
Tensor Regression Meets Gaussian Processes

Rose Yu*

*California Institute of Technology

Guangyu Li†

†University of Southern California

Yan Liu†

Abstract

Low-rank tensor regression, a new model class that learns high-order correlation from data, has recently received considerable attention. At the same time, Gaussian processes (GP) are well-studied machine learning models for structure learning. In this paper, we demonstrate interesting connections between the two, especially for multi-way data analysis. We show that low-rank tensor regression is essentially learning a multi-linear kernel in Gaussian processes, and the low-rank assumption translates to the constrained Bayesian inference problem. We prove the oracle inequality and derive the average case learning curve for the equivalent GP model. Our finding implies that low-rank tensor regression, though empirically successful, is highly dependent on the eigenvalues of covariance functions as well as variable correlations.

1 Introduction

High-order correlations are ubiquitous in modern data analytics. For instance, data generated from a sensor network contain measurements from different locations, time stamps, and variables. Accurate prediction requires models that can simultaneously capture correlations across time, space and variables. Low-rank tensor regression is a class of supervised learning models that aim to learn such high-order correlations. In recent years, low-rank tensor regression has been intensively studied, leading to successful applications in multi-task learning (Wimalawarne et al., 2014), deep learning (Novikov et al., 2015) and complex network analysis (Imaizumi and Hayashi, 2016).

In contrast to traditional unsupervised tensor decompo-

sition (Kolda and Bader, 2009), tensor regression (Zhou et al., 2013) learns a tensor model in a supervised fashion and imposes low-rank structure for dimension reduction. Tensor regression has several advantages over vector or matrix regression: from the modeling perspective, the model tensor provides an explicit parameterization for the multi-directional interdependence among variables. The low-rankness represents the shared latent space in the data. From the learning perspective, tensor model enjoys lower sample complexity. The tensor low-rank constraint regularizes the model to be more generalizable. However, a notable disadvantage of tensor regression is the absence of confidence intervals for the predictions, which calls for a probabilistic counterpart that can effectively represent the high-order correlations in the data.

Meanwhile, Gaussian processes (Rasmussen, 2006) are well-established techniques for modeling correlations structures. With versatile covariance design, GP remain popular in spatial statistics and time series analysis. A natural question then arises, “which method is better? And how are these two model classes related?” Known examples of similar connections include the Gaussian process latent variable model (Lawrence, 2004) for PCA, the multi-task Gaussian process model (Bonilla et al., 2007) for multi-task learning and the probabilistic Tucker model for Tucker tensor decomposition (Chu and Ghahramani, 2009). The probabilistic interpretation deepens the understanding of the regularized optimization approach, suggesting its generalization to non-Gaussian data with kernel methods.

In this paper, we make the first attempt at understanding this connection. We show that tensor regression is equivalent to learning a Gaussian process with multi-linear transformation kernel: multi-linear Gaussian process (MLGP). The low-rank assumption on the parameter tensor can be interpreted as a constrained Bayesian inference problem. We analyze the theoretical properties of MLGP by proving its oracle inequality and deriving the average case learning curve. We validate our theory with numerical simulations and provide a comparative analysis between different GP models. Finally, we showcase the model on three real-world

tensor regression applications: multi-linear multi-task learning, spatio-temporal forecasting, and multi-output regression. The model not only can achieve superior performance but also uncover interesting patterns from multi-way data.

Note that the goal of our work is fundamentally different from existing works on Bayesian estimator for tensor-variate regression (Guhaniyogi et al., 2015; Xu et al., 2015; Suzuki, 2015). For example, (Xu et al., 2015) propose a generative model for Bayesian tensor regression; (Suzuki, 2015) analyzes the minimax optimal rate of the estimator. These works emphasize probabilistic modeling instead of establishing the connections. And most existing theoretical analyses are asymptotic. In contrast, our work aims to provide deeper insights into the relationship between the optimizers of tensor regression and estimators for Gaussian process models.

2 Tensor Regression and Its Counterpart

2.1 Low-Rank Tensor Regression

Tensor regression exploits the high-order correlation in the data. It learns a multi-linear function whose parameters form a tensor. To represent shared latent spaces and address “the curse of dimensionality” issue, tensor regression usually constrains the mode tensor to be low-rank. Formally, given an input tensor \mathcal{X} , an output tensor \mathcal{Y} and a model parameter tensor \mathcal{W} , tensor regression aims to solve the following optimization problem:

$$\begin{aligned} \mathcal{W}^* &= \operatorname{argmin}_{\mathcal{W}} \hat{\mathcal{L}}(f(\mathcal{X}, \mathcal{W}); \mathcal{Y}) \\ \text{s.t.} \quad &\operatorname{rank}(\mathcal{W}) \leq R \end{aligned} \quad (1)$$

where $\hat{\mathcal{L}}$ denotes the loss function, and f represents a regression model (e.g. linear, logistic). The solution \mathcal{W}^* minimizes the empirical loss, subject to the tensor low-rank constraint $\operatorname{rank}(\mathcal{W}) \leq R$.

Low-rank tensor regression has many applications. One example is multi-linear multi-task learning¹ (MLMTL), which learns multiple tasks with a multi-level task hierarchy. For example, when forecasting the energy demand for multiple power plants, we can split the tasks by categories: coal, oil and natural gas. MLMTL improves the prediction by modeling the correlations within and across categories. We can encode such task hierarchy using a tensor, where the first dimension of the tensor represents features, and the rest to index the grouped tasks at each level.

¹Other applications can be re-formulated as special cases of multi-linear multi-task learning

Specifically, given T learning tasks with feature dimension T_1 , we can split them into T_2 groups, each of which contains $T_3 = T/T_2$ tasks. Assuming each task t contains n_t training data points $\{\mathbf{x}_{t,i}, \mathbf{y}_{t,i}\}_{i=1}^{n_t}$ and is parametrized by $\mathbf{w}_t \in \mathbb{R}^{T_1}$. We can form a tensor by concatenating all the parameters as a matrix $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_T]$ and folding along the feature dimension $\mathcal{W} = \operatorname{fold}_{(1)}(\mathbf{W}) \in \mathbb{R}^{T_1 \times T_2 \times T_3}$. The objective of MLMTL is to learn this parameter tensor subject to the low-rank constraint:

$$\begin{aligned} \mathcal{W}^* &= \operatorname{argmin}_{\mathcal{W}} \sum_{t=1}^T \sum_{i=1}^{n_t} \mathcal{L}(\langle \mathbf{x}_{t,i}, \mathbf{w}_t \rangle; \mathbf{y}_{t,i}) \\ \text{s.t.} \quad &\operatorname{rank}(\mathcal{W}) \leq R \end{aligned} \quad (2)$$

If the task hierarchy has two levels $T = T_2 \times T_3$, we obtain a third-order tensor. In general, one can use an $(m+1)$ -order tensor to represent an m -level task clustering hierarchy. Note that the definition of tensor rank is not unique (Kolda and Bader, 2009). One popular definition is Tucker rank due to its computational benefit. Tucker rank assumes that the tensor \mathcal{W} has a Tucker decomposition $\mathcal{W} = \mathcal{S} \times_1 \mathbf{U}_1 \times_2 \mathbf{U}_2 \times_3 \mathbf{U}_3$, with a core tensor $\mathcal{S} \in \mathbb{R}^{R_1 \times R_2 \times R_3}$ and orthonormal projection matrices $\{\mathbf{U}_m\}_{m=1}^3$. Tucker rank corresponds to the size of the core tensor \mathcal{S} .

