because $\tilde{\ell}_{1,m}^m$ is column-wise separable, PGD/PSGD-based algorithms for HOFMs with $\tilde{\ell}_{1,m}^m$ require to evaluate the following proximal operator:

$$\text{prox}_{\lambda \|\cdot\|_1^m}(\mathbf{p}, s) = \arg \min_{\mathbf{q} \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{p} - \mathbf{q}\|^2 + \lambda \|\mathbf{q}\|_1^m \quad (45)$$

For this proximal operator, we show a generalization of Theorem 7.

Theorem 10 *Assume that $\mathbf{p} \in \mathbb{R}^d$ is sorted in descending order by absolute value: $|p_1| \geq |p_2| \geq \dots \geq |p_d|$. Then, the solution to the proximal problem (45) $\mathbf{q}^* \in \mathbb{R}^d$ is*

$$q_j^* = \begin{cases} \text{sign}(p_j) [|p_j| - \lambda m S_\theta^{m-1}] & j \leq \theta, \\ 0 & \text{otherwise,} \end{cases} \quad (46)$$

$$S_j \in \left[0, \sum_{i=1}^j |p_i| \right] \quad \text{s.t.} \quad \lambda m j S_j^{m-1} + S_j - \sum_{i=1}^j |p_i| = 0 \quad (47)$$

and $\theta = \max\{j : |p_j| - \lambda m S_j^{m-1} \geq 0\}$.

Like (30), (45) can be evaluated in $O(d \log d)$ time or $O(d)$ time in expectation if S_j can be computed in $O(1)$. Unfortunately, S_j in (47) cannot be computed analytically when $m > 5$. However, HOFMs with $M = 3$ is typically recommended (Blondel et al., 2016a) and then it can be computed analytically. Even if $m > 5$, one can approximately compute S_j by using a numerical method, e.g., Newton's method.

6.1.2 EXTENSION OF Ω_{CS} AND $\ell_{2,1}^2$ FOR HOFMS

We next extend Ω_{CS} for higher-order feature interactions as

$$\Omega_{\text{CS}}^m(\mathbf{P}) := \sum_{j_m > \dots > j_1} \|\mathbf{p}_{j_1}\|_2 \|\mathbf{p}_{j_2}\|_2 \cdots \|\mathbf{p}_{j_m}\|_2 = K_A^m ((\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2)^\top, \mathbf{1}), \quad (48)$$

and we propose using the following regularization for (order-wise) feature selection in M -order HOFMs:

$$\Omega_{\text{CS}}^2(\mathbf{P}^{(2)}) + \dots + \Omega_{\text{CS}}^M(\mathbf{P}^{(M)}), \quad (49)$$

Clearly, $\Omega_{\text{CS}} = \Omega_{\text{CS}}^2$ holds. By using the multi-linearity, we can write it as

$$\Omega_{\text{CS}}^m(\mathbf{P}) = \|\mathbf{p}_j\|_2 \frac{\partial}{\partial \|\mathbf{p}_j\|_2} K_A^m ((\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2)^\top, \mathbf{1}) + \text{const.} \quad (50)$$

Therefore, $\Omega_{\text{CS}}^m(\mathbf{P}^{(m)})$ is also regarded as $\ell_{2,1}$ regularization for one row vector $\mathbf{p}_j^{(m)}$ like Ω_{CS} . Thus, the PBCD algorithm for HOFMs with (49) can be extended similarly as the PCD algorithm for HOFMs with (43).

Like the extension of $\tilde{\ell}_{1,2}^2$, we simply propose using $\ell_{2,1}^m$ as an extension of $\ell_{2,1}^2$. Then, PSGD/PGD-based algorithm requires to evaluate the following proximal operator for $\mathbf{P}^{(m)}$:

$$\text{prox}_{\lambda \|\cdot\|_{2,1}^m}(\mathbf{P}) = \arg \min_{\mathbf{Q} \in \mathbb{R}^{d \times k}} \frac{1}{2} \|\mathbf{Q} - \mathbf{P}\|_2^2 + \lambda \|\mathbf{Q}\|_{2,1}^m. \quad (51)$$

The following is a generalization of Theorem 8 and states that (51) can be evaluated analytically in $O(dk)$ when $m < 6$.

Theorem 11 *The solution to the proximal problem (51) $\mathbf{Q}^* \in \mathbb{R}^{d \times k}$ is*

$$\mathbf{q}_j^* = \begin{cases} \frac{c_j^*}{\|\mathbf{p}_j\|_2} \mathbf{p}_j & \|\mathbf{p}_j\|_2 \neq 0, \\ \mathbf{0} & \|\mathbf{p}_j\|_2 = 0, \end{cases} \quad (52)$$

where $\mathbf{c}^* = \text{prox}_{\lambda \|\cdot\|_1^m} \left((\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2)^\top \right)$.

Therefore, (51) can be evaluated by (i) computing the ℓ_2 norm of each row vector ($O(dk)$), (ii) evaluating $\mathbf{c}^* = \text{prox}_{\lambda \|\cdot\|_1^m} \left((\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2)^\top \right)$ ($O(d)$ in expectation), and (iii) computing \mathbf{q}_j^* as (52) for all $j \in [d]$ ($O(dk)$). When $m > 5$, (51) cannot be evaluated analytically since the evaluation of (45) is required.

6.2 All-subsets Model

For using HOFMs, a machine learning user must determine the maximum order of interactions, M . A machine learning user might want to consider all feature interactions. Although d -order HOFMs use all feature interactions, they require $O(dk(d^2)) = O(d^3k)$ computational cost and it might be prohibitive for a high-dimensional case. To overcome this problem, Blondel et al. (2016a) also proposed the all-subsets model, which uses all feature interactions efficiently. The output of the all-subsets model is defined by

$$f_{\text{all}}(\mathbf{x}; \mathbf{P}) := \sum_{s=1}^k K_{\text{all}}(\mathbf{x}, \mathbf{p}_{:,s}) = \sum_{m=0}^d \sum_{s=1}^k K_A^m(\mathbf{p}_{:,s}, \mathbf{x}), \quad (53)$$

where $\mathbf{P} \in \mathbb{R}^{d \times k}$ is the learnable parameter, and $K_{\text{all}} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is the all-subsets kernel (Blondel et al., 2016a):

$$K_{\text{all}}(\mathbf{x}, \mathbf{p}) := \prod_{j=1}^d (1 + x_j p_j) = \sum_{S \in 2^{[d]}} \prod_{j \in S} x_j p_j. \quad (54)$$

Clearly, the all-subsets kernel (54) can be evaluated in $O(d)$ (strictly speaking, $O(\text{nnz}(\mathbf{x}))$) time, so the all-subsets model (53) can be evaluated in $O(dk)$ (strictly speaking, $O(\text{nnz}(\mathbf{x}))k$) time. Blondel et al. (2016a) also proposed efficient CD and SGD-based algorithms for the all-subsets model.

6.2.1 EXTENSION OF Ω_{TI} AND $\tilde{\ell}_{1,2}^2$ FOR THE ALL-SUBSETS MODEL

The output of the all-subsets model (53) can be rearranged as

$$f_{\text{all}}(\mathbf{x}; \mathbf{P}) = \sum_{S \in 2^{[d]}} w_S \prod_{j \in S} x_j, \quad \text{where } w_S = \sum_{s=1}^k \prod_{j \in S} p_{j,s}. \quad (55)$$

Thus, we propose $\Omega_{\text{TI}}^{\text{all}}$, which is an extension of Ω_{TI} to the all-subsets model:

$$\Omega_{\text{TI}}^{\text{all}}(\mathbf{P}) := \sum_{S \in 2^{[d]}} \sum_{s=1}^k \prod_{j \in S} |p_{j,s}| = \sum_{s=1}^k \sum_{S \in 2^{[d]}} \prod_{j \in S} |p_{j,s}| \cdot 1 = \sum_{s=1}^k K_{\text{all}}(\text{abs}(\mathbf{p}_{:,s}), \mathbf{1}) = \sum_{m=1}^d \Omega_{\text{TI}}^m(\mathbf{P}). \quad (56)$$

Since the output of the all-subsets model is multi-linear and $\Omega_{\text{TI}}^{\text{all}}(\mathbf{P})$ can be regarded as a ℓ_1 regularization for one parameter, the all-subsets model with this regularization can be optimized efficiently by using the PCD algorithm. For $p_{j,s}$, $\Omega_{\text{TI}}^{\text{all}}(\mathbf{P})$ is written as $\{\partial K_{\text{all}}(\text{abs}(\mathbf{p}_{:,s}), \mathbf{1}) / \partial |p_{j,s}|\} \cdot |p_{j,s}| + \text{const}$, and $\partial K_{\text{all}}(\text{abs}(\mathbf{p}_{:,s}), \mathbf{1}) / \partial |p_{j,s}|$ can be computed in $O(1)$ time if $K_{\text{all}}(\text{abs}(\mathbf{p}_{:,s}), \mathbf{1})$ is given (Blondel

et al., 2016a). Thus, we can extend the PCD algorithm for the canonical all-subsets model (Blondel et al., 2016a) to the PCD algorithm for the all-subsets model with $\Omega_{\text{TI}}^{\text{all}}$ regularization similarly as for FMs and HOFMs.

We next extend $\tilde{\ell}_{1,2}^2$ to the all-subsets model. Based on Theorem 9 and (56), we extend it as $\sum_{m=1}^d \|\mathbf{P}^\top\|_{1,m}^m$. We leave the development of an efficient algorithm for evaluating the corresponding proximal problem for future work. Because the all-subsets model is multi-linear w.r.t $\mathbf{p}_1, \dots, \mathbf{p}_d$ (and of course $p_{1,1}, \dots, p_{d,k}$), it is also optimized by using the CD algorithm efficiently (Blondel et al., 2016a).

6.2.2 EXTENSION OF Ω_{CS} AND $\ell_{2,1}^2$ FOR THE ALL-SUBSETS MODEL

Next, we extend Ω_{CS} to all-subsets model as

$$\Omega_{\text{CS}}^{\text{all}}(\mathbf{P}) := \sum_{m=1}^d \Omega_{\text{CS}}^m(\mathbf{P}) = K_{\text{all}}((\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2)^\top, \mathbf{1}). \quad (57)$$

The all-subsets model with $\Omega_{\text{CS}}^{\text{all}}$ is efficiently optimized by using the PBCD algorithm in a similar manner because the all-subsets model is multi-linear w.r.t \mathbf{p}_j and $\Omega_{\text{CS}}^{\text{all}}$ is also multi-linear w.r.t $\|\mathbf{p}_j\|_2$ for all $j \in [d]$.

We also extend $\ell_{2,1}^2$ to the all-subsets model as $\sum_{m=1}^d \|\mathbf{P}\|_{2,1}^m$ but we leave the development of an efficient algorithm for evaluating the corresponding proximal problem for future work.

7. Related Work

Because FMs are equivalent to the QR with low-rank factorized \mathbf{W} and the QR is essentially a linear model, one can naïvely use any feature selection method (Tibshirani, 1996; Liu et al., 2017; Mallat and Zhang, 1993) for feature interaction selection in the QR. Especially, the QR with the ℓ_1 norm and the trace (nuclear) norm $\|\cdot\|_{\text{tr}}$ regularization is one of the natural choice for learning a low-rank \mathbf{W} with feature interaction selection. Formally, the corresponding objective function is

$$L_{\text{QR}}(\mathbf{w}, \mathbf{W}; \lambda_w) + \lambda_{\text{tr}} \|\mathbf{W}\|_{\text{tr}} + \tilde{\lambda} \|\mathbf{W}\|_1, \quad (58)$$

where $\lambda_w, \lambda_{\text{tr}}$, and $\tilde{\lambda} > 0$ are regularization-strength hyperparameters. We call the QR learned by minimizing (58) the sparse and low-rank QR (SLQR). Richard et al. (2012) firstly proposed proximal algorithms for convex objective functions with $\|\cdot\|_1$ and $\|\cdot\|_{\text{tr}}$ such as (58) for estimating a simultaneously sparse and low-rank matrix. The incremental PGD algorithm proposed by Richard et al. (2012) updates \mathbf{W} as

$$\mathbf{W} \leftarrow \text{prox}_{\eta\tilde{\lambda}} \left(\text{prox}_{\eta\lambda_{\text{tr}}\|\cdot\|_{\text{tr}}} (\mathbf{W} - \eta\nabla L_{\text{QR}}) \right). \quad (59)$$

Because the objective function (58) is convex, its all local minima are global minima unlike the FM-based existing methods. It is a great advantage of the QR-based methods compared to the FM-based methods. However, the QR (and of course the SLQR) requires $O(d^2)$ memory and $O(\text{nnz}(\mathbf{x})^2)$ time for evaluation, so it is hard to use QR-based methods for a high-dimensional case. Moreover, the evaluation of (59) takes $O(d^3)$ computational cost (Parikh and Boyd, 2014). Therefore, it is harder to use the SLQR for a high-dimensional case.

Agrawal et al. (2019); Yang et al. (2019); Morvan and Vert (2018), and Suzumura et al. (2017) proposed feature interaction selection methods in the QR and/or QR-like models. They can be more efficient than above-mentioned naïve methods (Tibshirani, 1996; Liu et al., 2017; Mallat and Zhang, 1993). However, the methods of Agrawal et al. (2019) and Yang et al. (2019) require super linear computational cost w.r.t d or N , and those of Morvan and Vert (2018) and Suzumura et al. (2017)

can be used only when $\mathbf{x} \in [0, 1]^d$. Moreover, as described in Section 2, the QR cannot estimate the weights for interactions that are not observed from the training data.

Cheng et al. (2014) proposed a greedy (forward) feature interaction selection algorithm in FMs for context-aware recommendation. They call FMs with their algorithm gradient boosting FMs (GBFMs). In general, greedy selection algorithms produce a sub-optimal solution and are often inferior to shrinkage methods (e.g., methods based on sparse regularization) (Hastie et al., 2009). Moreover, at each greedy interaction selection step, their algorithm sequentially selects each feature that constructs the interaction, namely, their greedy step is approximately greedy.

Chen et al. (2019b) proposed a Bayesian feature interaction selection method in FMs for personalized recommendation. Assume that there are U users and I items, and $\mathbf{x}_i, \dots, \mathbf{x}_I \in \mathbb{R}^d$ are feature vectors of items. Then, they proposed (Bayesian) personalized FMs, which predict the preference of each user for each item. The output of personalized FMs for u -th user $f_u : \mathbb{R}^d \rightarrow \mathbb{R}$ is defined as

$$f_u(\mathbf{x}; \pi_u, \boldsymbol{\mu}, \mathbf{P}, \phi) = \langle \bar{\mathbf{w}}_u, \mathbf{x} \rangle + \sum_{j_2 > j_1} \bar{w}_{u, j_1, j_2} x_{j_1} x_{j_2}, \quad (60)$$

where

$$\begin{aligned} \bar{w}_{u, j} &= \pi_{u, j} \mu_j, \quad \bar{w}_{u, j_1, j_2} = \pi_{u, j_1, j_2} \mu_{j_1, j_2}, \\ \pi_{u, j_1, j_2} &= \pi_{u, j_1} \pi_{u, j_2} [1 + (1 - \pi_{u, j_1} \pi_{u, j_2})(\pi_{u, j_1} + \pi_{u, j_2})], \end{aligned}$$

where $\boldsymbol{\mu} \in \mathbb{R}^d$, $\pi_u \in [0, 1]^d$ and $\mathbf{P}_u \in \mathbb{R}^{d \times k}$ are learnable parameters and $\mu_{j_1, j_2} \in \mathbb{R}^d$ is computed defined by using $\boldsymbol{\mu}$, \mathbf{p}_{j_1} , and \mathbf{p}_{j_2} (for more detail, please see (Chen et al., 2019b)). Intuitively, $\pi_{u, j}$ represents the selection probability of j -th feature. Similarly, π_{u, j_1, j_2} represents that of the interaction between j_1 -th and j_2 -th features. Unfortunately, their method cannot actually select feature interactions without selecting features because $\pi_{u, j_1, j_2} = 0$ if and only if $\pi_{u, j_1} = 0$ or $\pi_{u, j_2} = 0$ (on $\pi_{u, j} \in [0, 1]$), namely, their method actually is for feature selection.

Several researchers proposed FMs and deep-neural-network-extension of FMs that adapt feature interaction weights depending on an input feature vector \mathbf{x} (Xiao et al., 2017; Song et al., 2019; Hong et al., 2019; Xue et al., 2020). While such methods outperformed FMs on some recommender system tasks, they also require $O(\text{nnz}(\mathbf{x})^2)$ time for evaluation. Moreover, their feature interaction weights cannot be completely zero.

8. Experiments

In this section, we experimentally demonstrate the advantages of the proposed methods for feature interaction selection (and additionally feature selection) compared with existing methods. Firstly, we show the results for synthetic datasets. These results indicate that (i) TI-sparse FMs are useful for feature interaction selection in FMs while existing methods are not and (ii) CS-sparse FMs are more useful than existing methods for feature selection in FMs. Secondly, we show the results for some real-world datasets on an interpretability constraint setting. Finally, we compares the convergence speeds of some optimization algorithms for proposed methods on both synthetic and real-world datasets.

We implemented the existing and proposed methods in Nim ³ and ran the experiments on an Arch Linux desktop with an Intel Core i7-4790 (3.60 GHz) CPU and 16 GB RAM. The FM implementation was as fast as libFM (Rendle, 2012), which is the implementation of FMs in C++. Our implementation is available at <https://github.com/neonnnnn/nimfm>.




Figure 2: Comparison of proximal operators associated with **TI**, **CS**, **L21**, and **L1** regularizer. We evaluated the proximal operators at a randomly sampled \mathbf{P} with various λ . Left graph shows the number of used feature interactions and right graph shows the number of used features in $\mathbf{Q}^*(\mathbf{Q}^*)^\top$, where \mathbf{Q}^* is the output of the proximal operator: $\mathbf{Q}^* = \text{prox}_{\lambda\Omega(\cdot)}(\mathbf{P})$ and $\Omega = \ell_1$ (**L1** (Pan et al., 2016)), $\ell_{2,1}$ (**L21** (Xu et al., 2016; Zhao et al., 2017)), $\tilde{\ell}_{1,2}^2$ (**TI**, proposed in Section 4), or $\ell_{2,1}^2$ (**CS**, proposed in Section 5).

8.1 Comparison of Proximal Operators

Firstly, we compared the outputs of proximal operators of the existing methods and the proposed methods in the same way as in Section 3.1. We evaluated proximal operators of not only **L1** (Pan et al., 2016) and **L21** (Xu et al., 2016; Zhao et al., 2017) but also $\tilde{\ell}_{1,2}^2$ (**TI**) and $\ell_{2,1}^2$ (**CS**). Their corresponding proximal operators are (8), (6), (30), and (35), respectively.

