
Polynomial Networks and Factorization Machines: New Insights and Efficient Training Algorithms

Mathieu Blondel
Masakazu Ishihata
Akinori Fujino
Naonori Ueda

MATHIEU.BLONDEL@LAB.NTT.CO.JP
ISHIHATA.MASAKAZU@LAB.NTT.CO.JP
FUJINO.AKINORI@LAB.NTT.CO.JP
UEDA.NAONORI@LAB.NTT.CO.JP

NTT Communication Science Laboratories, 2-4 Hikaridai Seika-cho Soraku-gun, Kyoto 619-0237, Japan

Abstract

Polynomial networks and factorization machines are two recently-proposed models that can efficiently use feature interactions in classification and regression tasks. In this paper, we revisit both models from a unified perspective. Based on this new view, we study the properties of both models and propose new efficient training algorithms. Key to our approach is to cast parameter learning as a low-rank symmetric tensor estimation problem, which we solve by multi-convex optimization. We demonstrate our approach on regression and recommender system tasks.

1. Introduction

Interactions between features play an important role in many classification and regression tasks. One of the simplest approach to leverage such interactions consists in *explicitly* augmenting feature vectors with products of features (monomials), as in polynomial regression. Although fast linear model solvers can be used (Chang et al., 2010; Sonnenburg & Franc, 2010), an obvious drawback of this kind of approach is that the number of parameters to estimate scales as $O(d^m)$, where d is the number of features and m is the order of interactions considered. As a result, it is usually limited to second or third-order interactions.

Another popular approach consists in using a polynomial kernel so as to *implicitly* map the data via the kernel trick. The main advantage of this approach is that the number of parameters to estimate in the model is actually independent of d and m . However, the cost of storing and evaluating the model is now proportional to the number of training instances. This is sometimes called the curse of kernelization (Wang et al., 2010). Common ways to address the issue in-

clude the Nyström method (Williams & Seeger, 2001), random features (Kar & Karnick, 2012) and sketching (Pham & Pagh, 2013; Avron et al., 2014).

In this paper, in order to leverage feature interactions in possibly very high-dimensional data, we consider models which predict the output $y \in \mathbb{R}$ associated with an input vector $\mathbf{x} \in \mathbb{R}^d$ by

$$\hat{y}_{\mathcal{K}}(\mathbf{x}; \boldsymbol{\lambda}, \mathbf{P}) := \sum_{s=1}^k \lambda_s \mathcal{K}(\mathbf{p}_s, \mathbf{x}), \quad (1)$$

where $\boldsymbol{\lambda} = [\lambda_1, \dots, \lambda_k]^T \in \mathbb{R}^k$, $\mathbf{P} = [\mathbf{p}_1, \dots, \mathbf{p}_k] \in \mathbb{R}^{d \times k}$, \mathcal{K} is a kernel and k is a hyper-parameter. More specifically, we focus on two specific choices of \mathcal{K} which allow us to use feature interactions: the homogeneous polynomial and the ANOVA kernels. Our contributions are as follows. We show (Section 3) that choosing one kernel or the other allows us to recover **polynomial networks** (PNs) (Livni et al., 2014) and, surprisingly, **factorization machines** (FMs) (Rendle, 2010; 2012). Based on this new view, we show important properties of PNs and FMs. Notably, we show for the first time that the objective function of arbitrary-order FMs is multi-convex (Section 4). Unfortunately, the objective function of PNs is not multi-convex. To remedy this problem, we propose a *lifted* approach, based on casting parameter estimation as a low-rank tensor estimation problem (Section 5.1). Combined with a symmetrization trick, this approach leads to a multi-convex problem, for *both* PNs and FMs (Section 5.2). We demonstrate our approach on regression and recommender system tasks.

Notation. We denote vectors, matrices and tensors using lower-case, upper-case and calligraphic bold, e.g., \mathbf{w} , \mathbf{W} and \mathcal{W} . We denote the set of $\overbrace{d \times \dots \times d}^{m \text{ times}}$ real tensors by \mathbb{R}^{d^m} and the set of symmetric real tensors by \mathbb{S}^{d^m} . We use $\langle \cdot, \cdot \rangle$ to denote vector, matrix and tensor inner product. Given \mathbf{x} , we define a symmetric rank-one tensor by

$\mathbf{x}^{\otimes m} := \mathbf{x} \otimes \dots \otimes \mathbf{x} \in \mathbb{S}^{d^m}$, where $(\mathbf{x}^{\otimes m})_{j_1, j_2, \dots, j_m} = x_{j_1} x_{j_2} \dots x_{j_m}$. We use $[d]$ to denote the set $\{1, \dots, d\}$.

2. Related work

2.1. Polynomial networks

Polynomial networks (PNs) (Livni et al., 2014) of degree $m = 2$ predict the output $y \in \mathbb{R}$ associated with $\mathbf{x} \in \mathbb{R}^d$ by

$$\hat{y}_{\text{PN}}(\mathbf{x}; \mathbf{w}, \boldsymbol{\lambda}, \mathbf{P}) := \langle \mathbf{w}, \mathbf{x} \rangle + \langle \sigma(\mathbf{P}^\top \mathbf{x}), \boldsymbol{\lambda} \rangle, \quad (2)$$

where $\mathbf{w} \in \mathbb{R}^d$, $\mathbf{P} \in \mathbb{R}^{d \times k}$, $\boldsymbol{\lambda} \in \mathbb{R}^k$ and $\sigma(u) := u^2$ is evaluated element-wise. Intuitively, the right-hand term can be interpreted as a feedforward neural network with one hidden layer of k units and with activation function $\sigma(u)$. Livni et al. (2014) also extend (2) to the case $m = 3$ and show theoretically that PNs can approximate feedforward networks with sigmoidal activation. A similar model was independently shown to perform well on dependency parsing (Chen & Manning, 2014). Unfortunately, the objective function of PNs is non-convex. In Section 5, we derive a *multi-convex* objective based on low-rank symmetric tensor estimation, suitable for training arbitrary-order PNs.

2.2. Factorization machines

One of the simplest way to leverage feature interactions is polynomial regression (PR). For example, for second-order interactions, in this approach, we compute predictions by

$$\hat{y}_{\text{PR}}(\mathbf{x}; \mathbf{w}, \mathbf{W}) := \langle \mathbf{w}, \mathbf{x} \rangle + \sum_{j' > j} \mathbf{W}_{j, j'} x_j x_{j'},$$

where $\mathbf{w} \in \mathbb{R}^d$ and $\mathbf{W} \in \mathbb{R}^{d^2}$. Obviously, model size in PR does not scale well w.r.t. d . The main idea of (second-order) factorization machines (FMs) (Rendle, 2010; 2012) is to replace \mathbf{W} with a factorized matrix $\mathbf{P}\mathbf{P}^\top$:

$$\hat{y}_{\text{FM}}(\mathbf{x}; \mathbf{w}, \mathbf{P}) := \langle \mathbf{w}, \mathbf{x} \rangle + \sum_{j' > j} (\mathbf{P}\mathbf{P}^\top)_{j j'} x_j x_{j'},$$

where $\mathbf{P} \in \mathbb{R}^{d \times k}$. FMs have been increasingly popular for efficiently modeling feature interactions in high-dimensional data, see (Rendle, 2012) and references therein. In Section 4, we show for the first time that the objective function of arbitrary-order FMs is *multi-convex*.

3. Polynomial and ANOVA kernels

In this section, we show that the prediction functions used by polynomial networks and factorization machines can be written using (1) for a specific choice of kernel.

