

# Rethinking Bayesian Learning for Data Analysis: The Art of Prior and Inference in Sparsity-Aware Modeling

Lei Cheng<sup>‡</sup>, Feng Yin<sup>§</sup>, Sergios Theodoridis<sup>\*</sup>, Sotirios Chatzis<sup>†</sup>, and Tsung-Hui Chang<sup>§</sup>

<sup>‡</sup> College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, China

<sup>§</sup> School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, China

<sup>\*</sup> Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Greece  
and Department of Electronics Systems, Aalborg University, Denmark

<sup>†</sup> Department of Electrical & Computer Engineering and Informatics, Cyprus University of Technology, Cyprus

Correspondence Author: Feng Yin, yinfeng@cuhk.edu.cn

## Abstract

Sparse modeling for signal processing and machine learning, in general, has been at the focus of scientific research for over two decades. Among others, supervised sparsity-aware learning comprises two major paths paved by: a) discriminative methods that establish direct input-output mapping based on a regularized cost function optimization, and b) generative methods that learn the underlying distributions. The latter, more widely known as Bayesian methods, enable uncertainty evaluation with respect to the performed predictions. Furthermore, they can better exploit related prior information and also, in principle, can naturally introduce robustness into the model, due to their unique capacity to marginalize out uncertainties related to the parameter estimates. Moreover, hyper-parameters (tuning parameters) associated with the adopted priors, which correspond to cost function regularizers, can be learnt via the training data and not via costly cross-validation techniques, which is, in general, the case with the discriminative methods. To implement sparsity-aware learning, the crucial point lies in the choice of the function regularizer for discriminative methods and the choice of the prior distribution for Bayesian learning. Over the last decade or so, due to the intense research on deep learning, emphasis has been put on discriminative techniques. However, a come back of Bayesian methods is taking place that sheds new light on the design of deep neural networks, which also establish firm links with Bayesian models, such

Lei Cheng and Feng Yin contribute equally.

as Gaussian processes, and, also, inspire new paths for unsupervised learning, such as Bayesian tensor decomposition.

The goal of this article is two-fold. First, to review, in a unified way, some recent advances in incorporating sparsity-promoting priors into three highly popular data modeling/analysis tools, namely deep neural networks, Gaussian processes, and tensor decomposition. Second, to review their associated inference techniques from different aspects, including: evidence maximization via optimization and variational inference methods. Challenges such as small data dilemma, automatic model structure search, and natural prediction uncertainty evaluation are also discussed. Typical signal processing and machine learning tasks are considered, such as time series prediction, adversarial learning, social group clustering, and image completion. Simulation results corroborate the effectiveness of the Bayesian path in addressing the aforementioned challenges and its outstanding capability of matching data patterns automatically.

## I. INTRODUCTION

Over the past three decades or so, machine learning has been gradually established as the umbrella name to cover methods whose goal is to extract valuable information and knowledge from data, and then use it to make predictions [1]. Machine learning has been extensively applied to a wide range of disciplines, such as signal processing, data mining, communications, finance, bio-medicine, robotics, to name but a few. The majority of the machine learning methods first rely on adopting a parametric model to describe the data at hand, and then an inference/estimation technique to derive estimates that describe the unknown model parameters. In the discriminative methods, point estimates of the involved parameters are obtained via cost function optimization. In contrast, by practicing the Bayesian philosophy, one can infer the underlying statistical distributions that describe the unknown parameters given the observed data, and, thus, provide a generative mechanism that models the random process that generates the data.

For the newcomers to machine learning, the discriminative (also referred to as cost function optimization) perspective might be more straightforward. It first formulates a task that quantifies the overall deviation between the observed target data and the model predictions, and then solves it for the point parameter estimates via an optimization algorithm. On the contrary, the generative (Bayesian) perspective, which aims to reveal the generative process and the statistical properties of the observed data, sounds more complicated due to some “jargon” terms such as prior, likelihood, posterior, and evidence. Nevertheless, machine learning under the Bayesian perspective is gaining in popularity recently due to the *comparative advantages* that spring from the nature of the statistical modeling and the extra information returned by the posterior distributions. This article aims at demystifying the philosophy that underlies the Bayesian techniques, and then review, in a unified way, recent advances of *Bayesian sparsity-aware learning* for three analysis tools of high current interest. In the Bayesian framework, model sparsity is implemented

via sparsity-promoting priors that lead to automatic model determination by optimally sparsifying an, originally, over-parameterized model. The goal is to optimally predict the order of the system that corresponds to the best trade-off between accuracy and complexity, with the aim to combat overfitting, in line with the general concept of regularization. However, in Bayesian learning, all the associated (hyper-)parameters, which control the degree of regularization, can be optimally obtained via the training set during the learning phase. It is hoped that this article can help the newcomers grasp the essence of Bayesian learning, and at the same time, provide the experts with an update of some recent advances developed for different data modeling and analysis tasks.

In particular, we will focus on Bayesian sparsity-aware learning for three popular data modeling and analysis tools, namely the deep neural networks (DNNs), Gaussian processes (GPs), and tensor decomposition, that have promoted intelligent signal processing applications. Some typical examples are as follows.

In the supervised learning front with over-parameterized DNNs, novel data-driven mechanisms have been proposed in [2]–[6] to intelligently prune redundant neuron connections without human assistance. In a similar vein, in [7]–[9], sparsity-promoting priors have been used in the context of the GPs that give rise to optimal and interpretable kernels that are capable of identifying a sparse subset of effective frequency components automatically. In the unsupervised learning front, some advanced works on tensor decomposition, e.g., [10]–[15], have shown that sparsity-promoting priors are able to unravel the few underlying interpretable components in a completely tuning-free fashion. Such techniques have found various signal processing applications, including data classification [2], [5], [6], adversarial learning [3], [4], time-series prediction [7]–[9], [16], blind source separation [10], [13], [17], image completion [12], [14], [15], and wireless communications [18].

The aforementioned references address two state-of-the-art challenges: 1) *The art of prior*: how should the fundamental sparsity-promoting priors be chosen and tailored to fit modern data-driven models with complex structures? 2) *The art of inference*: how can recent optimization theory and stochastic approximation techniques be leveraged to design fast, accurate, and scalable inference algorithms? This tutorial-style article aims to give a unified treatment on the underlying common ideas and techniques to offer concrete answers to the above questions. It is yet the goal of this article to provide a comprehensive review of such sparsity-promoting techniques. On the one hand, we will introduce some newly proposed sparsity-promoting priors, as well as various salient ones that, although being powerful, had never been used before in our target models. On the other hand, we will showcase some recent developments of the associated inference algorithms. For readers with different backgrounds and familiarity with Bayesian statistics, we provide a roadmap in Fig. 1 to facilitate their reading.

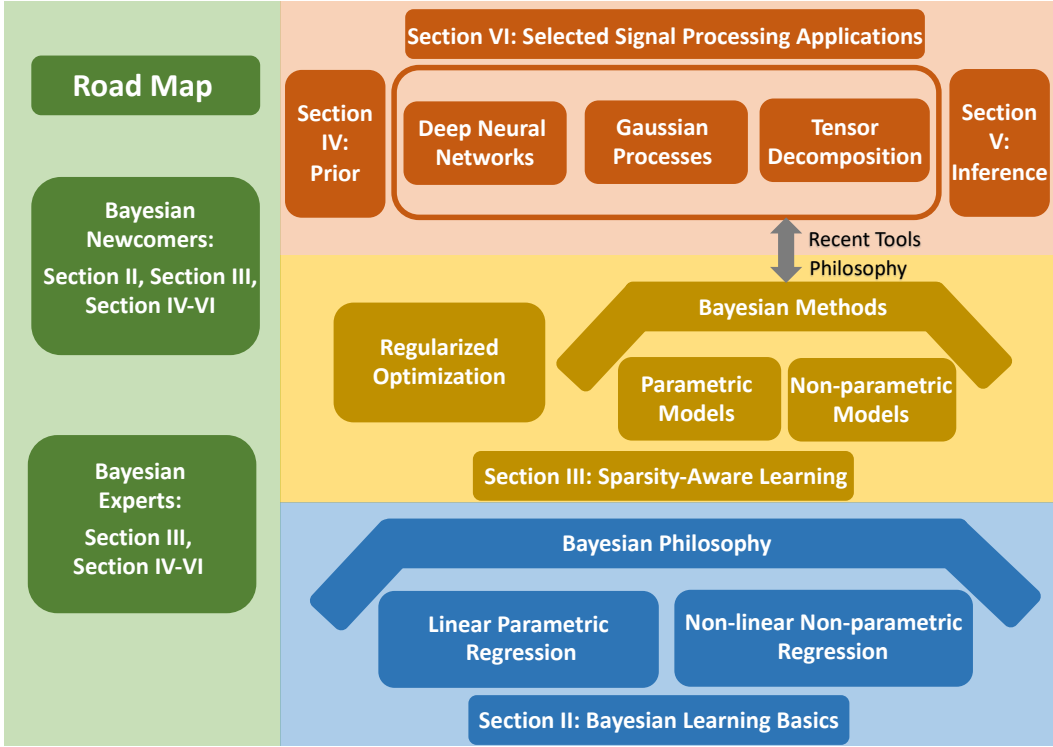


Fig. 1. The organization of this article and road map for the readers.

The remaining sections of this article are organized as follows. In Section II, we introduce some Bayesian learning basics, aiming to let the readers easily follow the main concepts, jargon terms, and math notations. In Section III, we first review two different paths (the regularized optimization and Bayesian paths), and further introduce some sparsity-promoting priors along the Bayesian path. In Section IV, we demonstrate how to integrate the introduced sparsity-promoting priors into three prevailing data analysis tools, i.e., the DNNs, GPs, and tensor decomposition. For the reviewed sparsity-aware learning models, we further introduce their associated inference methods in Section V. Various signal processing applications of high current interests enabled by the aforementioned models are exemplified in Section VI. Finally, we conclude the article and bring up some potential future research directions in Section VII.

## II. BAYESIAN LEARNING BASICS

In this section, we first provide some touches on the philosophy of Bayesian learning in Section II-A, and use *Bayesian linear regression* as an example to elucidate different symbol notations, terminology and unique features of Bayesian learning in Section II-B. Then, we discuss extensions to the non-linear and

non-parametric cases, shedding light on the connections between simple linear regression and advanced Gaussian process regression in Section II-C.

### A. Bayesian Philosophy Basics

1) *Bayes' Theorem*: Let  $\mathcal{D}$  be the observed (training) dataset and  $\mathcal{M}$  be the underlying model that is assumed to generate the data. For simplicity, we start our treatment with models that are parameterized in terms of a set of unknown parameters  $\boldsymbol{\theta} \in \mathbb{R}^{L \times 1}$ , where  $\mathbb{R}$  is the set of real numbers. By the definition of parametric models, the dimension  $L$  is pre-selected and fixed [1]. According to the Bayesian philosophy, these parameters are treated as random variables. Their randomness does not imply a random nature of these parameters, but essentially encodes our uncertainty with respect to their true (yet unknown) values, see related discussions in, e.g., [1]. First, in Bayesian modeling, we assume that the set of unknown random parameters is described by a prior distribution, i.e.,  $\boldsymbol{\theta} \sim p_{\mathcal{M}}(\boldsymbol{\theta}; \boldsymbol{\eta}_p)$ , which encodes our *prior* belief in  $\boldsymbol{\theta}$ ; that is, it encodes our uncertainty prior to receiving the dataset  $\mathcal{D}$ . As we are going to see soon, this corresponds to regularizing the learning task, since it will bias the solution that we seek towards certain regions in the parameter space. The prior  $p_{\mathcal{M}}(\boldsymbol{\theta}; \boldsymbol{\eta}_p)$  is specified via a set of deterministic yet unknown *hyper-parameters* (tuning parameters) stacked together in a vector and denoted by  $\boldsymbol{\eta}_p$ . The second quantity that is assumed to be known is the conditional distribution that describes the data given the values of the parameters,  $\boldsymbol{\theta}$ , which for the specific observed dataset  $\mathcal{D}$  is known as the *likelihood*  $p_{\mathcal{M}}(\mathcal{D}|\boldsymbol{\theta})$ .<sup>1</sup>

Having selected the likelihood and the prior distribution function, the goal of Bayesian inference is to infer (estimate) the *posterior* distribution of the parameters given the observations, i.e.,  $p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta})$ , that comprises the update of the prior assumption encoded in  $p_{\mathcal{M}}(\boldsymbol{\theta}; \boldsymbol{\eta}_p)$  after digesting the dataset  $\mathcal{D}$ . This process can be elegantly described by the celebrated *Bayes' theorem*, e.g., the one given in [1]:

$$p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta}) = \frac{p_{\mathcal{M}}(\mathcal{D}|\boldsymbol{\theta})p_{\mathcal{M}}(\boldsymbol{\theta}; \boldsymbol{\eta}_p)}{p_{\mathcal{M}}(\mathcal{D}; \boldsymbol{\eta})}. \quad (1)$$

Note that  $\boldsymbol{\eta}$  includes both the hyper-parameters associated with the prior,  $\boldsymbol{\eta}_p$ , and some extra hyper-parameters involved in the likelihood function  $p_{\mathcal{M}}(\mathcal{D}|\boldsymbol{\theta})$ , which are omitted for notation brevity.

