

# Which Factorization Machine Modeling Is Better: A Theoretical Answer with Optimal Guarantee

Ming Lin,<sup>1</sup> Shuang Qiu,<sup>2</sup> Jieping Ye,<sup>2</sup> Xiaomin Song,<sup>1</sup> Qi Qian,<sup>1</sup> Liang Sun,<sup>1</sup>  
Shenghuo Zhu,<sup>1</sup> Rong Jin<sup>1</sup>

<sup>1</sup>Alibaba Group, Bellevue, USA

<sup>2</sup>Department of Computational Medicine and Bioinformatics, University of Michigan, Ann Arbor, USA  
{ming.l, xiaomin.song, qi.qian, liang.sun, shenghuo.zhu, jinrong.jr}@alibaba-inc.com, {qiush, jpye}@umich.edu

## Abstract

Factorization machine (FM) is a popular machine learning model to capture the second order feature interactions. The optimal learning guarantee of FM and its generalized version is not yet developed. For a rank  $k$  generalized FM of  $d$  dimensional input, the previous best known sampling complexity is  $\mathcal{O}[k^3 d \cdot \text{polylog}(kd)]$  under Gaussian distribution. This bound is sub-optimal comparing to the information theoretical lower bound  $\mathcal{O}(kd)$ . In this work, we aim to tighten this bound towards optimal and generalize the analysis to sub-gaussian distribution. We prove that when the input data satisfies the so-called  $\tau$ -Moment Invertible Property, the sampling complexity of generalized FM can be improved to  $\mathcal{O}[k^2 d \cdot \text{polylog}(kd)/\tau^2]$ . When the second order self-interaction terms are excluded in the generalized FM, the bound can be improved to the optimal  $\mathcal{O}[kd \cdot \text{polylog}(kd)]$  up to the logarithmic factors. Our analysis also suggests that the positive semi-definite constraint in the conventional FM is redundant as it does not improve the sampling complexity while making the model difficult to optimize. We evaluate our improved FM model in real-time high precision GPS signal calibration task to validate its superiority.

## 1 Introduction

Factorization machine (FM) (Rendle 2010; Bayer et al. 2017; Juan, Lefortier, and Chapelle 2017; Juan et al. 2016; Zhao et al. 2017; Yamada et al. 2017; Luo et al. 2018; Lin et al. 2018) is a popular linear regression model to capture the second order feature interactions. It has been found effective in various applications, including recommendation systems (Rendle 2010), CTR prediction (Juan et al. 2016), computational medicine (Lin et al. 2016), social network (Hong, Doumith, and Davison 2013) and so on. Intuitively speaking, the second order feature interactions consider the factors jointly affecting the output. On the theoretical side, FM is closely related to the symmetric matrix sensing (Kueng, Rauhut, and Terstiege 2017; Cai and Zhang 2015; Yurtsever et al. 2017) and phase retrieval (Candes et al. 2013). While the conventional FM only considers the second order feature interactions, it is possible to extend the conventional FM to the high order functional space which leads to the Polynomial Network model (Blondel et al. 2016). FM is the cornerstone in modern machine

learning research as it abridges the linear regression and high order polynomial regression. It is therefore important to understand the theoretical foundation of FM.

Given an instance  $\mathbf{x} \in \mathbb{R}^d$ , the conventional FM assumes that the label  $y \in \mathbb{R}$  of  $\mathbf{x}$  is generated by

$$y = \mathbf{x}^\top \mathbf{w}^* + \mathbf{x}^\top M^* \mathbf{x} \quad \text{rank}(M^*) \leq k \quad (1)$$

where  $\{\mathbf{w}^*, M^*\}$  are the first order and the second order coefficients respectively. In the original FM paper (Rendle 2010), the authors additionally assumed that  $M^*$  is generated from a low-rank positive semi-definite (PSD) matrix with all its diagonal elements subtracted. That is

$$M^* = U^* U^{*\top} - \text{diag}(U^* U^{*\top}). \quad (2)$$

Eq. (2) consists of two parts. We call the first part  $U^* U^{*\top}$  as the PSD constraint and the second part  $-\text{diag}(\cdot)$  as the diagonal-zero constraint. Our key question in this work is whether the FM model (1) can be learned by  $\mathcal{O}[kd \log(kd)]$  observations and how the two additional constraints help the generalization ability of FM.

Although the FM has been widely applied, there is little research exploring the theoretical properties of the FM to answer the above key question. A naive analysis directly following the sampling complexity of the linear model would suggest  $\mathcal{O}(d^2)$  samples to recover  $\{\mathbf{w}^*, M^*\}$  which is too loose. When  $\mathbf{w}^* = 0$  and  $M^*$  is symmetric, Eq. (1) is equal to the symmetric matrix sensing problem. (Cai and Zhang 2015) proved the sampling complexity of this special case on well-bounded sub-gaussian distribution using trace norm convex programming under the  $\ell_2/\ell_1$ -RIP condition. (Yurtsever et al. 2017) developed a conditional gradient descent solver to recover  $M^*$ . However, when  $\mathbf{w}^* \neq \mathbf{0}$  the above methods and the theoretical results are no longer applicable. (Blondel, Fujino, and Ueda 2015) considered a convexified formulation of FM. Their FM model requires solving a trace-norm regularized loss function which is computationally expensive. They did not provide statistical learning guarantees for the convexified FM.

To the best of our knowledge, the most recent research dealing with the theoretical properties of the FM is (Lin and Ye 2016). In their study, the authors argued that the two constraints proposed in (Rendle 2010) can be removed if  $\mathbf{x}$  is sampled from the standard Gaussian distribution. However, their analysis heavily relies on the rotation invariance

of the Gaussian distribution therefore cannot be generalized to non-Gaussian cases. Even limiting on the Gaussian distribution, the sampling complexity given by their analysis is  $\mathcal{O}[k^3 d \cdot \text{polylog}(kd)]$  which is worse than the information-theoretic lower bound  $\mathcal{O}(kd)$ . It is still an open question whether the FM can be learned with  $\mathcal{O}(kd)$  samples and whether both constraints in the original formulation are necessary to make the model learnable.