Low-rank tensor regression is a challenging problem mainly due to the subspace of low-rank tensors is non-convex, resulting in a high-dimensional non-convex problem. Recent developments have seen efficient algorithms for solving Equation 1 and 2, e.g., (Yu and Liu, 2016; Rabusseau and Kadri, 2016), demonstrating low-rank tensor regression as a scalable method for multi-way data analysis. However, one major drawback of such formulation is that it trades uncertainty for efficiency: there is no confidence interval for the prediction. Hence, it is difficult for the learned tensor model to reason with uncertainty. In seek of its probabilistic counterpart, we resort to another class of structured learning models: Gaussian processes.

2.2 Multi-linear Gaussian Processes

Gaussian process regression infers continuous values with a GP prior. Given input \mathbf{x} , output \mathbf{y} , and a regression model

$$\mathbf{y} = f(\mathbf{x}) + \epsilon, \quad f(\mathbf{x}) \sim \operatorname{GP}(m, k) \quad (3)$$

with ϵ as the Gaussian noise. GP characterizes a prior distribution over function $f(\mathbf{x})$ with a mean function m and a covariance function k . By definition, we have $\mathbb{E}[f(\mathbf{x})] = m(\mathbf{x})$, $\operatorname{cov}(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}, \mathbf{x}')$. The mean function is usually defined to be zero. The covariance function completely defines the process’s behavior.

Next, we develop a GP model to describe the generative process of the MLMTL problem. Given a total of

$N = \sum_{t=1}^T n_t$ training data points $\{\mathbf{x}_{t,i}, \mathbf{y}_{t,i}\}_{i=1}^{n_t}$ from T related tasks, we assume that each data point $(\mathbf{x}_{t,i}, \mathbf{y}_{t,i})$ is drawn i.i.d from the following probabilistic model:

$$\mathbf{y}_{t,i} = f(\mathbf{x}_{t,i}) + \epsilon_t, \quad f(\mathbf{x}_{t,i}) \sim \text{GP}(0, k) \quad (4)$$

where the task t has a Gaussian noise $\epsilon_t \sim N(0, \sigma_t^2)$ with zero mean and variance σ_t^2 . To model multiple tasks, we can concatenate the data from all tasks:

$$\mathbf{y} = \begin{bmatrix} \mathbf{y}_{t,1} \\ \mathbf{y}_{t,2} \\ \dots \\ \mathbf{y}_{T,n_T} \end{bmatrix}, \quad \mathbf{X} = \begin{bmatrix} \mathbf{X}_1 & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{X}_2 & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{X}_T \end{bmatrix},$$

$$\mathbf{D} = \begin{bmatrix} \sigma_1^2 \otimes I_{n_1} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \sigma_2^2 \otimes I_{n_2} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \sigma_T^2 \otimes I_{n_T} \end{bmatrix}$$

where $\mathbf{X}_t = [\mathbf{x}_{t,1}; \mathbf{x}_{t,2}; \dots; \mathbf{x}_{t,n_t}]$ is the vectorization of the inputs for task t . In matrix form, the probabilistic model generalizes Equation 4 into:

$$\mathbf{y} = f(\mathbf{X}) + \mathbf{e}, \quad f(\mathbf{X}) \sim \text{GP}(\mathbf{0}, \mathbf{K}), \quad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{D})$$

with \mathbf{X} as the inputs, \mathbf{K} as the input covariance matrix and \mathbf{D} as the noise covariance.

To represent the multi-level task hierarchy $T = T_2 \times T_3$, we define the kernel matrix \mathbf{K} with Kronecker products:

$$\mathbf{K} = \phi(\mathbf{X})\mathbf{K}_3 \otimes \mathbf{K}_2 \otimes \mathbf{K}_1\phi(\mathbf{X})^\top$$

where \mathbf{K}_1 models the feature correlations, \mathbf{K}_2 models the correlations across groups, and \mathbf{K}_3 represents the dependences of tasks within the group. $\phi(\cdot)$ maps the inputs to a T_1 dimensional feature space.² This multi-linear kernel provides a multi-resolution compositional representation. It is expressive yet efficient. Figure 1 shows several examples of such construction with three kernel functions: Linear $k(x, x') = a + b(x - c)(x' - c)$, Squared Exponential $k(x, x') = a \exp \frac{-(x-x')^2}{2c}$ and Periodic $k(x, x') = a \exp -\frac{\sin^2(\pi|x-x'|)}{c}$ in different orders. We name this class of GP model multi-linear Gaussian processes (MLGP) as the kernel matrix encodes multi-linear structure.

2.3 Connection Between Two Models

In the following section, we connect low-rank tensor regression with multi-linear Gaussian processes by examining the common structures that the two models aim to learn.

²We want to clarify that the use of $\phi(\cdot)$ limits the model to a finite feature space. And the model itself is parametric, which is the same as the tensor regression formulation.