The Bayes' theorem solves for the *inverse problem* that is associated with any machine learning task. The *forward problem* is an easy one. Given the model  $\mathcal{M}$  and the values of the associated parameters  $\boldsymbol{\theta}$ , one can easily generate the output observations  $\mathcal{D}$  from the conditional distribution  $p_{\mathcal{M}}(\mathcal{D}|\boldsymbol{\theta})$ . The

<sup>1</sup>Throughout this paper, we use “;”, i.e.,  $p(x; \eta)$  if  $\eta$  is a deterministic parameter to be optimized or pre-selected by the user; and we use “|”, i.e.,  $p(x|\eta)$  if  $\eta$  is a random variable or a hyper-parameter treated as a random variable; that is, if the distribution is conditional on another random variable.

task of machine learning is the opposite and a more difficult one. Given the observed data  $\mathcal{D}$ , the task is to estimate/infer the model  $\mathcal{M}$ . This is known as the *inverse problem*, and Bayes' theorem applied to the machine learning task does exactly that. It relates the inverse problem (posterior) to the forward one (likelihood). All one needs for this “update” is to assume a prior and also to obtain an estimate of the distribution associated with the data, which comprises the denominator in (1). The latter term and the related information are neglected in the discriminative models, hence important information is not taken into account, see the discussion in, e.g., [1], [19].

Occasionally, we may need a point estimate of the model parameters as the intermediate result, and there are two commonly used estimates that can be computed from the posterior distribution,  $p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta})$ . Assuming that  $\boldsymbol{\theta}$  is known or an estimate is available, the first one is known as the *maximum-a-posteriori* (MAP) estimate and the other one as the *minimum-mean-squared-error* (MMSE) estimate, concretely [1],

$$\hat{\boldsymbol{\theta}}_{\text{MAP}} = \arg \max_{\boldsymbol{\theta}} p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta}), \quad (2)$$

$$\hat{\boldsymbol{\theta}}_{\text{MMSE}} = \int \boldsymbol{\theta} \cdot p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta}) d\boldsymbol{\theta}. \quad (3)$$

2) *Evidence Maximization for Hyper-parameter Learning*: In the prior distribution,  $p_{\mathcal{M}}(\boldsymbol{\theta}; \boldsymbol{\eta}_p)$ , the hyper-parameters,  $\boldsymbol{\eta}_p$ , could be either pre-selected according to the side information at hand, or learnt from the observed dataset  $\mathcal{D}$ . In Bayesian learning, the latter path is followed favorably. One popular alternative is to select the full set of hyper-parameters  $\boldsymbol{\eta}$  to be the most compatible with the observed dataset  $\mathcal{D}$ , which can be naturally formulated as the following so-called *evidence maximization*:

$$\max_{\boldsymbol{\eta}} \log p_{\mathcal{M}}(\mathcal{D}; \boldsymbol{\eta}), \quad (4)$$

where

$$p_{\mathcal{M}}(\mathcal{D}; \boldsymbol{\eta}) = \int p_{\mathcal{M}}(\mathcal{D}|\boldsymbol{\theta})p_{\mathcal{M}}(\boldsymbol{\theta}; \boldsymbol{\eta}_p)d\boldsymbol{\theta} \quad (5)$$

is known as the model *evidence*, since it measures the plausibility of the dataset  $\mathcal{D}$  given the hyper-parameters  $\boldsymbol{\eta}$ . Note that the evidence depends on the model itself and not on any specific value of the parameters  $\boldsymbol{\theta}$ , which have been integrated out (marginalized). *This is a crucial difference compared with the discriminative methods*. As it can be shown, the evidence maximization problem (4) involves a trade-off between accuracy (the achieved likelihood value) and model complexity, in line with Occam's razor rule [20], [1]. This allows computation of the model hyper-parameters  $\boldsymbol{\eta}$  directly from the observed dataset  $\mathcal{D}$ . At this point, recall that one of the major difficulties associated with machine learning, and the inverse problems in general, is *overfitting*. That is, if the model is too complex with respect to the number of training data samples, then the estimated models learn the specificities of the given training data and cannot generalize well when dealing with new unseen (test) data.

The use of regularization in the discriminative methods and priors in the Bayesian ones try to achieve the best trade-off between accuracy (fitting to the observed data) and generalization that heavily depends on the complexity of the model, see, e.g., [1], [19] for further discussions. Furthermore, note that in the Bayesian context, *model complexity* is interpreted from a broader view, since it depends not only on the number of parameters but also on the shape (e.g., variance and skewness) of the involved distributions of  $\theta$ , see e.g., [1], [19] for in-depth discussions. For example, under a broad enough Gaussian prior for the model parameters,  $\theta$ , and some limiting properties, it can be shown that the evidence in (5) results in the well-known Bayesian information criterion (BIC) for model selection [21], [1], which has the form:

$$\log p_{\mathcal{M}}(\mathcal{D}; \boldsymbol{\eta}) = \log p_{\mathcal{M}}(\mathcal{D}|\hat{\boldsymbol{\theta}}_{\text{MAP}}) - \frac{L}{2} \log N, \quad (6)$$

where the first term on the right-hand side is the accuracy (best likelihood fit) term and the second is the complexity term that “competes” in a trade-off fashion while maximizing the evidence, see, e.g., [22], [1] for further discussion. In (6),  $L$  denotes the number of unknown parameters in  $\theta$  and  $N$  is the size of the training data. A more recent interpretation of this trade-off, in the context of over-parameterized DNNs, is provided in [23], where the prior is viewed as the *inductive bias* that favors certain datasets.

3) *Marginalization for Prediction:* The learnt posterior  $p_{\mathcal{M}}(\theta|\mathcal{D}; \boldsymbol{\eta})$  provides uncertainty information about  $\theta$ , i.e., the plausibility of each possible  $\theta$  to be endorsed by the observed dataset  $\mathcal{D}$ , and it can be used to forecast an unseen dataset,  $\mathcal{D}_{\text{new}}$ , via *marginalization*:

$$p_{\mathcal{M}}(\mathcal{D}_{\text{new}}|\mathcal{D}; \boldsymbol{\eta}) = \int p_{\mathcal{M}}(\mathcal{D}_{\text{new}}|\boldsymbol{\theta})p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta})d\boldsymbol{\theta}. \quad (7)$$

From (7), Bayesian prediction can be interpreted as the weighted average of the predicted probability  $p_{\mathcal{M}}(\mathcal{D}_{\text{new}}|\boldsymbol{\theta})$  among all possible model configurations, each of which is specified by different model parameters,  $\boldsymbol{\theta}$ , and weighted by the respective posterior  $p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta})$ .<sup>2</sup> In other words, prediction *does not* depend on a *specific* point estimate of the unknown parameters, which equips Bayesian methods with great potential for more robust predictions against the estimation error of  $\boldsymbol{\theta}$ , see, e.g., [1], [19].

In summary, in light of the Bayes’ theorem, the four quantities (i.e., prior  $p_{\mathcal{M}}(\boldsymbol{\theta}; \boldsymbol{\eta}_p)$ , likelihood  $p_{\mathcal{M}}(\mathcal{D}|\boldsymbol{\theta})$ , posterior  $p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta})$  and evidence  $p_{\mathcal{M}}(\mathcal{D}; \boldsymbol{\eta})$ ) give a new perspective on the *inverse problem*. The resulting method combines the strength of the selected priors and the likelihood of the observed data to provide a corresponding posterior. The success of such an inference process strongly relies on the following three steps. First, incorporating a prior for each unknown model parameter/function enables one to naturally encode a desired structure into Bayesian learning. As it will be demonstrated

<sup>2</sup>In this paper, the unseen dataset  $\mathcal{D}_{\text{new}}$  is assumed to be statistically independent of the training dataset  $\mathcal{D}$ . Therefore,  $p_{\mathcal{M}}(\mathcal{D}_{\text{new}}|\boldsymbol{\theta}) = p_{\mathcal{M}}(\mathcal{D}_{\text{new}}|\mathcal{D}, \boldsymbol{\theta})$ .

in the rest of this article, a prior can be imposed to both parametric models with a fixed number of unknown parameters and non-parametric models that comprise unknown functions and/or an unknown set of parameters whose number is not fixed but it varies with the size of the dataset. Second, through evidence maximization, one can optimize the set of hyper-parameters that is associated with the selected Bayesian learning model to obtain enhanced generalization performance. Finally, marginalization ensures robust prediction and generalization performance by averaging over an ensemble of predictions using all possible parameter/function estimates weighted by the corresponding posterior probability. These three aspects will be discussed in detail in the following sections.

### B. Bayesian Linear Parametric Regression: A Pedagogic Example

Before moving to our next topics on more advanced Bayesian data analysis, we introduce the *Bayesian linear regression* model as an example to further elaborate the terminology and concepts discussed previously. It also serves as the cornerstone for the two recent supervised learning tools, namely the Bayesian neural networks and GP models to be elaborated in the following subsections.

1) *Linear Regression*: In statistics, the term “regression” refers to seeking the relationship between a *dependent* random variable,  $y$ , which is usually considered as the response of a system, and the associated *input/independent* variables,  $\mathbf{x} = [x_1, x_2, \dots, x_L]^T$ . When the system is modeled as a linear combiner with an additive disturbance or noise term  $v_n$ , the relationship between  $y_n$  and  $\mathbf{x}_n$  of the  $n$ -th data sample can be expressed as:

$$y_n = \boldsymbol{\theta}^T \mathbf{x}_n + v_n, \quad \forall n \in \{1, 2, \dots, N\}, \quad (8)$$

which specifies the linear regression task. For simplicity, we assume that the additive noise terms  $\{v_n\}$  are independently and identically distributed (i.i.d.) Gaussian with zero mean and variance  $\beta^{-1}$ , i.e.,  $\{v_n\} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(v_n; 0, \beta^{-1})$ , where  $\beta$  (i.e., inverse of the variance) is called “precision” in statistics and machine learning. The task of linear regression is to learn the weight parameters  $\boldsymbol{\theta} = [\theta_1, \theta_2, \dots, \theta_L]^T$  from the training/observed dataset  $\mathcal{D} \triangleq \{\mathbf{X}, \mathbf{y}\}$ , where the input matrix  $\mathbf{X} \triangleq [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]^T \in \mathbb{R}^{N \times L}$ , and the output vector  $\mathbf{y} \triangleq [y_1, y_2, \dots, y_N]^T$ .

2) *Bayesian Learning*: For the linear regression task, we take a Bayesian perspective by treating the unknown parameters  $\boldsymbol{\theta}$  as a random vector. As introduced in Section II-A, the inverse problem can be solved via the Bayes’ theorem after specifying the following four quantities.

■ Likelihood. The easiest one to derive is the likelihood function, which describes the forward problem of linear regression. Owing to the Gaussian and independence properties of the noise terms  $\{v_n\}$ , the



following Gaussian likelihood function can be easily obtained:

$$p_{\mathcal{M}}(\mathcal{D}|\boldsymbol{\theta}) = \prod_{n=1}^N \mathcal{N}(y_n; \boldsymbol{\theta}^T \mathbf{x}_n, \beta^{-1}). \quad (9)$$

■ Prior. Then, we specify a prior on the unknown parameters  $\boldsymbol{\theta}$ . For mathematical tractability, we adopt an i.i.d. Gaussian distribution as the prior:

$$p_{\mathcal{M}}(\boldsymbol{\theta}; \boldsymbol{\eta}_p) = \prod_{l=1}^L \mathcal{N}(\theta_l; 0, \alpha_l^{-1}), \quad (10)$$

where  $\alpha_l$  is the precision associated with each  $\theta_l$ , and  $\boldsymbol{\eta}_p = \boldsymbol{\alpha} \triangleq [\alpha_1, \alpha_2, \dots, \alpha_L]^T$  represents the hyper-parameters associated with the prior.

■ Evidence. After substituting the prior (10) and the likelihood (9) into (5), and performing the integration, we can derive the following Gaussian evidence:

$$p_{\mathcal{M}}(\mathcal{D}; \boldsymbol{\eta}) = \mathcal{N}(\mathbf{y}; \mathbf{0}, \beta^{-1} \mathbf{I} + \mathbf{X} \mathbf{A}^{-1} \mathbf{X}^T), \quad (11)$$

where the diagonal matrix  $\mathbf{A} \triangleq \text{diag}\{\boldsymbol{\alpha}\}$  and  $\mathbf{I}$  denotes the identity matrix. Here, we have  $\boldsymbol{\eta} = [\boldsymbol{\eta}_p^T, \beta]^T$ .

■ Posterior. Inserting the prior (10), the likelihood (9) and the evidence (5) into the Bayes' theorem (1), the posterior can be shown to be the Gaussian distribution:

$$p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta}) = \mathcal{N}(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (12)$$

where

$$\boldsymbol{\mu} = \beta \boldsymbol{\Sigma} \mathbf{X}^T \mathbf{y}, \quad (13a)$$

$$\boldsymbol{\Sigma} = (\mathbf{A} + \beta \mathbf{X}^T \mathbf{X})^{-1}. \quad (13b)$$

Once again, taking the above linear regression as a concrete example, we further demonstrate the merits of Bayesian learning in general.