In this work, we answer the above questions affirmatively. We show that when the data is sampled from sub-gaussian distribution and satisfies the so-called  $\tau$ -Moment Invertible Property (MIP), the generalized FM (without constraints) can be learned by  $\mathcal{O}[k^2 d / \tau^2 \cdot \text{polylog}(kd)]$  samples. The PSD constraint is not necessary to achieve this sharp bound. Actually the PSD constraint is harmful as it introduces asymmetric bias on the value  $y$  (see Experiment section). The optimal sampling complexity  $\mathcal{O}[kd \cdot \text{polylog}(kd)]$  is achievable if we further constrain that the diagonal elements of  $M^*$  are zero. This is not an artificial constraint but there is information-theoretic limitation prevents us recovering the diagonal elements of  $M^*$  on sub-gaussian distribution. Finally inspired by our theoretical results, we propose an improved version of FM, called iFM, which removes the PSD constraint and inherits the diagonal-zero constraint from the conventional modeling. Unlike the generalized FM, the sampling complexity of the iFM does not depend on the MIP constant of the data distribution.

The remainder of this paper is organized as follows. We revisit the modeling of the FM in Section 2 and show that the conventional modeling is sub-optimal when considered in a more general framework. In Section 3 we present the learning guarantee of the generalized FM on sub-gaussian distribution. We propose the high order moment elimination technique to overcome a difficulty in our convergence analysis. Based on our theoretical results, we propose the improved model iFM. Section 4 conducts numerical experiments on synthetic and real-world datasets to validate the superiority of iFM over the conventional FM. Section 5 encloses this work.

## 2 A Revisit of Factorization Machine

In this section, we revisit the modeling design of the FM and its variants. We briefly review the original formulation of the conventional FM to raise several questions about its optimality. We then highlight previous studies trying to establish the theoretical foundation of the FM modeling. Based on the above survey, we motivate our study and present our main results in the next section.

In their original paper, (Rendle 2010) assumes that the feature interaction coefficients in the FM can be embedded in a  $k$ -dimensional latent space. That is,

$$y = \sum_{i=1}^d w_i x_i + \sum_{i=1}^d \sum_{j=i+1}^d \langle \mathbf{u}_i, \mathbf{u}_j \rangle x_i x_j \quad (3)$$

where  $\mathbf{u}_i$  is a  $k \times 1$  vector. The original formulation Eq. (3) is equivalent to Eq. (1) with constraint Eq. (2). While the low-rank assumption is standard in the matrix sensing literature, the PSD and the diagonal-zero constraints are not.

A critical question is whether the two additional constraints are necessary or removable. Indeed we have strong reasons to remove both constraints.

The reasons to remove the diagonal-zero constraint are straightforward. First there is no theoretical result so far to motivate this constraint. Secondly subtracting the diagonal elements will make the second order derivative w.r.t.  $U$  non-PSD. This will raise many technical difficulties in optimization and learning theory as many research works assume convexity in their analysis.

The PSD constraint in the original FM modeling is the second term we wish to remove. Let us temporally forget about the diagonal-zero constraint and focus on the PSD constraint only. Obviously relaxing  $UU^\top$  with  $UV^\top$  will make the model more flexible. A more serious problem of the PSD constraint is that it implicitly assumes that the label  $y$  is more likely to be “positive”. This will introduce asymmetric bias about the distribution of  $y$ . To see this, suppose  $\hat{M} = U^* U^{*\top} = \bar{U} \Sigma \bar{U}^\top$  where  $\bar{U}$  is the eigenvector matrix of  $\hat{M}$ . We call  $\bar{U}$  the second order feature mapping matrix induced by  $\hat{M}$  since  $\mathbf{x}^\top \hat{M} \mathbf{x} = (\bar{U}^\top \mathbf{x})^\top \Sigma (\bar{U}^\top \mathbf{x})$ . The eigenvalue matrix  $\Sigma$  is the weights for the mapped features  $\bar{U}^\top \mathbf{x}$ . As  $\hat{M}$  is constrained to be PSD, the weights of  $\bar{U}^\top \mathbf{x}$  cannot be negative. In other words, the PSD constraint prevents the model learning patterns from negative class. Please check the Experiment section for more concrete examples.

Another issue of the PSD constraint raised is the difficulty in optimization. Suppose we choose least square as the loss function in FM. By enforcing  $\hat{M} = UU^\top$ , the loss function is a fourth order polynomial of  $U$ . This makes the initialization of  $U$  difficult since the scale of the initial  $U^{(0)}$  will affect the convergence rate. Clearly we cannot initialize  $U^{(0)} = 0$  since the gradient w.r.t.  $U$  will be zero. On the other hand, we cannot initialize  $U^{(0)}$  to have a large norm otherwise the problem will be ill-conditioned. This is because the spectral norm of the second order derivative w.r.t.  $U$  will be proportional to  $\|U^{(0)}\|_2^2$  therefore the (local) condition number depends on  $\|U^{(0)}\|_2$ . In practice, it is usually difficult to figure out the optimal scale of  $\|U^{(0)}\|_2$  resulting vanishing or explosion gradient norms. If we decouple the  $U$  and  $V$ , we can initialize  $\|U^{(0)}\|_2 = 1$  and  $V^{(0)} = 0$ . Then by alternating gradient descent, the decoupled FM model is easy to optimize.

In summary, the theoretical foundation of the FM is still not well-developed. On one hand, it is unclear whether the conventional FM modeling is optimal and on the other hand, there is strong motivation to modify the conventional formulation based on heuristic intuition. This inspires our study of the optimal modeling of the FM driven by theoretical analysis which is presented in the next section.

## 3 Main Results

In this section, we present our main results on the theoretical guarantees of the FM and its improved version iFM. We first give a sharp complexity bound for the generalized FM on sub-gaussian distribution. We show that the recovery error of the diagonal elements of  $M^*$  depends on a so-called  $\tau$ -MIP

condition of the data distribution. The sampling complexity bound can be improved to optimal by the diagonal-zero constraint.