Results are shown in Fig. 2. Unlike the existing methods (**L1** and **L21**), the proposed methods (**TI** and **CS**) could produce sparse but moderately sparse feature interaction matrices. Moreover, the number of used features in **TI** was 200 for all $\lambda \in \{2^{-7}, \dots, 2^7\}$. It means that **TI** successfully selects feature interactions in FMs. The number of used features in **CS** decreased gradually as λ increased. It indicates that **CS** can be more useful than **TI** and existing methods for feature selection.

8.2 Synthetic Datasets

We next evaluated the performance of the proposed methods and existing methods on feature interaction selection and feature selection problems using synthetic datasets. We ran the experiments on an Ubuntu 18.04 server with an AMD EPYC 7402P 24-Core Processor (2.80 GHz) and 128 GB RAM.

8.2.1 SETTINGS

Datasets. To evaluate the proposed TI regularizer and CS regularizer in the feature interaction selection and feature selection scenarios, we created datasets such that true models used partial second-order feature interactions. We defined the true model as the QR without the linear term with the block diagonal matrix \mathbf{W} . We defined each block diagonal matrix as a $(d_{\text{true}}/b, d_{\text{true}}/b)$ all-ones matrix, where b is the number of blocks (we set b such that d_{true} was dividable by b). Intuitively, there were b distinct groups of features, and $w_{j_1, j_2} = 1$ if j_1 and j_2 were in the same group of features and equaled zero otherwise. Precisely, if $\lceil j_1/(d_{\text{true}}/b) \rceil = \lceil j_2/(d_{\text{true}}/b) \rceil$, j_1 and j_2

3. <https://nim-lang.org/>

were in the same group of features. For the distribution of feature vector \mathbf{x} , we used a Gaussian distribution, $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. We set $\boldsymbol{\mu} = \mathbf{0}$ and $\Sigma_{j,j} = 1$ for all $j \in [d_{\text{true}}]$ and set $\Sigma_{j_1, j_2} = 0.2$ if $j_1 \neq j_2$ were in the same feature group and zero otherwise. Moreover, we concatenated d_{noise} -dimensional noise features to feature vector \mathbf{x} and used the concatenated vector as the observed feature vector (namely, the dimension of the observed feature vectors $d = d_{\text{true}} + d_{\text{noise}}$). We set the distribution of each noise feature to $\mathcal{N}(0, 1)$ (the noise features were independent of each other). Furthermore, we added noise to the observation of target $f_{\text{true}}(\mathbf{x})$. We used $\mathcal{N}(0, 0.1^2)$ for the target noises. We considered two settings.

- **Feature interaction selection setting:** $d_{\text{true}} = 80$, $b = 8$, and $d_{\text{noise}} = 20$. In this setting, there were eight groups of features, so the methods that perform only feature selection in FMs like CS-sparse FMs and $\ell_{2,1}$ -sparse FMs were not useful. Again, our main goal is to develop sparse FMs that are useful in this setting.
- **Feature selection setting:** $d_{\text{true}} = 20$, $b = 1$, and $d_{\text{noise}} = 80$. In this setting, there was only one group of features, so the methods that perform only feature selection in FMs were also useful.

Evaluation Metrics. We mainly used three metrics. They were computed using the parameters \mathbf{W} of the true models.

- **Estimation error:** $\|\mathbf{W} - \hat{\mathbf{P}}\hat{\mathbf{P}}^\top\|_{2,>} / \|\mathbf{W}\|_{2,>}$, where $\|\cdot\|_{2,>}$ is the ℓ_2 norm for only the strictly upper triangular elements, \mathbf{W} is the true feature interaction matrix, and $\hat{\mathbf{P}}$ is the learned parameter in FMs and sparse FMs. Lower is better.
- **F1-score:** the F1-score of the support prediction problem. To be more precise, we regarded (j_1, j_2) as a positive instance if $w_{j_1, j_2} \neq 0$ and as a negative instance otherwise, and we regarded (j_1, j_2) as a positive predicted instance if $\langle \hat{\mathbf{p}}_{j_1}, \hat{\mathbf{p}}_{j_2} \rangle \neq 0$ and as a negative predicted instance otherwise for all $j_2 > j_1$. Higher is better.
- **Percentage of successful support recovery (PSSR)** (Liu et al., 2017): the percentages of the results such that $\{(j_1, j_2) : w_{j_1, j_2} \neq 0, j_2 > j_1\} = \{(j_1, j_2) : \langle \hat{\mathbf{p}}_{j_1}, \hat{\mathbf{p}}_{j_2} \rangle \neq 0, j_2 > j_1\}$ among the different datasets. Higher is better.

Methods Compared. We compared the following eight methods.

- **TI:** $\tilde{\ell}_{1,2}^2$ -sparse FMs (TI-sparse FMs) optimized using PCD algorithm.
- **CS:** $\ell_{2,1}^2$ -sparse FMs (CS-sparse FMs) optimized using PBCD algorithm.
- **L21:** $\ell_{2,1}$ -sparse FMs (Xu et al., 2016; Zhao et al., 2017) optimized using PBCD algorithm.
- **L1:** ℓ_1 -sparse FMs (Pan et al., 2016) optimized using PCD algorithm.
- **Ω_* -nmAPGD:** Ω_* -sparse FMs optimized using non-monotone accelerated inexact PGD algorithm (Li and Lin, 2015).
- **Ω_* -SubGD:** Ω_* -sparse FMs optimized using SubGD algorithm.
- **SLQR:** the SLQR optimized using the incremental PGD algorithm (Richard et al., 2012) with acceleration (Nesterov, 1983; Beck and Teboulle, 2009) and restart (Giselsson and Boyd, 2014). The PGD algorithm with acceleration is known as a fast iterative shrinkage thresholding algorithm (FISTA) (Beck and Teboulle, 2009).
- **FM:** canonical FMs (Rendle, 2010) optimized using CD algorithm.

For all methods, we omitted the linear term $\langle \mathbf{w}, \mathbf{x} \rangle$ since the true models do not use it. Since the target values are in \mathbb{R} , we used the squared loss for $\ell(\cdot, \cdot)$.

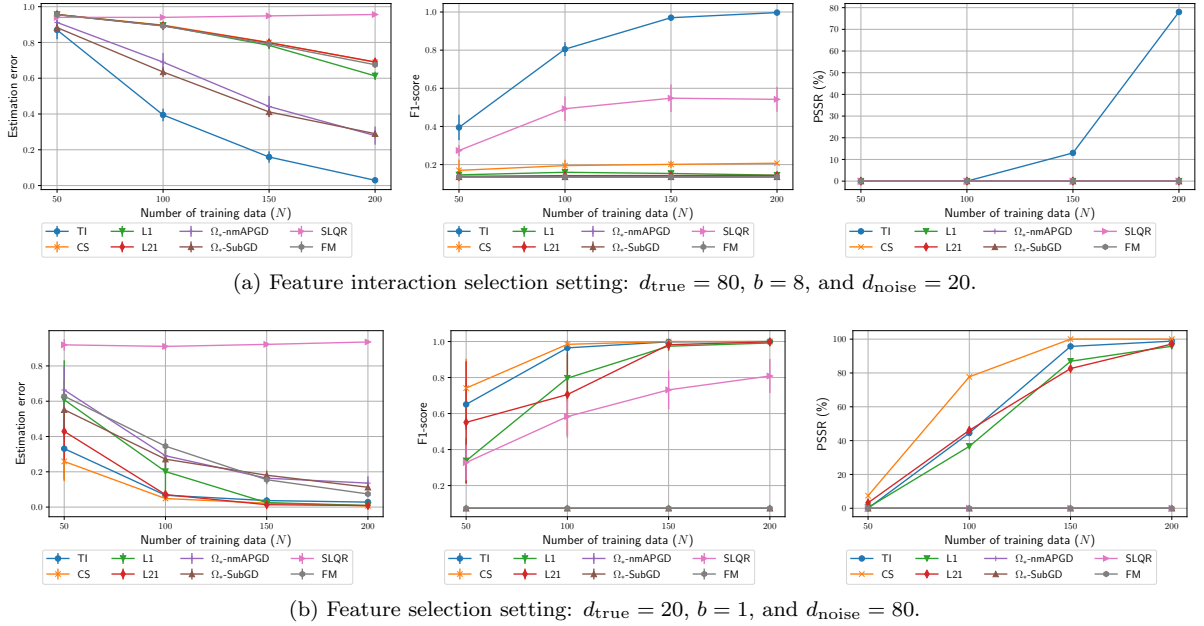


Figure 3: Experimental results with tuned parameters for **TI**, **CS**, **L21**, **L1**, Ω_* -**nmAPGD**, Ω_* -**SubGD**, **SLQR**, and **FM** methods on synthetic datasets using different numbers of training examples: (a) feature interaction selection setting datasets; (b) feature selection setting datasets. Left graphs show estimation error (lower is better), center graphs show F1-score (higher is better), and right graphs show PSSR (higher is better). In (a), the plots of estimation errors of **CS** and **L21**, those of F1-scores of **L1**, **L21**, Ω_* -**nmAPGD**, and Ω_* -**SubGD**, and those of PSSRs except for **TI** overlap, respectively. In (b), similarly, the plots of F1-scores of Ω_* -**nmAPGD** and Ω_* -**SubGD** and those of PSSRs of **SLQR**, Ω_* -**nmAPGD**, and Ω_* -**SubGD** overlap, respectively.

8.2.2 RESULTS WITH TUNED HYPERPARAMETERS

We compared the above-mentioned three metrics among the eight methods. Since these metrics use the parameter of the true model, \mathbf{W} , and the results clearly depended on the hyperparameter settings, we followed Liu et al. (2017) for evaluation: created 150 datasets (not the number of instances in a dataset but the number of datasets), divided them into 50 validation datasets and 100 test datasets, tuned the hyperparameters on the validation datasets (which also used \mathbf{W}), and finally learned models on the test datasets with tuned hyperparameter and evaluated them. We tuned hyperparameters λ_p and $\tilde{\lambda}_p$. For the sparse FM methods (i.e., **TI**, **CS**, **L21**, **L1**, Ω_* -**nmAPGD**, and Ω_* -**SubGD**), we set them to $10^{-2}, 10^{-1}, \dots, 10^2$. Since the **FM** method had only λ_p , we tuned it more carefully than the sparse FM methods: $10^{-3+0.7/24}, 10^{-3+1.7/24}, \dots, 10^{-3+24.7/24} = 10^4$. We set rank hyperparameter k to 30 for **FM** and the sparse FM methods and used $\mathcal{N}(0, 0.01^2)$ to initialize \mathbf{P} . We used line search techniques for computing step sizes in Ω_* -**SubGD** and Ω_* -**nmAPGD**. We used the SubGD method for solving the proximal operator (14) in Ω_* -**nmAPGD**. In this SubGD method, we used a diminishing step size: $\eta = 0.1/\sqrt{t}$ at the t -th iteration, and set the number of iteration to 20. In **SLQR**, we set λ_{tr} and $\tilde{\lambda}$ to $10^{-2}, 10^{-1}, \dots, 10^2$. We ran the experiment 10 times with different initial random seeds for FMs and sparse FMs since their learning results depend on the initial value of \mathbf{P} (we thus learned and evaluated above-mentioned methods $100 \times 10 = 1,000$ times). We also ran it with different numbers of instances in each dataset: $N = 50, 100, 150$, and 200. For fair comparison, we set the time budgets for optimization to $N/50$ second (CPU time) for all methods. However, we stopped optimization if the absolute difference





Figure 4: Sensitivity to regularization-strength hyperparameter $\tilde{\lambda}_p$ for sparse regularization for **TI**, **CS**, **L21**, and **L1** methods on feature selection setting synthetic datasets with $d_{\text{true}} = 20$, $b = 1$, $d_{\text{noise}} = 80$, and $N = 200$. Left graph shows estimation error (lower is better), center graph shows F1-score (higher is better), and right graph shows PSSR (higher is better). The plots of F1-scores of Ω_* -nmAPGD and Ω_* -SubGD and those of PSSRs of Ω_* -nmAPGD, Ω_* -SubGD, and SLQR overlap, respectively.

between the current parameter and previous parameter was less than 10^{-3} for **FM**, **L1**, **L21**, **TI**, and **CS**, and 10^{-7} for **SLQR**, Ω_* -SubGD, and Ω_* -nmAPGD.


As shown in Fig. 3a, **TI** performed the best on the feature interaction selection setting datasets for all metrics. Note that the plots of estimation errors of **CS** and **L21**, those of F1-scores of **L1**, **L21**, Ω_* -nmAPGD, and Ω_* -SubGD, and those of PSSRs except for **TI** overlap, respectively. Only **TI** successfully selected feature interactions in this setting: the F1-score and PSSR of **TI** increased with N , and **TI** achieved about 80% of PSSR when $N = 200$. The F1-scores and PSSRs of **CS**, **L21**, **L1**, Ω_* -nmAPGD, Ω_* -SubGD, and **FM** did not increase with N . Although **SLQR** achieved the better F1-scores than the other methods except for **TI**, its estimation errors were worst and its PSSRs were zero for all N . We observed that **SLQR** could achieve as low estimation errors as Ω_* -nmAPGD and Ω_* -SubGD if we set time budgets more than $10 \times N/50$ second. Namely, **SLQR** could select feature interactions and learn a good \mathbf{W} but it was inefficient compared to **TI**. Ω_* -nmAPGD and Ω_* -SubGD achieved lower estimation errors than the existing methods but their F1-scores and PSSRs were as low as those of the existing methods. This is because the nmAPGD algorithm with an inexact proximal operator and the SubGD algorithm do not produce a sparse \mathbf{P} and a sparse \mathbf{W} , and the PSSR and F1-score of a dense \mathbf{W} are low by definition. In contrast, not only **TI** but also **CS**, **L21**, and **L1** selected features correctly for the feature selection setting datasets (Fig. 3b). Note that the plots of F1-scores of Ω_* -nmAPGD and Ω_* -SubGD and those of PSSRs of **SLQR**, Ω_* -nmAPGD, and Ω_* -SubGD overlap, respectively. **CS** performed the best on all metrics for the feature selection setting datasets and it seems reasonable because **CS** used the upper bound of Ω_* and **CS** produces a row-wise sparse \mathbf{P} (i.e., **CS** is more preferable for feature selection than **TI**).

8.2.3 SENSITIVITY TO REGULARIZATION-STRENGTH HYPERPARAMETER

In this experiment, we investigated sensitivity to the regularization-strength hyperparameter for sparse regularization in the existing and proposed sparse FM methods. We evaluated and compared the estimation errors, F1-scores, and PSSRs for various $\tilde{\lambda}_p$, which is the regularization strength for sparse regularization. In this experiment, we used 50 feature selection datasets with $N = 200$ since the results on the feature interaction selection datasets were bad except **TI** even with tuned $\tilde{\lambda}_p$, as discussed in Section 8.2.2. We set $\tilde{\lambda}_p$ and $\tilde{\lambda}$ to $2^{-10}, 2^{-9}, \dots, 2^6$ and set $\lambda_p = \lambda_{\text{tr}} = 0.1$. The other settings (rank-hyper parameter, initialization, and stopping criterion) were the same as those described in Section 8.2.2. Again, we ran the experiment 10 times with different initial random seeds for FM-based methods.



(a) Feature interaction selection setting: $d_{\text{true}} = 80$, $b = 8$ and $d_{\text{noise}} = 20$ with $\tilde{\lambda}_p = 0.1$ (left) and 1.0 (right).



(b) Feature selection setting: $d_{\text{true}} = 20$, $b = 1$ and $d_{\text{noise}} = 80$ with $\tilde{\lambda}_p = 0.1$ (left) and 1.0 (right).

Figure 5: Trajectories of $\|\mathbf{W}\|_1$ regularized objective value for **TI**, **CS**, **L21**, **L1**, Ω_* -**nmAPGD**, Ω_* -**SubGD**, and **SLQR** methods on (a) feature interaction selection setting synthetic datasets and (b) feature selection setting synthetic datasets with $N = 200$.

As shown in Fig. 4, although the regions of an adequate $\tilde{\lambda}_p$ differed among methods, that of **CS** was wider than those of the other methods for all metrics. The regions of an adequate $\tilde{\lambda}_p$ of **TI** for the F1-score and PSSR were also wider than those of other methods. Thus, our proposed methods are less sensitive to $\tilde{\lambda}_p$ than the other methods. This is important in real-world applications, especially large-scale applications that require high computational costs for hyperparameter tuning.

8.2.4 EFFICIENCY AND SCALABILITY

In the third experiment, we evaluated the efficiency of the proposed and existing methods. We first compared convergences of the $\|\mathbf{W}\|_1$ (it is equivalent to $2\Omega_*(\mathbf{P}) + \|\mathbf{P}\|_2^2$ in FMs as shown in Section 3.2) regularized objective function value among **TI**, **CS**, **L21**, **L1**, Ω_* -**nmAPGD**, Ω_* -**SubGD**, and **SLQR** methods. We tracked the value in the optimization processes using 50 feature interaction selection setting datasets and feature selection setting datasets with $N = 200$. Since the appropriate $\tilde{\lambda}_p$ differed among methods, as mentioned in Section 8.2.3, we ran the experiment with $\tilde{\lambda}_p = 0.1$ and 1.0 . For both settings, we set $\lambda_p = \lambda_{\text{tr}} = 0.1$ and the other settings were the same as those described in Section 8.2.2. Note that we show the results of this experiment on some real-world datasets in Appendix C.1.

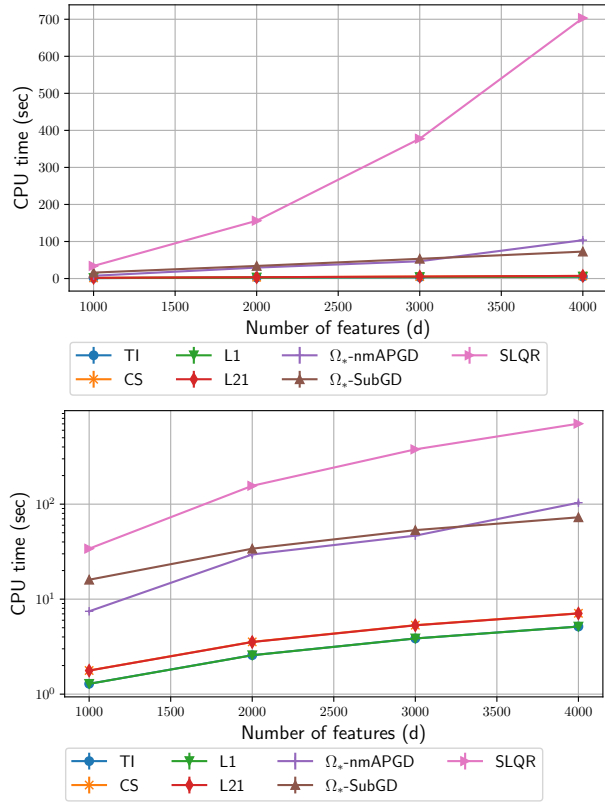


Figure 6: Running times for one epoch in **TI**, **CS**, **L21**, **L1**, Ω_* -**nmAPGD**, Ω_* -**SubGD**, and **SLQR** methods on synthetic datasets. The figure on the right is the same as the figure on the left, except that the y-axis is on a logarithmic scale. The plots of **TI** and **L1** overlap and those of **CS** and **L21** overlap.