■ Merit 1: Parameter Learning with Uncertainty Quantification. Using the Bayes' theorem, the posterior in (12) not only provides a point estimate  $\boldsymbol{\mu}$  in (13a) for the unknown parameters  $\boldsymbol{\theta}$ , but also provides a covariance matrix  $\boldsymbol{\Sigma}$  in (13b) that describes to which extent the posterior distribution is centered around the point estimate  $\boldsymbol{\mu}$ . In other words, it quantifies our uncertainty about the parameter estimate, which cannot be naturally obtained in any discriminative method. For the above example, we have  $\hat{\boldsymbol{\theta}}_{\text{MAP}} = \hat{\boldsymbol{\theta}}_{\text{MMSE}}$  because the posterior distribution  $p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta})$  follows a unimodal Gaussian distribution. Of course, Frequentist methods can also construct uncertainty region/confidence intervals by taking a few extra steps once the parameter estimates have been obtained. However, the Bayesian method provides, in one go, the posterior distribution of the model parameters, from which both a point estimate as well as the uncertainty region can be optimally derived via the learning optimization step.

■ Merit 2: Robust Prediction via Marginalization. After substituting (12) and (9) tailored to new observations into (7), the posterior/predictive distribution for a novel input  $\mathbf{x}_*$  is:

$$\begin{aligned} p_{\mathcal{M}}(y_*|\{\mathbf{X}, \mathbf{y}\}) &= \int \mathcal{N}(y_*; \boldsymbol{\theta}^T \mathbf{x}_*, \beta^{-1}) \mathcal{N}(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) d\boldsymbol{\theta} \\ &= \mathcal{N}(y_*; \boldsymbol{\mu}^T \mathbf{x}_*, \beta^{-1} + \mathbf{x}_*^T \boldsymbol{\Sigma} \mathbf{x}_*). \end{aligned} \quad (14)$$

The predicted value of  $y_*$  can be acquired via  $\boldsymbol{\mu}^T \mathbf{x}_*$ , and the posterior variance,  $\beta^{-1} + \mathbf{x}_*^T \boldsymbol{\Sigma} \mathbf{x}_*$ , quantifies the uncertainty about this point prediction. Rather than providing a point prediction like in the discriminative methods, Bayesian methods advocate averaging all possible predicted values via marginalization, and are thus more robust against erroneous parameter estimates.

### C. Bayesian Nonlinear Nonparametric Model: GP Regression Example

In order to improve the data representation power of Bayesian linear parametric models, a lot of efforts have been invested on designing *non-linear* and *non-parametric* models. A direct non-linear generalization of (8) is given by

$$y = f(\mathbf{x}) + v, \quad (15)$$

where instead of the linearity of (8), we employ a non-linear functional dependence  $f(\mathbf{x})$  and let  $v$  be the noise term like before. Moreover, the randomness associated with the weight parameters  $\boldsymbol{\theta}$  in (8) is now embedded into the function  $f(\mathbf{x})$  itself, which is assumed to be a *random process*. That is, the outcome/realization of each random experiment is a *function* instead of a single value/vector. Thus, in this case, we have to deal with priors related to non-linear functions directly, rather than indirectly, i.e., by specifying a family of non-linear parametric functions and placing priors over the associated weight parameters.

1) *GP Model:* In the sequel, we introduce one representative model that adopts this rationale, namely the GP model for non-linear regression. The GP models constitute a special family of random processes where the outcome of each experiment is a function or a sequence. For instance, in signal processing, this can be a continuous-time signal  $f(t)$  as a function of time  $t$  or a discrete-time signal  $f(n)$  in terms of the sequence index  $n$ . In this article, we treat the GP model as a data analysis tool whose input that acts as the argument in  $f(\cdot)$  is a vector, i.e., vector,  $\mathbf{x} = [x_1, x_2, \dots, x_L]^T$  [1], [24].

For clarity, we give the definition of GP as follows:

*Definition of GP* [1], [24]: A random process,  $f(\mathbf{x})$ , is called a GP if and only if for any finite number of points,  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$ , the associated joint probability density function (pdf),  $p(f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_N))$  is Gaussian.

A GP can be considered as an infinitely long vector of jointly Gaussian distributed random variables, so that it can be fully described by its mean function and covariance function, defined as follows:

$$m(\mathbf{x}) \triangleq \mathbb{E}[f(\mathbf{x})], \quad (16)$$

$$\text{cov}(\mathbf{x}, \mathbf{x}') \triangleq \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]. \quad (17)$$

A GP is said to be stationary if the mean function,  $m(\mathbf{x})$ , is a constant mean, and moreover its covariance function has the following simplified form:  $\text{cov}(\mathbf{x}, \mathbf{x}') = \text{cov}(\boldsymbol{\tau})$  with  $\boldsymbol{\tau} \triangleq \mathbf{x} - \mathbf{x}'$ .

When a GP is adopted for data modeling and analysis, we need to specify the mean function and the covariance function in order to make the model match the underlying data patterns. The mean function is often set to zero especially when there is no prior knowledge available. The data representation power of the non-parametric GP models is determined overwhelmingly by the covariance function, which is also known as the *kernel function* due to the positive semi-definite nature of a covariance function. In the following, we use  $k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\eta}_p) \triangleq \text{cov}(\mathbf{x}, \mathbf{x}')$  to represent a pre-selected kernel function with an explicit set of tuning kernel hyper-parameters,  $\boldsymbol{\eta}_p$ , for the observed data. Finally, we say that a function realization is drawn from the GP prior, and we write

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\eta}_p)). \quad (18)$$

The consequent GP regression model follows (15), where  $f(\mathbf{x})$  is represented by a GP model defined in (18) and the noise term  $v$  is assumed to be Gaussian distributed with zero mean and variance  $\beta^{-1}$ , like in the previous simple Bayesian linear regression example.

2) *GP Kernel Functions*: As mentioned before, the kernel function plays a crucial role in determining a GP model's representation power. To shed more light on the kernel function, especially on how it represents random functions as well as its good physical interpretations, we demonstrate the most widely used (but not necessarily optimal) squared-exponential (SE) kernel.

■ SE Kernel: The form of this widely used kernel function is given below:

$$k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\eta}_p) = \sigma_s^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{2\ell^2}\right), \quad (19)$$

where the hyper-parameter  $\sigma_s^2$  determines the magnitude of fluctuation of  $f(\mathbf{x})$  and the other hyper-parameter  $\ell$ , called *length-scale*, determines the statistical correlation between two points,  $f(\mathbf{x})$  and  $f(\mathbf{x}')$ , separated by a (Euclidean) distance  $d \triangleq \|\mathbf{x} - \mathbf{x}'\|_2$ . Thus, we have the kernel hyper-parameters,  $\boldsymbol{\eta}_p = [\sigma_s^2, \ell]^T$ , specifically for this kernel.

In Fig. 2, we show some sample functions generated from a GP (for one-dimensional input,  $x$ ) involving the SE kernel with different hyper-parameter configurations. From these illustrations, we can clearly spot

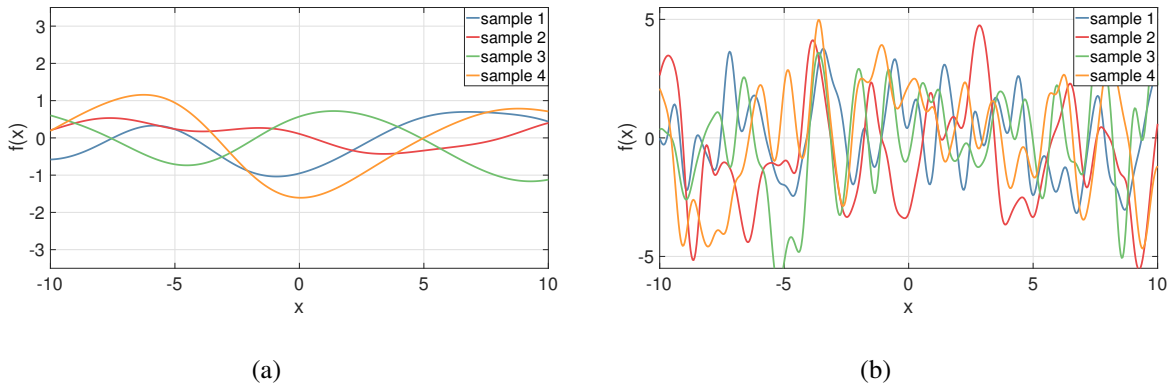


Fig. 2. Sample functions generated from a GP model using the SE kernel with two different hyper-parameter configurations. (a) GP with hyper-parameters,  $\sigma_s^2 = 1, \ell = 5$ , generates low-peaked and smooth sample functions; (b) GP with hyper-parameters,  $\sigma_s^2 = 5, \ell = 0.5$ , generate high-peaked and fast varying sample functions.

the physical meaning of the SE kernel hyper-parameters. There are many other classic kernel functions, such as the Ornstein-Uhlenbeck kernel, rational quadratic kernel, periodic kernel, locally periodic kernel as introduced in, e.g., [24]. They can even be combined, for instance in the form of a linearly-weighted sum, to enrich the overall modeling capacity [24], [25].

Designing a competent stationary kernel function for the GP model can also be considered in the frequency domain owing to the famous *Wiener-Khinchine Theorem* [1], [24]. The theorem states that the Fourier transform of a stationary kernel function,  $k(\boldsymbol{\tau})$ , and the associated *spectral density* of the process,  $S(\mathbf{s})$ , are Fourier duals:

$$k(\boldsymbol{\tau}) = \int S(\mathbf{s}) e^{2\pi i \mathbf{s}^T \boldsymbol{\tau}} d\mathbf{s}, \quad S(\mathbf{s}) = \int k(\boldsymbol{\tau}) e^{-2\pi i \mathbf{s}^T \boldsymbol{\tau}} d\boldsymbol{\tau}. \quad (20)$$

Here, it is noteworthy to mention that  $i$  is the imaginary unit and the operation  $\mathbf{s}^T \boldsymbol{\tau}$  refers to the inner product of the generalized frequency parameters,  $\mathbf{s}$ , and the time difference parameters,  $\boldsymbol{\tau}$ . In Section IV, we will introduce some optimal kernel design methods that were first built based on the spectral density in the frequency domain and then transformed back to the original input domain.

3) *GP for Regression*: In contrast to the Bayesian linear regression model, we set a GP prior directly on the underlying function in the GP regression model, namely,  $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\eta}_p))$ . Given the observed dataset,  $\mathcal{D} = \{\mathbf{X}, \mathbf{y}\}$  as defined before, the main goal of GP-based Bayesian non-linear regression is to compute the evidence,  $p(\mathbf{y}; \boldsymbol{\eta})$ , for optimizing the model hyper-parameters,  $\boldsymbol{\eta}$ , and to compute the posterior distribution,  $p(\mathbf{y}_* | \mathbf{y})$ , of  $\mathbf{y}_* = [y_{*,1}, y_{*,2}, \dots, y_{*,N_*}]^T$  evaluated at  $n_*$  novel test inputs  $\mathbf{X}_* = [\mathbf{x}_{*,1}, \mathbf{x}_{*,2}, \dots, \mathbf{x}_{*,N_*}]^T$ .

■ Evidence: This can be obtained in a straightforward way due to the regression model  $\mathbf{y} = \mathbf{f}(\mathbf{X}) + \mathbf{v}$ , where  $\mathbf{v}$  is independent of the GP model,  $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\eta}_p))$ , and we let  $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \beta^{-1}\mathbf{I})$ . As a consequence, it is easy to derive, see e.g., [1], [24], that

$$p(\mathbf{y}; \boldsymbol{\eta}) = \mathcal{N}(\mathbf{y}; \mathbf{0}, \mathbf{K}(\mathbf{X}, \mathbf{X}; \boldsymbol{\eta}_p) + \beta^{-1}\mathbf{I}), \quad (21)$$

where  $\boldsymbol{\eta} = [\boldsymbol{\eta}_p^T, \beta]^T$  and  $\mathbf{K}(\mathbf{X}, \mathbf{X}; \boldsymbol{\eta}_p)$  is the  $N \times N$  kernel matrix of  $\mathbf{f}(\mathbf{X}) \triangleq [f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_n)]^T$  evaluated for the training samples. Note that the kernel matrix is a square matrix whose  $(ij)$ -th entry is the pairwise covariance between  $f(\mathbf{x}_i)$  and  $f(\mathbf{x}_j)$ , computed according to (17), for any  $\mathbf{x}_i$  and  $\mathbf{x}_j$  in the training dataset. The covariance matrix,  $\mathbf{X}\mathbf{A}^{-1}\mathbf{X}$ , of the Bayesian linear regression function,  $f(\mathbf{x}) = \boldsymbol{\theta}^T \mathbf{x}$ , given in (11) can be regarded as one instance of the kernel matrix,  $\mathbf{K}(\mathbf{X}, \mathbf{X}; \boldsymbol{\eta}_p)$ . The latter can provide increased representation power through choosing more appropriate kernel forms and tuning the associated kernel hyper-parameters. As it will be shown in Section V, we will maximize this evidence function to get an optimal set of the model hyper-parameters.

■ Posterior Distribution: It turns out, see e.g. [1], [19], that the joint distribution of the training output  $\mathbf{y}$  and the test output  $\mathbf{y}_*$  is a Gaussian, of the following form:

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}_* \end{bmatrix} \sim \mathcal{N} \left( \begin{bmatrix} \mathbf{y} \\ \mathbf{y}_* \end{bmatrix}; \mathbf{0}, \begin{bmatrix} \mathbf{K}(\mathbf{X}, \mathbf{X}) + \beta^{-1}\mathbf{I} & \mathbf{K}(\mathbf{X}, \mathbf{X}_*) \\ \mathbf{K}(\mathbf{X}_*, \mathbf{X}) & \mathbf{K}(\mathbf{X}_*, \mathbf{X}_*) + \beta^{-1}\mathbf{I} \end{bmatrix} \right), \quad (22)$$

where  $\mathbf{K}(\mathbf{X}, \mathbf{X}_*)$  stands for the  $N \times N_*$  kernel matrix between the training inputs and test inputs and  $\mathbf{K}(\mathbf{X}_*, \mathbf{X}_*)$  for the  $N_* \times N_*$  kernel matrix among the test inputs. Here, we let  $\mathbf{K}(\mathbf{X}, \mathbf{X})$  be a short form of  $\mathbf{K}(\mathbf{X}, \mathbf{X}; \boldsymbol{\eta})$ .