The **polynomial kernel** is a popular kernel for using combinations of features. The kernel is defined as

$$\mathcal{P}_\gamma^m(\mathbf{p}, \mathbf{x}) := (\gamma + \langle \mathbf{p}, \mathbf{x} \rangle)^m,$$

where $m \in \mathbb{N}$ is the degree and $\gamma > 0$ is a hyper-parameter. We define the **homogeneous polynomial kernel** by

$$\mathcal{H}^m(\mathbf{p}, \mathbf{x}) := \mathcal{P}_0^m(\mathbf{p}, \mathbf{x}) = \langle \mathbf{p}, \mathbf{x} \rangle^m.$$

Let $\mathbf{p} = [p_1, \dots, p_d]^\top$ and $\mathbf{x} = [x_1, \dots, x_d]^\top$. Then,

$$\mathcal{H}^m(\mathbf{p}, \mathbf{x}) = \sum_{j_1=1}^d \dots \sum_{j_m=1}^d p_{j_1} x_{j_1} \dots p_{j_m} x_{j_m}.$$

We thus see that \mathcal{H}^m uses *all* monomials of degree m (i.e., all combinations of features *with* replacement).

A much lesser known kernel is the **ANOVA kernel** (Stitson et al., 1997; Vapnik, 1998). Following (Shawe-Taylor & Cristianini, 2004, Section 9.2), the ANOVA kernel of degree m , where $2 \leq m \leq d$, can be defined as

$$\mathcal{A}^m(\mathbf{p}, \mathbf{x}) := \sum_{j_m > \dots > j_1} p_{j_1} x_{j_1} \dots p_{j_m} x_{j_m}. \quad (3)$$

As a result, \mathcal{A}^m uses only monomials composed of *distinct* features (i.e., feature combinations *without* replacement). For later convenience, we also define $\mathcal{A}^0(\mathbf{p}, \mathbf{x}) := 1$ and $\mathcal{A}^1(\mathbf{p}, \mathbf{x}) := \langle \mathbf{p}, \mathbf{x} \rangle$.

With \mathcal{H}^m and \mathcal{A}^m defined, we are now in position to state the following lemma.

Lemma 1 Expressing PNs and FMs using kernels

Let $\hat{y}_{\mathcal{K}}(\mathbf{x}; \boldsymbol{\lambda}, \mathbf{P})$ be defined as in (1). Then,

$$\begin{aligned} \hat{y}_{\text{PN}}(\mathbf{x}; \mathbf{w}, \boldsymbol{\lambda}, \mathbf{P}) &= \langle \mathbf{w}, \mathbf{x} \rangle + \hat{y}_{\mathcal{H}^2}(\mathbf{x}; \boldsymbol{\lambda}, \mathbf{P}) \\ \hat{y}_{\text{FM}}(\mathbf{x}; \mathbf{w}, \mathbf{P}) &= \langle \mathbf{w}, \mathbf{x} \rangle + \hat{y}_{\mathcal{A}^2}(\mathbf{x}; \mathbf{1}, \mathbf{P}). \end{aligned}$$

The relation easily extends to higher orders. This new view allows us to state results that will be very useful in the next sections. The first one is that \mathcal{H}^m and \mathcal{A}^m are homogeneous functions, i.e., they satisfy

$$\lambda^m \mathcal{K}(\mathbf{p}, \mathbf{x}) = \mathcal{K}(\lambda \mathbf{p}, \mathbf{x}) \quad \forall \lambda \in \mathbb{R}, \forall m \in \mathbb{N}_+.$$

Another key property of $\mathcal{A}^m(\mathbf{p}, \mathbf{x})$ is multi-linearity.¹

Lemma 2 Multi-linearity of $\mathcal{A}^m(\mathbf{p}, \mathbf{x})$ w.r.t. p_1, \dots, p_d

Let $\mathbf{p}, \mathbf{x} \in \mathbb{R}^d$, $j \in [d]$ and $1 \leq m \leq d$. Then,

$$\mathcal{A}^m(\mathbf{p}, \mathbf{x}) = \mathcal{A}^m(\mathbf{p}_{-j}, \mathbf{x}_{-j}) + p_j x_j \mathcal{A}^{m-1}(\mathbf{p}_{-j}, \mathbf{x}_{-j})$$

where \mathbf{p}_{-j} denotes the $(d-1)$ -dimensional vector with p_j removed and similarly for \mathbf{x}_{-j} .

That is, everything else kept fixed, $\mathcal{A}^m(\mathbf{p}, \mathbf{x})$ is an affine function of p_j , $\forall j \in [d]$. Proof is given in Appendix B.1.

¹A function $f(\theta_1, \dots, \theta_k)$ is called multi-linear (resp. multi-convex) if it is linear (resp. convex) w.r.t. $\theta_1, \dots, \theta_k$ separately.

Assuming \mathbf{p} is dense and \mathbf{x} sparse, the cost of naively computing $\mathcal{A}^m(\mathbf{p}, \mathbf{x})$ by (3) is $O(n_z(\mathbf{x})^m)$, where $n_z(\mathbf{x})$ is the number of non-zero features in \mathbf{x} . To address this issue, we will make use of the following lemma for computing \mathcal{A}^m in nearly $O(mn_z(\mathbf{x}))$ time when $m \in \{2, 3\}$.

Lemma 3 *Efficient computation of ANOVA kernel*

$$\begin{aligned} \mathcal{A}^2(\mathbf{p}, \mathbf{x}) &= \frac{1}{2} [\mathcal{H}^2(\mathbf{p}, \mathbf{x}) - \mathcal{D}^2(\mathbf{p}, \mathbf{x})] \\ \mathcal{A}^3(\mathbf{p}, \mathbf{x}) &= \frac{1}{6} [\mathcal{H}^3(\mathbf{p}, \mathbf{x}) - 3\mathcal{D}^{2,1}(\mathbf{p}, \mathbf{x}) + 2\mathcal{D}^3(\mathbf{p}, \mathbf{x})] \end{aligned}$$

where we defined $\mathcal{D}^m(\mathbf{p}, \mathbf{x}) := \sum_{j=1}^d (p_j x_j)^m$ and $\mathcal{D}^{m,n}(\mathbf{p}, \mathbf{x}) := \mathcal{D}^m(\mathbf{p}, \mathbf{x})\mathcal{D}^n(\mathbf{p}, \mathbf{x})$.

See Appendix B.2 for a derivation.

4. Direct approach

Let us denote the training set by $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$ and $\mathbf{y} = [y_1, \dots, y_n]^T \in \mathbb{R}^n$. The most natural approach to learn models of the form (1) is to directly choose λ and \mathbf{P} so as to minimize some error function

$$D_{\mathcal{K}}(\lambda, \mathbf{P}) := \sum_{i=1}^n \ell(y_i, \hat{y}_{\mathcal{K}}(\mathbf{x}_i; \lambda, \mathbf{P})), \quad (4)$$

where $\ell(y_i, \hat{y}_i)$ is a convex loss function. Note that (4) is a convex objective w.r.t. λ regardless of \mathcal{K} . However, it is in general non-convex w.r.t. \mathbf{P} . Fortunately, when $\mathcal{K} = \mathcal{A}^m$, we can show that (4) is multi-convex.

Theorem 1 *Multi-convexity of (4) when $\mathcal{K} = \mathcal{A}^m$*

$D_{\mathcal{A}^m}$ is convex in λ and in each row of \mathbf{P} separately.