Table 2: Datasets used to demonstrate usefulness of proposed methods.

Dataset	Task	d	N_{train}	N_{valid}	N_{test}
ML100K (Harper and Konstan, 2016)	Regression	2,703	64,000	16,000	20,000
a9a (Chang and Lin, 2011)	Classification	123	26,048	6,513	16,281
RCV1 (Chang and Lin, 2011)	Classification	47,236	16,193	4,049	677,399
FD (Ke et al., 2017)	Classification	696	16,000	4,000	80,000

As shown in Fig. 5a, the proposed **TI** method achieved the lowest $\|\mathbf{W}\|_1$ regularized objective value on the feature interaction selection datasets. The difference was remarkable for $\tilde{\lambda}_p = 1.0$. As shown in Fig. 5b), the proposed **TI** and **CS** methods achieved lower $\|\mathbf{W}\|_1$ regularized objective values on the feature selection datasets than the other methods for all parameter settings. Moreover, the objective values converged faster with **TI** and **CS** than with the other methods all parameter settings. Thus, our proposed sparse FMs are more attractive alternatives to Ω_* -sparse FMs than the existing sparse FMs.

We next compared the scalability of the existing and proposed methods w.r.t the number of features d . We created synthetic datasets with varying d and compared the running time for one epoch among **TI**, **CS**, **L21**, **L1**, Ω_* -**nmAPGD**, Ω_* -**SubGD**, and **SLQR** methods. We changed d_{true} as 1000, 2000, 3000, and 4000. We set $b = d_{\text{true}}/100$, $d_{\text{noise}} = 0$, and $N = 2500$. We created ten datasets with different random seeds for all d and report the average running times. The other settings were the same as those described in Section 8.2.2.

As shown in Fig. 6, the running time of **TI**, **CS**, **L1** and **L21** linearly increased w.r.t d . On the other hand, that of Ω_* -**nmAPGD** and Ω_* -**SubGD** increased quadratically, and that of **SLQR** increased cubically w.r.t d . When $d = 4000$, **SLQR** ran more than 100 times slower than **TI**, **CS**, **L1**, and **L21**. Thus, the proposed **TI** and **CS** are better than Ω_* -**nmAPGD**, Ω_* -**SubGD**, and **SLQR** for a high-dimensional case.

8.3 Real-world Datasets

Next, we used real-world datasets to demonstrate the usefulness of the proposed methods.

Settings and Datasets. We compared existing and proposed sparse FMs on an interpretability-constraint setting: the number of interactions in sparse FMs were constrained to be (about) 1,000. We used one regression dataset, the MovieLens 100K (ML100K) (Harper and Konstan, 2016) dataset, and three binary classification datasets, a9a, RCV1 (Chang and Lin, 2011), and Flight Delay (FD) (Ke et al., 2017). Table 2 summarizes the details of these datasets. With the ML100K dataset, which is used for movie recommendation, we considered a regression problem: predicting the score given to a movie by a user. Possible scores were $1, 2, \dots, 5$. We created feature vectors by following the method of Blondel et al. (2016a). We divided the 100,000 user-item scores in the dataset into sets of 64,000, 16,000, and 20,000 for training, validation, and testing, respectively. For the a9a and RCV1 datasets, we used feature vectors and targets that are available from the LIB-SVM (Chang and Lin, 2011) datasets repository.⁴ Both a9a and RCV1 have already been divided into training and testing datasets; we used 20% of the training dataset as the validation dataset and the remaining 80% as the training dataset in this experiment. For the FD dataset, we considered the classification task to predict whether a flight will be delayed by more than 15 minutes. We used the scripts provided in <https://github.com/szilard/benchm-ml:ran-2-gendata.txt> and randomly sampled train, valid, and test datasets from `train-1m.csv` and `test.csv`. Then, each instance had eight attributes and we encoded them to one-hot features except `DepTime` and `Distance`. `DepTime` represents an actual departure time and consists of four integers, HHMM. We split such `DepTime` values to HH and MM, and encoded them to one-hot features (MM values were encoded to six-

4. <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

dimensional one-hot features based on their tens digit). For **Distance** values, we regarded them as numerical values and used their logarithm values.

Evaluation Metrics. As metrics, we used the root mean squared error (RMSE) for the ML100K dataset and the area under the receiver operating characteristic curve (ROC-AUC) for the a9a, RCV1, and FD datasets. Lower is better for the RMSE, and higher is better for the ROC-AUC.

Hyperparameter Settings. As in the experiments using synthetic datasets, we compared the **TI**, **CS**, **L21**, **L1**, and **FM** methods. We used the squared error as loss function ℓ . We set rank-hyper parameter k to 30. We used the linear term $\langle \mathbf{w}, \mathbf{x} \rangle$ and introduced a bias (intercept) term. We initialized each element in \mathbf{P} by using $\mathcal{N}(0, 0.01^2)$ as in the experiments using synthetic datasets. We initialized \mathbf{w} to $\mathbf{0}$ and the bias term to 0. We ran the experiment five times with different random seeds and calculated the average values of the evaluation metrics. For the **FM** method, we chose λ_w and λ_p from $0.5 \times 10^{-7}, 0.5 \times 10^{-6}, \dots, 0.5 \times 10^{-2}$. For the **TI**, **CS**, **L21**, and **L1** methods, we set $\lambda_w = \lambda_w^{\text{FM}}$ and $\lambda_p = \lambda_p^{\text{FM}}/10$, where λ_w^{FM} and λ_p^{FM} are the tuned λ_w and λ_p in the **FM** method. Since sparse FMs have additional regularizers, we set λ_p to $\lambda_p^{\text{FM}}/10$, not λ_p^{FM} . As described above, because we constrained the number of used interactions in sparse FMs to be about 1,000, the method for tuning $\tilde{\lambda}_p$ was complicated. We searched for the appropriate $\tilde{\lambda}_p$ by binary search since the number of used interactions (i.e., number of non-zero elements among the strictly upper triangular elements in $\mathbf{P}\mathbf{P}^\top$) tended to be monotonically non-decreasing in λ_p . For each sparse FM, the initial range (i.e., upper bound and lower bound) of the binary search was chosen from $10^{-7}, 10^{-6}, \dots, 10^{-2}$. After the initial range was chosen, we searched for the appropriate $\tilde{\lambda}_p$ by binary search. Since it was hard to achieve the number of used interactions to be exactly 1,000, we accepted the models with the number of used interactions to be in [990, 1,035]. The reason why we set the acceptable range to be [990, 1,035] is that **CS** and **L21** achieve only feature selection, and 990 and 1,035 are the nearest binomial coefficients to 1,000: $990 = \binom{45}{2}$ and $1,035 = \binom{46}{2}$. Moreover, if the gap between an upper and a lower bound in binary search was lower than 10^{-12} , we gave up tuning λ_p for such models and set their scores to be N/A. Although **FM** did not select feature interactions, we showed results of it for comparison. We also show the results of **TI**, **CS**, **L1** and **L21** with tuned $\tilde{\lambda}_p \in \{10^{-7}, 10^{-6}, \dots, 10^{-2}\}$ (i.e., the best results before doing binary search) for comparison although the numbers of used interactions in them were not close to 1,000.

Results. As shown in Table 3, although the differences were not remarkable, our proposed methods achieved the best performance for each dataset. For ML100K and a9a datasets, **TI** achieved the best performance. This results match the experimental results in Section 8.2.2. Among **TI**, **CS**, **L21**, and **L1**, only **TI** can perform feature interaction selection. For the RCV1 dataset, only **CS** succeeded in learning models using approximately 1,000 feature interactions. For the FD dataset, **TI** and **CS** succeeded in learning such models but **L21** and **L1** did not. This result matches the experimental results in Section 8.2.2 and Section 8.2.3. **TI** and **CS**, especially **CS**, are less sensitive to regularization-strength hyperparameter and can perform better feature selection than **L21** and **L1**. Moreover, as shown in Table 3b, **TI** achieved the best performance for a9a, RCV1, and FD datasets and **CS** achieved the best performance for ML100K dataset on no interpretability constraint setting. They used not all but partial feature interactions except **CS** for FD dataset. On the other hand, **FM** used all feature interactions for ML100K, a9a, and FD datasets and used no feature interactions for RCV1 dataset (in our training-validation-test splitting, the training ML100K dataset contained only 2,640 features and the training FD dataset contained only 643 features). **L1** and **L21** also used no feature interactions for RCV1 and FD datasets. From these results, we conclude that our proposed methods could select better (important) features and feature interactions in terms of the prediction performance.

9. Conclusion

In this paper, we have presented new sparse regularizers for feature interaction selection and feature selection in FMs, the TI regularizer and the CS regularizer, respectively, as well as efficient proximal

Table 3: Comparison of test RMSE for ML100K dataset and test ROC-AUC for a9a, RCV1, and FD datasets. (a) Results on interpretability constraint setting. (b) Results on no interpretability constraint setting ($\tilde{\lambda}_p \in \{10^{-7}, 10^{-6}, \dots, 10^{-2}\}$). We report not only RMSE and ROC-AUC (upper value in each cell) but also the number of used interactions (lower value in each cell). Lower is better for RMSE, and higher is better for ROC-AUC.

(a) Interpretability constraint setting.				
Method	ML100K (RMSE)	a9a (ROC-AUC)	RCV1 (ROC-AUC)	FD (ROC-AUC)
TI	0.93018 991.4	0.90301 1,001.6	N/A	0.71373 1,004.4
CS	0.93127 1,026.0	0.90259 1,035.0	0.99197 1,035.0	0.71437 1,035.0
L1	N/A	0.90197 1,026.0	N/A	N/A
L21	0.93302 1,030.2	0.90259 990.0	N/A	N/A

(b) No interpretability constraint setting ($\tilde{\lambda}_p$ was tuned from $\{10^{-7}, 10^{-6}, \dots, 10^{-2}\}$).				
Method	ML100K (RMSE)	a9a (ROC-AUC)	RCV1 (ROC-AUC)	FD (ROC-AUC)
TI	0.91848 186,600.4	0.90288 1,408.2	0.99258 663,793.8	0.71378 2,471.4
CS	0.91610 7264.0	0.90250 903.0	0.99243 170,820.0	0.71057 0.0
L1	0.92612 1,616,087.2	0.90193 946.0	0.99195 0.0	0.71057 0.0
L21	0.91975 1,775,679.2	0.90269 1,081.0	0.99195 0.0	0.71057 0.0
FM	0.91734 3,483,480.0	0.90280 7,503.0	0.99195 0.0	0.71334 206,403.0

optimization methods for these proposed methods. Our basic idea is the use of ℓ_1 regularizer for feature interaction weight matrix computed from the parameter matrix of FMs. This regularization seems appropriate for feature interaction selection in FMs because it is reported as one of the most promising sparse regularizers and selecting feature interactions necessarily means making \mathbf{W} sparse. Unfortunately, the associated objective function is hard to optimize w.r.t the parameter matrix of FMs. To overcome this difficulty, we have proposed the use of squares of sparsity-inducing (quasi-)norms as an upper bound of the ℓ_1 norm for \mathbf{W} , and we have presented such regularizers concretely, the TI regularizer and the CS regularizer. The TI enables feature interaction selection without feature selection, and the CS can select better (more important) features. Fortunately, the associated objective functions are now easy to optimize, because TI-sparse FMs and CS-sparse FMs can be optimized at the same computational cost as canonical FMs. We have demonstrated the effectiveness of the proposed methods on synthetic and real-world datasets.

As future work, we would like to (i) develop more efficient PSGD-based algorithms for TI/CS-sparse FMs and (ii) investigate the theoretical properties of the proposed methods.

Acknowledgments

This work was partially supported by JSPS KAKENHI Grants JP20J13620 and JP18H03337, and by the Global Station for Big Data and Cybersecurity, a project of the Global Institution for Collaborative Research and Education at Hokkaido University.

Appendix A. Proofs

Additional Notation. For given two matrices $\mathbf{P}, \mathbf{Q} \in \mathbb{R}^{n \times n}$, $\mathbf{P} =_{>} \mathbf{Q}$ means $p_{i,j} = q_{i,j}$ for all $0 < i < j \leq n$, that is, \mathbf{P} and \mathbf{Q} have same values in their strictly upper triangular elements. We use $\mathbf{e}_j^d \in \{0, 1\}^d$ for the d -dimensional standard basis vector whose j -th element is one and the others are zero. For given a scalar $\alpha \in \mathbb{R}$ and a subset C of \mathbb{R}^d , we define $\alpha C := \{\alpha \mathbf{c} : \mathbf{c} \in C\}$. For given a d -dimensional vector $\mathbf{x} \in \mathbb{R}^d$ and a subset C of \mathbb{R}^d , we use $\mathbf{x} + C$ as $\{\mathbf{x} + \mathbf{c} : \mathbf{c} \in C\}$.

A.1 Subdifferentials of Powers of Norms

For Theorem 10 and Theorem 11 (and of course Theorem 8), we first derive the subdifferentials of powers of norms.

Lemma 12 *Let $\mathbf{x} \in \mathbb{R}^d$ and $\|\cdot\|$ be a norm on \mathbb{R}^d . Then, for all $m \geq 1$, the subdifferential of $\|\cdot\|^m$ at \mathbf{x} , $\partial \|\mathbf{x}\|^m$ is*

$$\partial \|\mathbf{x}\|^m = m \|\mathbf{x}\|^{m-1} \partial \|\mathbf{x}\| = m \|\mathbf{x}\|^{m-1} \{ \mathbf{z} \in \mathbb{R}^d : \langle \mathbf{z}, \mathbf{x} \rangle = \|\mathbf{x}\|, \|\mathbf{z}\|_* \leq 1 \}, \quad (61)$$

where $\|\cdot\|_* : \mathbb{R}^d \rightarrow \mathbb{R}$ is the dual norm of $\|\cdot\| : \|\mathbf{x}\|_* = \sup_{\mathbf{z} \in \mathbb{R}^d, \|\mathbf{z}\| \leq 1} \langle \mathbf{z}, \mathbf{x} \rangle$.

Proof (i) We first show that $\partial \|\mathbf{x}\|^m \supseteq m \|\mathbf{x}\|^{m-1} \partial \|\mathbf{x}\|$. For all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ and $\mathbf{z} \in \partial \|\mathbf{x}\|$, from the definition of $\partial \|\mathbf{x}\|$ and the Hölder's inequality ($\langle \mathbf{z}, \mathbf{y} \rangle \leq \|\mathbf{z}\|_* \|\mathbf{y}\|$), we have

$$\|\mathbf{x}\|^m + \langle m \|\mathbf{x}\|^{m-1} \mathbf{z}, \mathbf{y} - \mathbf{x} \rangle = \|\mathbf{x}\|^m + m \|\mathbf{x}\|^{m-1} (\langle \mathbf{z}, \mathbf{y} \rangle - \langle \mathbf{z}, \mathbf{x} \rangle) \quad (62)$$

$$= \|\mathbf{x}\|^m + m \|\mathbf{x}\|^{m-1} (\langle \mathbf{z}, \mathbf{y} \rangle - \|\mathbf{x}\|) \leq \|\mathbf{x}\|^m + m \|\mathbf{x}\|^{m-1} \|\mathbf{z}\|_* \|\mathbf{y}\| - m \|\mathbf{x}\|^m \quad (63)$$

$$\leq (1 - m) \|\mathbf{x}\|^m + m \|\mathbf{x}\|^{m-1} \|\mathbf{y}\|. \quad (64)$$

Moreover, from the convexity of x^m on $x \geq 0$ ($m \geq 1$) and the non-negativity of the norm, we have

$$(\|\mathbf{y}\|)^m \geq (\|\mathbf{x}\|)^m + m(\|\mathbf{x}\|)^{m-1}(\|\mathbf{y}\| - \|\mathbf{x}\|) = (1 - m) \|\mathbf{x}\|^m + m \|\mathbf{x}\|^{m-1} \|\mathbf{y}\|. \quad (65)$$

Combining them, we obtain

$$\|\mathbf{y}\|^m \geq \|\mathbf{x}\|^m + \langle m \|\mathbf{x}\|^{m-1} \mathbf{z}, \mathbf{y} - \mathbf{x} \rangle. \quad (66)$$

The above inequality means $(m \|\mathbf{x}\|^{m-1} \mathbf{z}) \in \partial \|\mathbf{x}\|^m$. Therefore $\partial \|\mathbf{x}\|^m \supseteq m \|\mathbf{x}\|^{m-1} \partial \|\mathbf{x}\|$.

(ii) We next prove that $\partial \|\mathbf{x}\|^m \subseteq m \|\mathbf{x}\|^{m-1} \partial \|\mathbf{x}\|$ by contradiction.

1. $\langle \mathbf{z}, \mathbf{x} \rangle = m \|\mathbf{x}\|^{m-1} \|\mathbf{x}\| = m \|\mathbf{x}\|^m$ for all $\mathbf{z} \in \partial \|\mathbf{x}\|^m$. It clearly holds when $\mathbf{x} = \mathbf{0}$, so we consider the case where $\mathbf{x} \neq \mathbf{0}$. Assume that $\langle \mathbf{z}, \mathbf{x} \rangle = m \|\mathbf{x}\|^m + \varepsilon$, where $\varepsilon > 0$, and we show that then there exists $\mathbf{y} \in \mathbb{R}^d$ such that $\|\mathbf{y}\|^m < \|\mathbf{x}\|^m + \langle \mathbf{z}, \mathbf{y} - \mathbf{x} \rangle$ (this contradicts to the definition of $\partial \|\mathbf{x}\|^m$: if $\mathbf{z} \in \partial \|\mathbf{x}\|^m$, then $\|\mathbf{y}\|^m \geq \|\mathbf{x}\|^m + \langle \mathbf{z}, \mathbf{y} - \mathbf{x} \rangle$ for all $\mathbf{y} \in \mathbb{R}^d$). Let $\mathbf{y} = c\mathbf{x}$, $c > 0$, and define $h : \mathbb{R} \rightarrow \mathbb{R}$ as

$$h(c) := \|\mathbf{x}\|^m + \langle \mathbf{z}, \mathbf{y} - \mathbf{x} \rangle - \|\mathbf{y}\|^m = (1 - c^m) \|\mathbf{x}\|^m - (1 - c)(m \|\mathbf{x}\|^m + \varepsilon). \quad (67)$$

$h(1) = 0$, $h'(c) = -mc^{m-1} \|\mathbf{x}\|^m + (m \|\mathbf{x}\|^m + \varepsilon)$, and $h'(c) > 0$ on $c^{m-1} \in (0, 1 + \varepsilon/(m \|\mathbf{x}\|^m))$, so $h(c) > 0$ on $c \in (1, \{1 + \varepsilon/(m \|\mathbf{x}\|^m)\}^{1/(m-1)})$. This means that there exists $\mathbf{y} \in \mathbb{R}^d$ such that $\|\mathbf{y}\|^m < \|\mathbf{x}\|^m + \langle \mathbf{z}, \mathbf{y} - \mathbf{x} \rangle$, and this contradicts to the definition of the subdifferential. The contradiction can be derived under the assumption $\langle \mathbf{z}, \mathbf{x} \rangle = m \|\mathbf{x}\|^m - \varepsilon$ in a similar manner. Thus, $\langle \mathbf{z}, \mathbf{x} \rangle = m \|\mathbf{x}\|^m$ for all $\mathbf{z} \in \partial \|\mathbf{x}\|^m$.