By applying some classic conditional Gaussian results, see e.g., [1], we can derive the posterior distribution from the joint distribution in (22) as:

$$p(\mathbf{y}_* | \mathbf{y}) \sim \mathcal{N}(\mathbf{y}_*; \bar{\mathbf{m}}, \bar{\mathbf{V}}), \quad (23)$$

where the posterior mean (vector) and the posterior covariance (matrix) are respectively,

$$\bar{\mathbf{m}} = \mathbf{K}(\mathbf{X}_*, \mathbf{X}) [\mathbf{K}(\mathbf{X}, \mathbf{X}) + \beta^{-1}\mathbf{I}]^{-1} \mathbf{y}, \quad (24)$$

$$\bar{\mathbf{V}} = \mathbf{K}(\mathbf{X}_*, \mathbf{X}_*) + \beta^{-1}\mathbf{I} - \mathbf{K}(\mathbf{X}_*, \mathbf{X}) [\mathbf{K}(\mathbf{X}, \mathbf{X}) + \beta^{-1}\mathbf{I}]^{-1} \mathbf{K}(\mathbf{X}, \mathbf{X}_*). \quad (25)$$

The above posterior mean gives a point prediction, while the posterior covariance defines the uncertainty region of such prediction. A leading benefit of using the GP models over discriminative methods, such as the kernel ridge regression, lies in the natural uncertainty quantification given by (25).

A graphical illustration of GP working on a toy regression example is shown in Fig. 3. As we can see from the figures, the uncertainty in the GP prior is constantly large, reflecting our crude prior belief

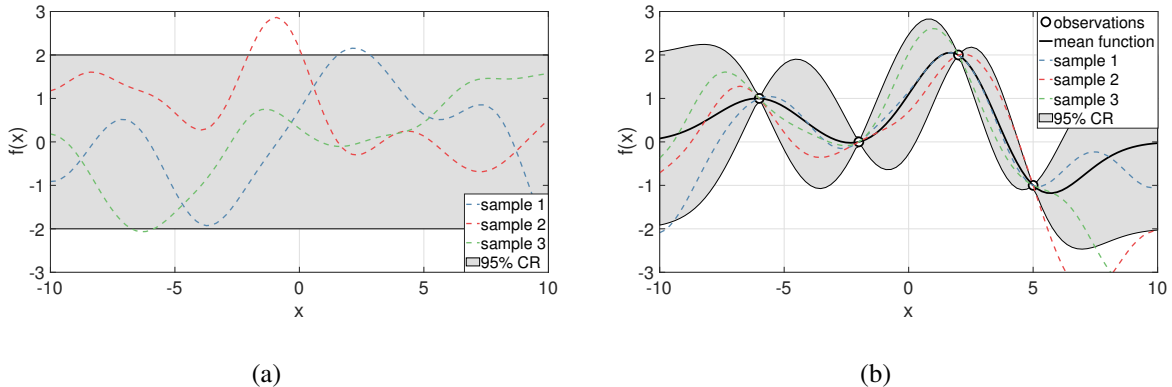


Fig. 3. Subfigure (a) shows three sample functions drawn randomly from a GP prior (refer to (18)) with an SE kernel (refer to (19)). Subfigure (b) shows three sample functions drawn from the GP posterior (refer to (23), (24), (25)) computed based on the prior shown in (a) as well as four noisy observations indicated by the black circles. The corresponding posterior mean function is depicted by the dark black curve. The grey shaded area represents the uncertainty region, taken as the 95% confidence region (CR) for both the prior and the posterior herein.

in the underlying function. While it has been significantly reduced in the neighborhood of the observed data points in the GP posterior, but still remains comparably large in regions where the observed data points are scarce.

Apart from the representation power of the GP model, it also connects to various other salient machine learning models, including for instance the relevance vector machine, support vector machine [24]. Also, it has been shown that a neural network, with one or multiple hidden layers, asymptotically approaches a GP, see e.g. [26], [27].

### III. SPARSITY-AWARE LEARNING: REGULARIZATION FUNCTIONS AND PRIOR DISTRIBUTIONS

In modern big data analysis, there is a trend to employ sophisticated models that involve an excessive number of parameters (sometimes even more than the number of data samples, e.g., in *over-parameterized models*). This makes the learning systems vulnerable to overfitting to the observed data. Thus, the obvious question concerns the right model size given the data sample. *Sparsity-aware learning* (SAL) that promotes sparsity on the structure of the learnt model comprises a major path in dealing with such models in a data-adaptive fashion. The term *sparsity* implies that most of the unknown parameters are pushed to (almost) zero values. This can be achieved either via the combination of a discriminative method and appropriate regularizers, or via the Bayesian path by adopting sparsity-promoting priors. The major difference between the two paths lies in the way “sparsity” is interpreted and embedded into the models, as it is explained in the following subsections.

In the sequel, we will first introduce the first path that leads to SAL via regularized optimization methods in Section III-A, followed by the SAL via Bayesian methods for parametric models in Section III-B and non-parametric models in Section III-C. Note that the aim of this article is not on comparing the two paths, but rather to rethink the Bayesian philosophy.

### A. SAL via Regularized Optimization Methods

Following the regularized optimization way, “sparsity” information is embedded through *regularization* functions. Using the linear regression task as an example, the regularized parameter optimization problem is formulated as:

$$\min_{\boldsymbol{\theta}} \underbrace{\frac{1}{2} \sum_{n=1}^N (y_n - \boldsymbol{\theta}^T \mathbf{x}_n)^2}_{\text{data fitting cost}} + \underbrace{\lambda}_{\text{regularization parameter}} \times \underbrace{r(\boldsymbol{\theta})}_{\text{regularization function}}, \quad (26)$$

where the regularization function  $r(\boldsymbol{\theta})$  steers the solution towards a preferred sparse structure, and the regularization parameter  $\lambda$  is to balance the trade-off between the data fitting cost and the regularization function for sparse structure embedding. In SAL, it is assumed that the unknown parameters  $\boldsymbol{\theta}$  have a majority of zero entries, and thus the adopted regularization function  $r(\boldsymbol{\theta})$  should help the optimization process unveil such zeros. Such regularization functions include the family of  $l_p$  norm functions with  $0 \leq p \leq 1$ , among which the  $l_1$  norm is most popular, since it retains the computationally attractive property of convexity. Furthermore, strong theoretical results have been derived, see e.g., [1], [28]. In recent years, SAL advances via regularized cost optimization prevail in the context of machine learning using data analysis tools. The literature is very rich and fairly well documented with many sparsity-promoting regularization functions. Although the resulting regularized cost function might be non-convex and/or non-smooth, efficient learning algorithms exist and have been built on solid theoretical foundations in optimization theory, see e.g., [29].

### B. SAL via Bayesian Methods For Parametric Models

Before we entail into a more formal presentation of a family of probability density functions (pdfs) that promote sparsity, let us first view sparsity from a slightly different angle. It is well known that there is a bridge between the estimate obtained from problem (26) and the MAP estimate (see Section II-A). For example, it is not difficult to see that if  $r(\boldsymbol{\theta})$  is the squared Euclidean norm (i.e., the  $l_2$  norm that gives rise to the so-called ridge regression), the resulting estimate corresponds to the MAP one when assuming the noise to be i.i.d. Gaussian and the prior on  $\boldsymbol{\theta}$  to be also of a Gaussian form. If, on the other hand,  $r(\boldsymbol{\theta})$  takes the  $l_1$  norm, this corresponds to imposing a Laplacian prior, instead of a Gaussian





Table I: Examples of GSM prior. Abbreviations: Ga = Gamma, IG = inverse Gamma, GIG = generalized inverse Gaussian,  $C^+$  = Half Cauchy.

GSM prior $p(\theta_l)$	Mixing distribution $p(\zeta_l)$
Student's $t$	Inverse Gamma: $p(\zeta_l; \boldsymbol{\eta}_p = [a, b]) = \text{IG}(\zeta_l; a, b)$
Normal-Jefferys	Log-uniform: $p(\zeta_l; \boldsymbol{\eta}_p = [ \ ]) \propto \frac{1}{ \zeta_l }$
Laplacian	Gamma: $p(\zeta_l; \boldsymbol{\eta}_p = [a, b]) = \text{Ga}(\zeta_l; a, b)$
Generalized hyperbolic	Generalized inverse Gaussian: $p(\zeta_l; \boldsymbol{\eta}_p = [a, b, \lambda]) = \text{GIG}(\zeta_l; a, b, \lambda)$
Horseshoe	$\zeta_l = \tau_l v_l, \boldsymbol{\eta}_p = [a, b]$ Half Cauchy: $p(\tau_l) = C^+(0, a)$ $p(v_l) = C^+(0, b)$

associated with the prior. Thus, the GSM prior for each  $\theta_l$  is expressed as

$$p(\theta_l; \boldsymbol{\eta}_p) = \int \mathcal{N}(\theta_l; 0, \zeta_l) p(\zeta_l; \boldsymbol{\eta}_p) d\zeta_l. \quad (27)$$

By varying the functional forms of  $p(\zeta_l; \boldsymbol{\eta}_p)$ , the marginalization (i.e., integrating out the dependence on  $\zeta_l$ ) performed in light of (27) induces different prior distributions of  $\boldsymbol{\theta}$ . For example, if  $p(\zeta_l; \boldsymbol{\eta}_p)$  is an inverse Gamma distribution, (27) induces a Student's  $t$  distribution [30]; if  $p(\zeta_l; \boldsymbol{\eta}_p)$  is a Gamma distribution, (27) induces a Laplacian distribution [30]. For clarity, Table I summarizes different heavy-tail distributions, including Normal-Jefferys, generalized hyperbolic, and horseshoe distributions, among others. To illustrate graphically the sparsity-promoting property endowed by their heavy-tail nature, in addition to the Laplacian distribution plotted in Fig. 4, we further depict two representative GSM prior distributions, namely the Student's  $t$  distribution and the horseshoe distribution, in Fig. 5. In Section IV-A and IV-C, we will show the use of GSM prior in modeling Bayesian neural networks and low-rank tensor decomposition models, respectively.

Besides the aforementioned families of sparsity-promoting priors, another path that has been followed to impose sparsity exploits the underlying property of the evidence function to provide a trade-off between the fitting accuracy and the model complexity, at its maximum value, as it has already been discussed in Section II-A. To this end, one imposes an individual Gaussian prior  $\mathcal{N}(0, \zeta_l)$  on each one of the unknown parameters, which are assumed to be mutually independent, and then treats the respective variances,  $\zeta_l$ ,  $l = 1, 2, \dots, L$ , as hyper-parameters that are obtained via the evidence function optimization. Due to the accuracy-complexity trade-off, the variances of the parameters that need to be pushed to zero (i.e., do not contribute much to the accuracy-likelihood term) get very large values and their corresponding means get values close to zero, see e.g., [19], [31], where a theoretical justification is provided. The key

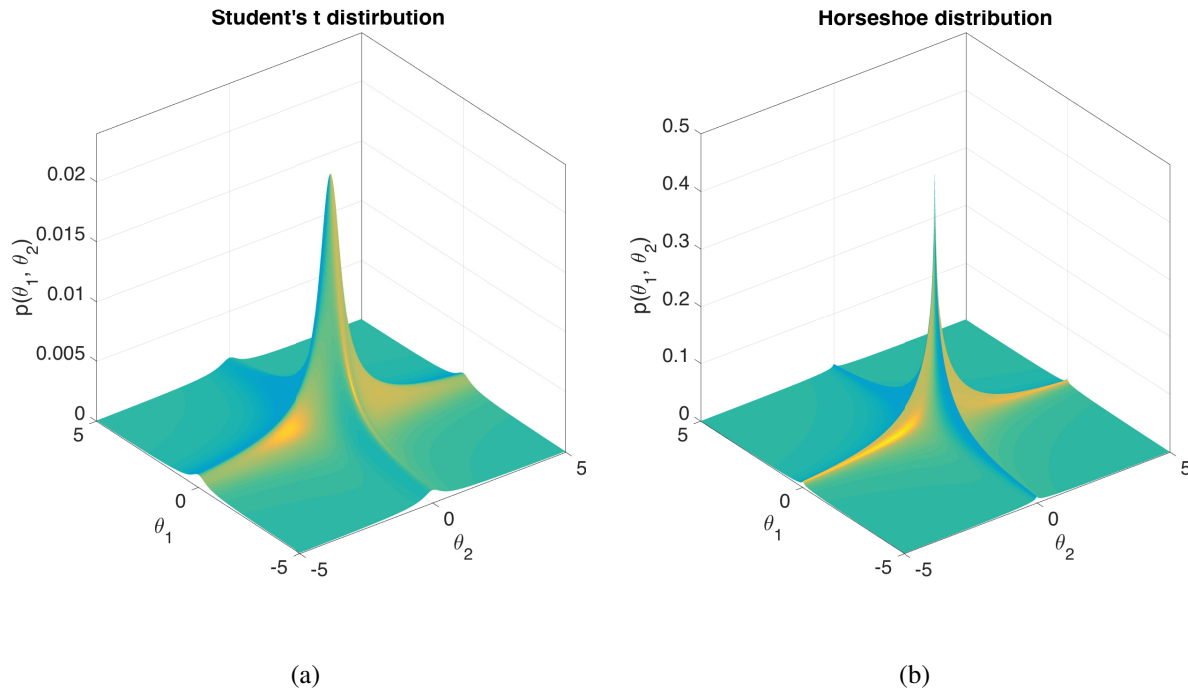


Fig. 5. Representative GSM prior distributions in two-dimensional space. Subfigure (a) and Subfigure (b) show the Student’s  $t$  distribution and the horseshoe distribution, respectively. It can be seen that these two distributions show different heavy-tail profiles and are both sparsity-promoting.

point here is that allowing the parameters to vary independently, with different variances, unveils specific relevance of every individual parameter to the observed data, and the “irrelevant” ones are pushed to zero with high probability. Such methods are also known as *Automatic Relevance Determination* (ARD). In Section IV-B, we demonstrate the use of ARD philosophy for designing recent sparse kernels.