We introduce a few more notations needed in this section. Suppose  $\mathbf{x}$  is sampled from coordinate sub-gaussian with zero mean and unit variance. The element-wise third order moment of  $\mathbf{x}$  is denoted as  $\kappa^* \triangleq \mathbb{E}\mathbf{x}^3$  and the fourth order moment is  $\phi^* \triangleq \mathbb{E}\mathbf{x}^4$ . All training instances are sampled identically and independently (i.i.d.). Denote the feature matrix  $X = [\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}] \in \mathbb{R}^{d \times n}$  and the label vector  $\mathbf{y} = [y_1, \dots, y_n]^\top \in \mathbb{R}^n$ .  $\mathcal{D}(\cdot)$  denotes the diagonal function. For any two matrices  $A$  and  $B$ , we denote their Hadamard product as  $A \circ B$ . The element-wise squared matrix is defined by  $A^2 \triangleq A \circ A$ . For a non-negative real number  $\xi \geq 0$ , the symbol  $O(\xi)$  denotes some perturbation matrix whose spectral norm is upper bounded by  $\xi$ . The  $i$ -th largest singular value of matrix  $M$  is  $\sigma_i(M)$ . We abbreviate  $\sigma_i^* \triangleq \sigma_i(M^*)$ . To abbreviate our high probability bounds, given a probability  $\eta$ , we use the symbol  $C_\eta$  and  $c_\eta$  to denote some polynomial logarithmic factors in  $1/\eta$  and any other necessary variables that do not change the polynomial order of the upper bounds.

### 3.1 Limitation of The Generalized FM

In order to derive the theoretical optimal FM models, we begin with the most general formulation of FM, that is, with no constraint except low-rank:

$$\mathbf{y} = \mathbf{x}^\top \mathbf{w}^* + \mathbf{x}^\top M^* \mathbf{x} \quad \text{s.t. } M^* = U^* V^{*\top}. \quad (4)$$

Clearly  $M^*$  must be symmetric but for now this does not matter. Eq. (4) is called the generalized FM (Lin and Ye 2016). It is proved that when  $\mathbf{x}$  is sampled from the Gaussian distribution, Eq. (4) can be learned by  $\mathcal{O}(k^3 d)$  training samples. Although this bound is not optimal, (Lin and Ye 2016) showed the possibility to remove Eq. (2) on the Gaussian distribution. However, their result no longer holds true on non-Gaussian distributions. In the following, we will show that the learning guarantee for the generalized FM on sub-gaussian distribution is much more complex than the Gaussian one.

Our first important observation is that model (4) is not always learnable on all sub-gaussian distributions.

**Proposition 1.** *When  $\mathbf{x} \in \{-1, +1\}^d$ , the generalized FM is not learnable.*

The above observation is easy to verify since  $\mathbf{x}^\top M^* \mathbf{x} = \text{tr}(M^*)$  when  $\mathbf{x} \in \{-1, +1\}^d$ . Therefore at least the diagonal elements of  $M^*$  cannot be recovered at all. Proposition 1 shows that there is information-theoretic limitation to learn the generalized FM on sub-gaussian distribution. In our analysis, we find that such limitation is related to a property of the data distribution which we call the *Moment Invertible Property* (MIP).

**Definition 2** (Moment Invertible Property). A zero mean unit variance sub-gaussian distribution  $\mathbb{P}(\mathbf{x})$  is called  $\tau$ -Moment Invertible if  $|\phi - 1 - \kappa^2| \geq \tau$  for some constants  $\tau \geq 0$ ,  $\phi \triangleq \mathbb{E}\mathbf{x}^4$ ,  $\kappa \triangleq \mathbb{E}\mathbf{x}^3$ .

With the MIP condition, the following theorem shows that the generalized FM is learnable via alternating gradient descent.

**Theorem 3.** *Suppose  $\mathbf{x}$  is sampled from a  $\tau$ -MIP sub-gaussian distribution with  $\mathbf{y}$  generated by Eq. (4). Then with probability at least  $1 - \eta$ , there is an alternating gradient descent based method which can achieve the recovery error after  $t$  iteration such that*

$$\|\mathbf{w}^{(t)} - \mathbf{w}^*\|_2 + \|M^{(t)} - M^*\|_2 \leq [(2\sqrt{5}\sigma_1^*/\sigma_k^* + 2)\delta]^t (\|\mathbf{w}^*\|_2 + \|M^*\|_2),$$

provided

$$\begin{aligned} n &\geq \frac{C_\eta}{\delta^2} (p+1)^2 \max\{p\tau^{-2}, (\sqrt{k} + |\text{tr}(M)|/\|M\|_2)^2 d\} \\ p &\triangleq \max\{1, \|\kappa^*\|_\infty, \|\phi^* - 3\|_\infty, \|\phi^* - 1\|_\infty\} \\ \delta &\leq \min\left\{\frac{1}{2\sqrt{5}\sigma_1^*/\sigma_k^* + 2}, \frac{1}{2\sqrt{5}}\sigma_k^*[\|\mathbf{w}^*\|_2 + \|M^*\|_2]^{-1}\right\}. \end{aligned}$$

Theorem 3 is our key result. In Theorem 3, we measure the quality of our estimation by the recovery error

$$\epsilon_t \triangleq \|\mathbf{w}^{(t)} - \mathbf{w}^*\|_2 + \|M^{(t)} - M^*\|_2.$$

The recovery error decreases linearly along the steps of alternating iteration with rate  $\delta \approx \mathcal{O}(1/\sqrt{n})$ . The sampling complexity is on order of  $\max\{\mathcal{O}(k^2 d), \mathcal{O}(1/\tau^2)\}$ . This bound delivers two messages. First when the distribution is close to the Gaussian distribution,  $\tau \approx 2$  and the bound is controlled by  $\mathcal{O}(k^2 d)$ . This result improves the previous  $\mathcal{O}(k^3 d)$  given by (Lin and Ye 2016) for the Gaussian distribution. Secondly when  $\tau$  is small, the sampling complexity is proportional to  $\mathcal{O}(1/\tau^2)$ . The sampling complexity will even trend to infinite when the data follows the binary Bernoulli distribution where  $\tau = 0$ . Therefore the  $\tau$ -MIP condition provides a sufficient condition to make the generalized FM learnable.

We have not given any detail about the alternating gradient descent algorithm mentioned in Theorem 3. We find that it is difficult to prove the convergence rate following the conventional alternating gradient descent framework. To address this difficulty, we use a high order moment elimination technique in the next subsection in the convergent analysis.

### 3.2 Alternating Gradient Descent with High Order Moment Elimination

In this subsection, we will construct an alternating gradient descent algorithm which achieves the convergence rate and the sampling complexity in Theorem 3. We first show that the conventional alternating gradient descent cannot be applied directly to prove Theorem 3. Then a high order moment elimination technique is proposed to overcome the difficulty.