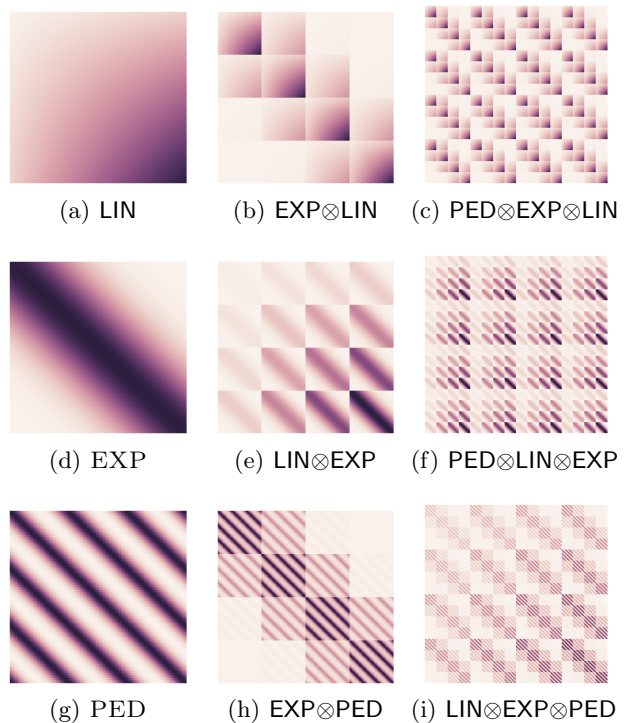


Figure 1: Multi-linear kernel construction by composing Linear (LIN) Squared Exponential (EXP) and Periodic (PED) kernels on a $(50 \times 50) \times (4 \times 4) \times (4 \times 4)$ grid.

When dealing with a large number of tasks and high dimensional data, learning $\{\mathbf{K}_m\}_{m=1}^3$ can be very expensive. To reduce the computational cost, we use the low-rank approximation for each correlation matrix:

$$\{\mathbf{K}_m = \mathbf{U}_m \mathbf{U}_m^\top \in \mathbb{R}^{T_m \times T_m}\}_{m=1}^3$$

where $\mathbf{U}_m \in \mathbb{R}^{T_m \times R_m}$ is an orthogonal matrix with dimension R_m much smaller than T_m .

The weight-space view of GP allows us to re-write the latent function: $f(\mathbf{X}) = \langle \text{vec}(\mathcal{W}), \phi(\mathbf{X}) \rangle$, where $\mathcal{W} \in \mathbb{R}^{T_1 \times T_2 \times T_3}$ is the regression model parameters with the following prior distribution:

$$\begin{aligned} \text{vec}(\mathcal{W}) &= (\mathbf{U}_1 \otimes \mathbf{U}_2 \otimes \mathbf{U}_3)^\top \text{vec}(\mathcal{S}) \\ \text{vec}(\mathcal{S}) &\sim \mathcal{N}(\mathbf{0}, \sigma_s^2 \mathbf{I}) \end{aligned}$$

Here $\mathcal{S} \in \mathbb{R}^{R_1 \times R_2 \times R_3}$ is a super-diagonal core tensor with i.i.d entries. $\{\mathbf{U}_m \in \mathbb{R}^{T_m \times R_m}\}$ is a set of orthogonal projection matrices.

Under the MLGP model, the prior distribution of the latent function follows Gaussian $p(f|\mathbf{X}) = \mathcal{N}(0, \mathbf{K})$, and the likelihood distribution is $p(\mathbf{y}|f) = \mathcal{N}(f, \mathbf{D})$. By integrating out the model parameters, we can obtain the marginal distribution of the outputs \mathbf{y} :

$$p(\mathbf{y}|\mathbf{X}) = \int_f p(\mathbf{y}, f, s|\mathbf{X}) df ds = \mathcal{N}(\mathbf{0}, \mathbf{K} + \mathbf{D})$$

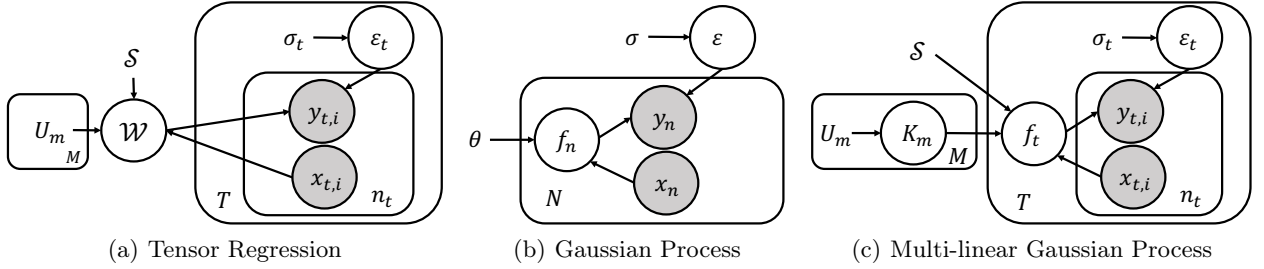


Figure 2: Graphical model for (a) tensor regression, (b) Gaussian process and (c) MLGP. The outer plate represents tasks, while the inner plate represents the repeated examples within a task.

where we omit the core tensor constant σ_s , which acts as a regularization term. The log-likelihood of the marginal distribution for MLGP is:

$$L = -\frac{1}{2} \log |\mathbf{K} + \mathbf{D}| - \frac{1}{2} \mathbf{y}^\top (\mathbf{K} + \mathbf{D})^{-1} \mathbf{y} + \text{const} \quad (5)$$

s.t. $\mathbf{K} = \phi(\mathbf{X}) \otimes_{m=1}^3 \mathbf{K}_m \phi(\mathbf{X})^\top$

Using the Kronecker product property $\otimes_{m=1}^3 \mathbf{U}_m \mathbf{U}_m^\top = (\otimes_{m=1}^3 \mathbf{U}_m)(\otimes_{m=1}^3 \mathbf{U}_m)^\top$, we can re-write the covariance matrix as:

$$\mathbf{K} = (\phi(\mathbf{X}) \otimes_{m=1}^3 \mathbf{U}_m)(\phi(\mathbf{X}) \otimes_{m=1}^3 \mathbf{U}_m)^\top$$

Denote $\tilde{\mathbf{U}} = \phi(\mathbf{X}) \otimes_{m=1}^3 \mathbf{U}_m$ and let the singular value decomposition of $\tilde{\mathbf{U}}$ be $\tilde{\mathbf{U}} = \mathbf{U}_x \boldsymbol{\Sigma}_x \mathbf{V}_x^\top$. We can maximize the log-likelihood by taking derivatives over L with respect to $\tilde{\mathbf{U}}$ and set it to zero, which gives the stationary point condition:

$$\mathbf{y} \mathbf{y}^\top (\mathbf{K} + \mathbf{D})^{-1} \tilde{\mathbf{U}} = \tilde{\mathbf{U}}$$

With some manipulation, we can obtain an equivalent eigenvalue problem. Detailed derivation can be found in Appendix A.1.

$$\mathbf{y} \mathbf{y}^\top \mathbf{U}_x = \mathbf{U}_x (\boldsymbol{\Sigma}_x^2 + \mathbf{D})$$