*Remark 1:* In practice, the choice of a specific prior depends on the trade-off between the expressive power of the prior and the difficulty of inference. As shown in Table I, advanced sparsity-promoting priors, e.g., the generalized hyperbolic prior and the horseshoe prior, come with more complicated mathematical expressions. These endue the priors flexibility to adapt to different levels of sparsity, while also pose difficulty in deriving efficient inference algorithm. Typically, when the noise power is known to be small, and/or the side information about the sparsity level is available, sparsity-promoting priors with simple mathematical forms, e.g., the Student’s  $t$ -prior, are recommended. Otherwise, one might consider the adoption of more complex members in the family of GSM priors, see e.g., [6], [10].

### C. SAL via Bayesian Methods for Non-Parametric Models

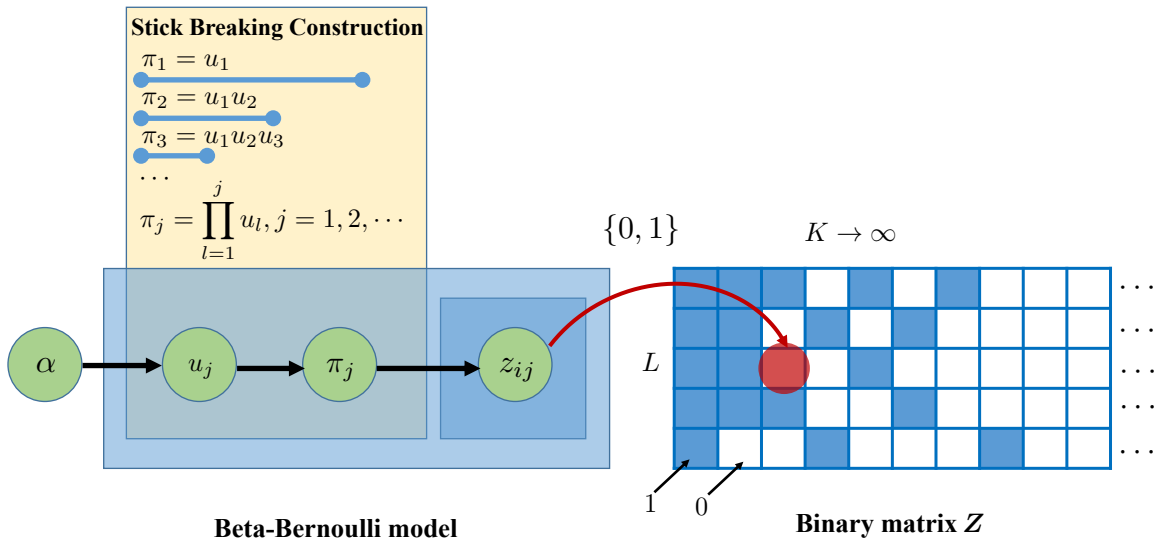


Fig. 6. Illustrative implementation of IBP via stick breaking construction.

In this subsection, we turn our focus on non-parametric models, where the number of the involved parameters in an adopted model is not considered to be known and fixed by the user but, in contrast, it has to be learnt from the data during training. A common path to this direction is to assume that the involved number of parameters is infinite (practically a very large number) and then leave it to the algorithm to recover a finite set of parameters out of the, initially, infinite ones. To this end, one has to involve priors that deal with infinite many parameters. The “true” number of parameters is recovered via the associated posteriors.

1) *Indian Buffet Process (IBP) Prior:* We will introduce the IBP in a general formulation, and then we will see how to adapt it to fit our needs in the context of designing DNNs. Let us first assume that an Indian restaurant offers a very large number,  $K$ , of dishes and let  $K \rightarrow \infty$ . There are  $L$  customers. The first one selects a number of dishes, with some probability. The second customer, selects some of the previously selected dishes with some probability and some new ones with another probability, and so on till all,  $L$ , customers have been considered. In the context of designing NNs, customers are replaced by the dimensions of the input (to each one of the layers) vector, and the infinite many dishes by the number of nodes in a layer. Since we have assumed that the architecture is unknown, that is, the number of nodes (neurons) in a layer, we consider infinite many of those. Then, following the rationale of IBP, the first dimension, say,  $x_1$  is linked to some of the infinite nodes, with certain probabilities, respectively. Then, the second dimension, say,  $x_2$  is linked to some of the previously linked nodes and to some new ones, according to certain probabilities. This reasoning carries on, till the last dimension of the input

vector,  $x_L$ , has been considered. As we will see soon, the IBP is a sparsity promoting prior, because out of the infinite many nodes, only a small number of those is *probabilistically* selected.

In a more formal way, we adopt a binary random variable,  $z_{ij} \in \{0, 1\}$ ,  $i = 1, 2, \dots, L$  and  $j = 1, 2, \dots$ . If  $z_{ij} = 1$ , the  $i$ -th customer ( $i$ -th dimension) selects the  $j$ -th dish (is linked to the  $j$ -th node). On the contrary, if  $z_{ij} = 0$ , the dish is not selected, (the  $i$ -th dimension is not linked to the  $j$ -th node). The binary matrix  $\mathbf{Z}$  that is defined from the elements  $z_{ij}$  is an infinite dimensional one and the IBP is a prior that promotes zeros in such binary matrices.

One way to implement the IBP is via the so-called *stick breaking construction*. The goal is to populate an infinite binary matrix,  $\mathbf{Z}$ , with each element being zero or one. To this end, we first generate *hierarchically* a sequence of, theoretically, infinite probability values,  $\pi_j$ ,  $j = 1, 2, \dots$ . To achieve this, the Beta distribution is mobilized. The Beta distribution, e.g., [1], is defined in terms of two parameters. For the IBP, we fix one of them to be equal to 1 and the other one,  $\alpha$ , is left as a (hyper-)parameter, which can either be pre-selected or learnt during training. Then, the following steps are in order:

$$u_j \sim \text{Beta}(u_j|\alpha, 1), \quad \pi_j = \prod_{l=1}^j u_l, \quad j = 1, 2, \dots, \quad (28)$$

where, the notation  $\sim$  indicates the sample drawn from a distribution. Then, the generated probabilities,  $\pi_j$ , are used to populate the matrix  $\mathbf{Z}$ , by drawing samples from a Bernoulli distribution, see e.g., [1], that generates an one, with probability  $\pi_j$  and a zero with probability  $1 - \pi_j$ , as

$$z_{ij} \sim \text{Bernoulli}(z_{ij}|\pi_j), \quad (29)$$

for each  $i = 1, 2, \dots, L$ , as illustrated in Fig. 6. The Beta distribution generates numbers between  $[0, 1]$ , and from the above construction it is obvious that the sequence of probabilities  $\{\pi_j\}$  goes rapidly to zero, due to the product of quantities  $\{u_l\}$  being less than one in magnitude. How fast this takes place is controlled by  $\alpha$ , which is known as the innovation or strength parameter of the IBP, see e.g., [1].

#### IV. THE ART OF PRIOR: SPARSITY-AWARE MODELING FOR THREE CASE STUDIES

In the previous section, we have introduced the indispensable ingredients for obtaining sparsity-aware modeling under the Bayesian learning framework, namely the priors. In this section, we will demonstrate how these priors can be incorporated into some popular data modeling and analysis tools to achieve sparsity-promoting properties. Concretely, we will introduce sparsity-aware modeling for *Bayesian deep neural networks* in Section IV-A, for Gaussian processes in Section IV-B, and for tensor decompositions in Section IV-C.

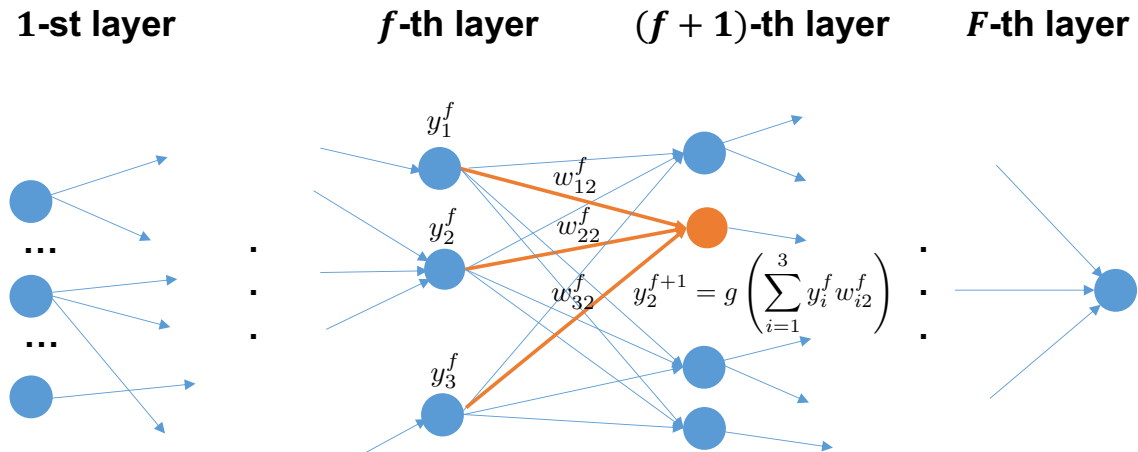


Fig. 7. Illustration of a deep fully connected network.

### A. Sparsity-Aware Modeling for Bayesian Deep Neural Networks

Our focus in this section is to deal with sparsity-promoting techniques in order to prune DNNs. That is, starting from a network with a large number of nodes, to optimally remove nodes and/or links. We are going to follow both paths, namely the parametric one via the GSM priors and the non-parametric one via the IBP prior.

1) *Fundamentals of DNNs*: Neural networks are learning machines that comprise a large number of neurons, which are connected in a layer-wise fashion. After 2010 or so, neural networks with many (more than three) hidden layers, known as *deep neural networks*, have dominated the field of machine learning due to their remarkable representation power and the outstanding prediction performance for various learning tasks. Since their introduction, one of the major tasks associated with their design has been the so-called *pruning*. That is, to remove redundant nodes and links so that their size and, hence, the number of the involved parameters is reduced. Of course, this is another name of what we have called “sparsification” of the network. Over the years, a number of rather ad-hoc techniques have been proposed, see e.g., [1], [19] for a review. In the most recent years, Bayesian techniques have been employed in a more formal and theoretically pleasing way. These techniques comprise our interest in this article. We will focus on the vanilla deep fully connected networks, yet, such techniques have been extended and can be used for the case of, e.g., convolutional networks [2], [3].

The name “deep fully connected networks” stresses out that each node in any of the layers is directly connected to every node of the previous layer. To state it in a more formal way, without loss of generality, we consider a deep fully connected network consisting of  $F$  layers. The number of nodes in the  $f$ -th

( $1 \leq f \leq F - 1$ ) layer is  $a^f$ .<sup>3</sup> For the  $i$ -th node in the  $f$ -th layer and the  $j$ -th node in the  $(f + 1)$ -th layer, the link between them has a weight  $w_{ij}^f$ , as illustrated in Fig. 7. The input vector to the  $(f + 1)$ -th layer consists of the outputs from the previous layer, denoted as  $\mathbf{y}^f = [y_1^f, y_2^f, \dots, y_{a^f}^f]^T$ , where  $y_i^f$  is the output at the  $i$ -th node. Denote  $\mathbf{w}_j^f = [w_{1j}^f, w_{2j}^f, \dots, w_{a^f j}^f]^T$  the vector that collects all the link weights associated with the  $j$ -th node. Then the output of the  $j$ -th node is:

$$y_j^{f+1} = g \left( \sum_{i=1}^{a^f} w_{ij}^f y_i^f \right) = g \left( [\mathbf{w}_j^f]^T \mathbf{y}^f \right), \quad (30)$$

where  $g(\cdot)$  is a non-linear transformation function (also called activation function), and the most widely used ones include the rectified linear unit (ReLU) function, the sigmoid function, and the hyperbolic tanh function, see e.g., [1].

2) *Sparsity-Aware Modeling Using GSM Priors*: The basic idea of this approach can be traced back to the pioneering work [32] of D. J. MacKay in 1995. He pointed out that for a neural network with a single hidden layer, the weights, each associated with a link between two nodes, can be treated as *random variables*. The connection weights are associated with zero-mean Gaussian priors, typically with a shared variance hyper-parameter. Then, appropriate (e.g., Gaussian) posteriors are learnt over the connection weights, which can be used for inference at test time. The variance hyper-parameters of the imposed Gaussian priors can be selected to be low enough, so that the corresponding connection weights exhibit *a priori* the tendency of being concentrated around the postulated zero mean. This induces a sparsity “bias” to the network. The major differences between the recent works [5], [6] and the early work [32] lies in their adopted priors.