The generalized FM defined in Eq. (4) can be written in the matrix form

$$\mathbf{y} = X^\top \mathbf{w}^* + \mathcal{A}(M^*) \quad (5)$$

where the operator  $\mathcal{A}(\cdot) : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^d$  is defined by  $\mathcal{A}(M) \triangleq [\mathbf{x}^{(1)\top} M \mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)\top} M \mathbf{x}^{(n)}]$ . The adjoint

---

**Algorithm 1** Alternating Gradient Descent with High Order Moment Elimination

---

**Require:** The mini-batch size  $n$ ; number of total update  $T$ ; training instances  $X^{(t)} \triangleq [\mathbf{x}^{(t,1)}, \mathbf{x}^{(t,2)}, \dots, \mathbf{x}^{(t,n)}]$ ,  $\mathbf{y}^{(t)} \triangleq [y^{(t,1)}, y^{(t,2)}, \dots, y^{(t,n)}]^\top$ ; rank  $k \geq 1$ .

**Ensure:**  $\mathbf{w}^{(T)}, U^{(T)}, V^{(T)}, M^{(t)} \triangleq U^{(t)}V^{(t)\top}$ .

- 1: Retrieve  $n$  training instances to estimate the third and fourth order moments  $\kappa$  and  $\phi$ .
  - 2: Compute  $G$  and  $H$  in Eq. (8) and (9).
  - 3: Initialize  $\mathbf{w}^{(0)} = \mathbf{0}$ ,  $V^{(0)} = \mathbf{0}$ .  $\bar{U}^{(0)} = \text{SVD}(\mathcal{M}^{(0)}(\mathbf{y}^{(0)}), k)$ , that is the top- $k$  singular vectors.
  - 4: **for**  $t = 1, 2, \dots, T$  **do**
  - 5: Retrieve  $n$  training instances  $X^{(t)}, \mathbf{y}^{(t)}$ , compute  $\mathcal{M}^{(t)}$  in Eq. (7) and update  $U^{(t)}$  and  $V^{(t)}$  as:
$$\hat{\mathbf{y}}^{(t)} = X^{(t)\top} \mathbf{w}^{(t-1)} + \mathcal{A}^{(t)}(U^{(t-1)}V^{(t-1)\top})$$

$$U^{(t)} = V^{(t-1)} - \mathcal{M}^{(t)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)})\bar{U}^{(t-1)}$$

$$\{\bar{U}^{(t)}, R^{(t)}\} = \text{QR}(U^{(t)})$$

$$V^{(t)} = V^{(t-1)}\bar{U}^{(t-1)\top}\bar{U}^{(t)} - \mathcal{M}^{(t)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)})\bar{U}^{(t)}$$
  - 6: Compute  $\mathcal{W}^{(t)}$  in Eq. (10) and update  $\mathbf{w}^{(t)} = \mathbf{w}^{(t-1)} - \mathcal{W}^{(t)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)})$ .
  - 7: **end for**
  - 8: **Output:**  $\mathbf{w}^{(T)}, \bar{U}^{(T)}, V^{(T)}$ .
- 

operator of  $\mathcal{A}$  is  $\mathcal{A}'$ . To recover  $\{\mathbf{w}^*, M^*\}$ , we minimize the square loss function

$$\min_{\mathbf{w}, U, V} \mathcal{L}(\mathbf{w}, U, V) \triangleq \frac{1}{2n} \|X^\top \mathbf{w} + \mathcal{A}(UV^\top) - \mathbf{y}\|^2. \quad (6)$$

A straightforward idea to prove Theorem 3 is to show that the alternating gradient descent will converge. However, we find that this is difficult in our problem. To see this, let us compute the expected gradient of  $\mathcal{L}(\mathbf{w}^{(t)}, U^{(t)}, V^{(t)})$  with respect to  $V^{(t)}$  at step  $t$ .

$$\mathbb{E} \nabla_V \mathcal{L}(\mathbf{w}^{(t)}, U^{(t)}, V^{(t)}) = 2(M^{(t)} - M^*)U^{(t)} + F^{(t)}U^{(t)}$$

where

$$F^{(t)} \triangleq \text{tr}(M^{(t)} - M^*)I + \mathcal{D}(\phi - 3)\mathcal{D}(M^{(t)} - M^*) + \mathcal{D}(\kappa)\mathcal{D}(\mathbf{w}^{(t)} - \mathbf{w}^*).$$

In previous studies, one expects  $\mathbb{E} \nabla \mathcal{L} \approx I$ . However, this is no longer the case in our problem. Clearly  $\|\frac{1}{2}\mathbb{E} \nabla \mathcal{L} - I\|_2$  is dominated by  $\|\kappa\|_\infty$  and  $\|\phi - 3\|_\infty$ . For non-Gaussian distributions, these two perturbation terms could dominate the gradient norm. Similarly the gradient of  $\mathbf{w}$  is biased by  $O(\|\kappa\|_\infty)$ .

The difficulty to follow the conventional gradient descent analysis inspires us to look for a new convergence analysis technique. The perturbation term  $F^{(t)}$  consists of high order moments of the sub-gaussian variable  $\mathbf{x}$ . It might be possible to construct a sequence of another high order moments to eliminate these perturbation terms. We call this idea the high order moment elimination method. The next question

is whether the desired moments exist and how to construct them efficiently. Unfortunately, this is impossible in general. A sufficient condition to ensure the existence of the elimination sequence is that the data distribution satisfies the  $\tau$ -MIP condition.

To construction an elimination sequence, for any  $\mathbf{z} \in \mathbb{R}^n$  and  $M \in \mathbb{R}^{d \times d}$ , define functions

$$\begin{aligned} \mathcal{P}^{(t,0)}(\mathbf{z}) &\triangleq \mathbf{1}^\top \mathbf{z}/n, & \mathcal{P}^{(t,1)}(\mathbf{z}) &\triangleq X^{(t)} \mathbf{z}/n \\ \mathcal{P}^{(t,2)}(\mathbf{z}) &\triangleq (X^{(t)})^2 \mathbf{z}/n - \mathcal{P}^{(t,0)}(\mathbf{z}) \\ \mathcal{A}^{(t)}(M) &\triangleq \mathcal{D}(X^{(t)\top} M X^{(t)}) \\ \mathcal{H}^{(t)}(\mathbf{z}) &\triangleq \mathcal{A}^{(t)\top} \mathcal{A}^{(t)}(\mathbf{z})/(2n). \end{aligned}$$