Further perform eigen-decomposition of the output covariance $\mathbf{y} \mathbf{y}^\top = \mathbf{U}_y \boldsymbol{\Lambda}_y \mathbf{U}_y^{-1}$, we have $\mathbf{U}_x = \mathbf{U}_y$, $\boldsymbol{\Sigma}_x = (\boldsymbol{\Lambda}_y - \mathbf{D})^{\frac{1}{2}}$. Therefore, the likelihood of the MLGP model is maximized when the solution satisfies

$$\phi(\mathbf{X}) \otimes_{m=1}^3 \mathbf{U}_m = \mathbf{U}_y (\boldsymbol{\Lambda}_y - \mathbf{D})^{\frac{1}{2}} \mathbf{V}_x^\top \quad (6)$$

which suggests that the maximum likelihood estimator of MLGP correspond to a multi-linear transformation from the feature space $\phi(\mathbf{X})$ to the principal subspace of the output. Recall that for tensor regression in Equation 1, the model parameter tensor \mathcal{W} also maps features to the output space with principal subspace projection using the Tucker decomposition of \mathcal{W} . Hence MLGP and tensor regression are essentially learning the same latent feature representations.

If we further consider the low-rank structure in the projection matrices, GP becomes degenerate. Degenerate GP has been shown in (Quinero-Candela and Rasmussen, 2005) to be equivalent to finite sparse linear models. Alternatively, we can interpret the low-rankness in MLGP and tensor regression using a constrained Bayesian inference approach (Koyejo and Ghosh, 2013). By minimizing the Kullback-Leibler (KL) divergence of the Bayesian posterior $\mathcal{N}(\mathbf{0}, \mathbf{K} + \mathbf{D})$ from any constructed GP prior $\mathcal{N}(\mathbf{0}, \mathbf{S})$, and assuming \mathbf{K} is low-rank, we have the following problem:

$$\min_{\mathbf{K}: \mathbf{K} \succeq \mathbf{0}, \text{rank}(\mathbf{K}) < R} \log \det[(\mathbf{K} + \mathbf{D}) \mathbf{S}^{-1}] + \text{tr}[(\mathbf{K} + \mathbf{D})^{-1} \mathbf{S}]$$

It turns out that the log-det of $\mathbf{K} + \mathbf{D}$ is a *smooth surrogate* for the rank of \mathbf{K} , which simultaneously minimizes the rank of \mathcal{W} . Therefore, the estimator for MLGP with low-rank kernel provides an approximate solution to the low-rank tensor regression problem. To this end, we have established the connections between tensor regression and Gaussian processes. Figure 2 depicts the graphical models of tensor regression, GP, and MLGP. It is evident that the parameter tensor in tensor regression maps to the covariance of the MLGP model. Latent tensor components become parameters of the covariance function.

We employ gradient-based optimization for Equation 5 to learn the hyper-parameters of MLGP. (see Appendix A.2 for details) Note that gradient-based optimization does not guarantee the orthonormality of the projection matrices. However, with a good initialization, we can still obtain reasonable approximations. As \mathbf{K} contains the Kronecker product and the low-rank structure, we can apply Woodbury matrix identity and exploit Kronecker properties to speedup the inference. The predictive distribution for the test data follows the standard GP regression procedure.

2.4 Theoretical Analysis

We study the theoretical properties of MLGP, which also shed light on the properties of existing tensor regression frameworks.

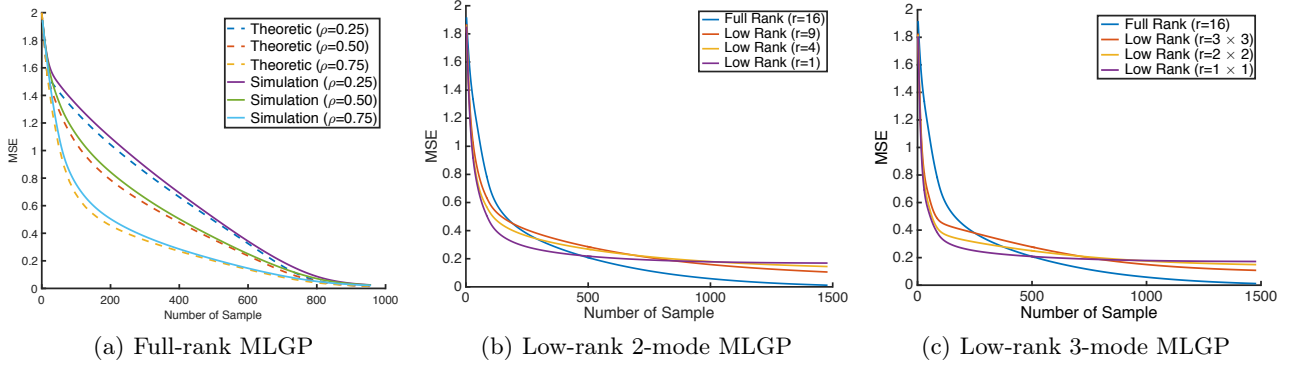


Figure 3: (a) Theoretical and numerically simulated learning curve for task correlation $\rho = 0.25, 0.25, 0.75$. (b) Learning curve for 2-mode MLGP with low-rank approximation $r = 9, 4, 1$. (c) Learning curve for 3-mode MLGP with low-rank approximation $r = 9, 4, 1$.

We first bound the excess risk of MLGP and derive the oracle inequality. Consider a tensor of functionals \mathcal{W} and define a space \mathcal{C}_N with sample size N :

$$\mathcal{C}_N = \{\mathcal{W} : \mathcal{W} = \mathcal{S} \times_1 \mathbf{U}_1 \times_2 \mathbf{U}_2 \times_3 \mathbf{U}_3, \|\mathcal{S}_{(1)}\|_* = \mathcal{O}\left(\frac{N}{T_2 T_3 + \log(T_1 T_2 T_3)}\right)^{1/4}\}$$

where $\|\cdot\|_*$ denotes the matrix nuclear norm. The following proposition states the oracle inequality:

Proposition 2.1. *Let $\hat{\mathcal{W}}$ be the estimator that minimizes the empirical risk $\hat{\mathcal{L}}(f(\mathcal{X}, \mathcal{W}); \mathcal{Y})$ over the space of functional tensors $\mathcal{W} \in \mathcal{C}_N$, then the excess risk, defined as \mathcal{L} satisfies:*

$$\mathcal{L}(\hat{\mathcal{W}}) - \inf_{\mathcal{W} \in \mathcal{C}_N} \mathcal{L}(\mathcal{W}) \xrightarrow{P} 0$$