2. $\|\mathbf{z}\|_* \leq m \|\mathbf{x}\|^{m-1}$ for all $\mathbf{z} \in \partial \|\mathbf{x}\|^m$. Assume that $\|\mathbf{z}\|_* = m \|\mathbf{x}\|^{m-1} + \varepsilon$, $\varepsilon > 0$. From the Hölder's inequality, $\langle \mathbf{z}, \mathbf{y} \rangle \leq \|\mathbf{z}\|_* \|\mathbf{y}\|$ for all $\mathbf{y} \in \mathbb{R}^d$. We can take $\mathbf{y} \in \mathbb{R}^d$ such that $\langle \mathbf{z}, \mathbf{y} \rangle = \|\mathbf{z}\|_* \|\mathbf{y}\|$. Then,

$$\|\mathbf{x}\|^m + \langle \mathbf{z}, \mathbf{y} - \mathbf{x} \rangle = \|\mathbf{x}\|^m + \|\mathbf{z}\|_* \|\mathbf{y}\| - m \|\mathbf{x}\|^m \quad (68)$$

$$= (1 - m) \|\mathbf{x}\|^m + (m \|\mathbf{x}\|^{m-1} + \varepsilon) \|\mathbf{y}\|. \quad (69)$$

We first consider the case where $\mathbf{x} \neq \mathbf{0}$. If we choose \mathbf{y} such that $\|\mathbf{y}\| = \|\mathbf{x}\|$, then

$$(1 - m) \|\mathbf{x}\|^m + (m \|\mathbf{x}\|^{m-1} + \varepsilon) \|\mathbf{y}\| = (1 + \varepsilon) \|\mathbf{y}\| > \|\mathbf{y}\|. \quad (70)$$

This contradicts to the definition of the subdifferential. We next consider the case where $\mathbf{x} = \mathbf{0}$. If we choose \mathbf{y} such that $\varepsilon > \|\mathbf{y}\|^{m-1}$, then

$$(1 - m) \|\mathbf{x}\|^m + (m \|\mathbf{x}\|^{m-1} + \varepsilon) \|\mathbf{y}\| = \varepsilon \|\mathbf{y}\| > \|\mathbf{y}\|^{m-1} \|\mathbf{y}\| = \|\mathbf{y}\|^m, \quad (71)$$

and this also contradicts to the definition of the subdifferential. Thus $\|\mathbf{z}\|_* \leq m \|\mathbf{x}\|^{m-1}$ for all $\mathbf{z} \in \partial \|\mathbf{x}\|^m$.

These results imply $\partial \|\mathbf{x}\|^m \subseteq m \|\mathbf{x}\|^{m-1} \partial \|\mathbf{x}\|$.

From (i) and (ii), we have $\partial \|\mathbf{x}\|^m = m \|\mathbf{x}\|^{m-1} \partial \|\mathbf{x}\|$. ■

From Theorem 12, we have the following corollaries.

Corollary 13 For all $m \geq 1$, the subdifferential of $\|\cdot\|_1^m$ at $\mathbf{p} \in \mathbb{R}^d$, $\partial \|\mathbf{p}\|_1^m$ is

$$\partial \|\mathbf{p}\|_1^m = m \|\mathbf{p}\|_1^{m-1} \partial \|\mathbf{p}\|_1 = m \|\mathbf{p}\|_1^{m-1} \prod_{i=1}^d J(p_i), \quad \text{where } J(p) = \begin{cases} \{1\} & p > 0, \\ \{-1\} & p < 0, \\ [-1, 1] & p = 0. \end{cases} \quad (72)$$

Corollary 14 For all $m \geq 1$, the subdifferential of $\|\cdot\|_{2,1}^m$ at $\mathbf{P} \in \mathbb{R}^{d \times k}$, $\partial \|\mathbf{P}\|_{2,1}^m$ is

$$\partial \|\mathbf{P}\|_{2,1}^m = m \|\mathbf{P}\|_{2,1}^{m-1} \partial \|\mathbf{P}\|_{2,1} = m \|\mathbf{P}\|_{2,1}^{m-1} \{\mathbf{Z} \in \mathbb{R}^{d \times k} : \mathbf{z}_i \in \partial \|\mathbf{p}_i\|_2, i \in [d]\}. \quad (73)$$

A.2 Proximal Operator for $\tilde{\ell}_{1,m}^m$

In this section, we prove Theorem 10. We first present some properties of $\mathbf{q}^* = \text{prox}_{\lambda \|\cdot\|_1^m}(\mathbf{p})$.

Proposition 15 Let $\mathbf{q}^* = \text{prox}_{\lambda \|\cdot\|_1^m}(\mathbf{p}) \in \mathbb{R}^d$. Then, the followings hold for all $i, j \in [d]$:

(i) $p_i = 0 \rightarrow q_i^* = 0$,

(ii) $p_i > 0 \rightarrow q_i^* \geq 0$,

(iii) $p_i < 0 \rightarrow q_i^* \leq 0$,

(iv) $|p_i| \geq |p_j| \rightarrow |q_i^*| \geq |q_j^*|$,

Proof (i) is trivial.

(ii) Assume $p_i > 0$. We prove (ii) by showing that \mathbf{q} is not an optimal value if $q_i < 0$. We consider the objective function in the proximal operation: $g_{\lambda \|\cdot\|_1^m}(\mathbf{q}; \mathbf{p}) = \|\mathbf{p} - \mathbf{q}\|_2^2 / 2 + \lambda \|\mathbf{q}\|_1^m$. By construction $\mathbf{q}^* = \arg \min_{\mathbf{q}} g_{\lambda \|\cdot\|_1^m}(\mathbf{q}; \mathbf{p})$. Let \mathbf{q} be the d -dimensional vector with $q_i < 0$ and \mathbf{q}' be the d -dimensional vector such that $q'_i = |q_i|$ and $q'_j = q_j$ for all $j \in [d] \setminus \{i\}$ (i.e, the sign of i -th

element is reversed compared to \mathbf{q}). Then, $g_{\lambda\|\cdot\|_1^m}(\mathbf{q}'; \mathbf{p}) < g_{\lambda\|\cdot\|_1^m}(\mathbf{q}; \mathbf{p})$ since $\|\mathbf{q}'\|_1^m = \|\mathbf{q}\|_1^m$ but $\|\mathbf{p} - \mathbf{q}'\|_2^2 < \|\mathbf{p} - \mathbf{q}\|_2^2$. That is, when $p_i > 0$, \mathbf{q} with $q_i < 0$ is not an optimal solution. It implies (ii).

(iii) can be derived as in (ii).

(iv) Assume $|p_i| > |p_j|$ and we prove (iv) as in the proof of (ii). Let \mathbf{q} be the d -dimensional vector such that $|q_i| < |q_j|$ and $\text{sign}(p_i) \cdot \text{sign}(q_i) \geq 0$, $\text{sign}(p_j) \cdot \text{sign}(q_j) \geq 0$ (i.e., q_i and q_j satisfy (i)-(iii)). Moreover, let \mathbf{q}' be the d -dimensional vector such that $q'_i = \text{sign}(p_i)|q_j|$, $q'_j = \text{sign}(p_j)|q_i|$, and $q'_{j'} = q_{j'}$ for all $j' \in [d] \setminus \{i, j\}$ (namely, \mathbf{q}' is defined by exchanging the absolute value of i -th and j -th element in \mathbf{q}). Then, $g_{\lambda\|\cdot\|_1^m}(\mathbf{q}'; \mathbf{p}) < g_{\lambda\|\cdot\|_1^m}(\mathbf{q}; \mathbf{p})$ like (ii): the values of second terms are same ($\|\mathbf{q}'\|_1^m = \|\mathbf{q}\|_1^m$) but

$$\|\mathbf{p} - \mathbf{q}'\|_2^2 - \|\mathbf{p} - \mathbf{q}\|_2^2 = -2|p_i||q'_i| - 2|p_j||q'_j| + 2|p_i||q_i| + 2|p_j||q_j| \quad (74)$$

$$= -2|p_i||q_j| - 2|p_j||q_i| + 2|p_i||q_i| + 2|p_j||q_j| = -2(|p_i| - |p_j|)(|q_j| - |q_i|) < 0. \quad (75)$$

That is, when $|p_i| > |p_j|$, \mathbf{q} with $|q_i| < |q_j|$ is not an optimal solution, and it implies (iv). \blacksquare

Finally, we prove Theorem 10.

Theorem 10 Assume that $\mathbf{p} \in \mathbb{R}^d$ is sorted in descending order by absolute value: $|p_1| \geq |p_2| \geq \dots \geq |p_d|$. Then, the solution to the proximal problem (45) $\mathbf{q}^* \in \mathbb{R}^d$ is

$$q_j^* = \begin{cases} \text{sign}(p_j) [|p_j| - \lambda m S_\theta^{m-1}] & j \leq \theta, \\ 0 & \text{otherwise,} \end{cases} \quad (46)$$

$$S_j \in \left[0, \sum_{i=1}^j |p_i| \right] \text{ s.t. } \lambda m j S_j^{m-1} + S_j - \sum_{i=1}^j |p_i| = 0 \quad (47)$$

and $\theta = \max\{j : |p_j| - \lambda m S_j^{m-1} \geq 0\}$.

Proof We first prove that there exists $\theta \in [d]$ such that Equation (46) holds. $\mathbf{q}^* = \text{prox}_{\lambda\|\cdot\|_1^m}(\mathbf{p})$ means the subdifferential of $g_{\lambda\|\cdot\|_1^m}$ at \mathbf{q}^* includes $\mathbf{0}$. From $\partial g_{\lambda\|\cdot\|_1^m}(\mathbf{q}^*) = \mathbf{q}^* - \mathbf{p} + \lambda \partial \|\mathbf{q}^*\|_1^m$ and Theorem 13, we have

$$q_j^* = \begin{cases} p_j - \lambda m \|\mathbf{q}^*\|_1^{m-1} & q_j^* > 0, \\ p_j + \lambda m \|\mathbf{q}^*\|_1^{m-1} & q_j^* < 0. \end{cases} \quad (76)$$

Moreover, from Theorem 15, \mathbf{q}^* is also sorted by absolute value and there exists $\theta \in [d]$ such that $|q_j| > 0$ for all $j \leq \theta$ and $|q_j| = 0$ for all $j > \theta$. Therefore, Equation (76) can be rewritten as

$$q_j^* = \begin{cases} \text{sign}(p_j) \left[|p_j| - \lambda m \left(\sum_{i=1}^\theta |q_i^*| \right)^{m-1} \right] & j \leq \theta, \\ 0 & \text{otherwise,} \end{cases} \quad (77)$$

Here, about the summation of $|q_1^*|, \dots, |q_\theta^*|$, we have

$$\sum_{i=1}^\theta |q_i^*| = \sum_{i=1}^\theta \text{sign}(q_i^*) q_i^* = \sum_{i=1}^\theta \text{sign}(p_i) q_i^* = \sum_{i=1}^\theta |p_i| - \lambda m \theta \left(\sum_{i=1}^\theta |q_i^*| \right)^{m-1}. \quad (78)$$

Thus, (77) is rewritten as

$$q_j^* = \begin{cases} \text{sign}(p_j) [|p_j| - \lambda m S_\theta^{m-1}] & j \leq \theta, \\ 0 & \text{otherwise,} \end{cases} \quad (45)$$

where

$$S_j \in \left[0, \sum_{i=1}^j |p_i| \right] \text{ s.t. } \lambda m j S_j^{m-1} + S_j - \sum_{i=1}^j |p_i| = 0. \quad (47)$$

Such S_j always exists uniquely for all $j \in [d]$ because $\lambda m j \cdot 0^{m-1} + 0 - \sum_{i=1}^j |p_i| \leq 0$, $\lambda m j (\sum_{i=1}^j |p_i|)^{m-1} + \sum_{i=1}^j |p_i| - \sum_{i=1}^j |p_i| \geq 0$, and $\lambda m j S^{m-1} + S$ is monotonically increasing w.r.t S when $S \geq 0$.

It is worth noting that S_j is monotonically non-decreasing for all j such that $|p_j| - \lambda m S_j^{m-1} \geq 0$. Assume that $|p_{j'}| - \lambda m S_{j'}^{m-1} \geq 0$. Then, for all $j \in [j']$, we have

$$\lambda m j S_{j'}^{m-1} + S_{j'} = \lambda m (j + j' - j') S_{j'}^{m-1} + S_{j'} = \sum_{i=1}^j |p_i| + \left[\sum_{i=j+1}^{j'} |p_i| + \lambda m (j - j') S_{j'}^{m-1} \right] \quad (79)$$

$$= \sum_{i=1}^j |p_i| + \left[\sum_{i=j+1}^{j'} (|p_i| - \lambda m S_{j'}^{m-1}) \right] \geq \sum_{i=1}^j |p_i| = \lambda m j S_j^{m-1} + S_j, \quad (80)$$

where the inequality follows from the assumption $|p_j| \geq |p_{j'}|$ for all $j \in [j']$. Because $\lambda m j S^{m-1} + S$ is monotonically increasing w.r.t S (when $S \geq 0$), (80) implies $S_{j'} \geq S_j$. Obviously, it also implies $|p_j| - \lambda m S_j^{m-1} \geq |p_{j'}| - \lambda m S_{j'}^{m-1} \geq 0$ since $|p_j| \geq |p_{j'}|$. Therefore, S_j is monotonically non-decreasing for all j such that $|p_j| - \lambda m S_j^{m-1} \geq 0$. Moreover, if $|p_{j+1}| - \lambda m S_{j+1}^{m-1} = 0$, $S_{j+1} = S_j$. Assume that $|p_{j+1}| - \lambda m S_{j+1}^{m-1} = 0$. Then, $\lambda m (j+1) S_{j+1}^{m-1} + S_{j+1} - \sum_{i=1}^{j+1} |p_i| = \lambda m j S_{j+1}^{m-1} + S_{j+1} - \sum_{i=1}^j |p_i| = 0$ and thus $S_{j+1} = S_j$.

Finally, we prove $\theta = \max\{j : |p_j| - \lambda m S_j^{m-1} \geq 0\}$ by contradiction. Since $|p_1| - \lambda m S_1^{m-1} = S_1$, the maximum value exists. Let $\theta' = \max\{j : |p_j| - \lambda m S_j^{m-1} \geq 0\}$. First suppose $\theta > \theta'$ (of course we assume that $\theta' < d$). Then, from the assumption $|p_\theta| - \lambda m S_\theta^{m-1} < 0$, we have $\text{sign}(p_\theta) \text{sign}(q_\theta) = -1$. This contradicts (i)-(iii) in Theorem 15. Next suppose $\theta < \theta'$ (we assume $\theta' > 1$). Then, the subdifferential of $g_{\lambda \|\cdot\|_1^m}$ at \mathbf{q}^* is

$$\partial g_{\lambda \|\cdot\|_1^m}(\mathbf{q}^*) = \mathbf{q}^* - \mathbf{p} + \lambda \partial \|\mathbf{q}^*\|_1^m = \mathbf{q}^* - \mathbf{p} + \lambda m \|\mathbf{q}^*\|_1^{m-1} \prod_{i=1}^d J_i(q_i^*) \quad (81)$$

$$= \mathbf{q}^* - \mathbf{p} + \lambda m S_\theta^{m-1} \prod_{i=1}^d J_i(q_i^*). \quad (82)$$

If $S_\theta = S_{\theta'}$, we can clearly replace θ with θ' and thus we assume $S_\theta < S_{\theta'}$ (S_j is monotonically non-decreasing for all $j \in [\theta']$). Then, $|p_{\theta'}| \geq \lambda m S_{\theta'}^{m-1} > \lambda m S_\theta^{m-1}$ and hence for all $\alpha \in [-1, 1]$

$$q_{\theta'}^* - p_{\theta'} + \lambda m S_\theta^{m-1} \alpha = 0 - p_{\theta'} + \lambda m S_\theta^{m-1} \alpha \neq 0. \quad (83)$$

It implies $\mathbf{0} \notin \partial g_{\lambda \|\cdot\|_1^m}(\mathbf{q}^*; \mathbf{p}) \iff \mathbf{q}^* \neq \arg \min g_{\lambda \|\cdot\|_1^m}(\mathbf{q}; \mathbf{p})$, and this is a contradiction. \blacksquare

A.3 Proximal Operator for $\ell_{2,1}^m$

Theorem 11 *The solution to the proximal problem (51) $\mathbf{Q}^* \in \mathbb{R}^{d \times k}$ is*

$$\mathbf{q}_j^* = \begin{cases} \frac{c_j^*}{\|\mathbf{p}_j\|_2} \mathbf{p}_j & \|\mathbf{p}_j\|_2 \neq 0, \\ \mathbf{0} & \|\mathbf{p}_j\|_2 = 0, \end{cases} \quad (52)$$

where $\mathbf{c}^* = \text{prox}_{\lambda \|\cdot\|_1^m} \left((\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2)^\top \right)$.

Proof If $\mathbf{p}_j = \mathbf{0}$, clearly $\mathbf{q}_j^* = \mathbf{0}$. Then, we can eliminate j -th row vector from the proximal problem (51). Hence, with out loss of generality, we assume that $\mathbf{p}_j \neq \mathbf{0}$ for all $j \in [d]$.