Notice when  $n \rightarrow \infty$ ,

$$\begin{aligned} \mathcal{P}^{(t,0)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}) &\approx \text{tr}(M^{(t)} - M^*) \\ \mathcal{P}^{(t,1)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}) &\approx D(M^{(t)} - M^*)\kappa + \mathbf{w}^{(t)} - \mathbf{w}^* \\ \mathcal{P}^{(t,2)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}) &\approx D(M^{(t)} - M^*)(\phi - 1) \\ &\quad + D(\kappa)(\mathbf{w}^{(t)} - \mathbf{w}^*). \end{aligned}$$

This inspires us to find a linear combination of  $\mathcal{P}^{(t,\cdot)}$  to eliminate  $F^{(t)}$ . The solution for this linear combination equation is

$$\begin{aligned} \mathcal{M}^{(t)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}) &\triangleq \mathcal{H}^{(t)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}) \\ &\quad - \frac{1}{2} \mathcal{D}\left(G_1 \circ \mathcal{P}^{(t,1)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)})\right) \\ &\quad - \frac{1}{2} \mathcal{D}\left(G_2 \circ \mathcal{P}^{(t,2)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)})\right) \end{aligned} \quad (7)$$

where

$$G_{j,:}^\top = \begin{bmatrix} 1 & \kappa_j \\ \kappa_j & \phi_j - 1 \end{bmatrix}^{-1} \begin{bmatrix} \kappa_j \\ \phi_j - 3 \end{bmatrix} \quad (8)$$

$$H_{j,:}^\top = \begin{bmatrix} 1 & \kappa_j \\ \kappa_j & \phi_j - 1 \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \quad (9)$$

Similarly to eliminate the high order moments in the gradient of  $\mathbf{w}^{(t)}$ , we construct

$$\begin{aligned} \mathcal{W}^{(t)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}) &\triangleq H_1 \circ \mathcal{P}^{(t,1)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}) \\ &\quad + H_2 \circ \mathcal{P}^{(t,2)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}). \end{aligned} \quad (10)$$

The overall construction is given in Algorithm 1.

We briefly prove that the construction in Algorithm 1 will eliminate the high order moments in  $F^{(t)}$  by which a global linear convergence rate is immediately followed. Please check appendix for details. We will omit the superscript  $(t)$  in  $X^{(t)}$  and  $\mathcal{P}^{(t,\cdot)}$  when not raising confusion.

First we show that  $\frac{1}{n}\mathcal{A}'\mathcal{A}$  is conditionally independent re-strict isometric after shifting its expectation (Shift CI-RIP). The proof can be found in appendix.

**Theorem 4** (Shift CI-RIP). *Suppose  $d \geq (2 + \|\phi^* - 3\|_\infty)^2$ . Fixed a rank- $k$  matrix  $M$ , with probability at least  $1 - \eta$ ,*

$$\begin{aligned} \frac{1}{n} \mathcal{A}' \mathcal{A}(M) &= 2M + \text{tr}(M)I + \mathcal{D}(\phi^* - 3)\mathcal{D}(M) \\ &\quad + O(\delta \|M\|_2) \end{aligned}$$

*provided  $n \geq c_\eta(\sqrt{k} + |\text{tr}(M)|)^2 d / \delta^2$ .*

Theorem 4 is the main theorem in our analysis. The key ingredient of our proof is to apply the matrix Bernstein's inequality with an improved version of sub-gaussian Hanson-Wright inequality proved by (Rudelson and Vershynin 2013). Please check appendix for more details.

Based on the shifted CI-RIP condition of operator  $\mathcal{A}$ , we prove the following perturbation bounds.

**Lemma 5.** *For  $n \geq C_\eta(\sqrt{k} + |\text{tr}(M)|)^2 d / \delta^2$ , with probability at least  $1 - \eta$ ,*

$$\begin{aligned} \frac{1}{n} \mathcal{A}'(X^\top \mathbf{w}) &= \mathcal{D}(\boldsymbol{\kappa}^*) \mathbf{w} + O(\delta \|\mathbf{w}\|_2) \\ \mathcal{P}^{(0)}(\mathbf{y}) &\triangleq \frac{1}{n} \mathbf{1}^\top \mathbf{y} = \text{tr}(M) + O[\delta(\|\mathbf{w}\|_2 + \|M\|_2)] \\ \mathcal{P}^{(1)}(\mathbf{y}) &\triangleq \frac{1}{n} X \mathbf{y} = \mathcal{D}(M) \boldsymbol{\kappa}^* + \mathbf{w} + O[\delta(\|\mathbf{w}\|_2 + \|M\|_2)] \\ \mathcal{P}^{(2)}(\mathbf{y}) &\triangleq \frac{1}{n} X^2 \mathbf{y} - \mathcal{P}^{(0)}(\mathbf{y}) = \mathcal{D}(M)(\boldsymbol{\phi}^* - 1) \\ &\quad + \mathcal{D}(\boldsymbol{\kappa}^*) \mathbf{w} + O[\delta(\|\mathbf{w}\|_2 + \|M\|_2)]. \end{aligned}$$

Lemma 5 shows that  $\mathcal{A}'X^\top$  and  $\mathcal{P}^{(t,\cdot)}$  are all concentrated around their expectations with no more than  $O(C_\eta k^2 d)$  samples. To finish our construction, we need to bound the deviation of  $G$  and  $H$  from their expectation  $G^*$  and  $H^*$ . This is done in the following lemma.

**Lemma 6.** *Suppose the distribution of  $\mathbf{x}$  is  $\tau$ -MIP with  $\tau > 0$ . Then in Algorithm 1,*

$$\|G - G^*\|_\infty \leq \delta, \quad \|H - H^*\|_\infty \leq \delta,$$

provided

$$n \geq C_\eta(1 + \tau^{-1} \sqrt{\|\boldsymbol{\kappa}^*\|_\infty^2 + \|\boldsymbol{\phi}^* - 3\|_\infty^2}) / (\tau \delta^2).$$

Lemma 6 shows that  $G \approx G^*$  as long as  $n \geq O(1/\tau^2)$ . The matrix inversion in the definition of  $G$  requires that the  $\tau$ -MIP condition must be satisfied with  $\tau > 0$ .

We are now ready to show that  $\mathcal{M}^{(t)}$  and  $\mathcal{W}^{(t)}$  are almost isometric.