Proof. Denote $\mathbb{E}[\text{cov}(\mathcal{Y}, \mathbf{U}_1(\mathcal{X}))] = \mathbf{\Sigma}(\mathbf{U}_1)$, we first bound the difference:

$$\mathcal{L}(\mathcal{W}) - \hat{\mathcal{L}}(\mathcal{W}) \leq C \max\{2, \|\mathcal{S}_{(1)}\|_*^2\} \|\mathbf{\Sigma}(\mathbf{U}_1) - \hat{\mathbf{\Sigma}}(\mathbf{U}_1)\|_2$$

The empirical risk is:

$$\begin{aligned} \mathcal{L}(\hat{\mathcal{W}}) - \mathcal{L}(\mathcal{W}^*) &\leq [\mathcal{L}(\hat{\mathcal{W}}) - \hat{\mathcal{L}}(\hat{\mathcal{W}})] - [\mathcal{L}(\mathcal{W}^*) - \hat{\mathcal{L}}(\mathcal{W}^*)] \\ &\leq \mathcal{O}\left(\|\mathcal{S}_{(1)}\|_*^2 \|\mathbf{\Sigma}(\mathbf{U}_1) - \hat{\mathbf{\Sigma}}(\mathbf{U}_1)\|_2\right) \end{aligned}$$

if we assume $\|\mathcal{S}_{(1)}\|_* = \mathcal{O}\left(\left(\frac{N}{T_2 T_3 + \log(T_1 T_2 T_3)}\right)^{1/4}\right)$, then $\mathcal{L}(\hat{\mathcal{W}}) - \mathcal{L}(\mathcal{W}^*) \leq \mathcal{O}(1)$. Details of the derivation are deferred to Appendix A.3. \square

This shows the estimation error tending to zero under a scaling assumption on the sample size N and the dimensions $\{T_m\}$. However, asymptotic results can only capture the large N regime and will not apply for finite sample sizes in practice. The following theorem states the explicit form of the non-asymptotic learning curve for the MLGP model under full-rank and low-rank scenarios:

Theorem 2.2. *Assume the eigenfunction decomposition for the data-dependent part of covariance $\phi(\mathbf{x})\mathbf{K}_1\phi(\mathbf{x}')^\top = \sum_i \lambda_i \psi_i(\mathbf{x})\psi_i(\mathbf{x}')^\top$, denote $\mathbf{\Lambda}$ as the diagonal matrix of $\{\delta_{i,j}\lambda_i\}$, the average case learning curve for MLGP of single task t satisfies*

$$\epsilon(N)_t = \text{tr} \mathbf{P}_{t_1, \dots, t_M} \left(\mathbf{\Lambda}'^{-1} + \sum_{s=1}^T \text{diag}\left(\frac{n_s}{\sigma_s^2 + \epsilon_s}\right) \mathbf{P}_{s_m} \right)^{-1}$$

when $\mathbf{\Lambda}'$ is full-rank

$$\epsilon(N)_t = \text{tr} \mathbf{P} \left(\mathbf{\Lambda}' - \left(\sum_{s=1}^T \text{diag}\left(\frac{\sigma_s^2 + \epsilon_s}{n_s}\right) \mathbf{P}_{s_m} + \mathbf{\Lambda}' \right)^{-1} \mathbf{\Lambda}'^2 \right)$$

when $\mathbf{\Lambda}'$ is rank-deficient, where $\mathbf{P}_{t_1, \dots, t_M}$, or \mathbf{P} is the linear operator that maps index t to a set of indices $\{t_m\}$, and $\mathbf{\Lambda}' = \otimes_{m=2}^M \mathbf{K}_m \otimes \mathbf{\Lambda}$.

Proof. The Bayes error, defined as $\hat{\epsilon} = \mathbb{E}_{\mathbf{x}}[(\mathbf{w} - \hat{\mathbf{w}})^2]$, has the following form for the low-rank case:

$$\hat{\epsilon} = \text{tr}(\mathbf{\Lambda}) - \text{tr}(\mathbf{D} + \mathbf{\Psi}\mathbf{\Lambda}\mathbf{\Psi}^\top)^{-1} \mathbf{\Psi}\mathbf{\Lambda}^2\mathbf{\Psi}^\top \quad (7)$$

and

$$\hat{\epsilon} = \text{tr}(\mathbf{\Lambda}^{-1} + \mathbf{\Psi}^\top \mathbf{D}^{-1} \mathbf{\Psi})^{-1} \quad (8)$$

for the full-rank case. And $\mathbf{\Lambda}$ and $\mathbf{\Psi}$ are the eigen-components of the covariance. The size of $\mathbf{\Lambda}$ is equal to the number of kernel eigenfunctions. When the GP has a non-degenerate kernel, $\mathbf{\Lambda}$ is full-rank. We can apply the Woodbury lemma to Equation 7, which yields a simplified version as in Equation 8.

Using method of characteristics (Sollich and Halees, 2002), we can obtain a corresponding lower bound for the average case learning curve:

$$\epsilon(N) = \text{tr}(\mathbf{\Lambda}) - \text{tr}\left(\frac{\sigma^2 + \epsilon}{N} \mathbf{I} + \mathbf{\Lambda}\right)^{-1} \mathbf{\Lambda}^2 \quad (9)$$

$$\epsilon(N) = \text{tr}\left(\mathbf{\Lambda}^{-1} + \frac{N}{\sigma^2 + \epsilon} \mathbf{I}\right)^{-1} \quad (10)$$

For MLGP, due to the task hierarchy, a task index t is projected to a set of indexes $\{t_m\}$ along different modes of a tensor. Define the projection on m th mode as $\mathbf{P}_{t_m} = \mathbf{e}_{t_m} \mathbf{e}_{t_m}^\top$, where \mathbf{e}_{t_m} is a unit vector with all zero but t_m th entry as one. Assume eigenfunction decomposition for the data-dependent part of covariance $\phi(\mathbf{x}) \mathbf{K}_1 \phi(\mathbf{x})^\top = \sum_i \lambda_i \psi_i(\mathbf{x}) \psi_i(\mathbf{x})^\top$, we have

$$\mathbf{K}_{jk} = \prod_{m=2}^M \mathbf{K}_{m,(\tau_j, \tau_k)} \sum_i \lambda_i \delta_{\tau_j, t} \psi_i(\mathbf{x}_j) \delta_{\tau_k, t} \psi_i(\mathbf{x}_k)^\top$$

$$\mathbf{K} = \Psi (\otimes_{m=2}^M \mathbf{K}_m \otimes \Lambda) \Psi^\top = \Psi \Lambda' \Psi^\top$$