We first show that

$$\exists C_j^* \geq 0 \text{ s.t. } \mathbf{q}_j^* = C_j^* \mathbf{p}_j \quad \forall j \in [d]. \quad (84)$$

If $\mathbf{q}_j^* = \mathbf{0}$, then $\mathbf{q}_j^* = C_j^* \mathbf{p}_j$ holds with $C_j^* = 0$. Thus, we next consider the case where $\mathbf{q}_j^* \neq \mathbf{0}$. Let $g_{\lambda \|\cdot\|_{2,1}^m}(\mathbf{Q}; \mathbf{P}) = \|\mathbf{P} - \mathbf{Q}\|_2^2 / 2 + \lambda \|\mathbf{Q}\|_{2,1}^m$. By construction, $\mathbf{Q}^* = \arg \min_{\mathbf{Q}} g_{\lambda \|\cdot\|_{2,1}^m}(\mathbf{Q}; \mathbf{P})$ and $\mathbf{0} \in \partial g_{\lambda \|\cdot\|_{2,1}^m}(\mathbf{Q}^*; \mathbf{P})$. From Theorem 14, we have

$$\mathbf{0} \in -\mathbf{p}_j + \mathbf{q}_j^* + \lambda m \|\mathbf{Q}^*\|_{2,1}^{m-1} \partial \|\mathbf{q}_j^*\|_2, \quad \forall j \in [d]. \quad (85)$$

By the assumption $\mathbf{q}_j^* \neq \mathbf{0}$, $\partial \|\mathbf{q}_j^*\|_2 = \left\{ \mathbf{q}_j^* / \|\mathbf{q}_j^*\|_2 \right\}$. Therefore, we obtain

$$\mathbf{0} = -\mathbf{p}_j + \mathbf{q}_j^* + \lambda m \|\mathbf{Q}^*\|_{2,1}^{m-1} \frac{1}{\|\mathbf{q}_j^*\|_2} \mathbf{q}_j^* \rightarrow \left(1 + \frac{\lambda m \|\mathbf{Q}^*\|_{2,1}^{m-1}}{\|\mathbf{q}_j^*\|_2} \right) \mathbf{q}_j^* = \mathbf{p}_j. \quad (86)$$

It implies (84).

Since (84) holds, we consider the following optimization problem instead of (51):

$$\mathbf{C}^* = \arg \min_{\mathbf{C} \in \mathbb{R}_{\geq 0}^d} \frac{1}{2} \sum_{j=1}^d (\|\mathbf{p}_j\|_2 - C_j \|\mathbf{p}_j\|_2)^2 + \lambda \left\| (\|C_1 \mathbf{p}_1\|_2, \dots, \|C_d \mathbf{p}_d\|_2)^\top \right\|_1^m. \quad (87)$$

Let $c_j^* = C_j^* \|\mathbf{p}_j\|_2$ for all $j \in [d]$. Then, we have $\mathbf{q}_j^* = (c_j^* / \|\mathbf{p}_j\|_2) \cdot \mathbf{p}_j$, where

$$\mathbf{c}^* = \arg \min_{\mathbf{c} \in \mathbb{R}_{\geq 0}^d} \frac{1}{2} \sum_{j=1}^d (\|\mathbf{p}_j\|_2 - c_j)^2 + \lambda \|\mathbf{c}\|_1^m \quad (88)$$

$$= \arg \min_{\mathbf{c} \in \mathbb{R}^d} \frac{1}{2} \left\| (\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2)^\top - \mathbf{c} \right\|_2^2 + \lambda \|\mathbf{c}\|_1^m \quad (89)$$

$$= \text{prox}_{\lambda \|\cdot\|_1^m} \left((\|\mathbf{p}_1\|_2, \dots, \|\mathbf{p}_d\|_2)^\top \right), \quad (90)$$

which concludes the proof. ■

Note that Theorem 11 is a generalization of Theorem 8.

A.4 Regularization by Powers of Norms

Theorem 4 $\Omega : \mathbb{R}^{d \times k} \rightarrow \mathbb{R}_{\geq 0}$ is an m -homogeneous quasi-norm if and only if there exists a quasi-norm $\|\cdot\|'$ such that $\Omega(\cdot) = (\|\cdot\|')^m$.

Proof (\Rightarrow) Let $\Omega : \mathbb{R}^{d \times k} \rightarrow \mathbb{R}_{\geq 0}$ be an m -homogeneous quasi-norm with $\Omega(\mathbf{P} + \mathbf{Q}) \leq K(\Omega(\mathbf{P}) + \Omega(\mathbf{Q}))$ for all $\mathbf{P}, \mathbf{Q} \in \mathbb{R}^{d \times k}$. We show that $\sqrt[m]{\Omega}$ satisfies the axiom of quasi-norm. For all $\mathbf{P}, \mathbf{Q} \in \mathbb{R}^{d \times k}$, we have

- $\sqrt[m]{\Omega(\mathbf{P})} \geq 0$ and $\sqrt[m]{\Omega(\mathbf{P})} = 0 \iff \mathbf{P} = \mathbf{0}$ since $\Omega(\mathbf{P}) \geq 0$ and $\Omega(\mathbf{P}) = 0 \iff \mathbf{P} = \mathbf{0}$,
- $\sqrt[m]{\Omega(\alpha\mathbf{P})} = \sqrt[m]{|\alpha|^m \Omega(\mathbf{P})} = |\alpha| \sqrt[m]{\Omega(\mathbf{P})}$ for all $\alpha \in \mathbb{R}$, and
- $\sqrt[m]{\Omega(\mathbf{P} + \mathbf{Q})} \leq \sqrt[m]{K(\Omega(\mathbf{P}) + \Omega(\mathbf{Q}))} \leq \sqrt[m]{K} \left\{ \sqrt[m]{\Omega(\mathbf{P})} + \sqrt[m]{\Omega(\mathbf{Q})} \right\}$.

Thus $\sqrt[m]{\Omega}$ is a quasi-norm.

(\Leftarrow) Let $\|\cdot\|'$ be a quasi-norm with $\|\mathbf{P} + \mathbf{Q}\|' \leq K(\|\mathbf{P}\|' + \|\mathbf{Q}\|')$ for all $\mathbf{P}, \mathbf{Q} \in \mathbb{R}^{d \times k}$. We show that $(\|\cdot\|')^m$ satisfies the axiom of m -homogeneous quasi-norm. For all $\mathbf{P}, \mathbf{Q} \in \mathbb{R}^{d \times k}$, the following holds:

- $(\|\mathbf{P}\|')^m \geq 0$ and $(\|\mathbf{P}\|')^m = 0 \iff \mathbf{P} = \mathbf{0}$ since $\|\mathbf{P}\|' \geq 0$ and $\|\mathbf{P}\|' = 0 \iff \mathbf{P} = \mathbf{0}$,
- $(\|\alpha\mathbf{P}\|')^m = (|\alpha| \|\mathbf{P}\|')^m = |\alpha|^m (\|\mathbf{P}\|')^m$ for all $\alpha \in \mathbb{R}$, and
- since $\|\mathbf{P}\|' \geq 0$ for all $\mathbf{P} \in \mathbb{R}^{d \times k}$ and x^m is convex and monotone for all $x \geq 0$ and $m \geq 1$,

$$(\|\mathbf{P} + \mathbf{Q}\|')^m \leq \{K(\|\mathbf{P}\|' + \|\mathbf{Q}\|')\}^m = (2K)^m \left\{ \frac{1}{2} \|\mathbf{P}\|' + \frac{1}{2} \|\mathbf{Q}\|' \right\}^m \quad (91)$$

$$\leq (2K)^m \left\{ \frac{1}{2} (\|\mathbf{P}\|')^m + \frac{1}{2} (\|\mathbf{Q}\|')^m \right\} \quad (92)$$

$$= (2^{\frac{m-1}{m}} K)^m \{(\|\mathbf{P}\|')^m + (\|\mathbf{Q}\|')^m\}, \quad (93)$$

where the last inequality follows from Jensen's inequality. ■

Before proving Theorem 9, we show that all quasi-norms on a finite-dimensional vector space (in this paper we consider $\mathbb{R}^{d \times k}$) are equivalent to each other.

Lemma 16 *For given two quasi-norms on $\mathbb{R}^{d \times k}$, $\Omega, \Omega' : \mathbb{R}^{d \times k} \rightarrow \mathbb{R}$, there exist $c, C \in \mathbb{R}_{>0}$ and for all \mathbf{P} such that*

$$c\Omega(\mathbf{P}) \leq \Omega'(\mathbf{P}) \leq C\Omega(\mathbf{P}). \quad (94)$$

Proof We consider only $\mathbf{P} \neq \mathbf{0}$ since it trivially holds for $\mathbf{P} = \mathbf{0}$. We can prove it like the equivalence of norms on a finite-dimensional vector space. Let $\Omega_\infty(\mathbf{P}) := \max_{j \in [d], s \in [k]} |p_{j,s}|$. Ω_∞ is clearly (quasi-)norm and it is sufficient to prove that any quasi-norm is equivalent to Ω_∞ : if for given two quasi-norms $\Omega, \Omega' : \mathbb{R}^{d \times k} \rightarrow \mathbb{R}$ there exist $c, C, c', C' > 0$ such that $c\Omega_\infty(\mathbf{P}) \leq \Omega(\mathbf{P}) \leq C\Omega_\infty(\mathbf{P})$ and $c'\Omega_\infty(\mathbf{P}) \leq \Omega'(\mathbf{P}) \leq C'\Omega_\infty(\mathbf{P})$ for all $\mathbf{P} \in \mathbb{R}^{d \times k}$, then (94) holds: $C/c'\Omega(\mathbf{P}) \leq \Omega'(\mathbf{P}) \leq C'/c\Omega(\mathbf{P})$.

We first show that for any quasi-norm Ω there exists $C > 0$ such that $\Omega(\mathbf{P}) \leq C\Omega_\infty(\mathbf{P})$ for all \mathbf{P} . Assume that Ω is quasi-norm with $\Omega(\mathbf{P} + \mathbf{Q}) \leq K(\Omega(\mathbf{P}) + \Omega(\mathbf{Q}))$ ($K \geq 1$). Let $\mathbf{E}_{j,s}^{d,k}$ be the $d \times k$ matrix such that its (j, s) element is 1 and others are 0. Then, we have

$$\Omega(\mathbf{P}) = \Omega \left(\sum_{j \in [d], s \in [k]} p_{j,s} \mathbf{E}_{j,s}^{d,k} \right) \leq K^{dk} \sum_{j \in [d], s \in [k]} |p_{j,s}| \Omega(\mathbf{E}_{j,s}^{d,k}) \quad (95)$$

$$\leq \left(K^{dk} \sum_{j \in [d], s \in [k]} \Omega(\mathbf{E}_{j,s}^{d,k}) \right) \max_{j \in [d], s \in [k]} |p_{j,s}| = \left(K^{dk} \sum_{j \in [d], s \in [k]} \Omega(\mathbf{E}_{j,s}^{d,k}) \right) \Omega_\infty(\mathbf{P}). \quad (96)$$

Since $K^{dk} \sum_{j \in [d], s \in [k]} \Omega(\mathbf{E}_{j,s}^{d,k})$ does not depend on \mathbf{P} , setting C to be $K^{dk} \sum_{j \in [d], s \in [k]} \Omega(\mathbf{E}_{j,s}^{d,k})$ produces $\Omega(\mathbf{P}) \leq C\Omega_\infty(\mathbf{P})$ for all \mathbf{P} .

We next show that for any quasi-norm Ω there exists $c > 0$ such that $c\Omega_\infty(\mathbf{P}) \leq \Omega(\mathbf{P})$ for all \mathbf{P} . We prove it by contradiction: suppose that there does not exist $c > \mathbb{R}$ such that $c\Omega_\infty(\mathbf{P}) \leq \Omega(\mathbf{P})$. This implies for all $c > 0$ there exists \mathbf{P} such that $\Omega_\infty(\mathbf{P})/\Omega(\mathbf{P}) > 1/c$, i.e., $\sup \Omega_\infty(\mathbf{P})/\Omega(\mathbf{P}) = \infty$. Thus, for all $n \in \mathbb{N}$, there exists $\mathbf{P}^{(n)}$ such that $\Omega_\infty(\mathbf{P}^{(n)})/\Omega(\mathbf{P}^{(n)}) > n$, and without loss of generality we can assume that $\Omega_\infty(\mathbf{P}^{(n)}) = 1$ and $\Omega(\mathbf{P}^{(n)}) < 1/n$ (if $\Omega_\infty(\mathbf{Q}^{(n)})/\Omega(\mathbf{Q}^{(n)}) > n$, then define $\mathbf{P}^{(n)}$ as $\mathbf{Q}^{(n)}/\Omega_\infty(\mathbf{Q}^{(n)})$). We show the contradiction by constructing a convergent sequence $\lim_{n \rightarrow \infty} \mathbf{P}^{(n)} = \mathbf{0}$.

Since $\Omega_\infty(\mathbf{P}^{(n)}) = \sup_{j \in [d], s \in [k]} |p_{j,s}|$ for all $n \in \mathbb{N}$, there exists $(j', s') \in [d] \times [k]$ such that $p_{j',s'}^{(n)} = 1$ for infinitely many $n \in \mathbb{N}$. We assume $(j', s') = (1, 1)$ without loss of generality and $\{\mathbf{P}^{(n)}\}$ be the sequence such that $\Omega_\infty(\mathbf{P}^{(n)}) = 1$, $\Omega(\mathbf{P}^{(n)}) < 1/n$, and $p_{1,1}^{(n)} = 1$. Then, $\{\mathbf{P}^{(n)}\}$ is a bounded sequence in $\mathbb{R}^{d \times k}$ and hence it has a convergent subsequence on the normed space $(\mathbb{R}^{d \times k}, \|\cdot\|_2)$ (Bolzano–Weierstrass theorem). Let \mathbf{P} be the limit of that subsequence and we redefine $\{\mathbf{P}^{(n)}\}$ as the corresponding convergent subsequence. Then, $\Omega_\infty(\mathbf{P} - \mathbf{P}^{(n)}) \rightarrow 0$ as $n \rightarrow \infty$ since Ω_∞ is a norm and we have

$$\Omega(\mathbf{P}) \leq K(\Omega(\mathbf{P} - \mathbf{P}^{(n)}) + \Omega(\mathbf{P}^{(n)})) \quad (97)$$

$$\leq K \left(C\Omega_\infty(\mathbf{P} - \mathbf{P}^{(n)}) + \frac{1}{n} \right). \quad (98)$$

Since n is arbitrary, $\Omega(\mathbf{P}) = 0$ and hence $\mathbf{P} = \mathbf{0}$. However, $p_{1,1}^{(n)} = 1$ for all $n \in \mathbb{N}$. It implies $p_{1,1} = 1$ and $\mathbf{P} \neq \mathbf{0}$. This is a contradiction and thus the assumption is wrong. Therefore, there exists $c > \mathbb{R}$ such that $c\Omega_\infty(\mathbf{P}) \leq \Omega(\mathbf{P})$. \blacksquare

Finally, we prove Theorem 9.

Theorem 9 *For any m -homogeneous quasi-norm Ω^m , there exists $C > 0$ such that $\Omega_*^m(\mathbf{P}) \leq C\Omega^m(\mathbf{P})$ for all $\mathbf{P} \in \mathbb{R}^{d \times k}$.*

Proof From Theorem 4 and Theorem 16, it is sufficient to prove there exists a quasi-norm $\|\cdot\|$ such that $\Omega_*^m(\mathbf{P}) \leq C\|\mathbf{P}\|^m$ for all $\mathbf{P} \in \mathbb{R}^{d \times k}$, and here $\|\cdot\|_1$ corresponds to such a norm: $\Omega_*^m(\mathbf{P}) \leq \Omega_{\text{TI}}^m(\mathbf{P}) \leq \|\mathbf{P}\|_1^m$. \blacksquare

A.5 Justification for TI Regularizer

Theorem 5 *Let $L_{\text{QR}}(\mathbf{w}, \mathbf{W}; \lambda_w)$ be the objective function of the QR with ℓ_2^2 regularization for \mathbf{w} :*

$$L_{\text{QR}}(\mathbf{w}, \mathbf{W}; \lambda_w) := \sum_{n=1}^N \ell(f_{\text{QR}}(\mathbf{x}_n), y_n)/N + \lambda_w \|\mathbf{w}\|_2^2. \quad (22)$$

Then, for any $\lambda_w, \lambda_p, \tilde{\lambda}_p \geq 0$, there exists $k' \leq d(d-1)/2$ such that for all $k \geq k'$,

$$\begin{aligned} & \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{P} \in \mathbb{R}^{d \times k}} L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \tilde{\lambda}_p \Omega_{\text{TI}}(\mathbf{P}) \\ & \leq \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{W} \in \mathbb{R}^{d \times d}} L_{\text{QR}}(\mathbf{w}, \mathbf{W}; \lambda_w) + (\tilde{\lambda}_p + 2\lambda_p) \|\mathbf{W}\|_1. \end{aligned} \quad (23)$$

Moreover, if $\lambda_p = 0$, the equality holds, and $f_{\text{FM}}(\mathbf{x}; \mathbf{w}_{\text{TI}}^, \mathbf{P}_{\text{TI}}^*) = f_{\text{QR}}(\mathbf{x}; \mathbf{w}_{\text{QR}}^*, \mathbf{W}_{\text{QR}}^*)$ for all $\mathbf{x} \in \mathbb{R}^d$, where $\{\mathbf{w}_{\text{TI}}^*, \mathbf{P}_{\text{TI}}^*\}$ and $\{\mathbf{w}_{\text{QR}}^*, \mathbf{W}_{\text{QR}}^*\}$ are the solutions of the left- and right-hand sides, respectively, of Equation (23).*

Proof Without loss of generality, we can omit the linear term. We first consider the case $\lambda_p = 0$. We prove (23) with $\lambda_p = 0$ by showing that for any strictly upper triangular matrix $\mathbf{W} \in \mathbb{R}^{d \times d}$ there exists $\mathbf{P} \in \mathbb{R}^{d \times d(d-1)/2}$ such that

$$\mathbf{P}\mathbf{P}^\top =_{>} \mathbf{W}, \quad \Omega_{\text{TI}}(\mathbf{P}) = \Omega_*(\mathbf{W}). \quad (99)$$

It is sufficient for (23) since $\mathbf{P}\mathbf{P}^\top =_{>} \mathbf{W}$ implies $L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \mathcal{D}, \lambda_w, 0) = L_{\text{QR}}(\mathbf{w}, \mathbf{W}; \mathcal{D}, \lambda_w, 0)$ and \mathbf{W}_{QR}^* is always strictly upper triangular matrix since lower triangle elements are not used in f_{QR} . Fix \mathbf{W} be a strictly upper triangular matrix and let \mathbf{Q} be the $d \times d^2$ matrix with

$$\mathbf{q}_j = \text{vec} \left(\left(\underbrace{\left(\sqrt{|w_{1,j}|} \mathbf{e}_j^d, \dots, \sqrt{|w_{j-1,j}|} \mathbf{e}_j^d \right)}_{j-1}, \text{sign}(\mathbf{w}_j) \circ \sqrt{\text{abs}(\mathbf{w}_j)}, \underbrace{\mathbf{0}, \dots, \mathbf{0}}_{d-j} \right) \right) \in \mathbb{R}^{d^2}, \quad (100)$$

where $\sqrt{\cdot}$ for a vector is the element-wise square root. Then, for all $0 < j_1 < j_2 \leq d$,

$$\begin{aligned} \langle \mathbf{q}_{j_1}, \mathbf{q}_{j_2} \rangle &= \sum_{i=1}^{j_1-1} \left\langle \sqrt{|w_{i,j_1}|} \mathbf{e}_{j_1}^d, \sqrt{|w_{i,j_2}|} \mathbf{e}_{j_2}^d \right\rangle + \left\langle \text{sign}(\mathbf{w}_{j_1}) \circ \sqrt{\text{abs}(\mathbf{w}_{j_1})}, \sqrt{|w_{j_1,j_2}|} \mathbf{e}_{j_2}^d \right\rangle \\ &\quad + \sum_{i=j_1+1}^{j_2-1} \left\langle \mathbf{0}, \sqrt{|w_{i,j_2}|} \mathbf{e}_{j_2}^d \right\rangle + \left\langle \mathbf{0}, \text{sign}(\mathbf{w}_{j_2}) \circ \sqrt{\text{abs}(\mathbf{w}_{j_2})} \right\rangle + \sum_{i=j_2+1}^d \langle \mathbf{0}, \mathbf{0} \rangle \quad (101) \end{aligned}$$

$$= (d^2 - 1) \cdot 0 + \text{sign}(w_{j_1,j_2}) \sqrt{|w_{j_1,j_2}|} \sqrt{|w_{j_1,j_2}|} = w_{j_1,j_2} \iff \mathbf{Q}\mathbf{Q}^\top =_{>} \mathbf{W}, \quad (102)$$

$$\sum_{i=1}^{d^2} |q_{j_1,i}| |q_{j_2,i}| = |w_{j_1,j_2}| \iff \Omega_{\text{TI}}(\mathbf{Q}) = \Omega_*(\mathbf{W}). \quad (103)$$

This proves (23) with $\lambda_p = 0$ when $k \geq d^2$.