**Lemma 7.** *Under the same settings of Theorem 3, with probability at least  $1 - \eta$ ,*

$$\begin{aligned} \mathcal{M}^{(t)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}) &= M^{(t-1)} - M^* + O(\delta \epsilon_{t-1}) \\ \mathcal{W}^{(t)}(\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}) &= \mathbf{w}^{(t-1)} - \mathbf{w}^* + O(\delta \epsilon_{t-1}) \end{aligned}$$

provided

$$n \geq C_\eta(p+1)^2 / \delta^2 \max\{p/\tau^2, (\sqrt{k} + |\text{tr}(M)|/\|M\|_2)^2 d\}$$

where  $p \triangleq \max\{1, \|\boldsymbol{\kappa}^*\|_\infty, \|\boldsymbol{\phi}^* - 3\|_\infty, \|\boldsymbol{\phi}^* - 1\|_\infty\}$ .

Lemma 7 shows that  $\mathcal{M}^{(t)}$  and  $\mathcal{W}^{(t)}$  are almost isometric when the number of samples is larger than  $O(k^2 d)$  and  $O(1/\tau^2)$ . The proof of Lemma 7 consists of two steps. First we replace each operator or matrix with its expectation plus a small perturbation given in Lemma 5 and Lemma 6. Then Lemma 7 follows after simplification. Theorem 3 is obtained by combining Lemma 7 with alternating gradient descent analysis. Please check appendix for the complete proof.

### 3.3 Improved Factorization Machine

Theorem 3 shows that learning the generalized FM is hard on non-gaussian distribution. Especially, when the data distribution has a very small  $\tau$ -MIP constant, the sampling complexity to recover  $M^*$  in Eq. (4) will be as large as  $O(1/\tau^2)$ . The recovery is even impossible on  $\tau = 0$  distributions such as the Bernoulli distribution. Clearly, a well-defined learnable model should not depend on the  $\tau$ -MIP condition.

Indeed the bound given by Theorem 3 is quite sharp. It explains well why we cannot recover  $M^*$  on the Bernoulli distribution. Therefore it is unlikely to remove the  $\tau$  dependency by designing a better elimination sequence in Algorithm 1. After examining the proof of Theorem 1 carefully, we find that the only reason our bound contains  $\tau$  is that the diagonal elements of  $M^*$  are allowed to be non-zero. If we constrain  $\mathcal{D}(M^*) = 0$  and  $\mathcal{D}(M^{(t)}) = 0$ , the  $F^{(t)}$  in the expected gradient  $\mathbb{E} \nabla_V \mathcal{L}(\mathbf{w}^{(t)}, U^{(t)}, V^{(t)})$  will be zero and then we do not need to eliminate it during the alternating iteration. This greatly simplifies our convergence analysis as we only need Theorem 4 which now becomes

$$\frac{1}{n} \mathcal{A}' \mathcal{A}(M) = 2M + O(\delta \|M\|_2). \quad (11)$$

Eq. (4) already shows that  $\frac{1}{n} \mathcal{A}' \mathcal{A}$  is almost isometric that immediately implies the linear convergence rate of alternating gradient descent. As a direct corollary of Theorem 4, the sampling complexity could be improved to  $O(c_\eta k d)$  which is optimal up to some logarithmic constants  $c_\eta$ . Inspired by these observations, we propose to learn the following FM model

$$\begin{aligned} y &= \mathbf{x}^\top \mathbf{w}^* + \mathbf{x}^\top M^* \mathbf{x} \\ \text{s.t. } M^* &= U^* V^{*\top} - \mathcal{D}(U^* V^{*\top}). \end{aligned} \quad (12)$$

We called the above model the Improved Factorization Machine (iFM). The iFM model is a trade-off between the conventional FM model and the generalized FM model. It decouples the PSD constraint with  $U \neq V$  in the generalized FM model but keeps the diagonal-zero constraint as the conventional FM model. Unlike the conventional FM model, the iFM model is proposed in a theoretical-driven way. The decoupling of  $\{U, V\}$  makes the iFM easy to optimize while the diagonal-zero constraint makes it learnable with the optimal  $O(kd)$  sampling complexity. In the next section, we will verify the above discussion with numerical experiments.

## 4 Experiments

We first use synthetic data in subsection 4.1 to show the modeling power of iFM and the PSD bias of the conventional FM. In subsection 4.2 we apply iFM in a real-word problem, the vTEC estimation task, to demonstrate its superiority over baseline methods.

### 4.1 Synthetic Data

In this subsection, we construct numerical examples to support our theoretical results. To this end, we choose  $d = 100$ ,  $k = 5$ .  $\{\mathbf{w}^*, U^*, V^*, \mathbf{x}\}$  are all sampled from the Gaussian




Figure 1: RMSE Curve of iFM v.s. FM

distribution with variance  $1/d$ . We randomly sample  $30kd$  instances as training set and 10000 instances as testing set. In Figure 1, we report the convergence curve of iFM and FM on the testing set averaged over 10 trials. The x-axis is the iteration step and the y-axis is the Root Mean Square Error (RMSE) of  $y$ . In Figure (a), we generate label  $y$  following the conventional FM assumption. Both iFM and FM converge well. In Figure (b), we flip the sign of the label in (a). While iFM still converges well, FM cannot model the sign flip. This example shows why we should avoid to use the conventional FM both in theory and in practice: even a simple flipping operation can make the model under-fit the data. In Figure (c), we generate  $y$  from  $M^*$  with both positive eigenvalues and negative eigenvalues. Again the conventional FM cannot fit this data as the distribution of  $y$  is now symmetric in both direction.

## 4.2 vTEC Estimation

In this subsection, we demonstrate the superiority of iFM in a real-world application, the *vertical Total Electron Content* (vTEC) estimation. The vTEC is an important descriptive parameter of the ionosphere of the Earth. It integrates the total number of electrons integrated when traveling from space to earth with a perpendicular path. One important application is the real-time high precision GPS signal calibration. The accuracy of the GPS system heavily depends on the pseudo-range measurement between the satellite and the receiver. The major error in the pseudo-range measurement is caused by the vTEC which is dynamic.