Proof is given in Appendix B.3. As a corollary, the objective function of FMs of arbitrary order is thus multi-convex. Theorem 1 suggests that we can minimize (4) efficiently when $\mathcal{K} = \mathcal{A}^m$ by solving a succession of convex problems w.r.t. λ and the rows of \mathbf{P} . We next show that when m is odd, we can just fix $\lambda = \mathbf{1}$ without loss of generality.

Lemma 4 *When is it useful to fit λ ?*

Let $\mathcal{K} = \mathcal{H}^m$ or \mathcal{A}^m . Then

$$\min_{\lambda \in \mathbb{R}^k, \mathbf{P} \in \mathbb{R}^{d \times k}} D_{\mathcal{K}}(\lambda, \mathbf{P}) \leq \min_{\mathbf{P} \in \mathbb{R}^{d \times k}} D_{\mathcal{K}}(\mathbf{1}, \mathbf{P}) \quad \text{if } m \text{ is even}$$

$$\min_{\lambda \in \mathbb{R}^k, \mathbf{P} \in \mathbb{R}^{d \times k}} D_{\mathcal{K}}(\lambda, \mathbf{P}) = \min_{\mathbf{P} \in \mathbb{R}^{d \times k}} D_{\mathcal{K}}(\mathbf{1}, \mathbf{P}) \quad \text{if } m \text{ is odd.}$$

The result stems from the fact that \mathcal{H}^m and \mathcal{A}^m are homogeneous functions. If we define $\mathbf{v} := \text{sign}(\lambda) \sqrt[|m|]{|\lambda|} \mathbf{p}$, then we obtain $\lambda \mathcal{H}^m(\mathbf{p}, \mathbf{x}) = \mathcal{H}^m(\mathbf{v}, \mathbf{x}) \forall \lambda$ if m is odd, and similarly for \mathcal{A}^m . That is, λ can be absorbed into \mathbf{v} without loss of generality. When m is even, $\lambda < 0$ cannot be absorbed unless we allow complex numbers. Because FMs fix $\lambda = \mathbf{1}$, Lemma 4 shows that the class of functions that FMs can represent is possibly smaller than our framework.

5. Lifted approach

5.1. Conversion to low-rank tensor estimation problem

If we set $\mathcal{K} = \mathcal{H}^m$ in (4), the resulting optimization problem is *neither convex nor* multi-convex w.r.t. \mathbf{P} . In (Blondel et al., 2015), for $m = 2$, it was proposed to cast parameter estimation as a low-rank symmetric matrix estimation problem. A similar idea was used in the context of phase retrieval in (Candès et al., 2013). Inspired by these works, we propose to convert the problem of estimating λ and \mathbf{P} to that of estimating a low-rank *symmetric* tensor $\mathcal{W} \in \mathbb{S}^{d^m}$. Combined with a symmetrization trick, this approach leads to an objective that is multi-convex, for *both* $\mathcal{K} = \mathcal{A}^m$ and $\mathcal{K} = \mathcal{H}^m$ (Section 5.2).

We begin by rewriting the kernel definitions using rank-one tensors. For $\mathcal{H}^m(\mathbf{p}, \mathbf{x})$, it is easy to see that

$$\mathcal{H}^m(\mathbf{p}, \mathbf{x}) = \langle \mathbf{p}^{\otimes m}, \mathbf{x}^{\otimes m} \rangle. \quad (5)$$

For $\mathcal{A}^m(\mathbf{p}, \mathbf{x})$, we need to ignore irrelevant monomials. For convenience, we introduce the following notation:

$$\langle \mathcal{W}, \mathcal{X} \rangle_{>} := \sum_{j_m > \dots > j_1} \mathcal{W}_{j_1, \dots, j_m} \mathcal{X}_{j_1, \dots, j_m} \quad \forall \mathcal{W}, \mathcal{X} \in \mathbb{S}^{d^m}.$$

We can now concisely rewrite the ANOVA kernel as

$$\mathcal{A}^m(\mathbf{p}, \mathbf{x}) = \langle \mathbf{p}^{\otimes m}, \mathbf{x}^{\otimes m} \rangle_{>}. \quad (6)$$

Our key insight is described in the following lemma.

Lemma 5 *Link between tensors and kernel expansions*

Let $\mathcal{W} \in \mathbb{S}^{d^m}$ have a symmetric outer product decomposition (Comon et al., 2008)

$$\mathcal{W} = \sum_{s=1}^k \lambda_s \mathbf{p}_s^{\otimes m}. \quad (7)$$

Let $\lambda = [\lambda_1, \dots, \lambda_k]^T$ and $\mathbf{P} = [\mathbf{p}_1, \dots, \mathbf{p}_k]$. Then,

$$\langle \mathcal{W}, \mathbf{x}^{\otimes m} \rangle = \hat{y}_{\mathcal{H}^m}(\mathbf{x}; \lambda, \mathbf{P}) \quad (8)$$

$$\langle \mathcal{W}, \mathbf{x}^{\otimes m} \rangle_{>} = \hat{y}_{\mathcal{A}^m}(\mathbf{x}; \lambda, \mathbf{P}). \quad (9)$$

The result follows immediately from (5) and (6), and from the linearity of $\langle \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle_{>}$. Given $\mathcal{W} \in \mathbb{S}^{d^m}$, let us define the following objective functions

$$\begin{aligned} L_{\mathcal{H}^m}(\mathcal{W}) &:= \sum_{i=1}^n \ell(y_i, \langle \mathcal{W}, \mathbf{x}_i^{\otimes m} \rangle) \\ L_{\mathcal{A}^m}(\mathcal{W}) &:= \sum_{i=1}^n \ell(y_i, \langle \mathcal{W}, \mathbf{x}_i^{\otimes m} \rangle_{>}). \end{aligned}$$

If \mathcal{W} is decomposed as in (7), then from Lemma 5, we obtain $L_{\mathcal{K}}(\mathcal{W}) = D_{\mathcal{K}}(\lambda, \mathbf{P})$ for $\mathcal{K} = \mathcal{H}^m$ or \mathcal{A}^m . This

suggests that we can convert the problem of learning λ and \mathbf{P} to that of learning a symmetric tensor \mathcal{W} of (symmetric) rank k . Thus, the problem of finding a small number of bases $\mathbf{p}_1, \dots, \mathbf{p}_k$ and their associated weights $\lambda_1, \dots, \lambda_k$ is converted to that of learning a *low-rank* symmetric tensor. Following (Candès et al., 2013), we call this approach *lifted*. Intuitively, we can think of \mathcal{W} as a tensor that contains the weights for predicting y of monomials of degree m . For instance, when $m = 3$, $\mathcal{W}_{i,j,k}$ is the weight corresponding to the monomial $x_i x_j x_k$.

5.2. Multi-convex formulation

Estimating a low-rank symmetric tensor $\mathcal{W} \in \mathbb{S}^{d^m}$ for arbitrary integer $m \geq 2$ is in itself a difficult non-convex problem. Nevertheless, based on a symmetrization trick, we can convert the problem to a multi-convex one, which we can easily minimize by alternating minimization. We first present our approach for the case $m = 2$ to give intuitions then explain how to extend it to $m \geq 3$.