Here, we show that \mathbf{Q} has $d(d+1)/2$ all-zeros columns. Let $\mathbf{Q}^j = (\mathbf{q}_{\cdot, d(j-1)+1}, \dots, \mathbf{q}_{\cdot, d(j-1)+d}) \in \mathbb{R}^{d \times d}$, i.e., $\mathbf{Q} = (\mathbf{Q}^1, \dots, \mathbf{Q}^d)$. Then, $1, \dots, j$ -th columns in \mathbf{Q}^j are all-zeros vectors since the row vectors in \mathbf{Q}^j are

$$\mathbf{q}_{j_1}^j = (q_{j_1, d(j-1)+1}, \dots, q_{j_1, d(j-1)+d})^\top = \mathbf{0} \text{ for all } j_1 < j, \quad (104)$$

$$\mathbf{q}_j^j = \mathbf{w}_j = (0, \dots, 0, w_{j,j+1}, \dots, w_{j,d})^\top, \quad (105)$$

$$\mathbf{q}_{j_2}^j = \mathbf{e}_{j_2}^d = \underbrace{(0, \dots, 0, 1, 0, \dots, 0)}_{j_2-1}^\top \text{ for all } j_2 > j. \quad (106)$$

Thus, \mathbf{Q} has $1 + 2 + \dots + d = d(d+1)/2$ all-zeros columns and let $\mathbf{P} \in \mathbb{R}^{d(d-1)/2}$ be the sub-matrix of \mathbf{Q} such that its all-zeros columns are removed. Then $\mathbf{P}\mathbf{P}^\top =_{>} \mathbf{W}$ and $\Omega_{\text{TI}}(\mathbf{P}) = \|\mathbf{W}\|_1$. It proves (23) with $\lambda_p = 0$. Furthermore, since $\|\mathbf{W}\|_1 \leq \Omega_{\text{TI}}(\mathbf{P})$ for all $\mathbf{P}\mathbf{P}^\top =_{>} \mathbf{W}$, the inverse inequality clearly holds if $\lambda_p = 0$:

$$\min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{P} \in \mathbb{R}^{d \times k}} L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \mathcal{D}, \lambda_w, 0) + \tilde{\lambda}_p \Omega_{\text{TI}}(\mathbf{P}) \geq \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{W} \in \mathbb{R}^{d \times d}} L_{\text{QR}}(\mathbf{w}, \mathbf{W}; \mathcal{D}, \lambda_w) + \tilde{\lambda}_p \|\mathbf{W}\|_1. \quad (107)$$

It implies the equality holds in (23) and $f_{\text{FM}}(\mathbf{x}; \mathbf{w}_{\text{TI}}^*, \mathbf{P}_{\text{TI}}^*) = f_{\text{QR}}(\mathbf{x}; \mathbf{w}_{\text{QR}}^*, \mathbf{W}_{\text{QR}}^*)$ for all $\mathbf{x} \in \mathbb{R}^d$.

Next, we prove (23) with $\lambda_p \geq 0$. For \mathbf{P} as defined above, we have

$$\|\mathbf{P}\|_2^2 = \sum_{j=1}^d \|\mathbf{p}_j\|_2^2 = \sum_{j=1}^d \left\{ \sum_{i=1}^{j-1} |w_{i,j}| + \sum_{i=1}^d |w_{j,i}| \right\} = 2 \|\mathbf{W}\|_1. \quad (108)$$

Algorithm 2 Computation of the proximal operator (30) with sorting \mathbf{p} ($O(d \log d)$).

Input: $\mathbf{p} \in \mathbb{R}^d$, $\lambda \geq 0$

 1: $\tilde{\mathbf{p}} \leftarrow (p_{(1)}, \dots, p_{(d)})^\top$; $\triangleright O(d \log d)$

 2: $S_0 \leftarrow 0$;

 3: **for** $j = 1, \dots, d$ **do**

 4: $S_j \leftarrow S_{j-1} + |\tilde{p}_j|$;

 5: **end for**

 6: $S_j \leftarrow S_j / (1 + 2\lambda j)$ for all $j \in [d]$;

 7: $\theta \leftarrow \max\{j \in [d] : |\tilde{p}_j| - 2\lambda S_j \geq 0\}$;

 8: $q_j^* \leftarrow \text{sign}(p_j) \max\{|p_j| - 2\lambda S_\theta, 0\}$ for all $j \in [d]$;

Output: $\mathbf{q}^* (= \text{prox}_{\lambda \|\cdot\|_1^2}(\mathbf{p}))$

 Combining it with the above-mentioned proof for $\lambda_p = 0$ implies (23) for all $\lambda_p \geq 0$. ■

 We also obtain a similar relationship between TI-sparse FMs and Ω_* -sparse FMs.

Theorem 17 For any $\lambda_w, \lambda_p, \tilde{\lambda}_p \geq 0$ and $k_* \in \mathbb{N}_{>0}$, there exists $k' \leq d(d-1)/2$ such that for all $k \geq k'$,

$$\begin{aligned} & \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{P} \in \mathbb{R}^{d \times k}} L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, \lambda_p) + \tilde{\lambda}_p \Omega_{\text{TI}}(\mathbf{P}) \\ & \leq \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{P} \in \mathbb{R}^{d \times k_*}} L_{\text{FM}}(\mathbf{w}, \mathbf{P}; \lambda_w, (d-1)\lambda_p) + \tilde{\lambda}_p \Omega_*(\mathbf{P}). \end{aligned} \quad (109)$$

Proof Let $\mathbf{P}_{\Omega_*}^*$ be the optimal solution of the RHS in (109). Then, we easily obtain (109) with $\lambda_p = 0$ by substituting the strictly upper triangular elements of $\mathbf{P}_{\Omega_*}^* (\mathbf{P}_{\Omega_*}^*)^\top$ to those of \mathbf{W} in the proof of Theorem 5. Thus, for $\lambda_p \geq 0$, we show that $\|\mathbf{P}\|_2^2 \leq (d-1) \|\mathbf{P}_{\Omega_*}^*\|_2^2$, where $\mathbf{P} \in \mathbb{R}^{d \times d(d-1)/2}$ is constructed as in the proof of Theorem 5. It is sufficient for (109). From (108), we have

$$\|\mathbf{P}\|_2^2 = 2 \sum_{j_2 > j_1} |w_{j_1, j_2}| = 2 \sum_{j_2 > j_1} |\langle \mathbf{p}_{\Omega_*, j_1}^*, \mathbf{p}_{\Omega_*, j_2}^* \rangle| \leq 2 \sum_{j_2 > j_1} \|\mathbf{p}_{\Omega_*, j_1}^*\|_2 \|\mathbf{p}_{\Omega_*, j_2}^*\|_2 \quad (110)$$

$$\leq \sum_{j_2 > j_1} \|\mathbf{p}_{\Omega_*, j_1}^*\|_2^2 + \|\mathbf{p}_{\Omega_*, j_2}^*\|_2^2 = (d-1) \|\mathbf{P}_{\Omega_*}^*\|_2^2. \quad (111)$$

■

Appendix B. Implementation Details

In this section, we briefly review and show the implementation details of the existing and proposed algorithms.

$O(d \log d)$ *Time Algorithm for Proximal Operator* (30) (Martins et al., 2011). Algorithm 2 shows the procedure for solving (30) in $O(d \log d)$ time, where $p_{(i)}$ is the i -th largest value of $p_{(i)}$ in absolute value; i.e., $|p_{(1)}| \geq |p_{(2)}| \geq \dots \geq |p_{(d)}|$ and $\forall j \in [d] \exists i \in [d]$ s.t. $p_j = p_{(i)}$.

$O(d)$ *Time Algorithm for Proximal Operator* (30) (Martins et al., 2011). In fact, the proximal operator (30) can be computed in $O(d)$ time. Given $p_{(\theta)}$ (not the index θ , but the value $p_{(\theta)}$), one can compute S_θ in $O(d)$ time: $S_{G_{(\theta)}} := \sum_{j \in G_{(\theta)}} |p_j| / (1 + 2\lambda |G_{(\theta)}|) = S_\theta$, where $G_{(\theta)} := \{j \in [d] : |p_j| \geq |p_{(\theta)}|\}$ (clearly, $\theta = |G_{(\theta)}|$). Thus, even if \mathbf{p} is not sorted by absolute value, one can compute \mathbf{q}^* in $O(d)$ time if only $p_{(\theta)}$ is found. Fortunately, one can find $p_{(\theta)}$ in $O(d)$ time in expectation by using the randomized-median-finding-like algorithm (Duchi et al., 2008). Algorithm 3 shows the

Algorithm 3 Computation of the proximal operator (30) without sorting ($O(d)$)

Input: $\mathbf{p} \in \mathbb{R}^d$, $\lambda \geq 0$
 1: $C \leftarrow [d]$; ▷ Candidate index set
 2: $S \leftarrow 0$;
 3: $\theta \leftarrow 0$;
 4: **while** $C \neq \emptyset$ **do** ▷ Find $p_{(\theta)}$
 5: Pick $i \in C$ at random;
 6: $G \leftarrow \{j \in C : |p_j| \geq |p_i|\}$;
 7: $L \leftarrow \{j \in C : |p_j| < |p_i|\}$;
 8: $S_{G_i} \leftarrow (S + \sum_{j \in G} |p_j|)/(1 + 2\lambda(\theta + |G|))$;
 9: **if** $|p_i| - 2\lambda S_{G_i} \geq 0$ **then** ▷ $|p_i| \geq |p_{(\theta)}|$;
 10: $C \leftarrow L$, $S \leftarrow S + \sum_{j \in G} |p_j|$, $\theta \leftarrow \theta + |G|$;
 11: **else** ▷ $|p_i| < |p_{(\theta)}|$
 12: $C \leftarrow G \setminus \{j \in G : |p_j| = |p_i|\}$;
 13: **end if**
 14: **end while**
 15: $S_\theta \leftarrow S/(1 + 2\lambda\theta)$;
 16: $q_j^* \leftarrow \text{sign}(p_j) \max\{|p_j| - 2\lambda S_\theta, 0\}$ for all $j \in [d]$;
Output: $\mathbf{q}^* (= \text{prox}_{\lambda \|\cdot\|_1^2}(\mathbf{p}))$

procedure for computing the proximal operator (30) in $O(d)$ time. It finds $p_{(\theta)}$ and computes S_θ by repeating (i) randomly sampling p_i from $\{p_i : i \in C\}$, (ii) determining whether $|p_i| \geq |p_{(\theta)}|$ (although $p_{(\theta)}$ is unknown), and (iii) reducing C in accordance with whether $|p_i| \geq |p_{(\theta)}|$, where $C \subseteq [d]$ is the set of candidates of θ initialized as $[d]$. When $i \in C$ is sampled, the algorithm partitions the candidate index set C into $G := \{j \in C : |p_j| \geq |p_i|\}$ and $L := \{j \in C : |p_j| < |p_i|\}$ and discards one of them as follows. If $|p_i| \geq \lambda S_{G_i}$, where $G_i := \{j \in [d] : |p_j| \geq |p_i|\}$, $|p_{(\theta)}|$ is smaller than $|p_i|$, so the algorithm next searches for $p_{(\theta)}$ from L (i.e., it discards G). In this case, the algorithm updates S and θ as $S + \sum_{j \in G} |p_j|$ and $\theta + |G|$, respectively (both S and θ are initially 0 and are simply desired values when $C = \emptyset$). Maintaining S and θ reduces the computation cost of S_{G_i} : $S_{G_i} = (S + \sum_{j \in G} |p_j|)/(1 + 2\lambda(\theta + |G|))$. Otherwise (that is, if $|p_i| < 2\lambda S_{G_i}$), the algorithm discards L and sets $C = G \setminus \{j \in G : |p_j| = |p_i|\}$ since $|p_{(\theta)}|$ is larger than $|p_i|$.

CD Algorithm for Canonical FMs. Algorithm 4 shows the CD algorithm for objective function (2). The CD algorithm requires the predictions of all training instances $f_n = f_{\text{FM}}(\mathbf{x}_n; \mathbf{w}, \mathbf{P})$ for all $n \in \text{supp}(\mathbf{x}_{:,j})$ for updating $p_{j,s}$. The naïve computation of such f_n at each iteration is too costly, so a method for *updating (synchronizing)* predictions is essential for an efficient implementation. Typically, an efficient algorithm computes and caches f_n for all $n \in [N]$ before starting the optimization and synchronizes them every time a parameter is updated. Fortunately, FMs are multi-linear w.r.t w_1, \dots, w_d and $p_{1,1}, \dots, p_{d,k}$, so each prediction can be easily synchronized in $O(1)$. For w_j , the predictions are written as $f_n = x_{n,j}w_j + \text{const}$ for all $n \in [N]$, so f_n are synchronized as $f_n \leftarrow f_n - x_{n,j}\eta\delta$ after updating w_j as $w_j \leftarrow w_j - \eta\delta$. To be more precise, the algorithm synchronizes f_n for only $n \in \text{supp}(\mathbf{x}_{:,j})$ since f_n does not change for all $n \notin \text{supp}(\mathbf{x}_{:,j})$. For $p_{j,s}$, the predictions are written as $f_n = f'_{n,s}p_{j,s} + \text{const}$, where $f'_{n,s} = x_{n,j}(a_{n,s} - x_{n,j}p_{j,s})$ and $a_{n,s} := \langle \mathbf{x}_n, \mathbf{p}_s \rangle$. Thus the algorithm can synchronize f_n as in the case for w_j by caching $a_{n,s}$ for all $n \in [N]$, $s \in [k]$. Clearly, the gradient of the loss w.r.t $p_{j,s}$ can be efficiently computed by using $f'_{n,s}$: $\partial \ell(f_n, y_n) / \partial p_{j,s} = \partial \ell(f_n, y_n) / \partial f_n \cdot f'_{n,s}$. At each iteration, the algorithm requires $a_{n,s}$ for all $n \in [N]$ and therefore the additional space complexity for caches is $O(N)$ (the algorithm does not require $a_{n,s}$ for all $s \in [k]$ at the same time).

SGD for Canonical FMs. Algorithm 5 shows the SGD algorithm for objective function (2). Since it is almost straightforward, we only describe the *lazy update* technique (Duchi et al., 2011), which

Algorithm 4 CD algorithm for canonical FMs

Input: $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$, $k \in \mathbb{N}_{>0}$, $\lambda_w, \lambda_p \geq 0$
 1: Initialize $\mathbf{P} \in \mathbb{R}^{d \times k}$, $\mathbf{w} \in \mathbb{R}^d$;
 2: Compute predictions: $f_n \leftarrow f_{\text{FM}}(\mathbf{x}_n; \mathbf{w}, \mathbf{P})$ for all $n \in [N]$;
 3: **while** not convergence **do**
 4: **for** $j = 1, \dots, d$ **do** ▷ Update \mathbf{w}
 5: $\delta \leftarrow \sum_{i \in \text{supp}(\mathbf{x}_{:,j})} \ell'(f_n, y_n) x_{n,j} / N + 2\lambda_w w_j$;
 6: $\eta \leftarrow (\mu \|\mathbf{x}_{:,j}\|_2^2 / N + 2\lambda_w)^{-1}$;
 7: $w_j \leftarrow w_j - \eta \delta$;
 8: $f_n \leftarrow f_n - x_{n,j} \eta \delta$ for all $n \in \text{supp}(\mathbf{x}_{:,j})$;
 9: **end for**
 10: **for** $s = 1, \dots, k$ **do** ▷ Update \mathbf{P}
 11: $a_{n,s} \leftarrow \langle \mathbf{x}_n, \mathbf{p}_{:,s} \rangle$ for all $n \in [N]$; ▷ Cache for updating $\mathbf{p}_{:,s}$
 12: **for** $j = 1, \dots, d$ **do**
 13: $f'_{n,s} \leftarrow x_{n,j} (a_{n,s} - x_{n,j} p_{j,s})$ for all $n \in \text{supp}(\mathbf{x}_{:,j})$;
 14: $\delta \leftarrow \sum_{i \in \text{supp}(\mathbf{x}_{:,j})} \ell'(f_n, y_n) \cdot f'_{n,s} / N + 2\lambda_p p_{j,s}$;
 15: $\eta \leftarrow [\mu \sum_{n \in \text{supp}(\mathbf{x}_{:,j})} (f'_{n,s})^2 / N + 2\lambda_p]^{-1}$;
 16: $p_{j,s} \leftarrow p_{j,s} - \eta \delta$;
 17: $f_n \leftarrow f_n - f'_{n,s} \eta \delta$ for all $n \in \text{supp}(\mathbf{x}_{:,j})$;
 18: $a_{n,s} \leftarrow a_{n,s} - x_{n,j} \eta \delta$ for all $n \in \text{supp}(\mathbf{x}_{:,j})$;
 19: **end for**
 20: **end for**
 21: **end while**
Output: Learned \mathbf{P} and \mathbf{w}

improves the efficiency of the algorithm by leveraging the sparsity of a sampled instance. For all $j \notin \text{supp}(\mathbf{x}^t)$, where \mathbf{x}^t is the sampled instance at t -th iteration, the gradients of loss function w.r.t w_j and \mathbf{p}_j are $2\lambda_w w_j$ and $2\lambda_p \mathbf{p}_j$, and update rules are $w_j \leftarrow (1 - 2\eta_w^t) w_j$ and $\mathbf{p}_j \leftarrow (1 - 2\eta_p^t) \mathbf{p}_j$, respectively. The lazy update technique enables us to omit updating such parameters and makes the computational cost of each iteration $O(\text{nnz}(\mathbf{x}_n) + \text{nnz}(\mathbf{x}_n)k)$ while a naïve implementation takes $O(d + dk)$. Although we hereinafter consider for only \mathbf{w} for simplicity, the same holds for \mathbf{P} . Assume that the algorithm is at t -th iteration and $x_j^t \neq 0$, and also assume that there exists $t_j < t$ such that $x_j^{t_j} \neq 0$ and $x_j^{t'} = 0$ for all $t_j < t' < t$ (namely, j -th feature is 0 from $(t_j + 1)$ -th until $(t - 1)$ -th iteration. Then, the value of w_j at $(t - 1)$ -th iteration, w_j^{t-1} , is written as

$$w_j^{t-1} = (1 - 2\eta_w^{t-1} \lambda_w) w_j^{t-2} = \left\{ \prod_{t'=t_j+1}^{t-1} (1 - 2\eta_w^{t'} \lambda_w) \right\} w_j^{t_j} \quad (112)$$

$$= \left\{ \frac{\prod_{t'=1}^{t-1} (1 - 2\eta_w^{t'} \lambda_w)}{\prod_{t'=1}^{t_j} (1 - 2\eta_w^{t'} \lambda_w)} \right\} w_j^{t_j} =: \frac{\alpha_w^{t-1}}{\alpha_w^{t_j}} w_j^{t_j}. \quad (113)$$

Therefore, given α_w^{t-1} and $\alpha_w^{t_j}$, the algorithm can update w_j^t from $w_j^{t_j}$ in $O(1)$. Since $x_j^{t'} = 0$ for all $t' \in \{t_j + 1, \dots, t - 1\}$, omitting to update w_j from $(t_j + 1)$ -th until $(t - 1)$ -th iteration does not affect the result. $\alpha_{w,j}$ in Algorithm 5 corresponds to $\alpha_w^{t_j}$. In our implementation, if $\alpha_w < 10^{-9}$, algorithm updates all w_j and resets α_w and $\alpha_{w,j}$ for all $j \in [d]$ in order to avoid numerical errors. The additional space complexity for caches is $O(d)$ (scaling values $\alpha_{w,j}, \alpha_{p,j}$ for all $j \in [d]$).