In order to estimate the vTEC, we build a triangle mesh grid system in an anonymous region. Each node in the grid is a ground stations equipped with dual-frequency high precision GPS receiver. The distance between two nodes is around 75 kilometers. The station sends out resolved GPS data every second. Formally, our system solves an online regression problem. Our system receives  $N_t$  data points at every time step  $t$  measured in seconds. Each data point  $\mathbf{x}^{(t,i)} \in \mathbb{R}^4$  presents an ionospheric pierce point of a satellite-receiver pair. The first two dimensions are the latitude  $\alpha^{(t,i)}$  and the longitude  $\beta^{(t,i)}$  of the pierce point. The third dimension and the fourth dimension are the zenith angle  $\theta^{(t,i)}$  and the azimuth angle  $\gamma^{(t,i)}$  respectively. We will omit the superscript  $t$  below as we always stay within the

same time window  $t$ . In order to build a localized prediction model, we encode  $\{\alpha^{(i)}, \beta^{(i)}\}$  into a high dimensional vector. Suppose we have  $m$  satellites in total. First we collect data points for 60 seconds. Then we cluster the collected  $\{\alpha^{(i)}, \beta^{(i)}\}$  into  $m$  clusters via K-means algorithm. Denote the cluster center of K-means as  $\{\mathbf{c}^{(1)}, \mathbf{c}^{(2)}, \dots, \mathbf{c}^{(m)}\}$  and the  $i$ -th data point belongs to the  $g_i$ -th cluster. The first two dimensions of the  $i$ -th data point  $\{\alpha^{(i)}, \beta^{(i)}\}$  are then encoded as  $\mathbf{v}^{(i)} \in \mathbb{R}^{2m}$  where

$$\mathbf{v}_j^{(i)} = \begin{cases} \{\alpha^{(i)}, \beta^{(i)}\} - \mathbf{c}^{(g_i)} & j = g_i \\ 0 & \text{otherwise} \end{cases}$$

Finally, each data point  $\mathbf{x}^{(i)}$  is encoded into an  $\mathbb{R}^{2m+8}$  vector:

$$\text{Enc}(\mathbf{x}^{(i)}) \triangleq [\mathbf{v}^{(i)}, \sin(\theta^{(i)}), \cos(\theta^{(i)}), \theta^{(i)}, (\theta^{(i)})^2, \sin(\gamma^{(i)}), \cos(\gamma^{(i)}), \gamma^{(i)}, (\gamma^{(i)})^2].$$

**Evaluation** It is important to note that our problem does not fit the conventional machine learning framework. We only have validation set for model selection and evaluation set to evaluate the model performance. We introduce the training station and the testing station that correspond to the “training set” and “testing set” in the conventional machine learning framework. However please be advised that they are not exactly the same concepts.

To evaluate the performance of our models, we randomly select one ground station as the testing station. Around the testing station, we choose 16 ground stations as training stations to learn the online prediction model. Suppose the online prediction model is  $\mathcal{F}$  which maps  $\text{Enc}(\mathbf{x}^{(i)})$  to the corresponding vTEC value

$$\mathcal{F} : \text{Enc}(\mathbf{x}^{(i)}) \rightarrow \text{vTEC}(\mathbf{x}^{(i)}) \in \mathbb{R}.$$

In our system, we are only given the double difference of the vTEC values due to the signal resolving process. Suppose two satellites  $a, b$  and two ground stations  $c, d$  are connected in the data-link graph,  $\mathbf{x}^{(a,c)}$  denotes the ionospheric pierce point between  $a$  and  $c$ . The observed double difference  $y^{(a,b,c,d)}$  is given by

$$y^{(a,b,c,d)} \triangleq \text{vTEC}(\mathbf{x}^{(a,c)}) - \text{vTEC}(\mathbf{x}^{(a,d)}) - \text{vTEC}(\mathbf{x}^{(b,c)}) + \text{vTEC}(\mathbf{x}^{(b,d)}).$$

Table 1: RMSE &amp; RTK of iFM v.s. Baseline Methods

	Ridge	LASSO	ElasticNet	Kernel	FM	iFM
TestDay1	<b>0.02161</b>	0.02260	0.02222	0.02570	0.02178	0.02164
TestDay2	0.01430	0.01461	0.01454	0.01683	0.01439	<b>0.01404</b>
TestDay3	0.01508	0.01524	0.01496	0.01875	<b>0.01484</b>	<b>0.01484</b>
TestDay4	0.01449	0.01487	0.01460	0.01564	0.01432	<b>0.01423</b>
TestDay5	0.01610	0.01579	0.01612	0.01744	0.01606	<b>0.01567</b>
TestDay6	0.01487	0.01491	0.01483	0.01684	0.01470	<b>0.01459</b>
TestDay7	0.01828	0.01849	0.01827	0.02209	0.01830	<b>0.01805</b>
TestDay8	0.01461	0.01552	0.01519	0.01629	0.01466	<b>0.01453</b>
TestDay9	0.01657	0.01646	0.01639	0.02096	0.01646	<b>0.01636</b>
Average	0.01621	0.01650	0.01635	0.01895	0.01617	<b>0.01599</b>
RTK	62.67%	62.18%	63.07%	52.09%	63.10%	<b>64.44%</b>

Since  $v\text{TEC}(\cdot)$  is an unknown function, we need to approximate it by  $\mathcal{F}$ . Once  $\mathcal{F}$  is learned from training stations, we can apply it to predict the double difference  $y^{(a,b,c,d)}$  where either  $c$  or  $d$  is the testing station.

Once we get the  $v\text{TEC}$  estimation, we use it to calibrate the GPS signal and finally compute the geometric coordinate of the user. The RTK ratio measures the quality of the positioning service. It is a real number presenting the probability of successful positioning with accuracy at least one centimeter. The RTK ratio is computed from a commercial system that is much slower than the computation of RMSE.

**Dataset and Results** We select a ground station at the region center as testing station. Around the testing station 16 stations are selected as training stations. We collect 5 consecutive days’ data as validation set for parameter tuning. The following 9 days’ data are used as evaluation set. We update the prediction model per 60 seconds. The learned model is then used to predict the double differences relating to the testing station. We compare the predicted double differences to the true values detected by the testing station. The number of valid satellites in our experiment is around 10 to 20.

In Table 1, we report the root-mean-square error (RMSE) over 9 days period. The dates are denoted as TestDay1 to TestDay9 for anonymity. Five baseline methods are evaluated: Ridge Regression, LASSO, ElasticNet, Kernel Ridge Regression (Kernel) with RBF kernel and the conventional Factorization Machine (FM). More computational expensive models such as deep neural network are not feasible for our online system. For Ridge, LASSO, Kernel and ElasticNet, their parameters are tuned from  $1 \times 10^{-6}$  to  $1 \times 10^6$ . The regularizer parameters of FM and iFM are tuned from  $1 \times 10^{-6}$  to  $1 \times 10^6$ . The rank of  $M$  is tuned in set  $\{1, 2, \dots, 10\}$ . We use Scikit-learn (Pedregosa et al. 2011) and fastFM (Bayer 2016) to implement the baseline methods.