where τ_j is the task index for j th example, further projected to the mode-wise indexes. Augmented eigenfunction matrix $\Psi_{j,it} = \delta_{\tau_j, t} \psi_i(\mathbf{x}_j)$ accounts for missing data, where the column index of Ψ runs over all eigenfunctions and all tasks. For task t , denote $k_t(\mathbf{x}, \cdot) = k(\mathbf{x}_t, \cdot)$

$$\mathbb{E}_{\mathbf{x}}[k_t(\mathbf{x}, \mathbf{X}) k(\mathbf{X}, \mathbf{x}_t)] = \Psi (\otimes_{m=2}^M (\mathbf{K}_m \mathbf{P}_{t_m} \mathbf{K}_m) \otimes \Lambda^2) \Psi^\top$$

where \mathbf{P}_{t_m} is the m th mode index for task t . The Bayes error can be written as:

$$\hat{\epsilon}_t = \mathbb{E}_{\mathbf{x}}[k_t(\mathbf{x}, \mathbf{x})] - \mathbb{E}_{\mathbf{x}}[k_t(\mathbf{x}, \mathbf{X}) (\mathbf{K} + \mathbf{D})^{-1} k_t(\mathbf{X}, \mathbf{x})]$$

For the first term

$$\begin{aligned} \mathbb{E}_{\mathbf{x}}[k_t(\mathbf{x}, \mathbf{x})] &= \prod_{m=2}^M \mathbf{e}_{t_m}^\top \mathbf{K}_m \mathbf{e}_{t_m} \mathbb{E}_{\mathbf{x}}[\phi(\mathbf{x}) \mathbf{K}_1 \phi(\mathbf{x})^\top] \\ &= \text{tr} \otimes_{m=2}^M \mathbf{P}_{t_m} \mathbf{K}_m \otimes \Lambda \end{aligned}$$

For the second term

$$\begin{aligned} &\text{tr} (\mathbf{D} + \Psi \Lambda' \Psi^\top)^{-1} \Psi (\otimes_{m=2}^M (\mathbf{K}_m \mathbf{P}_{t_m} \mathbf{K}_m) \otimes \Lambda^2) \Psi^\top \\ &\text{tr} (\mathbf{D} + \Psi \Lambda' \Psi^\top)^{-1} \Psi (\otimes_{m=2}^M (\mathbf{K}_m \mathbf{P}_{t_m} \mathbf{K}_m) \otimes \Lambda^2) \Psi^\top \end{aligned}$$

With $\otimes_m \mathbf{P}_{t_m} = \mathbf{P}_{t_1, \dots, t_M}$, compare Equation 11 with Equation 7, we have

$$\begin{aligned} \hat{\epsilon}_t &= \mathbf{P}_{t_1, \dots, t_M} \left(\text{tr}(\Lambda') - \text{tr}(\mathbf{D} + \Psi \Lambda' \Psi^\top)^{-1} \Psi \Lambda'^2 \Psi^\top \right) \\ \Lambda' &= \otimes_{m=2}^M \mathbf{K}_m \otimes \Lambda \end{aligned}$$

The Bayes error of task t is that of all tasks projected to each of its mode-wise task indices. Using an analogous method of characteristic curves, we can obtain a set of self-consistency equations for the learning curve of MLGP (see Appendix A.4 for details). \square

Theorem 2.2 indicates the performance dependency of MLGP, hence tensor regression, on the eigenvalues of the covariance function as well as the task correlation matrix. When the number of examples for all tasks becomes large, the Bayes errors $\hat{\epsilon}_t$ will be small and eventually be negligible compared to the noise variances σ_t . This also reflects a commonly accepted claim for

the asymptotic useless of multi-task learning: when the number of samples becomes large, the learning curves would come close to single task learning, except for the fully corrected case.

We further conduct numerical simulations to better understand the derived learning curve. Consider the case with 16 identical tasks, and set the task correlation matrix $\otimes_{m=2}^M \mathbf{K}_m$ to have ρ everywhere except for the principal diagonal elements. Assuming all the tasks are identical, Figure 3(a) compares the theoretic learning curve with the numerically simulated learning curve for different task relatedness. The theoretical learning curves generally lay slightly below the actual learning curves, providing a tight lower bound. With a higher value of ρ , tasks share higher interdependence, resulting in faster convergence w.r.t. Bayes error.

Figure 3(b) shows the learning curve for 2-modes MLGP with different low-rank approximations with $R_m = [1, 4, 9, 16]$. The low-rankness alleviates the noise variance error, leading to a faster convergence rate but eventually converges to a solution with a larger approximation gap. Figure 3(c) displays the learning curves for the 3-modes MLGP model, with the similar low-rank approximation. We observe that under the same rank assumption, the 3-mode MLGP imposes a stronger prior, leading to superior performances over 2-model MLGP with sparse observations.

2.5 Relation to Other Methods

It turns out that for multi-output regression, where all the tasks share the same inputs $\mathbf{X}_0 \in \mathbb{R}^{n_0 \times D}$, we can write $\mathbf{X} = \mathbf{X}_0 \otimes \mathbf{I}_T$, and noise becomes $\mathbf{D} = \text{diag}([\sigma_1, \dots, \sigma_T]) \otimes \mathbf{I}_{n_0}$. The covariance $\mathbf{K} = (\otimes_{m=2}^M \mathbf{K}_m) \otimes \phi(\mathbf{X}_0) \mathbf{K}_1 \phi(\mathbf{X}_0)^\top = (\otimes_{m=2}^M \mathbf{K}_m) \otimes \mathbf{K}_x$, where $\otimes_{m=2}^M \mathbf{K}_m$ encodes task similarity and \mathbf{K}_x is the kernel matrix over inputs \mathbf{X}_0 . When the number of modes $M = 2$, the model reduces to the multi-task Gaussian process (MTGP) model with free-form parameters (Bonilla et al., 2007). Here we factorize over Kronecker product operands as the low-rank approximation while MTGP uses Nyström approximation.

The multi-linear kernel $\phi(\mathbf{X}) (\otimes_{m=1}^M \mathbf{K}_m) \phi(\mathbf{X})^\top$ allows us to compute $\{\mathbf{K}_m\}_{m=1}^M$ separately, which avoids inversion of the big covariance matrix \mathbf{K} . This property has also been exploited in (Wilson et al., 2014) for multidimensional pattern extrapolation (GPatt). In there, inputs are assumed to be on a multidimensional grid $\mathbf{x} \in \mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_M$, the covariance matrix has decomposition $\mathbf{K} = \otimes_{m=1}^M \mathbf{K}_m$ where each factor \mathbf{K}_m is a kernel matrix over the space \mathcal{X}_m . The difference is that we use Kronecker products to learn multi-directional task correlations while GPatt performs kernel learning for each dimension of the inputs.