PBCD Algorithm for Sparse FMs. Algorithm 6 shows the PBCD algorithm for sparse FMs with/without the line search (Tseng and Yun, 2009), where $\Omega : \mathbb{R}^{d \times k}$ is a sparsity-inducing regular-

Algorithm 5 SGD algorithm for canonical FMs

Input: $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$, $k \in \mathbb{N}_{>0}$, $\lambda_w, \lambda_p \geq 0$
 1: Initialize $\mathbf{P} \in \mathbb{R}^{d \times k}$, $\mathbf{w} \in \mathbb{R}^d$;
 2: $t \leftarrow 1$;
 3: $\alpha_w, \alpha_p, \alpha_{w,j}, \alpha_{p,j} \leftarrow 1$ for all $j \in [d]$; ▷ For lazy update
 4: **while** not convergence **do**
 5: Sample $(\mathbf{x}^t, \mathbf{y}^t) \in \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\}$;
 6: $w_j \leftarrow \alpha_w / \alpha_{w,j} w_j$, $\mathbf{p}_j \leftarrow \alpha_p / \alpha_{p,j} \mathbf{p}_j$ for all $j \in \text{supp}(\mathbf{x}^t)$; ▷ Lazy update for regularization term
 7: Compute the prediction: $f^t \leftarrow f_{\text{FM}}(\mathbf{x}^t; \mathbf{w}, \mathbf{P})$;
 8: $f'_{j,s} \leftarrow x_j^t [\langle \mathbf{x}^t, \mathbf{p}_{:,s} \rangle - x_j^t p_{j,s}]$ for all $j \in \text{supp}(\mathbf{x}^t)$, $s \in [k]$;
 9: $L'_{w,j} \leftarrow \ell'(f^t, \mathbf{y}^t) \cdot x_j^t + 2\lambda_w w_j$ for all $j \in \text{supp}(\mathbf{x}^t)$;
 10: $L'_{p,j,s} \leftarrow \ell'(f^t, \mathbf{y}^t) \cdot f'_{j,s} + 2\lambda_p p_{j,s}$ for all $j \in \text{supp}(\mathbf{x}^t)$, $s \in [k]$;
 11: Compute step size parameter η_w^t, η_p^t ;
 12: $w_j \leftarrow w_j - \eta_w^t L'_{w,j}$ for all $j \in \text{supp}(\mathbf{x}^t)$;
 13: $p_{j,s} \leftarrow p_{j,s} - \eta_p^t L'_{p,j,s}$; for all $j \in \text{supp}(\mathbf{x}^t)$, $s \in [k]$;
 14: $\alpha_w \leftarrow (1 - 2\eta\lambda_w)\alpha_w$, $\alpha_p \leftarrow (1 - 2\eta\lambda_p)\alpha_p$;
 15: $\alpha_{w,j} \leftarrow \alpha_w$, $\alpha_{p,j} \leftarrow \alpha_p$ for all $j \in \text{supp}(\mathbf{x}^t)$;
 16: $t \leftarrow t + 1$;
 17: **end while**
 18: $w_j \leftarrow \alpha_w / \alpha_{w,j} w_j$, $\mathbf{p}_j \leftarrow \alpha_p / \alpha_{p,j} \mathbf{p}_j$ for all $j \in [d]$; ▷ Finalize
Output: Learned \mathbf{P} and \mathbf{w}

izer. The operator $\text{prox}_\Omega(\cdot, j)$ in line 18 is the proximal operator for only j -th row vector. The PBCD algorithm for such sparse FMs updates not $p_{j,s}$ but \mathbf{p}_j at each iteration and it takes $O(\text{nnz}(\mathbf{x}_{:,j})k)$ computational cost (if the proximal operator can be evaluated in $O(k)$). In this algorithm, we show the two variants for choosing step size η : using the line search method proposed by Tseng and Yun (2009) and using the Lipschitz constant of the gradient. $\sigma, \rho \in (0, 1)$ are hyperparameters for the line search. In our experiments, we used the BCD algorithm without the line search. The additional space complexity for caches is $O(Nk)$ (caching $\mathbf{a}_n \in \mathbb{R}^k$ for all $n \in [N]$ requires $O(Nk)$ and caching $\mathbf{q}, \boldsymbol{\delta}, \mathbf{d} \in \mathbb{R}^k$ requires $O(k)$). Some sparse regularizers (e.g, CS regularizer) require some additional caches and how to compute/use caches depends on the regularizer. Moreover, in the case of using the line search, some regularizers might also require a non-obvious efficient incremental evaluation method (i.e., naïve computation of $\Omega(\cdot)$ might take a high computational cost that is prohibitive at each line search iteration).

PSGD Algorithm for Sparse FMs. The extension of Algorithm 5 to PSGD algorithm for sparse FMs is straightforward: evaluates a proximal operator for a sparse regularization after updating \mathbf{P} at each iteration. For $\ell_{1,2}^2$ -sparse FMs (TI-sparse FMs), all parameters must be the latest values at each iteration, i.e., a lazy update technique cannot be used.

Extension to HOFMs. Here we describe the extension of above described algorithms for M -order HOFMs. The extended algorithms update $\mathbf{w}, \mathbf{P}^{(2)}, \dots, \mathbf{P}^{(M)}$ sequentially. HOFMs are also multi-linear w.r.t $\mathbf{w}, \mathbf{p}_1^{(2)}, \dots, \mathbf{p}_d^{(M)}$ (and clearly $w_0, \dots, w_d, p_{1,1}^{(2)}, \dots, p_{d,k}^{(M)}$) (Blondel et al., 2016a) and thus the output of HOFMs are written as $f_{\text{HOFM}}^M(\mathbf{x}) = \langle \boldsymbol{\theta}, \nabla_{\boldsymbol{\theta}} f_{\text{HOFM}}^M(\mathbf{x}) \rangle + \text{const}$ for a parameter $\boldsymbol{\theta} \in \{\mathbf{w}, \mathbf{p}_1^{(2)}, \dots, \mathbf{p}_d^{(M)}\}$. Given $\nabla_{\boldsymbol{\theta}} f_{\text{HOFM}}^M(\mathbf{x})$, the parameter $\boldsymbol{\theta}$ and predictions can be updated efficiently in both CD (updates only one element in $\boldsymbol{\theta}$) and BCD algorithms. For each case, $\partial f_{\text{HOFM}}^M(\mathbf{x}) / \partial \theta_j$ is written as

$$\frac{\partial f_{\text{HOFM}}^M(\mathbf{x})}{\partial \theta_j} = \begin{cases} x_j & \text{if } \theta_j = w_j, j \in [d], \\ x_j K_A^{m-1}(\mathbf{x}_{-j}, (\mathbf{p}_{:,s})_{-j}) & \text{if } \theta_j = p_{j,s}^{(m)}, m \in \{2, \dots, M\}, j \in [k], \end{cases} \quad (114)$$

Algorithm 6 PBCD algorithm for sparse FMs

Input: $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$, $k \in \mathbb{N}_{>0}$, $\lambda_w, \lambda_p, \tilde{\lambda}_p \geq 0$, optional: $\sigma, \rho \in (0, 1)$ (for line search)

- 1: Initialize $\mathbf{P} \in \mathbb{R}^{d \times k}$, $\mathbf{w} \in \mathbb{R}^d$;
- 2: Compute predictions: $f_n \leftarrow f_{\text{FM}}(\mathbf{x}_n; \mathbf{w}, \mathbf{P})$ for all $n \in [N]$;
- 3: **while** not convergence **do**
- 4: Optimize \mathbf{w} and update caches as in canonical FMs;
- 5: $\mathbf{a}_n \leftarrow \mathbf{P}^\top \mathbf{x}_n \in \mathbb{R}^s$ for all $n \in [N]$; ▷ Cache for update \mathbf{P}
- 6: **if** Perform line search **then**
- 7: $L \leftarrow \sum_{n=1}^N \ell(f_n, y_n)/N + \lambda_p \|\mathbf{P}\|_2^2 + \tilde{\lambda}_p \Omega(\mathbf{P})$; ▷ Objective value used in line search
- 8: **end if**
- 9: **for** $j = 1, \dots, d$ **do**
- 10: $\mathbf{f}'_n \leftarrow x_{n,j}(\mathbf{a}_n - x_{n,j}\mathbf{p}_j)$ for all $n \in \text{supp}(\mathbf{x}_{:,j})$;
- 11: $\nabla_{\mathbf{p}_j} \ell_n \leftarrow \ell'(f_n, y_n) \mathbf{f}'_n$ for all $n \in \text{supp}(\mathbf{x}_{:,j})$;
- 12: $\mathbf{d} \leftarrow (\sum_{n \in \text{supp}(\mathbf{x}_{:,j})} \nabla_{\mathbf{p}_j} \ell_n)/N + 2\lambda_p \mathbf{p}_j$;
- 13: **if** Perform line search **then**
- 14: $\eta^{-1} \leftarrow \max_{s \in [k]} \{2\lambda_p + \sum_{n \in \text{supp}(\mathbf{x}_{:,j})} \ell''(f_n, y_n) (\partial f_{\text{FM}}(\mathbf{x}_n)/\partial p_{j,s})^2/N\}$;
- 15: **else**
- 16: $\eta^{-1} \leftarrow 2\lambda_p + \mu \sum_{n \in \text{supp}(\mathbf{x}_{:,j})} \|\nabla_{\mathbf{p}_j} \ell_n\|_2^2/N$;
- 17: **end if**
- 18: $\mathbf{q} \leftarrow \text{prox}_{\tilde{\lambda}_p \eta \Omega}(\mathbf{P} - \eta \mathbf{e}_j^d \mathbf{d}^\top; j)$; ▷ Apply for only j -th row vector
- 19: $\boldsymbol{\delta} \leftarrow \mathbf{p}_j - \mathbf{q}$;
- 20: **if** Perform line search **then**
- 21: $L_{\text{new}} \leftarrow L + \sum_{n \in \text{supp}(\mathbf{x}_{:,j})} [\ell(f_n - \langle \mathbf{f}'_n, \boldsymbol{\delta} \rangle, y_n) - \ell(f_n, y_n)]/N$;
- 22: $L_{\text{new}} \leftarrow L_{\text{new}} + \lambda_p (\|\mathbf{q}\|_2^2 - \|\mathbf{p}_j\|_2^2) + \tilde{\lambda}_p [\Omega(\mathbf{P} - \mathbf{e}_j^d \boldsymbol{\delta}^\top) - \Omega(\mathbf{P})]$;
- 23: $\alpha \leftarrow 1$;
- 24: **while** not $L_{\text{new}} - L \leq \sigma \alpha \{-\langle \mathbf{d}, \boldsymbol{\delta} \rangle + \tilde{\lambda}_p [\Omega(\mathbf{P} - \mathbf{e}_j^d \boldsymbol{\delta}^\top) - \Omega(\mathbf{P})]\}$ **do**
- 25: $L_{\text{new}} \leftarrow L_{\text{new}} + \sum_{n \in \text{supp}(\mathbf{x}_{:,j})} [\ell(f_n - \langle \mathbf{f}'_n, \alpha \rho \boldsymbol{\delta} \rangle, y_n) - \ell(f_n - \langle \mathbf{f}'_n, \alpha \boldsymbol{\delta} \rangle, y_n)]/N$;
- 26: $L_{\text{new}} \leftarrow L_{\text{new}} + \lambda_p [\|\mathbf{p}_j - \alpha \rho \boldsymbol{\delta}\|_2^2 - \|\mathbf{p}_j - \alpha \boldsymbol{\delta}\|_2^2]$;
- 27: $L_{\text{new}} \leftarrow L_{\text{new}} + \tilde{\lambda}_p [\Omega(\mathbf{P} - \alpha \rho \mathbf{e}_j^d \boldsymbol{\delta}^\top) - \Omega(\mathbf{P} - \alpha \mathbf{e}_j^d \boldsymbol{\delta}^\top)]$;
- 28: $\alpha \leftarrow \alpha \rho$;
- 29: **end while**
- 30: $L \leftarrow L_{\text{new}}$;
- 31: $\boldsymbol{\delta} \leftarrow \alpha \boldsymbol{\delta}$;
- 32: **end if**
- 33: $\mathbf{p}_j \leftarrow \mathbf{p}_j - \boldsymbol{\delta}$;
- 34: $\mathbf{a}_n \leftarrow \mathbf{a}_n - x_{n,j} \boldsymbol{\delta}$ for all $n \in \text{supp}(\mathbf{x}_{:,j})$;
- 35: $f_n \leftarrow f_n - \langle \mathbf{f}'_n, \boldsymbol{\delta} \rangle$ for all $n \in \text{supp}(\mathbf{x}_{:,j})$;
- 36: **end for**
- 37: **end while**

Output: Learned \mathbf{P} and \mathbf{w}

where $\mathbf{x}_{-j} \in \mathbb{R}^{d-1}$ is the $(d-1)$ -dimensional vector with x_j removed. Thus, the replacement of $f'_{n,s}$ in Algorithm 4 (\mathbf{f}'_n in Algorithm 6) with $\partial f_{\text{HOFM}}^m(\mathbf{x}_n)/\partial p_{j,s}^{(m)}$ ($\nabla_{\mathbf{p}_j^{(m)}} f_{\text{HOFM}}^m(\mathbf{x}_n)$) produces the CD (BCD) algorithm for HOFMs. Algorithm 5 can be also extended to the SGD algorithm for HOFMs similarly. Here, the issue is clearly how to compute $\partial f_{\text{HOFM}}^M(\mathbf{x}_n)/\partial p_{j,s}^{(m)}$ efficiently. Fortunately, Blondel et al. (2016a) and Atarashi et al. (2020) proposed efficient computation algorithms for $K_A^{m-1}(\mathbf{x}_{-j}, (\mathbf{p}_{:,s})_{-j})$. For more detail, please see (Blondel et al., 2016a; Atarashi et al., 2020).




 (a) ML100K Dataset with $\tilde{\lambda}_p = 10^{-4}$ (left) and 10^{-5} (right).

 (b) a9a Dataset with $\tilde{\lambda}_p = 10^{-4}$ (left) and 10^{-5} (right).

 Figure 7: Trajectories of $\|W\|_1$ regularized objective value for **TI**, **CS**, **L21**, **L1**, Ω_* -nmAPGD, Ω_* -SubGD, and **SLQR** methods on (a) ML100K dataset and (b) a9a dataset.

Appendix C. Additional Experiments

C.1 Efficiency Comparison on Real-world Datasets

We evaluated the efficiency of the proposed and existing methods on the ML100K dataset and the a9a dataset as in Section 8.2.4. We ran the experiment with $\tilde{\lambda}_p = 10^{-4}$ and 10^{-5} on both datasets. We set $\lambda_p = \lambda_{tr} = 10^{-4}$ on the ML100K dataset and $\lambda_p = \lambda_{tr} = 10^{-3}$ on the a9a dataset. The other settings are the same those in Section 8.2.4. Because we compared the convergence speeds of objective values, we didn't separate datasets to training, validation, and testing datasets, i.e., we used 100,000 and 48,642 instances for training on the ML100K and the a9a datasets, respectively.

As shown in Fig. 7a, **TI** and **CS** converged much faster than the other methods in terms of Ω_* ($\|W\|_1$) regularized objective function on both the ML100K dataset and the a9a dataset. These results indicate that the proposed regularizers can be good alternative to Ω_* for not only synthetic datasets but also some real-world datasets.

C.2 Optimization Methods Comparison

We compared the convergence speeds of some algorithms for TI-sparse FMs and CS-sparse FMs on both synthetic and real-world datasets. We ran this experiment on an Arch Linux desktop with an Intel Core i7-4790 (3.60 GHz) CPU and 16 GB RAM.

Methods Compared. We compared the following algorithms for TI-sparse FMs and CS-sparse FMs:

- **PCD**: the proximal coordinate descent algorithm (only TI-sparse FMs).
- **PBCD**: the proximal block coordinate descent algorithm (only CS-sparse FMs).
- **APGD**: FISTA with restart.
- **nmAPGD**: the non-monotone APGD algorithm (Li and Lin, 2015).
- **PSGD**: the PSGD algorithm.
- **MB-PSGD**: the mini-batch PSGD algorithm (only on real-world datasets).
- **Katyusha**: the Katyusha algorithm proposed by Allen-Zhu (2017). It is similar to an accelerated proximal stochastic variance reduction gradient algorithm (Nitanda, 2014) but introduces an additional moment term, which is called Katyusha momentum.
- **MB-Katyusha**: the mini-batch Katyusha algorithm (only on real-world datasets).