Intuition with the second-order case. For the case $m = 2$, we need to estimate a low-rank symmetric matrix $\mathbf{W} \in \mathbb{S}^{d^2}$. Naively parameterizing $\mathbf{W} = \mathbf{P} \text{diag}(\lambda) \mathbf{P}^\top$ and solving for λ and \mathbf{P} does not lead to a multi-convex formulation for the case $\mathcal{K} = \mathcal{H}^2$. This is due to the fact that $\langle \mathbf{P} \text{diag}(\lambda) \mathbf{P}^\top, \mathbf{x}^{\otimes 2} \rangle$ is *quadratic* in \mathbf{P} . Our key idea is to parametrize $\mathbf{W} = \mathcal{S}(\mathbf{U}\mathbf{V}^\top)$ where $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{d \times r}$ and $\mathcal{S}(\mathbf{M}) := \frac{1}{2}(\mathbf{M} + \mathbf{M}^\top) \in \mathbb{S}^{d^2}$ is the *symmetrization* of $\mathbf{M} \in \mathbb{R}^{d^2}$. We then minimize $L_{\mathcal{K}}(\mathcal{S}(\mathbf{U}\mathbf{V}^\top))$ w.r.t. \mathbf{U}, \mathbf{V} .

The main advantage is that both $\langle \mathcal{S}(\mathbf{U}\mathbf{V}^\top), \cdot \rangle$ and $\langle \mathcal{S}(\mathbf{U}\mathbf{V}^\top), \cdot \rangle_{\mathcal{K}}$ are *bi-linear* in \mathbf{U} and \mathbf{V} . This implies that $L_{\mathcal{K}}(\mathcal{S}(\mathbf{U}\mathbf{V}^\top))$ is *bi-convex* in \mathbf{U} and \mathbf{V} and can therefore be efficiently minimized by alternating minimization. Once we obtained $\mathbf{W} = \mathcal{S}(\mathbf{U}\mathbf{V}^\top)$, we can *optionally* compute its eigendecomposition $\mathbf{W} = \mathbf{P} \text{diag}(\lambda) \mathbf{P}^\top$, with $k = \text{rank}(\mathbf{W})$ and $r \leq k \leq 2r$, then apply (8) or (9) to obtain the model in kernel expansion form.

Extension to higher-order case. For $m \geq 3$, we now estimate a low-rank symmetric tensor $\mathcal{W} = \mathcal{S}(\mathcal{M}) \in \mathbb{S}^{d^m}$, where $\mathcal{M} \in \mathbb{R}^{d^m}$ and $\mathcal{S}(\mathcal{M})$ is the symmetrization of \mathcal{M} (cf. Appendix A.2). We decompose \mathcal{M} using m matrices of size $d \times r$. Let us call these matrices $\{\mathbf{U}^t\}_{t=1}^m$ and their columns $\mathbf{u}_s^t = [u_{1s}^t, \dots, u_{ds}^t]^\top$. Then the decomposition of \mathcal{M} can be expressed as a sum of rank-one tensors

$$\mathcal{M} = \sum_{s=1}^r \mathbf{u}_s^1 \otimes \dots \otimes \mathbf{u}_s^m. \quad (10)$$

Due to multi-linearity of (10) w.r.t. $\mathbf{U}^1, \dots, \mathbf{U}^m$, the objective function $L_{\mathcal{K}}$ is *multi-convex* in $\mathbf{U}^1, \dots, \mathbf{U}^m$.

Computing predictions efficiently. When $\mathcal{K} = \mathcal{H}^m$, predictions are computed by $\langle \mathcal{W}, \mathbf{x}^{\otimes m} \rangle$. To compute them efficiently, we use the following lemma.

Lemma 6 *Symmetrization does not affect inner product*

$$\langle \mathcal{S}(\mathcal{M}), \mathcal{X} \rangle = \langle \mathcal{M}, \mathcal{X} \rangle \quad \forall \mathcal{M} \in \mathbb{R}^{d^m}, \mathcal{X} \in \mathbb{S}^{d^m}, m \geq 2. \quad (11)$$

Proof is given in Appendix A.2. Using $\mathbf{x}^{\otimes m} \in \mathbb{S}^{d^m}$, $\mathcal{W} = \mathcal{S}(\mathcal{M})$ and (10), we then obtain

$$\langle \mathcal{W}, \mathbf{x}^{\otimes m} \rangle = \langle \mathcal{M}, \mathbf{x}^{\otimes m} \rangle = \sum_{s=1}^r \prod_{t=1}^m \langle \mathbf{u}_s^t, \mathbf{x} \rangle.$$

As a result, we never need to explicitly compute the symmetrized tensor. For the case $\mathcal{K} = \mathcal{A}^2$, cf. Appendix D.3.

6. Regularization

In some applications, the number of bases or the rank constraint are not enough for obtaining good generalization performance and it is necessary to consider additional form of regularization. For the lifted objective with $\mathcal{K} = \mathcal{H}^2$ or \mathcal{A}^2 , we use the typical Frobenius-norm regularization

$$\tilde{L}_{\mathcal{K}}(\mathbf{U}, \mathbf{V}) := L_{\mathcal{K}}(\mathcal{S}(\mathbf{U}\mathbf{V}^\top)) + \frac{\beta}{2} (\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2), \quad (12)$$

where $\beta > 0$ is a regularization hyper-parameter. For the direct objective, we introduce the new regularization

$$\tilde{D}_{\mathcal{K}}(\lambda, \mathbf{P}) := D_{\mathcal{K}}(\lambda, \mathbf{P}) + \beta \sum_{s=1}^k |\lambda_s| \|\mathbf{p}_s\|^2. \quad (13)$$

This allows us to regularize λ and \mathbf{P} with a single hyper-parameter. Let us define the following nuclear norm penalized objective:

$$\bar{L}_{\mathcal{K}}(\mathbf{M}) := L_{\mathcal{K}}(\mathcal{S}(\mathbf{M})) + \beta \|\mathbf{M}\|_*. \quad (14)$$

We can show that (12), (13) and (14) are equivalent in the following sense.

Theorem 2 *Equivalence of regularized problems*

Let $\mathcal{K} = \mathcal{H}^2$ or \mathcal{A}^2 , then

$$\min_{\mathbf{M} \in \mathbb{R}^{d^2}} \bar{L}_{\mathcal{K}}(\mathbf{M}) = \min_{\substack{\mathbf{U} \in \mathbb{R}^{d \times r} \\ \mathbf{V} \in \mathbb{R}^{d \times r}}} \tilde{L}_{\mathcal{K}}(\mathbf{U}, \mathbf{V}) = \min_{\substack{\lambda \in \mathbb{R}^k \\ \mathbf{P} \in \mathbb{R}^{d \times k}}} \tilde{D}_{\mathcal{K}}(\lambda, \mathbf{P})$$

where $\text{rank}(\mathbf{M}^*) \leq r = k$ and $\mathbf{M}^* \in \text{argmin}_{\mathbf{M} \in \mathbb{R}^{d^2}} \bar{L}_{\mathcal{K}}(\mathbf{M})$.

Proof is given in Appendix C. Our proof relies on the variational form of the nuclear norm and is thus limited to $m = 2$. One of the key ingredients of the proof is to show that the minimizer of (14) is always a symmetric matrix. In addition to Theorem 2, from (Abernethy et al., 2009), we also know that every local minimum \mathbf{U}, \mathbf{V} of (12) gives a global solution $\mathbf{U}\mathbf{V}^\top$ of (14) provided that $\text{rank}(\mathbf{M}^*) \leq r$. Proving a similar result for (13) is a future work. When $m \geq 3$, as used in our experiments, a squared Frobenius norm penalty on \mathbf{P} (direct objective) or on $\{\mathbf{U}^t\}_{t=1}^m$ (lifted objective) works well in practice, although we lose the theoretical connection with the nuclear norm.