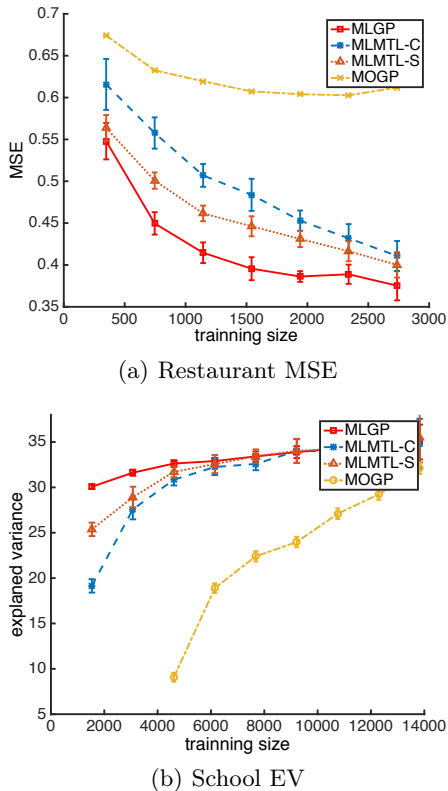


Figure 4: Multi-linear multi-task learning benchmark comparison (a) mean square error on the restaurant dataset. (b) expected variance on the school dataset. w.r.t sample size for MLGP and baselines.

3 Experiments

We conduct experiments for a series of tensor regression applications and demonstrate comparable prediction performances of MLGP with confidence intervals.

3.1 Multi-linear Multi-task learning

We evaluate on two benchmark datasets for MLMTL: school exam scores and restaurant ratings. School exam scores contain 15,362 students exam records with 21 features from 139 schools across 3 years. Restaurant ratings contain 3,483 rating records with 45 features from 138 consumers for 3 aspects.

We compare with the following baselines. (1) *MLMTC* (Romera-Paredes et al., 2013): latent trace norm optimization with alternating direction method of multipliers (ADMM) (2) *MLMTL-S* (Wimalawarne et al., 2014): scaled latent trace norm optimization with ADMM, and (3) *MOGP* (Alvarez and Lawrence, 2011): multi-output Gaussian process with DTC variational kernel. As all methods consider linear regression tasks, we use linear kernel MLGP as a fair comparison. For MOGP, we use 20 inducing points.

We randomly selected from a range of 10% to 80% of the entire data set as the training set. We selected 10% instances as the validation set and the rest was used as the test set. The regularization parameter for each norm was selected by minimizing the mean squared error on the validation set. We repeat the experiments for 10 times and average the results. All the baselines are the implementations of the original authors.

Figure 4(a) shows the restaurant rating prediction mean square error (MSE) for different methods over number of training samples. Figure 4(b) demonstrates the expected variance (EV) for the task of school exam score prediction. We observe superior performances of MLGP on restaurant data and comparable results for school data. In particular, when the size of the training data is small, MLGP shows significant advantages for both tasks. This justifies the benefit of MLGP for sparse observations.

3.2 Spatio-temporal Forecasting

Spatio-temporal forecasting has been shown to be a special case of tensor regression, with an additional spatial Laplacian matrix (Bahadori et al., 2014). We evaluate the spatio-temporal forecasting performance for 4 datasets reported in the original paper. For all the datasets, each variable is normalized by removing mean and dividing by variance. A third-order vector auto-regressive (VAR-3) model is employed for multi-variate time series modeling. We perform an 80/20 split along the time direction for training/testing and use validation to select the rank hyper-parameter.

Table 1 displays the forecasting MSE comparison. We compare with the reported best algorithm Greedy (Bahadori et al., 2014) for this task. We also include matrix multi-task learning with trace-norm regularization (MTL-Trace) to justify the benefit of the tensor-based approach. For all the 4 datasets, MLGP obtains similar prediction accuracy as Greedy. The predictive variance from MTGP directly provides empirical confidence intervals, which we append to the MSE.

To better understand the learned predictive distribution, we use a fine-grained USHCN dataset from California (Yu and Liu, 2016) and visualize the predictive variance of different locations on the map. We interpolate the variance values across locations and draw the contour plots. Figure 5 shows the contour plot for 54 locations of 5 climate variables. We observe interesting correlations between the predictive variance and geographical attributes. For example, precipitation (PRCP) and maximum temperature (TMAX) have relatively low-variance due to the subtropical climate in California. Snow depth (SNWD) demonstrates high variance along the mountains in Sierra Nevada.

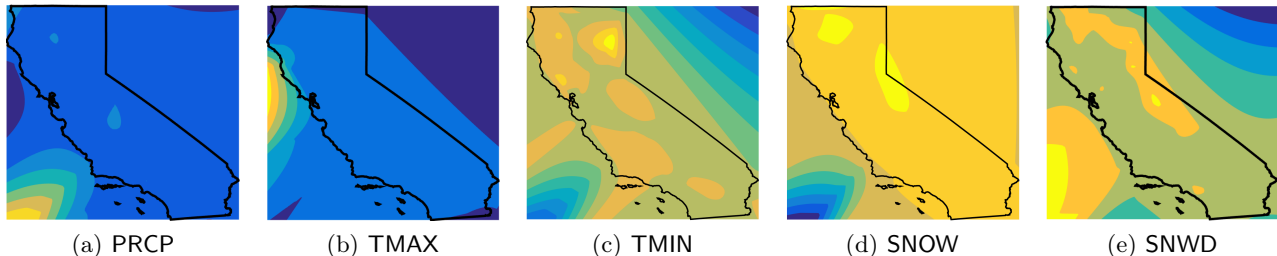


Figure 5: Contour plots for the MLGP predictive variance w.r.t precipitation PRCP, max temperate TMAX min temperate TMIN, snowfall SNOW and snow depth SNWD. Yellow is high variance and blue means low variance.