Let us consider a network with multiple hidden layers [5], [6], as illustrated in Fig. 8. For the  $i$ -th node in the  $f$ -th layer and the  $j$ -th node in the  $(f + 1)$ -th layer, their link has a weight  $w_{ij}^f$ . For each of the random weights, we can adopt a sparsity-promoting GSM prior so that

$$p(w_{ij}^f) = \int \mathcal{N}(w_{ij}^f; 0, \zeta_{ij}^f) p(\zeta_{ij}^f; \boldsymbol{\eta}) d\zeta_{ij}^f, \quad (31)$$

in which each functional form of  $p(\zeta_{ij}^f; \boldsymbol{\eta})$  in Table I corresponds to a GSM prior. Particularly, the Normal-Jeffreys prior and the horseshoe prior were used in [5], [6]. Next, we show how to conduct node-wise sparsity-aware modeling (for all the weights connected to that node). Inspired by the idea reported in [32], we group the weights  $\{w_{ij}^f\}_{j=1}^{a^{f+1}}$  connected to the  $i$ -th node, and assign a common scale parameter

<sup>3</sup>Note that  $f$  in  $a^f$  stands for the  $f$ -th layer and acts as a superscript. It does not denote  $a$  to the power of  $f$ .

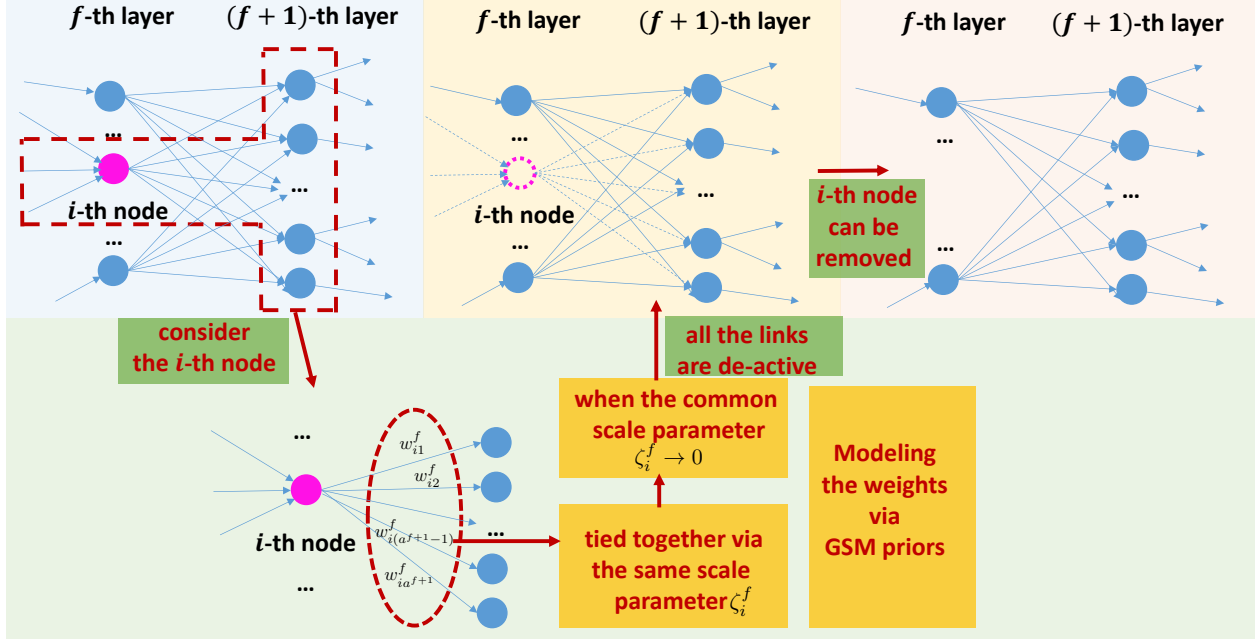


Fig. 8. Illustration of node-wise sparsity-aware modeling for DNNs using GSM priors.

$\zeta_i^f$  to their GSM priors, i.e.,  $\zeta_{ij}^f = \zeta_i^f, \forall j$ . Then we have the prior modeling for the  $i$ -th node related weights  $\{w_{ij}^f\}_{j=1}^{a^{f+1}}$ :

$$\begin{aligned}
 p(\{w_{ij}^f\}_{j=1}^{a^{f+1}}) &= \int p(\{w_{ij}^f\}_{j=1}^{a^{f+1}}, \zeta_i^f) d\zeta_i^f \\
 &= \int p(\{w_{ij}^f\}_{j=1}^{a^{f+1}} | \zeta_i^f) p(\zeta_i^f; \boldsymbol{\eta}) d\zeta_i^f \\
 &= \int \prod_{j=1}^{a^{f+1}} \mathcal{N}(w_{ij}^f; 0, \zeta_i^f) p(\zeta_i^f; \boldsymbol{\eta}) d\zeta_i^f.
 \end{aligned} \tag{32}$$

Furthermore, assuming that the nodes in the  $f$ -th layer are mutually independent, we obtain the prior modeling for all the weights  $\{\{w_{ij}^f\}_{i=1}^{a^f}\}_{j=1}^{a^{f+1}}$  forwarded from the  $f$ -th layer:

$$p(\{\{w_{ij}^f\}_{i=1}^{a^f}\}_{j=1}^{a^{f+1}}) = \prod_{i=1}^{a^f} p(\{w_{ij}^f\}_{j=1}^{a^{f+1}}) = \prod_{i=1}^{a^f} \int \prod_{j=1}^{a^{f+1}} \mathcal{N}(w_{ij}^f; 0, \zeta_i^f) p(\zeta_i^f; \boldsymbol{\eta}) d\zeta_i^f. \tag{33}$$

By this modeling strategy, the weights  $\{w_{ij}^f\}_{j=1}^{a^{f+1}}$  related to the  $i$ -th node are tied together in the sense that when  $\zeta_i^f$  (a single scalar value) goes to zero, the associated weights  $\{w_{ij}^f\}_{j=1}^{a^{f+1}}$  all become negligible. This makes the  $i$ -th node in the  $f$ -th layer disconnected to the  $(f+1)$ -th layer, and thus blocks the information flow. Together with the sparsity-promoting nature of GSM priors, the prior derived in (33) inclines a lot of nodes to be removed from the DNN without degrading the data fitting performance. This leads to the node-wise sparsity-aware modeling for deep fully connected neural networks. Of course, as

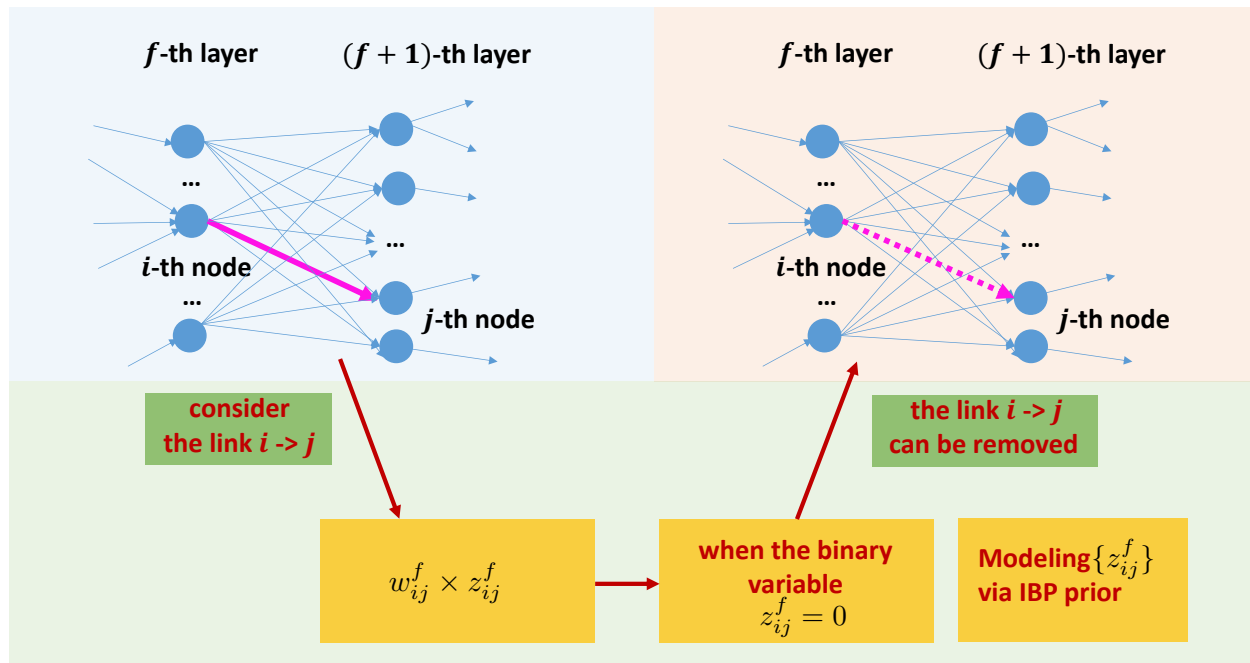


Fig. 9. Illustration of link-wise sparsity-aware modeling for deep neural networks using IBP prior.

it is always the case with Bayesian learning, the goal is to learn the corresponding posterior distributions, and node removal is based on the learnt values for the respective means and variances.

3) *Sparsity-Aware Modeling Using IBP Prior*: The previous approach on imposing sparsity inherits a major drawback that is shared by all techniques for designing DNNs. That is, the number of nodes per layer has to be specified and pre-selected. Of course, one may say that we can choose a very large number of nodes and then harness “sparsity” to prune the network. However, if one overdoes it, he/she soon runs into problems due to over-parameterization. In contrast, we are now going to turn our attention to non-parametric techniques. We are going to assume that the nodes per layer are theoretically infinite (in practice a large enough number) and then use the IBP prior to enforce sparsity.

In line with what has been said while introducing the IBP (Section III-C1), we multiply each weight, i.e.,  $w_{ij}^f$ , with a corresponding auxiliary (hidden) binary random variable,  $z_{ij}^f$ . The required priors for these variables,  $\{\{z_{ij}^f\}_{i=1}^{a^f}\}_{j=1}^{a^{f+1}}$  are generated via the IBP prior. In particular, we define a binary matrix  $\mathbf{Z}^f \in \mathbb{R}^{a^f \times a^{f+1}}$ , with its  $(ij)$ -th element being  $z_{ij}^f$  for the  $f$ -th layer. Due to the sparsity-promoting nature of IBP prior, most elements in  $\mathbf{Z}^f$  tend to be zero, nulling the corresponding weights in  $\{\{w_{ij}^f\}_{i=1}^{a^f}\}_{j=1}^{a^{f+1}}$ , due to the involved multiplication. This leads to an alternative sparsity-promoting modeling for DNNs [2], [3].

The *stick breaking construction* for the IBP prior was utilized since it turns out to be readily amenable to



variational inference. This is a desirable property that facilitates both training and inference through recent advances in black box variational inference, namely stochastic gradient variational Bayes, as explained in Section V-D. For each  $i$ , the considered hierarchical construction reads as follows:

$$u_j^f \sim \text{Beta}(u_j^f | \alpha, 1), \quad \pi_j^f = \prod_{l=1}^j u_l^f, \quad z_{ij}^f \sim \text{Bernoulli}(z_{ij}^f | \pi_j^f). \quad (34)$$

During training, *posterior* estimates of the respective probabilities are obtained, which then allow for a naturally-arising *component omission (link pruning) mechanism* by introducing a *cut-off threshold*  $\tau$ ; any link/weight with inferred posterior below this threshold value is deemed unnecessary and can be safely omitted from computations. This inherent capability renders the considered approach a *fully automatic, data-driven, principled paradigm* for sparsity-aware learning based on explicit inference of component utility based on *dedicated latent variables*.

By utilizing the aforementioned construction, we can easily incorporate the IBP mechanism in conventional ReLU-based networks and perform inference. However, the flexibility of the link-wise formulation allows us to go one step further.

In recent works, the stick-breaking IBP prior has been employed in conjunction with a radically different, biologically-inspired and competition-based activation, namely the stochastic local winner-takes-all (LWTA) [2], [3]. In the general LWTA context, neurons in a conventional hidden layer are replaced by LWTA blocks comprising *competing linear units*. In other words, each node comprises a set of linear (inner product) units. When presented with an input, each unit in each block computes its activation; the unit with the *strongest* activation is deemed to be the *winner* and passes its output to the next layer, while the rest are inhibited to silence, i.e., the zero value. This is how non-linearity is achieved.