7. Coordinate descent algorithms

We now describe how to learn the model parameters by coordinate descent, which is a state-of-the-art *learning-rate free* solver for multi-convex problems (e.g., Yu et al. (2012)). In the following, we assume that ℓ is μ -smooth.

Direct objective with $\mathcal{K} = \mathcal{A}^m$ for $m \in \{2, 3\}$. First, we note that minimizing (13) w.r.t. λ can be reduced to a standard ℓ_1 -regularized convex objective via a simple change of variable. Hence we focus on minimization w.r.t. \mathbf{P} .

Let us denote the elements of \mathbf{P} by p_{js} . Then, our algorithm cyclically performs the following update for all $s \in [k]$ and $j \in [d]$:

$$p_{js} \leftarrow p_{js} - \eta^{-1} \left[\sum_{i=1}^n \ell'(y_i, \hat{y}_i) \frac{\partial \hat{y}_i}{\partial p_{js}} + 2\beta |\lambda_s| p_{js} \right],$$

where $\eta := \mu \sum_{i=1}^n \left(\frac{\partial \hat{y}_i}{\partial p_{js}} \right)^2 + 2\beta |\lambda_s|$. Note that when ℓ is the squared loss, the above is equivalent to a Newton update and is the exact coordinate-wise minimizer.

The key challenge to use CD is computing $\frac{\partial \hat{y}_i}{\partial p_{js}} = \lambda_s \frac{\partial \mathcal{A}^m(\mathbf{p}_s, \mathbf{x}_i)}{\partial p_{js}}$ efficiently. Let us denote the elements of \mathbf{X} by x_{ji} . Using Lemma 3, we obtain $\frac{\partial \mathcal{A}^2(\mathbf{p}_s, \mathbf{x}_i)}{\partial p_{js}} = \langle \mathbf{p}_s, \mathbf{x}_i \rangle x_{ji} - p_{js} x_{ji}^2$ and $\frac{\partial \mathcal{A}^3(\mathbf{p}_s, \mathbf{x}_i)}{\partial p_{js}} = \mathcal{A}^2(\mathbf{p}_s, \mathbf{x}_i) x_{ji} - p_{js} x_{ji}^2 \langle \mathbf{p}_s, \mathbf{x}_i \rangle + p_{js}^2 x_{ji}^3$. If for all $i \in [n]$ and for s fixed, we maintain $\langle \mathbf{p}_s, \mathbf{x}_i \rangle$ and $\mathcal{A}^2(\mathbf{p}_s, \mathbf{x}_i)$ (i.e., keep in sync after every update of p_{js}), then computing $\frac{\partial \hat{y}_i}{\partial p_{js}}$ takes $O(m)$ time. Hence the cost of one epoch, i.e. updating all elements of \mathbf{P} once, is $O(mkn_z(\mathbf{X}))$. Complete details and pseudo code are given in Appendix D.1.

To our knowledge, this is the first CD algorithm capable of training third-order FMs. Supporting arbitrary $m \in \mathbb{N}$ is an important future work.

Lifted objective with $\mathcal{K} = \mathcal{H}^m$. Recall that we want to learn the matrices $\{\mathbf{U}^t\}_{t=1}^m$, whose columns we denote by $\mathbf{u}_s^t = [u_{1s}^t, \dots, u_{ds}^t]^\top$. Our algorithm cyclically performs the following update for all $t \in [m]$, $s \in [r]$ and $j \in [d]$:

$$u_{js}^t \leftarrow u_{js}^t - \eta^{-1} \left[\sum_{i=1}^n \ell'(y_i, \hat{y}_i) \frac{\partial \hat{y}_i}{\partial u_{js}^t} + \beta u_{js}^t \right],$$

where $\eta := \mu \sum_{i=1}^n \left(\frac{\partial \hat{y}_i}{\partial u_{js}^t} \right)^2 + \beta$. The main difficulty is computing $\frac{\partial \hat{y}_i}{\partial u_{js}^t} = \prod_{t' \neq t} \langle \mathbf{u}_{s'}^{t'}, \mathbf{x}_i \rangle x_{ji}$ efficiently. If for all $i \in [n]$ and for t and s fixed, we maintain $\xi_i := \prod_{t' \neq t} \langle \mathbf{u}_{s'}^{t'}, \mathbf{x}_i \rangle$, then the cost of computing $\frac{\partial \hat{y}_i}{\partial u_{js}^t}$ is $O(1)$. Hence the cost of one epoch is $O(mrn_z(\mathbf{X}))$, the same as SGD. Complete details are given in Appendix D.2.

Convergence. The above updates decrease the objective *monotonically*. Convergence to a stationary point is guar-

anteed following (Bertsekas, 1999, Proposition 2.7.1).

8. Inhomogeneous polynomial models

The algorithms presented so far are designed for homogeneous polynomial kernels \mathcal{H}^m and \mathcal{A}^m . These kernels only use monomials of the *same* degree m . However, in many applications, we would like to use monomials of *up to* some degree. In this section, we propose a simple idea to do so using the algorithms presented so far, unmodified. Our key observation is that we can easily turn homogeneous polynomials into inhomogeneous ones by augmenting the dimensions of the training data with dummy features.

We begin by explaining how to learn inhomogeneous polynomial models using \mathcal{H}^m . Let us denote $\tilde{\mathbf{p}}^\top := [\gamma, \mathbf{p}^\top] \in \mathbb{R}^{d+1}$ and $\tilde{\mathbf{x}}^\top := [1, \mathbf{x}^\top] \in \mathbb{R}^{d+1}$. Then, we obtain

$$\mathcal{H}^m(\tilde{\mathbf{p}}, \tilde{\mathbf{x}}) = \langle \tilde{\mathbf{p}}, \tilde{\mathbf{x}} \rangle^m = (\gamma + \langle \mathbf{p}, \mathbf{x} \rangle)^m = \mathcal{P}_\gamma^m(\mathbf{p}, \mathbf{x}).$$

Therefore, if we prepare the augmented training set $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n$, the problem of learning a model of the form $\sum_{s=1}^k \lambda_s \mathcal{P}_{\gamma_s}^m(\mathbf{p}_s, \mathbf{x})$ can be converted to that of learning a rank- k symmetric tensor $\mathcal{W} \in \mathbb{S}^{(d+1)^m}$ using the method presented in Section 5. Note that the parameter γ_s is automatically learned from data for each basis \mathbf{p}_s .

Next, we explain how to learn inhomogeneous polynomial models using \mathcal{A}^m . Using Lemma 2, we immediately obtain for $1 \leq m \leq d$:

$$\mathcal{A}^m(\tilde{\mathbf{p}}, \tilde{\mathbf{x}}) = \mathcal{A}^m(\mathbf{p}, \mathbf{x}) + \gamma \mathcal{A}^{m-1}(\mathbf{p}, \mathbf{x}). \quad (15)$$

For instance, when $m = 2$, we obtain

$$\mathcal{A}^2(\tilde{\mathbf{p}}, \tilde{\mathbf{x}}) = \mathcal{A}^2(\mathbf{p}, \mathbf{x}) + \gamma \mathcal{A}^1(\mathbf{p}, \mathbf{x}) = \mathcal{A}^2(\mathbf{p}, \mathbf{x}) + \gamma \langle \mathbf{p}, \mathbf{x} \rangle.$$

Therefore, if we prepare the augmented training set $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n$, we can easily learn a combination of linear kernel and second-order ANOVA kernel using methods presented in Section 4 or Section 5. Note that (15) only states the relation between two ANOVA kernels of consecutive degrees. Fortunately, we can also apply (15) recursively. Namely, by adding $m - 1$ dummy features, we can sum the kernels from \mathcal{A}^m down to \mathcal{A}^1 (i.e., linear kernel).