Table 1: Mean square error comparison of MLGP and baselines for spatio-temporal forecasting on 4 datasets with 10% testing set. Tensor regression models use VAR-3 with moving window.

| Dataset | MLGP | Greedy | MLMTL-C | MLMTL-S | MTL-Trace |
|----------|---------------------|--------|---------|---------|-----------|
| USHNC-US | 0.8973 \pm 0.0008 | 0.9069 | 0.9528 | 0.9543 | 0.9273 |
| CCDS | 0.8498 \pm 0.0013 | 0.8325 | 0.9105 | 0.8394 | 0.8632 |
| FSQ | 0.1248 \pm 0.0006 | 0.1223 | 0.1495 | 0.1243 | 0.1245 |
| YELP | 1.0725 \pm 0.0007 | NA | 1.0857 | 1.0876 | 1.0736 |

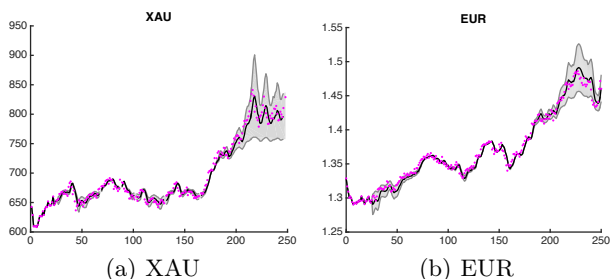


Figure 6: Predictive mean (solid line) and variance (shaded area) for foreign exchange rate of XAU and EUR using from MLGP for 50 time step ahead forecasting. Magenta points are observations.

3.3 Multi-output regression

Multiple output regression concerns with the case when predictor tensor is shared among all of the responses. One such application is the foreign exchange rate prediction task (Alvarez and Lawrence, 2011). The original dataset contains 3 precious metals and 12 international currencies. To show the benefit of exploiting multi-directional task interdependence, we select the foreign exchange rate of 6 international currencies (EUR, GBP, CHF, JPY, HKD, KRW) and 3 precious metals (gold, silver, and platinum), which forms three groups: precious metal, European currency and Asian currency. The dataset consists of all the data available for the 251 working days in the year of 2007.

We use the VAR-3 model for all the low-rank tensor regression baselines. MLGP achieves 0.0563 MSE while

best performance of low-rank tensor regression is 0.0657. These results are slightly worse than 0.0301 of MOGP with PITC approximation. However, since MLGP does not require all the responses to be of equal size, it runs much faster than MOGP, which involves a missing value imputation step to satisfy the size constraint. To further interpret the learned model, we plot out the predictive mean and variance together with observations in Figure 6. We observe high predictive variance whenever the time series encounters sharp changes.

4 Discussion and Conclusion

In this paper, we establish interesting connections between tensor regression and Gaussian processes. We develop a probabilistic counterpart: multi-linear Gaussian processes (MLGP). With the low-rank constraint, the Bayesian estimator of MLGP learns a smooth surrogate for the low-rank structure in tensor regression. Theoretical analysis shows its performance dependency on the eigenvalues of the covariance matrix and task correlation. Comparable (if not better) performance are observed in a series of real-world applications.

This relationship hints upon our choice of tools for multi-way data analysis. Tensor regression is fast and simple to implement. It is guaranteed to output orthonormal basis of the latent subspaces but does not generate confidence intervals. MLGP, on the other hand, can better handle sparse observations, and is more versatile with kernels. In terms of future directions, one interesting question is to study the robustness of both methods under adversarial corruptions.

5 Acknowledgment

We thank the anonymous reviewers for their helpful feedback and comments. The research was sponsored by NSF Award #1254206 and #1539608. The views and conclusions are those of the authors and should not be interpreted as representing the official policies of the funding agency, or the U.S. Government.

References

- Mauricio A Alvarez and Neil D Lawrence. Computationally efficient convolved multiple output gaussian processes. *The Journal of Machine Learning Research*, 12:1459–1500, 2011.
- Mohammad Taha Bahadori, Rose Yu, and Yan Liu. Fast multivariate spatio-temporal analysis via low rank tensor learning. In *Advances in Neural Information Processing Systems*, pages 3491–3499, 2014.
- Edwin V Bonilla, Kian M Chai, and Christopher Williams. Multi-task gaussian process prediction. In *Advances in neural information processing systems*, pages 153–160, 2007.
- Wei Chu and Zoubin Ghahramani. Probabilistic models for incomplete multi-dimensional arrays. In *International Conference on Artificial Intelligence and Statistics*, pages 89–96, 2009.
- Rajarshi Guhaniyogi, Shaan Qamar, and David B Dunson. Bayesian tensor regression. *arXiv preprint arXiv:1509.06490*, 2015.
- Masaaki Imaizumi and Kohei Hayashi. Doubly decomposing nonparametric tensor regression. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 727–736, 2016.
- Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- Oluwasanmi Koyejo and Joydeep Ghosh. Constrained bayesian inference for low rank multitask learning. In *Uncertainty in Artificial Intelligence*, page 341. Citeseer, 2013.
- Neil D Lawrence. Gaussian process latent variable models for visualisation of high dimensional data. In *NIPS*, 2004.
- Alexander Novikov, Dmitrii Podoprikin, Anton Osokin, and Dmitry P Vetrov. Tensorizing neural networks. In *Advances in Neural Information Processing Systems*, pages 442–450, 2015.
- Joaquin Quinonero-Candela and Carl Edward Rasmussen. Analysis of some methods for reduced rank gaussian process regression. In *Switching and Learning in Feedback Systems*, pages 98–127. Springer, 2005.
- Guillaume Rabusseau and Hachem Kadri. Low-rank regression with tensor responses. In *Advances in Neural Information Processing Systems*, pages 1867–1875, 2016.
- Carl Edward Rasmussen. Gaussian processes for machine learning. 2006.
- Bernardino Romera-Paredes, Hane Aung, Nadia Bianchi-Berthouze, and Massimiliano Pontil. Multilinear multitask learning. In *Proceedings of The 30th International Conference on Machine Learning*, pages 1444–1452, 2013.
- Peter Sollich and Anason Halees. Learning curves for gaussian process regression: Approximations and bounds. *Neural computation*, 14(6):1393–1428, 2002.
- Taiji Suzuki. Convergence rate of bayesian tensor estimator and its minimax optimality. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 1273–1282, 2015.
- Andrew Wilson, Elad Gilboa, John P Cunningham, and Arye Nehorai. Fast kernel learning for multidimensional pattern extrapolation. In *Advances in Neural Information Processing Systems*, pages 3626–3634, 2014.
- Kishan Wimalawarne, Masashi Sugiyama, and Ryota Tomioka. Multitask learning meets tensor factorization: task imputation via convex optimization. In *Advances in Neural Information Processing Systems*, pages 2825–2833, 2014.
- Zenglin Xu, Feng Yan, and Yuan Qi. Bayesian nonparametric models for multiway data analysis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 37(2):475–487, 2015.
- Rose Yu and Yan Liu. Learning from multiway data: Simple and efficient tensor regression. In *Proceedings of the 33rd International Conference on Machine Learning (ICML-16)*, 2016.
- Hua Zhou, Lexin Li, and Hongtu Zhu. Tensor regression with applications in neuroimaging data analysis. *Journal of the American Statistical Association*, 108(502):540–552, 2013.