This deterministic winner selection, known as *hard* LTWA, is the standard form of an LTWA. However in [2], a new variant was proposed to replace the hard LTWA by a novel *stochastic* adaptation of the competition mechanism implemented via a competitive random sampling procedure founded on Bayesian arguments. To be more specific, let a layer in the NN that comprises  $a^f = L$  inputs, i.e.,  $x_i$ ,  $i = 1, 2, \dots, L$ , where we use  $\mathbf{x}$  to denote the input to any layer in order to simplify the discussion. Also, assume that the number of LWTA blocks in the layer is  $a^{f+1} = K$ . We also relax the notation on the number of layer  $f$ , and our analysis refers to any node of any layer. Each LTWA block comprises  $J$  linear units, each one associated with a corresponding weight,  $w_{ikj}$ ,  $i = 1, 2, \dots, L$ ,  $k = 1, 2, \dots, K$ ,  $j = 1, 2, \dots, J$ . Consider the  $k$ -th LTWA block. We introduce an auxiliary *latent* variable,  $\xi_{kj}$ , and the output of the corresponding  $j$ -th linear unit in the  $k$ -th block is given by,

$$y_{kj} = \xi_{kj} \mathbf{w}_{kj}^T \mathbf{x} = \xi_{kj} \sum_{i=1}^L w_{ikj} x_i, \quad \xi_{kj} \in \{0, 1\}, \quad \sum_{j=1}^J \xi_{kj} = 1. \quad (35)$$

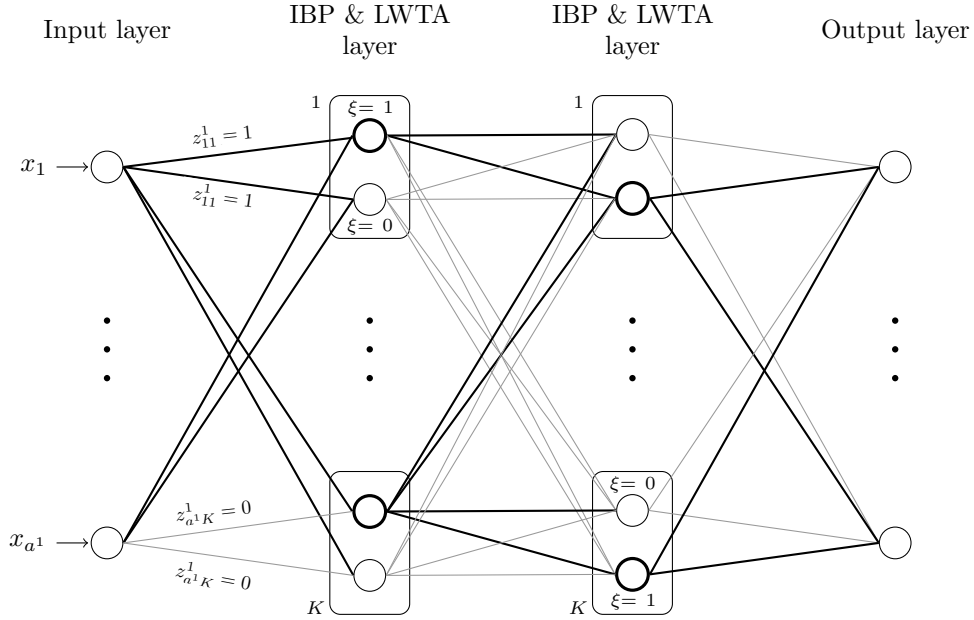


Fig. 10. A graphical illustration of the LWTA and IBP-based architecture. Bold edges denote active (effective) connections (with  $z_{ik}^f = 1$ ); nodes with bold contours denote winner units, i.e., that correspond to  $\xi = 1$  (we do not use  $\xi_{kj}^f$  to unclutter notation); rectangles denote LWTA blocks. For simplicity in the figure, each LWTA block comprises two ( $J = 2$ ) competing linear units,  $k = 1, 2, \dots, K$ .

In other words, the outputs of the linear units are either the respective inner product between the input vector and the associated weight vector or zero, depending on the value of  $\xi_{kj}$ , which can be either zero or one. Furthermore, only one of the  $\xi$ 's in a block can be one, and the rest are zero. Thus, we can associate with each LWTA block, a vector  $\xi_k \in \mathbb{R}^J$ , with only one of its elements being one and the rest being zero, see Fig. 10. In the ML jargon, this is known as one-hot vector and can be denoted as  $\xi_k \in \text{one\_hot}(J)$ . If we stack together all the  $\xi_k$ ,  $k = 1, 2, \dots, K$ , for the specific layer together, we can write,  $\xi \in \text{one\_hot}(J)^K$  (strictly speaking  $\xi^{a^{f+1}} \in \text{one\_hot}(J)^K$ ).

In the stochastic LTWA, all  $\xi$ 's are treated as binary random variables. The respective probabilities, which control the firing (corresponding  $\xi = 1$ ) of each linear unit within a single LTWA, are computed via a *softmax* type of operation, see e.g., [33], [1], that is,

$$P_{nkj} = \frac{\exp(h_{nkj})}{\sum_{j=1}^J \exp(h_{nkj})}, \quad h_{nkj} = \sum_{i=1}^L (z_{ik} w_{ikj}) x_{ni}. \quad (36)$$

Note that in the above equation, the *firing probability* of a linear unit depends on both the input,  $x_n$ ,

and on whether the link to the corresponding LWTA block is active or not (determined by the value of the corresponding utility variable  $z_{ik}$ ). Basically, the stochastic LWTA introduces a *lateral* competition among units in the same layer. How the  $w$ 's as well as the corresponding utility binary variables are learnt is provided in Section V-D. A graphical illustration of the considered approach is depicted in Fig. 10. Note that as the input changes, a *different subnetwork*, via different connected links, may be followed, to pass the input information to the output, with high probability. This is how nonlinearity is achieved in the context of the stochastic LWTA blocks.

### B. Sparsity-Aware Modeling for GPs

We already discussed in Section II-C that the kernel function determines, to a large extent, the expressive power of the GP model. More specifically, the kernel function profoundly controls the characteristics (e.g., smoothness and periodicity) of a GP. In order to provide a kernel function with more expressive power and adaptive to any given dataset, one way is to expand the kernel function as a linear combination of  $Q$  subkernels/basis kernels, i.e.,

$$k(\mathbf{x}, \mathbf{x}') = \sum_{i=1}^Q \alpha_i k_i(\mathbf{x}, \mathbf{x}'), \quad (37)$$

where the weights,  $\alpha_i$ ,  $i = 1, 2, \dots, Q$ , can either be set manually or be optimized. Each one of these subkernels can be any one of the known kernels or any function that admits the properties that define a kernel, see e.g., [19]. One may consider to construct such a kernel either in the original input domain or in the frequency domain. The most straightforward way is to linearly combine a set of elementary kernels, such as the SE kernel, rational quadratic kernel, periodic kernel, etc., with varying kernel hyperparameters in the original input domain, see e.g., [24], [25], [34]. For high-dimensional inputs, one can first detect pairwise interactions between the inputs and for each interaction pair adopt an elementary kernel or an advanced deep kernel [16]. Such resulting kernel belongs essentially to the Analysis-of-Variance (ANOVA) family as surveyed in [1], which has a hierarchical structure and *good interpretability*. Alternatively, one may perform optimal kernel design in the frequency domain by using the idea of sparse spectrum kernel representation. Due to its solid theoretical foundation, in this paper, we will focus on the sparse spectrum kernel representation and review some representative works, such as [7], [35], [36], at the end of this subsection.

1) *Rationale behind Sparsity-Awareness*: The corresponding GP model with the kernel form in (37) can be regarded as a linearly-weighted sum of  $Q$  independent GPs. In other words, we can assume that the underlying function takes the form  $f(\mathbf{x}) = \sum_{i=1}^Q f_i(\mathbf{x})$ , where  $f_i(\mathbf{x}) \sim \mathcal{GP}(0, \alpha_i k_i(\mathbf{x}, \mathbf{x}'))$ , for  $i = 1, 2, \dots, Q$ . In practice,  $Q$  is selected to be a large value compared to the “true” number of the

underlying effective components that generated the data, whose exact value is not known in practice. In the sequel, one can mobilize the ARD philosophy (see Section III-B), during the evidence function optimization, to *drive all unnecessary sub-kernels to zero*, namely promoting sparsity on  $\alpha$ . To this end, let us first establish a bridge between the non-parametric GP model and the Bayesian linear regression model that was considered in Section II-B.

From the theory of kernels, see e.g., [1], [19], [37], each one of the subkernel functions can be written as the inner product of the corresponding feature mapping function, namely,  $k_i(\mathbf{x}, \mathbf{x}') = \phi_i^T(\mathbf{x})\phi_i(\mathbf{x}')$ , where  $\phi_i(\mathbf{x}) : \mathbb{R}^L \mapsto \mathbb{R}^{L'}$  and it is often assumed  $L' \gg L$ . As a matter of fact, the feature mapping function results by fixing one of the arguments of the kernel function and making it a function of a single argument, i.e.,  $\phi_i(\mathbf{x}) = k_i(\mathbf{x}, \cdot)$ , where “ $\cdot$ ” denotes the free variable(s) of the function and is filled by  $\mathbf{x}'$  before. In general,  $\phi_i(\mathbf{x})$  is a function. However, in practice, if needed, this can be approximated by a very high dimensional vector constructed via the famous random Fourier feature approximation [38], [1]. Then, each independent GP process,  $f_i(\mathbf{x}) \sim \mathcal{GP}(0, \alpha_i k_i(\mathbf{x}, \mathbf{x}'))$ , by mobilizing the definition of the covariance matrix, can be equivalently interpreted as  $f_i(\mathbf{x}) \triangleq \boldsymbol{\theta}_i^T \phi_i(\mathbf{x})$ , where the weights,  $\boldsymbol{\theta}_i$ , of size  $L' \times 1$ , are assumed to follow a zero-mean Gaussian distribution, i.e.,  $\boldsymbol{\theta}_i \sim \mathcal{N}(\mathbf{0}, \alpha_i \mathbf{I})$ . Therefore, one can alternatively write  $f(\mathbf{x}) = \sum_{i=1}^Q \boldsymbol{\theta}_i^T \phi_i(\mathbf{x})$ , where  $\boldsymbol{\theta}_i$  and  $\boldsymbol{\theta}_j$  are assumed to be mutually independent for  $i \neq j$ . Essentially, a Gaussian process with such kernel configuration is a special case of the more general sparse linear model family, which can also incorporate, apart from a Gaussian prior, heavy-tailed priors to promote sparsity such as those surveyed in Section III-B. A more detailed presentation of sparse linear models can be found in some early references, such as [31], [39].

As we mentioned before, the GP model hyper-parameters can be optimized through maximizing the logarithm of the evidence function,  $L(\boldsymbol{\eta}) \triangleq \log p(\mathbf{y}; \boldsymbol{\eta})$ , and using (21), we obtain:

$$\begin{aligned} \hat{\boldsymbol{\eta}} &= \arg \max_{\boldsymbol{\eta}} \log \mathcal{N}(\mathbf{y}; \mathbf{0}, \beta^{-1} \mathbf{I} + \sum_{i=1}^Q \alpha_i \mathbf{K}_i(\mathbf{X}, \mathbf{X})) \\ &\equiv \arg \min_{\boldsymbol{\eta}} \left\{ \log \det \left( \beta^{-1} \mathbf{I} + \sum_{i=1}^Q \alpha_i \mathbf{K}_i(\mathbf{X}, \mathbf{X}) \right) + \mathbf{y}^T \left( \beta^{-1} \mathbf{I} + \sum_{i=1}^Q \alpha_i \mathbf{K}_i(\mathbf{X}, \mathbf{X}) \right)^{-1} \mathbf{y} \right\} \quad (38) \\ &\equiv \arg \min_{\boldsymbol{\eta}} \left\{ \log \det \left( \beta^{-1} \mathbf{I} + \sum_{i=1}^Q \alpha_i \boldsymbol{\Phi}_i(\mathbf{X}) \boldsymbol{\Phi}_i^T(\mathbf{X}) \right) + \mathbf{y}^T \left( \beta^{-1} \mathbf{I} + \sum_{i=1}^Q \alpha_i \boldsymbol{\Phi}_i(\mathbf{X}) \boldsymbol{\Phi}_i^T(\mathbf{X}) \right)^{-1} \mathbf{y} \right\}, \quad (39) \end{aligned}$$

where  $\boldsymbol{\eta} = [\boldsymbol{\alpha}^T, \beta]^T$ , and (38) in the second line corresponds to the original GP model, and (39) in the third line corresponds to the equivalent Bayesian linear model mentioned above. Note that  $\mathbf{K}_i(\mathbf{X}, \mathbf{X})$  represents the  $N \times N$  kernel matrix of  $k_i(\mathbf{x}, \mathbf{x}')$  evaluated for all the training input pairs, while  $\boldsymbol{\Phi}_i(\mathbf{X}) \triangleq$

$[\phi_i(\mathbf{x}_1), \phi_i(\mathbf{x}_2), \dots, \phi_i(\mathbf{x}_N)]^T$ , of size  $N \times L'$ , contains the explicit mapping vectors evaluated at the training data. In the sequel, they are denoted as  $\mathbf{K}_i$  and  $\Phi_i$  for brevity.