9. Experimental results

In this section, we present experimental results, focusing on regression tasks. Datasets are described in Appendix E. In all experiments, we set $\ell(y, \hat{y})$ to the squared loss.

9.1. Direct optimization: is it useful to fit λ ?

As explained in Section 4, there is no benefit to fitting λ when m is odd, since \mathcal{A}^m and \mathcal{H}^m can absorb λ into \mathbf{P} . This is however not the case when m is even: \mathcal{A}^m and \mathcal{H}^m

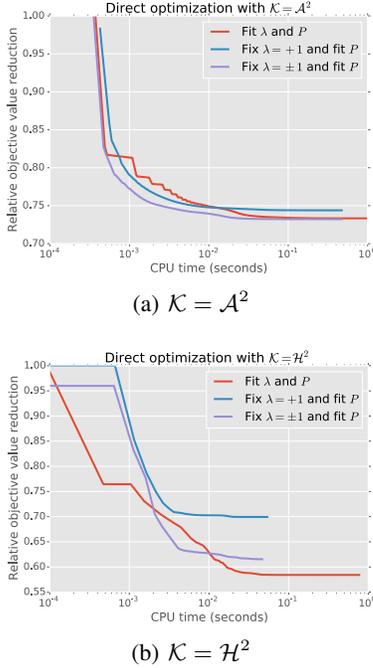


Figure 1. Effect of fitting λ when using direct optimization on the *diabetes* dataset with $m = 2$, $\beta = 10$ and $k = 4$. Objective values were computed by (13) and were normalized by the worst initialization’s objective value.

can absorb absolute values but not negative signs (unless complex numbers are allowed for parameters). Therefore, when m is even, the class of functions we can represent with models of the form (1) is possibly smaller if we fix $\lambda = 1$ (as done in FMs).

To check that this is indeed the case, on the *diabetes* dataset, we minimized (13) with $m = 2$ as follows:

- minimize w.r.t. both λ and P alternately,
- fix $\lambda_s = 1$ for $s \in [k]$ and minimize w.r.t. P ,
- fix $\lambda_s = \pm 1$ with proba. 0.5 and minimize w.r.t. P .

We initialized elements of P by $p_{js} \sim \mathcal{N}(0, 0.01)$ for all $j \in [d]$, $s \in [k]$. Our results are shown in Figure 1. For $\mathcal{K} = \mathcal{A}^2$, we use CD and for $\mathcal{K} = \mathcal{H}^2$, we use L-BFGS. Note that since (13) is convex w.r.t. λ , a) is insensitive to the initialization of λ as long as we fit λ before P . Not surprisingly, fitting λ allows us to achieve a smaller objective value. This is especially apparent when $\mathcal{K} = \mathcal{H}^2$. However, the difference is much smaller when $\mathcal{K} = \mathcal{A}^2$. We give intuitions as to why this is the case in Section 10.

We emphasize that this experiment was designed to confirm that fitting λ does indeed improve representation power of the model when m is even. In practice, it is possible that fixing $\lambda = 1$ reduces overfitting and thus improves *generalization error*. However, this highly depends on the data.

9.2. Direct vs. lifted optimization

In this section, we compare the direct and lifted optimization approaches on high-dimensional data when $m = 2$. To compare the two approaches fairly, we propose the following initialization scheme. Recall that, at the end of the day, both approaches are essentially learning a low rank symmetric matrix: $W = S(UV^T)$ for lifted and $W = P \text{diag}(\lambda)P^T$ for direct optimization. This suggests that we can easily convert the matrices $U, V \in \mathbb{R}^{d \times r}$ used for initializing lifted optimization to $P \in \mathbb{R}^{d \times k}$ and $\lambda \in \mathbb{R}^{d \times k}$ by computing the (reduced) eigendecomposition of $S(UV^T)$. Note that because we solve the lifted optimization problem by coordinate descent, UV^T is never symmetric and therefore the rank of $S(UV^T)$ is usually twice that of UV^T . Hence, in practice, we have that $r = k/2$. In our experiment, we compared four methods: lifted objective solved by CD, direct objective solved by CD, L-BFGS and SGD. For lifted optimization, we initialized the elements of U and V by sampling from $\mathcal{N}(0, 0.01)$. For direct optimization, we obtained P and λ as explained. Results on the *E2006-tfidf* high-dimensional dataset are shown in Figure 2. For $\mathcal{K} = \mathcal{A}^2$, we find that Lifted (CD) and Direct (CD) have similar convergence speed and both outperform Direct (L-BFGS). For $\mathcal{K} = \mathcal{H}^2$, we find that Lifted (CD) outperforms both Direct (L-BFGS) and Direct (SGD). Note that we did not implement Direct (CD) for $\mathcal{K} = \mathcal{H}^2$ since the direct optimization problem is not coordinate-wise convex, as explained in Section 5.

9.3. Recommender system experiment

To confirm the ability of the proposed framework to infer the weights of unobserved feature interactions, we conducted experiments on *Last.fm* and *MovieLens 1M*, two standard recommender system datasets. Following (Rendle, 2012), matrix factorization can be reduced to FMs by creating a dataset of (x_i, y_i) pairs where x_i contains the one-hot encoding of the user and item and y_i is the corresponding rating (i.e., number of training instances equals number of ratings). We compared four models:

- $\mathcal{K} = \mathcal{A}^2$ (augment): $\hat{y} = \hat{y}_{\mathcal{A}^2}(\tilde{x})$, with $\tilde{x}^T := [1, x^T]$,
- $\mathcal{K} = \mathcal{A}^2$ (linear combination): $\hat{y} = \langle w, x \rangle + \hat{y}_{\mathcal{A}^2}(x)$,
- $\mathcal{K} = \mathcal{H}^2$ (augment): $\hat{y} = \hat{y}_{\mathcal{H}^2}(\tilde{x})$ and
- $\mathcal{K} = \mathcal{H}^2$ (linear combination): $\hat{y} = \langle w, x \rangle + \hat{y}_{\mathcal{H}^2}(x)$,

where $w \in \mathbb{R}^d$ is a vector of first-order weights, estimated from training data. Note that b) and d) are exactly the same as FMs and PNs, respectively. Results are shown in Figure 3. We see that \mathcal{A}^2 tends to outperform \mathcal{H}^2 on these tasks. We hypothesize that this is the case because features are binary (cf., discussion in Section 10). We also see that simply augmenting the features as suggested in Section 8 is comparable or better than learning additional first-order feature weights, as done in FMs and PNs.

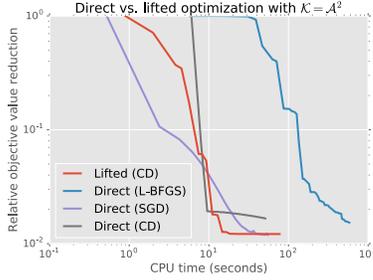
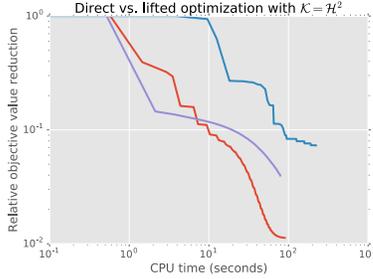
(a) $\mathcal{K} = \mathcal{A}^2$ (b) $\mathcal{K} = \mathcal{H}^2$

Figure 2. Comparison of the direct and lifted optimization approaches on the *E2006-tfidf* high-dimensional dataset with $m = 2$, $\beta = 100$ and $r = k/2 = 10$. In order to learn an inhomogeneous polynomial, we added a dummy feature to all training instances, as explained in Section 8. Objective values shown were computed by (13) and (12) and normalized by the initialization’s objective value.