In Table 1, we observe that iFM is uniformly better than the baseline methods. We average the root squared error over  $9 \times 24 \times 60 = 12960$  minutes in the last second row. The 95% confidence interval is within  $1 \times 10^{-5}$  in our experiment. In our experiment, the optimal rank of FM is 2 and the optimal rank of iFM is 6. We note that FM is better than the first order linear models since it captures the second order information.

This indicates that the second order information is indeed helpful.

In the last row of Table 1, we report the RTK ratio averaged over the 9 days. We find that the RTK ratio will improve a lot even with small improvement of  $v\text{TEC}$  estimation. This is because the error of  $v\text{TEC}$  estimation will be broadcasted and magnified in the RTK computation pipeline. The RTK ratio of iFM is about 1.77% better than that of Ridge regression and is more than 12% better than Kernel regression. Comparing to FM, it is 1.34% better. We conclude that iFM achieves overall better performance and the improvement is statistically significant.

## 5 Conclusion

We study the learning guarantees of the FM solved by alternating gradient descent on sub-gaussian distributions. We find that the conventional modeling of the factorization machine might be sub-optimal in capturing negative second order patterns. We prove that the constraints in the conventional FM can be removed resulting a generalized FM model learnable by  $\max\{\mathcal{O}(k^2d), \mathcal{O}(1/\tau^2)\}$  samples. The sampling complexity can be improved to the optimal  $\mathcal{O}(kd)$  with diagonal-zero constraint. Our theoretical analysis shows that the optimal modeling of high order linear model does not always agree with the heuristic intuition. We hope this work could inspire future researches of non-convex high order machines with solid theoretical foundation.

## 6 Acknowledgments

This research is supported in part by NSF (III-1539991). The high precision GPS dataset is provided by Qianxun Spatial Intelligence Inc. China. We appreciate Dr. Wotao Yin from University of California Los Angeles and anonymous reviewers for their insightful comments.

## References

Bayer, I.; He, X.; Kanagal, B.; and Rendle, S. 2017. A generic coordinate descent framework for learning from implicit feedback. In *Proceedings of the 26th International Conference on World Wide Web, WWW ’17*, 1341–1350.

- Bayer, I. 2016. fastFM: A Library for Factorization Machines. *Journal of Machine Learning Research* 17(184):1–5.
- Blondel, M.; Ishihata, M.; Fujino, A.; and Ueda, N. 2016. Polynomial Networks and Factorization Machines: New Insights and Efficient Training Algorithms. In *Proceedings of The 33rd International Conference on Machine Learning*, 850–858.
- Blondel, M.; Fujino, A.; and Ueda, N. 2015. Convex Factorization Machines. In *Machine Learning and Knowledge Discovery in Databases*, number 9285 in Lecture Notes in Computer Science.
- Cai, T. T., and Zhang, A. 2015. ROP: Matrix recovery via rank-one projections. *The Annals of Statistics* 43(1):102–138.
- Candes, E.; Eldar, Y.; Strohmer, T.; and Voroninski, V. 2013. Phase Retrieval via Matrix Completion. *SIAM Journal on Imaging Sciences* 6(1):199–225.
- Hong, L.; Doumith, A. S.; and Davison, B. D. 2013. Co-factorization Machines: Modeling User Interests and Predicting Individual Decisions in Twitter. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, 557–566.
- Juan, Y.; Zhuang, Y.; Chin, W.-S.; and Lin, C.-J. 2016. Field-aware Factorization Machines for CTR Prediction. In *Proceedings of the 10th ACM Conference on Recommender Systems*, 43–50.
- Juan, Y.; Lefortier, D.; and Chapelle, O. 2017. Field-aware factorization machines in a real-world online advertising system. In *Proceedings of the 26th International Conference on World Wide Web Companion*, WWW ’17 Companion, 680–688.
- Kueng, R.; Rauhut, H.; and Terstiege, U. 2017. Low rank matrix recovery from rank one measurements. *Applied and Computational Harmonic Analysis* 42(1):88–116.
- Lin, M., and Ye, J. 2016. A non-convex one-pass framework for generalized factorization machine and rank-one matrix sensing. In *Advances in Neural Information Processing Systems*, 1633–1641.
- Lin, K.; Xu, J.; Baytas, I. M.; Ji, S.; and Zhou, J. 2016. Multi-Task Feature Interaction Learning. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1735–1744.
- Lin, X.; Zhang, W.; Zhang, M.; Zhu, W.; Pei, J.; Zhao, P.; and Huang, J. 2018. Online compact convexified factorization machine. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, 1633–1642.
- Luo, L.; Zhang, W.; Zhang, Z.; Zhu, W.; Zhang, T.; and Pei, J. 2018. Sketched follow-the-regularized-leader for online factorization machine. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1900–1909.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12:2825–2830.
- Rendle, S. 2010. Factorization machines. In *IEEE 10th International Conference On Data Mining*, 995–1000.
- Roman Vershynin. 2017. *High-Dimensional Probability: An Introduction with Applications in Data Science*.
- Rudelson, M., and Vershynin, R. 2013. Hanson-Wright inequality and sub-gaussian concentration. *Electronic Communications in Probability* 18.
- Stewart, G. W., and Sun, J.-g. 1990. *Matrix Perturbation Theory*. Boston: Academic Press, 1 edition edition.
- Tao, T. 2012. *Topics in Random Matrix Theory*, volume 132. Amer Mathematical Society.
- Yamada, M.; Lian, W.; Goyal, A.; Chen, J.; Wimalawarne, K.; Khan, S. A.; Kaski, S.; Mamitsuka, H.; and Chang, Y. 2017. Convex factorization machine for toxicogenomics prediction. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1215–1224.
- Yurtsever, A.; Udell, M.; Tropp, J. A.; and Cevher, V. 2017. Sketchy Decisions: Convex Low-Rank Matrix Optimization with Optimal Storage. *arXiv:1702.06838 [math, stat]*.
- Zhao, H.; Yao, Q.; Li, J.; Song, Y.; and Lee, D. L. 2017. Meta-graph based recommendation fusion over heterogeneous information networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’17, 635–644.