We used line search techniques for the step size in **APGD** and **nmAPGD** (Beck and Teboulle, 2009; Li and Lin, 2015) but we did not use them in **PCD** and **PBCD** because their sub-problems are smooth. For the (mini-batch) PSGD-based algorithms, we ran the experiment using initial step size $\eta_0 = 1.0, 0.1$ and 0.01 but we show only the best results. In **PSGD** and **MB-PSGD**, rather than using a constant step size, we used a diminishing step size as proposed by Bottou (2012): $\eta = \eta_0(1 + \eta_0\lambda_p t)^{-1}$ at the t -th iteration. On the other hand, **Katyusha** and **MB-Katyusha** used a constant step size: $\eta = \eta_0$. In **MB-PSGD** and **MB-Katyusha**, we set the number of instances in one mini-batch N_b to be $Nd/\text{nnz}(\mathbf{X})$, which reduces the computational cost per epoch from $O(NdK)$ to $O(\text{nnz}(\mathbf{X})k)$. For the evaluation of the proximal operator (30), we used Algorithm 3, which runs in $O(d)$ time in expectation.

Datasets. We used feature interaction selection setting datasets, feature selection setting datasets, the ML100K dataset, and the a9a dataset. We set the number of instances of synthetic datasets 200 and 20,000 in order to compare scalabilities of algorithms w.r.t N . As in Appendix C.1, we used 100,000 and 48,642 instances for training on the ML100K and the a9a datasets, respectively. We ran the experiment ten times using different initial random seeds. On synthetic datasets, we set $\lambda_p = \tilde{\lambda}_p = 0.1$ for the batch algorithms (**PCD**, **PBCD**, **APGD**, and **nmAPGD**). For the stochastic algorithms (**PSGD** and **Katyusha**), we first scaled the feature vectors and targets: we used \mathbf{x}_n/\sqrt{d} and y_n/d as feature vectors and targets and set λ_p and $\tilde{\lambda}_p$ to $0.1/d^2$. Since we used the squared loss and $f_{\text{FM}}(\mathbf{x}/\sqrt{d}; \mathbf{0}, \mathbf{P}) = f_{\text{FM}}(\mathbf{x}; \mathbf{0}, \mathbf{P})/d$, the balance between the loss term and regularization term was the same as that in the batch algorithms. On the ML100K dataset, we set $\lambda_w = 0.5 \times 10^{-4}$, $\lambda_p = 0.5 \times 10^{-4}$, and $\tilde{\lambda}_p = 0.5 \times 10^{-4}$ for all methods. On the a9a dataset, we set $\lambda_w = 0.5 \times 10^{-2}$, $\lambda_p = 0.5 \times 10^{-3}$, and $\tilde{\lambda}_p = 0.5 \times 10^{-3}$ for all methods. They were chosen based on the results in Section 8.3.

Results: batch vs stochastic. As shown in Fig. 8, on synthetic datasets, when the number of training instances was 20,000, the stochastic algorithms (**PSGD** and **Katyusha**) were faster than the batch algorithms (**PCD**, **PBCD**, **APGD**, and **nmAPGD**). This indicates that stochastic algorithms can be more useful than batch algorithms for large-scale dense datasets, as described in Section 4.1 and Section 5.1. However, the stochastic algorithms take $O(dk)$ time at each iteration even if a sampled feature vector is sparse. On synthetic datasets, the feature vectors were dense since

they were generated from Gaussian distributions, so the stochastic algorithms might be relatively slower in some real-world applications. Indeed, as shown in Fig. 9, on real-world datasets, the completely stochastic algorithms (**PSGD** and **Katyusha**) were not faster than batch algorithms although although $N \gg d$ on both datasets. The use of appropriate size mini-batch improved the convergence speed as our expected: **MB-PSGD** and **MB-Katyusha** were faster than their completely stochastic versions. Nevertheless, such mini-batch algorithms were slower than **PCD** and **PBCD** on the ML100K dataset and the **PCD** on the a9a dataset. Moreover, performances of (mini-batch) stochastic algorithms were sensitive w.r.t the choice of the step size hyperparameter and the objective values usually diverged with $\eta_0 = 1.0$. Thus, as our analysis in Section 4.1, the PSGD-based algorithms should be used only when $\text{nnz}(\mathbf{X})/d$ is large. Note that $\text{nnz}(\mathbf{X})/d$ of the ML100K and the a9a datasets in this experiment are 338 and 5,506, respectively (clearly, on synthetic datasets $\text{nnz}(\mathbf{X})/d = N$).

Results: PCD/PBCD vs APGD/nmAPGD. On synthetic datasets (Fig. 8), **PCD/PBCD**, **APGD**, and **nmAPGD** tended to show similar results but **PCD** was much faster than **APGD** and **nmAPGD** on the feature interaction selecting dataset with $N = 200$. On real-world datasets (Fig. 9), **PCD** and **PBCD** were much faster than **APGD** and **nmAPGD**. Strictly speaking, not PCD-based algorithms but PGD/PSGD-based algorithms should be used since **TI (CS)** regularizer is not separable w.r.t each coordinate (each row vector) in \mathbf{P} . However, our results indicate that **PCD** and **PBCD** can work better than **nmAPGD** and **APGD** practically. Moreover, again, **PCD** and **PBCD** have some important practical advantages: (i) easy to implement, (ii) easy to extend to related models (as shown in Section 6), and (iii) having few hyperparameters.

C.3 Comparison of Algorithm 2 and Algorithm 3

We compared two algorithms for the proximal operator (30), Algorithm 2 (**Sort**) and Algorithm 3 (**Random**) proposed by Martins et al. (2011). For a d -dimensional vector, **Sort** and **Random** run in $O(d \log d)$ time and $O(d)$ time, respectively. We evaluated the runtime of two algorithms for a d -dimensional vector generated from a Gaussian distribution, $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_d)$ with varying $d \in \{2^3, 2^4, \dots, 2^{14}\}$. We set $\sigma = 1.0$ and $\sigma = 10.0$ and the regularization strength $\lambda = 0.001$, $\lambda = 0.1$, and $\lambda = 10.0$. We ran the experiment 100 times with different initial random seeds and report the average runtimes. For **Sort**, we used Nim’s standard **sort** procedure, which is an implementation of merge sort.

As shown in Fig. 10, **Random** outperformed **Sort** in most cases. **Sort** sometimes ran faster than **Random** only when d was very small ($d \leq 2^5 = 32$). Thus, we basically recommend **Random** (Algorithm 3) rather than **Sort** (Algorithm 2) for the evaluation of the proximal operator (30).


References

- Raj Agrawal, Brian Trippe, Jonathan Huggins, and Tamara Broderick. The kernel interaction trick: Fast Bayesian discovery of pairwise interactions in high dimensions. In *ICML*, pages 141–150, 2019.
- Zeyuan Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. *Journal of Machine Learning Research*, 18(1):8194–8244, 2017.
- Zeyuan Allen-Zhu. Natasha 2: Faster non-convex optimization than sgd. In *NeurIPS*, pages 2675–2686, 2018.
- Kyohei Atarashi, Satoshi Oyama, and Masahito Kurihara. Link prediction using higher-order feature combinations across objects. *IEICE Transactions on Information and Systems*, E103.D(8):1833–1842, 2020.


- Francis Bach, Rodolphe Jenatton, Julien Mairal, and Guillaume Obozinski. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 4(1):1–106, 2012.
- Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- Mathieu Blondel, Akinori Fujino, Naonori Ueda, and Masakazu Ishihata. Higher-order factorization machines. In *NeurIPS*, pages 3351–3359, 2016a.
- Mathieu Blondel, Masakazu Ishihata, Akinori Fujino, and Naonori Ueda. Polynomial networks and factorization machines: New insights and efficient training algorithms. In *ICML*, pages 850–858, 2016b.
- Léon Bottou. Stochastic gradient descent tricks. In *Neural Networks: Tricks of the Trade*, pages 421–436. Springer, 2012.
- Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
- Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):27, 2011.
- Xiaoshuang Chen, Yin Zheng, Jiaying Wang, Wenye Ma, and Junzhou Huang. Rafm: rank-aware factorization machines. In *International Conference on Machine Learning*, pages 1132–1140. PMLR, 2019a.
- Yifan Chen, Pengjie Ren, Yang Wang, and Maarten de Rijke. Bayesian personalized feature interaction selection for factorization machines. In *SIGIR*, pages 665–674, 2019b.
- Chen Cheng, Fen Xia, Tong Zhang, Irwin King, and Michael R Lyu. Gradient boosting factorization machines. In *RecSys*, pages 265–272, 2014.
- John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. Efficient projections onto the l_1 -ball for learning in high dimensions. In *ICML*, pages 272–279, 2008.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. A note on the group lasso and a sparse group lasso. *arXiv preprint arXiv:1001.0736*, 2010.
- Pontus Giselsson and Stephen Boyd. Monotonicity and restart in fast gradient methods. In *53rd IEEE Conference on Decision and Control*, pages 5058–5063. IEEE, 2014.
- Bin Gu, De Wang, Zhouyuan Huo, and Heng Huang. Inexact proximal gradient methods for non-convex and non-smooth optimization. In *AAAI*, 2018.
- Benjamin D Haeffele and René Vidal. Structured low-rank matrix factorization: Global optimality, algorithms, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(6):1468–1482, 2019.

- F Maxwell Harper and Joseph A Konstan. The movielens datasets: history and context. *ACM Transactions on Interactive Intelligent Systems*, 5(4):19, 2016.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2009.
- Fuxing Hong, Dongbo Huang, and Ge Chen. Interaction-aware factorization machines for recommender systems. In *AAAI*, pages 3804–3811, 2019.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. In *NeurIPS*, pages 3146–3154, 2017.
- Huan Li and Zhouchen Lin. Accelerated proximal gradient methods for nonconvex programming. In *NeurIPS*, pages 379–387, 2015.
- Xiao Li, Zhihui Zhu, Anthony Man-Cho So, and Rene Vidal. Nonconvex robust low-rank matrix recovery. *SIAM Journal on Optimization*, 30(1):660–686, 2020.
- Bo Liu, Xiao-Tong Yuan, Lezi Wang, Qingshan Liu, and Dimitris N Metaxas. Dual iterative hard thresholding: From non-convex sparse minimization to non-smooth concave maximization. In *ICML*, pages 2179–2187, 2017.
- Tie-Yan Liu. *Learning to Rank for Information Retrieval*. Springer, 2011.
- Stéphane G Mallat and Zhifeng Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, 1993.
- Andre Filipe Torres Martins, Noah Smith, Eric Xing, Pedro Aguiar, and Mario Figueiredo. Online learning of structured predictors with multiple kernels. In *ICML*, pages 507–515, 2011.
- Marine Le Morvan and Jean-Philippe Vert. WHInter: A working set algorithm for high-dimensional sparse second order interaction models. In *ICML*, pages 3635–3644, 2018.
- Yurii Nesterov. A method for unconstrained convex minimization problem with the rate of convergence $o(1/k^2)$. In *Doklady an ussr*, volume 269, pages 543–547, 1983.
- Yurii Nesterov et al. *Lectures on convex optimization*, volume 137. Springer, 2018.
- Atsushi Nitanda. Stochastic proximal gradient descent with acceleration techniques. In *NeurIPS*, pages 1574–1582, 2014.
- Zhen Pan, Enhong Chen, Qi Liu, Tong Xu, Haiping Ma, and Hongjie Lin. Sparse factorization machines for click-through rate prediction. In *ICDM*, pages 400–409, 2016.
- Neal Parikh and Stephen Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):127–239, 2014.
- Steffen Rendle. Factorization machines. In *ICDM*, pages 995–1000, 2010.
- Steffen Rendle. Factorization machines with libFM. *ACM Transactions on Intelligent Systems and Technology*, 3(3):57, 2012.
- Steffen Rendle, Zeno Gantner, Christoph Freudenthaler, and Lars Schmidt-Thieme. Fast context-aware recommendations with factorization machines. In *SIGIR*, pages 635–644, 2011.
- Emile Richard, Pierre-André Savalle, and Nicolas Vayatis. Estimation of simultaneously sparse and low rank matrices. In *ICML*, pages 51–58, 2012.


- Mark Schmidt, Nicolas Le Roux, and Francis Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. In *NeurIPS*, pages 1458–1466, 2011.
- John Shawe-Taylor and Nello Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- Weiping Song, Chence Shi, Zhiping Xiao, Zhijian Duan, Yewen Xu, Ming Zhang, and Jian Tang. AutoInt: Automatic feature interaction learning via self-attentive neural networks. In *CIKM*, pages 1161–1170, 2019.
- Shinya Suzumura, Kazuya Nakagawa, Yuta Umezu, Koji Tsuda, and Ichiro Takeuchi. Selective inference for sparse high-order interaction models. In *ICML*, pages 3338–3347, 2017.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Paul Tseng and Sangwoon Yun. A coordinate gradient descent method for nonsmooth separable minimization. *Mathematical Programming*, 117(1-2):387–423, 2009.
- Jun Xiao, Hao Ye, Xiangnan He, Hanwang Zhang, Fei Wu, and Tat-Seng Chua. Attentional factorization machines: Learning the weight of feature interactions via attention networks. In *IJCAI*, pages 3119–3125, 2017.
- Jianpeng Xu, Kaixiang Lin, Pang-Ning Tan, and Jiayu Zhou. Synergies that matter: Efficient interaction selection via sparse factorization machine. In *SDM*, pages 108–116, 2016.
- Niannan Xue, Bin Liu, Huifeng Guo, Ruiming Tang, Fengwei Zhou, Stefanos P Zafeiriou, Yuzhou Zhang, Jun Wang, and Zhenguo Li. Autohash: Learning higher-order feature interactions for deep ctr prediction. *IEEE Transactions on Knowledge and Data Engineering*, 2020.
- Shuo Yang, Yanyao Shen, and Sujay Sanghavi. Interaction hard thresholding: Consistent sparse quadratic regression in sub-quadratic time and space. In *NeurIPS*, pages 7926–7936, 2019.
- Quanming Yao, James T Kwok, Fei Gao, Wei Chen, and Tie-Yan Liu. Efficient inexact proximal gradient algorithm for nonconvex problems. In *IJCAI*, pages 3308–3314, 2017.
- Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- Huan Zhao, Quanming Yao, Jianda Li, Yangqiu Song, and Dik Lun Lee. Meta-graph based recommendation fusion over heterogeneous information networks. In *KDD*, pages 635–644, 2017.
- Dengyong Zhou, Sumit Basu, Yi Mao, and John Platt. Learning from the wisdom of crowds by minimax entropy. In *NeurIPS*, pages 2204–2212, 2012.



(a) Feature interaction selection setting: $d_{\text{true}} = 80$, $b = 8$ and $d_{\text{noise}} = 20$ with $N = 200$ (left) and 20,000 (right) for **TI** method.




(b) Feature selection setting: $d_{\text{true}} = 20$, $b = 1$ and $d_{\text{noise}} = 80$ with $N = 200$ (left) and 20,000 (right) for **TI** method.




(c) Feature selection setting: $d_{\text{true}} = 20$, $b = 1$ and $d_{\text{noise}} = 80$ with $N = 200$ (left) and 20,000 (right) for **CS** method.

Figure 8: Runtime comparisons among **PCD**, **PBCD**, **APGD**, **nmAPGD**, **PSGD**, and **Katyusha** algorithms on synthetic datasets using different amounts of training data: (a) feature interaction selection setting datasets for **TI** method; (b) feature selection setting datasets for **TI** method; (c) feature selection setting datasets for **CS** method. Left and right graphs show results for datasets with $N = 200$ and 20,000, respectively.



(a) ML100K dataset ($N = 100,000$): **TI** method (left) and **CS** method (right) with $\lambda_w = 5 \times 10^{-4}$, $\lambda_p = 5 \times 10^{-5}$, and $\tilde{\lambda}_p = 5 \times 10^{-5}$.



(b) a9a dataset ($N = 48,842$): **TI** method (left) and **CS** method (right) with $\lambda_w = 5 \times 10^{-2}$, $\lambda_p = 5 \times 10^{-4}$, and $\tilde{\lambda}_p = 5 \times 10^{-4}$.

Figure 9: Runtime comparisons among **PCD**, **PBCD**, **APGD**, **nmAPGD**, **PSGD**, **MB-PSGD**, **Katyusha**, and **MB-Katyusha** algorithms on (a) ML100K dataset and (b) a9a dataset. Left and right graphs show results for **TI** and **CS** method, respectively.

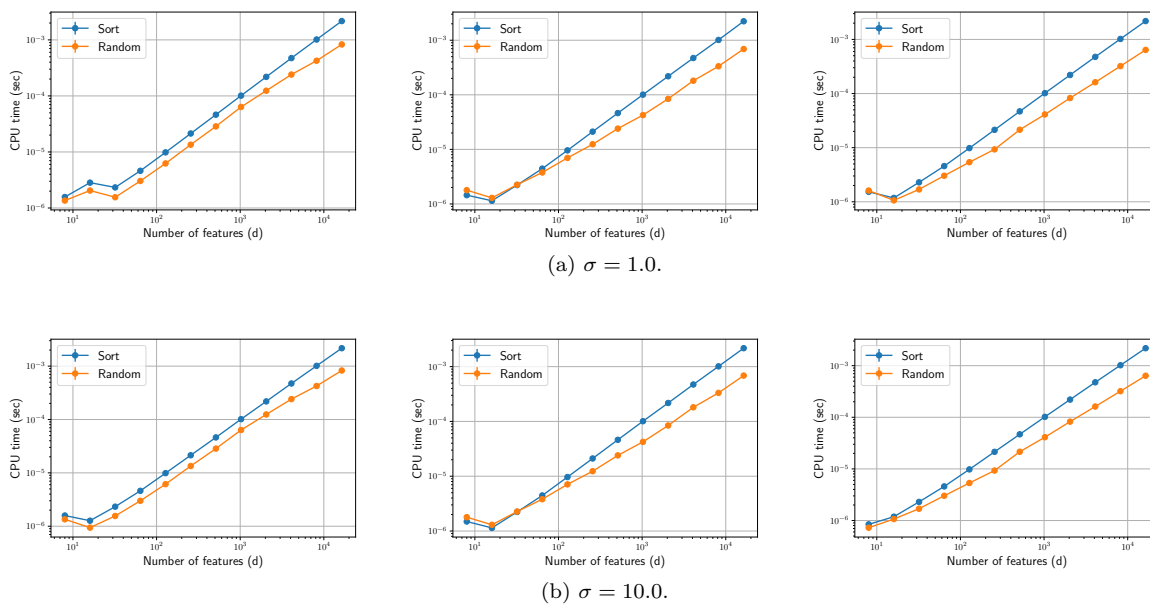


Figure 10: Comparison of Algorithm 2 (**Sort**) and Algorithm 3 (**Random**) with $\lambda = 0.001$ (left), $\lambda = 0.1$ (center), and $\lambda = 10.0$ (right).