9.4. Low-budget non-linear regression experiment

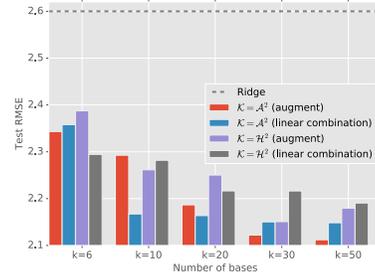
In this experiment, we demonstrate the ability of the proposed framework to reach good regression performance with a small number of bases k . We compared:

- Proposed with $\mathcal{K} = \mathcal{H}^3$ (with augmented features),
- Proposed with $\mathcal{K} = \mathcal{A}^3$ (with augmented features),
- Nyström method with $\mathcal{K} = \mathcal{P}_1^3$ and
- Random Selection: choose $\mathbf{p}_1, \dots, \mathbf{p}_k$ uniformly at random from training set and use $\mathcal{K} = \mathcal{P}_1^3$.

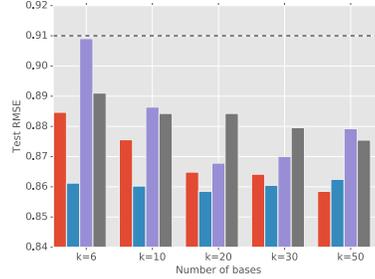
For a) and b) we used the lifted approach. For fair comparison in terms of model size (number of floats used), we set $r = k/3$. Results on the *abalone*, *cadata* and *cpusmall* datasets are shown in Figure 4. We see that *i*) the proposed framework reaches the same performance as kernel ridge regression with much fewer bases than other methods and *ii*) \mathcal{H}^3 tends to outperform \mathcal{A}^3 on these tasks. Similar trends were observed when using $\mathcal{K} = \mathcal{H}^2$ or \mathcal{A}^2 .

10. Discussion

Ability to infer weights of unobserved interactions. In our view, one of the strengths of PNs and FMs is their



(a) Last.fm



(b) MovieLens 1M

Figure 3. Predicted rating error on the *Last.fm* and *MovieLens 1M* datasets. The metric used is RMSE on the test set (lower is better). The hyper-parameter β was selected from 10 log-spaced values in the interval $[10^{-3}, 10^3]$ by 5-fold cross-validation.

ability to infer the weights of unobserved feature interactions, unlike traditional kernel methods. To see why, recall that in kernel methods, predictions are computed by $\hat{y} = \sum_{i=1}^n \alpha_i \mathcal{K}(\mathbf{x}_i, \mathbf{x})$. When $\mathcal{K} = \mathcal{H}^m$ or \mathcal{A}^m , by Lemma 5, this is equivalent to $\hat{y} = \langle \widetilde{\mathcal{W}}, \mathbf{x}^{\otimes m} \rangle$ or $\langle \widetilde{\mathcal{W}}, \mathbf{x}^{\otimes m} \rangle_>$ if we set $\widetilde{\mathcal{W}} := \sum_{i=1}^n \alpha_i \mathbf{x}_i^{\otimes m}$. Thus, in kernel methods, the weight associated with $x_{j_1} \dots x_{j_m}$ can be written as a linear combination of the training data’s monomials:

$$\widetilde{\mathcal{W}}_{j_1, \dots, j_m} = \sum_{i=1}^n \alpha_i x_{j_1 i} \dots x_{j_m i}.$$

Assuming binary features, the weights of monomials that were never observed in the training set are zero. In contrast, in PNs and FMs, we have $\mathcal{W} = \sum_{s=1}^k \lambda_s \mathbf{p}_s^{\otimes m}$ and therefore the weight associated with $x_{j_1} \dots x_{j_m}$ becomes

$$\mathcal{W}_{j_1, \dots, j_m} = \sum_{s=1}^k \lambda_s p_{j_1 s} \dots p_{j_m s}.$$

Because parameters are shared across monomials, PNs and FMs are able to *interpolate* the weights of monomials that were never observed in the training set. This is the key property which makes it possible to use them on recommender system tasks. In future work, we plan to apply PNs and FMs to biological data, where this property should be

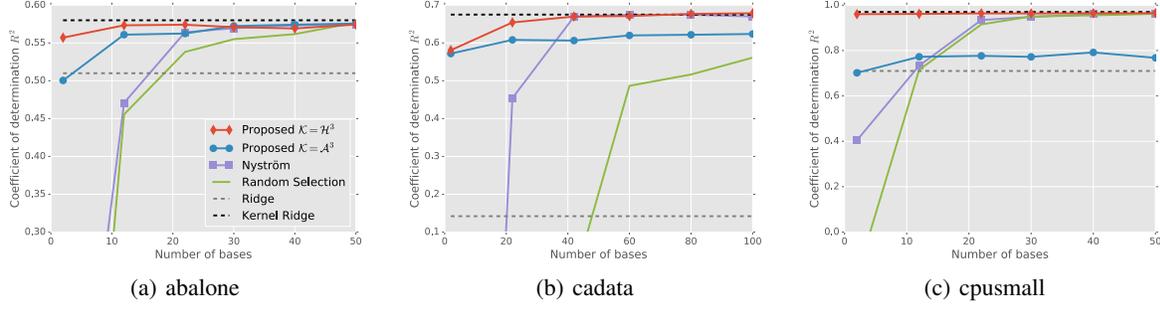


Figure 4. Regression performance as a function of the number of bases. The metric used is the coefficient of determination on the test set (higher is better). Regularization parameter was selected from 10 log-spaced values in $[10^{-4}, 10^4]$ by 5-fold cross-validation.

very useful, e.g., for inferring higher-order interactions between genes.

ANOVA kernel vs. polynomial kernel. One of the key properties of the ANOVA kernel $\mathcal{A}^m(\mathbf{p}, \mathbf{x})$ is multilinearity w.r.t. elements of \mathbf{p} (Lemma 2). This is the key difference with $\mathcal{H}^m(\mathbf{p}, \mathbf{x})$ which makes the direct optimization objective multi-convex when $\mathcal{K} = \mathcal{A}^m$ (Theorem 1). However, because we need to ignore irrelevant monomials, computing the kernel and its gradient is more challenging. Deriving efficient training algorithms for arbitrary $m \in \mathbb{N}$ is an important future work.

In our experiments in Section 9.1, we showed that fixing $\lambda = 1$ works relatively well when $\mathcal{K} = \mathcal{A}^2$. To see intuitively why this is the case, note that fixing $\lambda = 1$ is equivalent to constraining the weight matrix \mathbf{W} to be positive semidefinite, i.e., $\exists \mathbf{P}$ s.t. $\mathbf{W} = \mathbf{P}\mathbf{P}^\top$. Next, observe that we can rewrite the prediction function as

$$\hat{y}_{\mathcal{A}^2}(\mathbf{x}; \mathbf{1}, \mathbf{P}) = \langle \mathbf{P}\mathbf{P}^\top, \mathbf{x}^{\otimes 2} \rangle_{>} = \langle \mathcal{U}(\mathbf{P}\mathbf{P}^\top), \mathbf{x}^{\otimes 2} \rangle,$$

where $\mathcal{U}(\mathbf{M})$ is a mask which sets diagonal and lower-diagonal elements of \mathbf{M} to zero. We therefore see that when using $\mathcal{K} = \mathcal{A}^2$, we are learning a *strictly upper-triangular matrix*, parametrized by $\mathbf{P}\mathbf{P}^\top$. Importantly, the matrix $\mathcal{U}(\mathbf{P}\mathbf{P}^\top)$ is not positive semidefinite. This is what gives the model some degree of freedom, even though $\mathbf{P}\mathbf{P}^\top$ is positive semidefinite. In contrast, when using $\mathcal{K} = \mathcal{H}^2$, if we fix $\lambda = 1$, then we have that

$$\hat{y}_{\mathcal{H}^2}(\mathbf{x}; \mathbf{1}, \mathbf{P}) = \langle \mathbf{P}\mathbf{P}^\top, \mathbf{x}^{\otimes 2} \rangle = \mathbf{x}^\top \mathbf{P}\mathbf{P}^\top \mathbf{x} \geq 0$$

and therefore the model is unable to predict negative values.

Empirically, we showed in Section 9.4 that \mathcal{H}^m outperforms \mathcal{A}^m for low-budget non-linear regression. In contrast, we showed in Section 9.3 that \mathcal{A}^m outperforms \mathcal{H}^m for recommender systems. The main difference between the two experiments is the nature of the features used: continuous for the former and binary for the latter. For binary features, squared features x_1^2, \dots, x_d^2 are redundant with

x_1, \dots, x_d and are therefore not expected to help improve accuracy. On the contrary, they might introduce bias towards first-order features. We hypothesize that the ANOVA kernel is in general a better choice for binary features, although this needs to be verified by more experiments, for instance on natural language processing (NLP) tasks.

Direct vs. lifted optimization. The main advantage of direct optimization is that we only need to estimate $\lambda \in \mathbb{R}^k$ and $\mathbf{P} \in \mathbb{R}^{d \times k}$ and therefore the number of parameters to estimate is independent of the degree m . Unfortunately, the approach is neither convex nor multi-convex when using $\mathcal{K} = \mathcal{H}^m$. In addition, the regularized objective (13) is non-smooth w.r.t. λ . In Section 5, we proposed to reformulate the problem as one of low-rank symmetric tensor estimation and used a symmetrization trick to obtain a multi-convex smooth objective function. Because this objective involves the estimation of m matrices of size $d \times r$, we need to set $r = k/m$ for fair comparison with the direct objective in terms of model size. When $\mathcal{K} = \mathcal{A}^m$, we showed that the direct objective is readily multi-convex. However, an advantage of our lifted objective when $\mathcal{K} = \mathcal{A}^m$ is that it is convex w.r.t. larger block of variables than the direct objective.

11. Conclusion

In this paper, we revisited polynomial networks (Livni et al., 2014) and factorization machines (Rendle, 2010; 2012) from a unified perspective. We proposed direct and lifted optimization approaches and showed their equivalence in the regularized case for $m = 2$. With respect to PNs, we proposed the first CD solver with support for arbitrary integer $m \geq 2$. With respect to FMs, we made several novel contributions including making a connection with the ANOVA kernel, proving important properties of the objective function and deriving the first CD solver for third-order FMs. Empirically, we showed that the proposed algorithms achieve excellent performance on non-linear regression and recommender system tasks.

Acknowledgments

This work was partially conducted as part of “Research and Development on Fundamental and Applied Technologies for Social Big Data”, commissioned by the National Institute of Information and Communications Technology (NICT), Japan. We also thank Vlad Niculae, Olivier Grisel, Fabian Pedregosa and Joseph Salmon for their valuable comments.

References

- Abernethy, Jacob, Bach, Francis, Evgeniou, Theodoros, and Vert, Jean-Philippe. A new approach to collaborative filtering: Operator estimation with spectral regularization. *J. Mach. Learn. Res.*, 10:803–826, 2009.
- Avron, Haim, Nguyen, Huy, and Woodruff, David. Subspace embeddings for the polynomial kernel. In *Advances in Neural Information Processing Systems 27*, pp. 2258–2266. 2014.
- Bertsekas, Dimitri P. *Nonlinear programming*. Athena scientific Belmont, 1999.
- Blondel, Mathieu, Fujino, Akinori, and Ueda, Naonori. Convex factorization machines. In *Proceedings of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, 2015.
- Candès, Emmanuel J., Eldar, Yonina C., Strohmer, Thomas, and Voroninski, Vladislav. Phase retrieval via matrix completion. *SIAM Journal on Imaging Sciences*, 6(1):199–225, 2013.
- Chang, Yin-Wen, Hsieh, Cho-Jui, Chang, Kai-Wei, Ringgaard, Michael, and Lin, Chih-Jen. Training and testing low-degree polynomial data mappings via linear svm. *Journal of Machine Learning Research*, 11:1471–1490, 2010.
- Chen, Danqi and Manning, Christopher D. A fast and accurate dependency parser using neural networks. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, volume 1, pp. 740–750, 2014.
- Comon, Pierre, Golub, Gene, Lim, Lek-Heng, and Mourrain, Bernard. Symmetric tensors and symmetric tensor rank. *SIAM Journal on Matrix Analysis and Applications*, 30(3):1254–1279, 2008.
- Kar, Purushottam and Karnick, Harish. Random feature maps for dot product kernels. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pp. 583–591, 2012.
- Livni, Roi, Shalev-Shwartz, Shai, and Shamir, Ohad. On the computational efficiency of training neural networks. In *Advances in Neural Information Processing Systems*, pp. 855–863, 2014.
- Mazumder, Rahul, Hastie, Trevor, and Tibshirani, Robert. Spectral regularization algorithms for learning large incomplete matrices. *The Journal of Machine Learning Research*, 11:2287–2322, 2010.
- 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. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Pham, Ninh and Pagh, Rasmus. Fast and scalable polynomial kernels via explicit feature maps. In *Proceedings of the 19th KDD conference*, pp. 239–247, 2013.
- Rendle, Steffen. Factorization machines. In *Proceedings of International Conference on Data Mining*, pp. 995–1000. IEEE, 2010.
- Rendle, Steffen. Factorization machines with libfm. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(3):57–78, 2012.
- Shawe-Taylor, John and Cristianini, Nello. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- Sonnenburg, Sören and Franc, Vojtech. Coffin: A computational framework for linear svms. In *Proceedings of the 27th International Conference on Machine Learning*, pp. 999–1006, 2010.
- Stitson, Mark, Gammernan, Alex, Vapnik, Vladimir, Vovk, Volodya, Watkins, Chris, and Weston, Jason. Support vector regression with anova decomposition kernels. Technical report, Royal Holloway University of London, 1997.
- Vapnik, Vladimir. *Statistical learning theory*. Wiley, 1998.
- Wang, Z., Crammer, K., and Vucetic, S. Multi-class pegasos on a budget. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, pp. 1143–1150, 2010.
- Williams, Christopher K. I. and Seeger, Matthias. Using the nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems 13*, pp. 682–688. 2001.
- Yu, Hsiang-Fu, Hsieh, Cho-Jui, Si, Si, and Dhillon, Inderjit S. Scalable coordinate descent approaches to parallel matrix factorization for recommender systems. In *ICDM*, pp. 765–774, 2012.