Mathematical proof of the sparsity property follows that of the *relevance vector machine* (RVM) [40] for the classic sparse linear model. Let us focus on the last expression in (39), which involves both the log-determinant and the inverse of the overall covariance matrix,  $\mathbf{C}$ , and mathematically,

$$\mathbf{C} = \beta^{-1}\mathbf{I} + \sum_{m=1, m \neq i}^Q \alpha_m \Phi_m \Phi_m^T + \alpha_i \Phi_i \Phi_i^T = \mathbf{C}_{-i} + \alpha_i \Phi_i \Phi_i^T, \quad (40)$$

where we have separated the  $i$ -th subkernel from the rest. For clarity, let us focus on the kernel hyper-parameters,  $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \dots, \alpha_Q]^T$ , namely we regard  $\beta$  as known and remove it from  $\boldsymbol{\eta}$ . Applying the classic matrix identities [24], and inserting the results back to (39), we get  $L(\boldsymbol{\alpha}) = L(\boldsymbol{\alpha}_{-i}) + \gamma(\alpha_i)$ , where  $L(\boldsymbol{\alpha}_{-i})$  is simply the evidence function with the  $i$ -th subkernel removed, and the newly introduced quantity  $\gamma(\alpha_i)$  is

$$\gamma(\alpha_i) \triangleq -\frac{1}{2} \ln(\alpha_i) - \frac{1}{2} \log \left| \mathbf{I} + \alpha_i \Phi_i^T \mathbf{C}_{-i}^{-1} \Phi_i \right| + \frac{1}{2} \mathbf{y}^T \mathbf{C}_{-i}^{-1} \Phi_i \left( \alpha_i^{-1} \mathbf{I} + \Phi_i^T \mathbf{C}_{-i}^{-1} \Phi_i \right)^{-1} \Phi_i^T \mathbf{C}_{-i}^{-1} \mathbf{y}. \quad (41)$$

It is not difficult to verify that the evidence maximization problem in (39) boils down to maximizing the  $\gamma(\alpha_i)$  when fixing the rest of the parameters to their previous estimates. This means that we can solve for the hyper-parameters in a sequential manner. Taking the gradient of  $\gamma(\alpha_i)$  with respect to the scalar parameter  $\alpha_i$  and setting it to zero gives the global maximum as can be verified by its second order derivative. Interestingly, the solution to  $\alpha_i$  is either zero or a positive value, mainly depending on the relevance between the  $i$ -th subkernel function and the observed data [19]. Only if their relevance is high enough,  $\alpha_i$  will take a non-zero, positive value. This explains the sparsity-promoting rationale behind the method. In Section V, we will provide an advanced numerical method for solving the hyper-parameters that maximize the evidence function.

2) *Sparse Spectrum Kernel Representation*: At the beginning of this subsection, we expressed kernel expansion in terms of a number of subkernels and introduced two major paths (either in the original input domain or in the frequency domain). In the sequel, we turn our focus to the frequency domain representation of a kernel function and on techniques that promote sparsity in the frequency domain, leading to the sparse spectrum kernel representation.

To start with, it is assumed that the underlying function only has a few effective frequency bands/points in the real physical world. Second, the kernel function takes a linearly-weighted sum of basis functions, similar to the ARD method for linear parametric models, thus only a small number of which are supposed to be relevant to the given data from the algorithmic point of view. Sparse solutions can be obtained from maximizing the associated evidence function as will be introduced in Section V.

For ease of narration, we constrain ourselves to one-dimensional input space, namely,  $x \in \mathbb{R}^1$ , but the idea can be easily extended to the multi-dimensional input space. Often, we have  $x = t$  for one-dimensional time series modeling.

The earliest sparse spectrum kernel representation was proposed in [35] and developed upon a Bayesian linear regression model with trigonometric basis functions, namely,

$$f(x) = \sum_{i=1}^Q a_i \cos(2\pi\omega_i x) + b_i \sin(2\pi\omega_i x), \quad (42)$$

where  $\{\cos(2\pi\omega_i x), \sin(2\pi\omega_i x)\}$  constitute one pair of basis functions parameterized in terms of the center frequencies  $\omega_i$ ,  $i \in \{1, 2, \dots, Q\}$ , and the random weights,  $a_i$  and  $b_i$  are independent and follow the same Gaussian distribution,  $\mathcal{N}(0, \sigma_0^2/Q)$ .<sup>4</sup> Under such assumptions,  $f(x)$  can be regarded as a GP according to Section II-C, and the corresponding covariance/kernel function can be easily derived as

$$k(x, x') = \frac{\sigma_0^2}{Q} \phi^T(x) \phi(x') = \frac{\sigma_0^2}{Q} \sum_{i=1}^Q \cos(2\pi\omega_i(x - x')), \quad (43)$$

where the feature mapping vector  $\phi(x)$  contains all  $Q$  pairs of trigonometric basis functions. Usually, we favor a large value of  $Q$ , well exceeding the expected number of effective components. If the frequency points are *randomly* sampled from the underlying spectral density, denoted by  $\tilde{S}(\omega)$ , then (43) is equivalent to the random Fourier feature approximation of a stationary kernel function [38]. However in [35], the center frequencies are optimized through maximizing the evidence function. As it is claimed in the paper, such additional flexibility of the kernel obtained through optimization can significantly improve the fitting performance as it enables automatic learning of the best kernel function for any specific problem. The resulting spectral density of (43) is a set of *sparse Dirac deltas* for approximating the underlying spectral density.

In [36], the Dirac deltas are replaced with a *mixture of Gaussian* basis functions in the frequency domain, leading to the so-called *spectral mixture* (SM) kernel. The SM kernel can approximate any stationary kernel arbitrarily well in the  $\ell_1$  norm sense due to the Wiener's theorem of approximation [41]. Concretely, the underlying spectral density is approximated by a Gaussian mixture as

$$S(\omega) = \frac{1}{2} \sum_{i=1}^Q \frac{\alpha_i}{\sqrt{2\pi\sigma_i^2}} \left\{ \exp \left[ \frac{-(\omega - \mu_i)^2}{2\sigma_i^2} \right] + \exp \left[ \frac{-(\omega + \mu_i)^2}{2\sigma_i^2} \right] \right\}, \quad (44)$$

where  $Q$  is a fixed number of mixture components, and  $\alpha_i$ ,  $\mu_i$ ,  $\sigma_i^2$  are the weight, mean and variance parameters of the  $i$ -th mixture component, respectively. It is noteworthy that the sum of the two exponential functions on the right-hand-side of (44) ensure the symmetry of the spectral density. For illustration

<sup>4</sup>It is noteworthy that  $\omega \in [0, 1/2)$  represents a normalized frequency, namely the physical frequency over the sampling frequency.

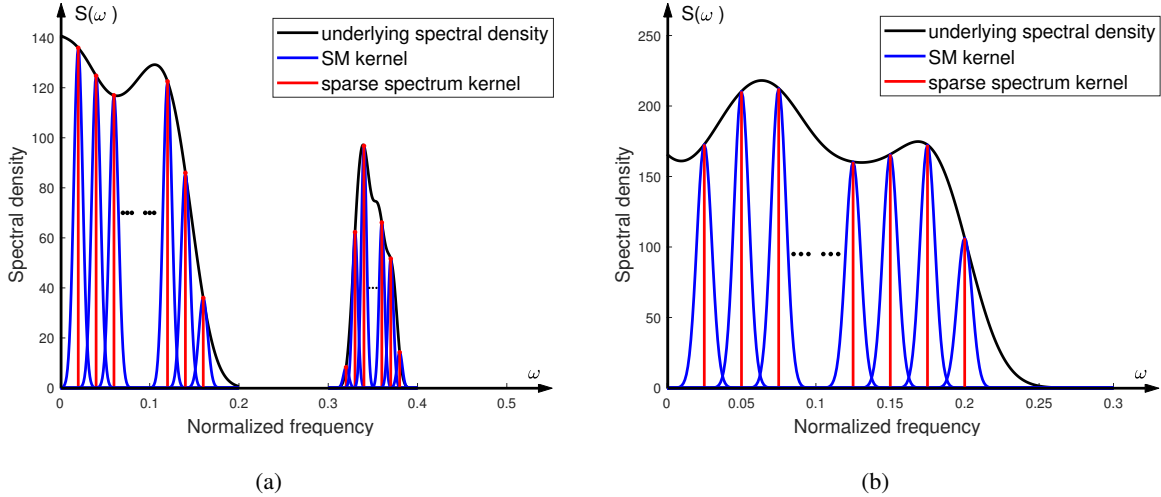


Fig. 11. Comparison between the SM kernel and the original sparse spectrum kernel in (43) for approximating the underlying spectral density. Herein, the SM kernel employs a mixture of Gaussian basis functions (see the blue curves), while the sparse spectrum kernel employs a mixture of Dirac deltas (see the red vertical lines).

purpose, we draw the comparison between the original sparse spectrum kernel [35] and the SM kernel [36] in Fig. 11.

Taking the inverse Fourier transform of  $S(\omega)$  yields a stationary kernel in the time-domain as follows:

$$k(t, t'; \boldsymbol{\eta}_p) = k(\tau; \boldsymbol{\eta}_p) = \sum_{i=1}^Q \alpha_i \exp[-2\pi^2 \tau^2 \sigma_i^2] \cos(2\pi \tau \mu_i), \quad (45)$$

where  $\boldsymbol{\eta}_p = [\alpha_1, \alpha_2, \dots, \alpha_Q, \mu_1, \mu_2, \dots, \mu_Q, \sigma_1^2, \sigma_2^2, \dots, \sigma_Q^2]^T$  denotes the hyper-parameters of the SM kernel to be optimized and  $\tau \triangleq x - x'$  owing to the stationary assumption. For accurate approximation, however, we need to choose a large  $Q$ , which potentially leads to an over-parameterized model with many redundant localized Gaussian components. Besides, optimizing the frequency and variance parameters is numerically difficult as a non-convex problem, and often incurs bad local minimal.

To remedy the aforementioned numerical issue, in [7], it was proposed to fix the frequency and variance parameters,  $\mu_1, \mu_2, \dots, \mu_Q, \sigma_1^2, \sigma_2^2, \dots, \sigma_Q^2$ , in the original SM kernel to some known grids and focus solely on the weight parameters,  $\alpha_1, \alpha_2, \dots, \alpha_Q$ . The resulting kernel is called *grid spectral mixture* (GridSM) kernel. By fixing the frequency and variance parameters, the above GridSM kernel can be regarded as a linear multiple kernel with  $Q$  basis subkernels,  $k_i(\tau) \triangleq \exp[-2\pi^2 \tau^2 \sigma_i^2] \cos(2\pi \tau \mu_i)$ ,  $i = 1, 2, \dots, Q$ . In [7], it was shown that for sufficiently small variance, each subkernel matrix has a low-rank smaller than  $N/2$ , namely half of the data size. Therefore, it falls under the formulation in (38). The corresponding weight parameters of such over-parameterized kernel can be obtained effectively via

optimizing the evidence function in (38), and the solution turns out to be sparse, as it is demonstrated in Section V.

### C. Sparsity-Aware Modeling for Tensor Decompositions

In the previous subsections, we elucidate the sparsity-aware modeling for the two recent supervised data analysis tools, namely the DNNs and GPs. The underlying idea of employing an over-parameterized model and embedding sparsity via an appropriate prior has inspired recent sparsity-aware modeling for the unsupervised learning tools in the context of tensor decomposition, see e.g., [10]–[15].

For a pedagogical purpose, we first introduce the basics of tensors and tensor *canonical polyadic decomposition* (CPD), the most fundamental tensor decomposition model in *unsupervised learning*.

1) *Tensors and CPD*: Tensors are regarded as *multi-dimensional generalization of matrices*, thus providing a natural representation for any multi-dimensional dataset. Specifically, a  $P$ -dimensional ( $P$ -D) dataset can be represented by a  $P$ -D tensor  $\mathcal{D} \in \mathbb{R}^{J_1 \times J_2 \times \dots \times J_P}$  [42]. Given a tensor-represented dataset  $\mathcal{D}$ , the unsupervised learning considered in this article aims to identify the underlying source signals that generate the observed data. In different fields, this task gets different names, such as “clustering” in social network analysis [43], “blind source separation” in electroencephalogram (EEG) and functional magnetic resonance imaging (fMRI) data analysis [44], [45], and “blind signal estimation” in radar/sonar signal processing [46]. In these applications, tensor CPD has been proven to be a powerful tool with good interpretability.

Formally, the tensor CPD is defined as follows:

*Definition of Tensor CPD* [42]: Given a  $P$ -D tensor  $\mathcal{D} \in \mathbb{R}^{J_1 \times J_2 \times \dots \times J_P}$ , CPD seeks to find the vectors  $\{\mathbf{a}_r^{(1)}, \mathbf{a}_r^{(2)}, \dots, \mathbf{a}_r^{(P)}\}_{r=1}^R$  such that

$$\begin{aligned} \mathcal{D} &= \sum_{r=1}^R \underbrace{\mathbf{a}_r^{(1)} \circ \mathbf{a}_r^{(2)} \circ \dots \circ \mathbf{a}_r^{(P)}}_{\text{rank-1 tensor}}, \\ &\triangleq \llbracket \mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \dots, \mathbf{A}^{(P)} \rrbracket, \end{aligned} \quad (46)$$

where  $\circ$  denotes vector outer product;  $\mathbf{A}^{(p)} \triangleq [\mathbf{a}_1^{(p)}, \mathbf{a}_2^{(p)}, \dots, \mathbf{a}_R^{(p)}] \in \mathbb{R}^{J_p \times R}, \forall p$ , is called the factor matrix. The minimal number  $R$  that yields the above expression is termed as the tensor rank.

From this definition, it is readily seen that the tensor CPD is a multi-dimensional generalization of a matrix decomposition in terms of rank-1 representation. Particularly, when  $P = 2$ , Eq. (46) reduces to decomposing a matrix  $D \in \mathbb{R}^{J_1 \times J_2}$  into the summation of  $R$  rank-1 matrices, i.e.,  $D = \sum_{r=1}^R \mathbf{a}_r^{(1)} \circ \mathbf{a}_r^{(2)}$ .



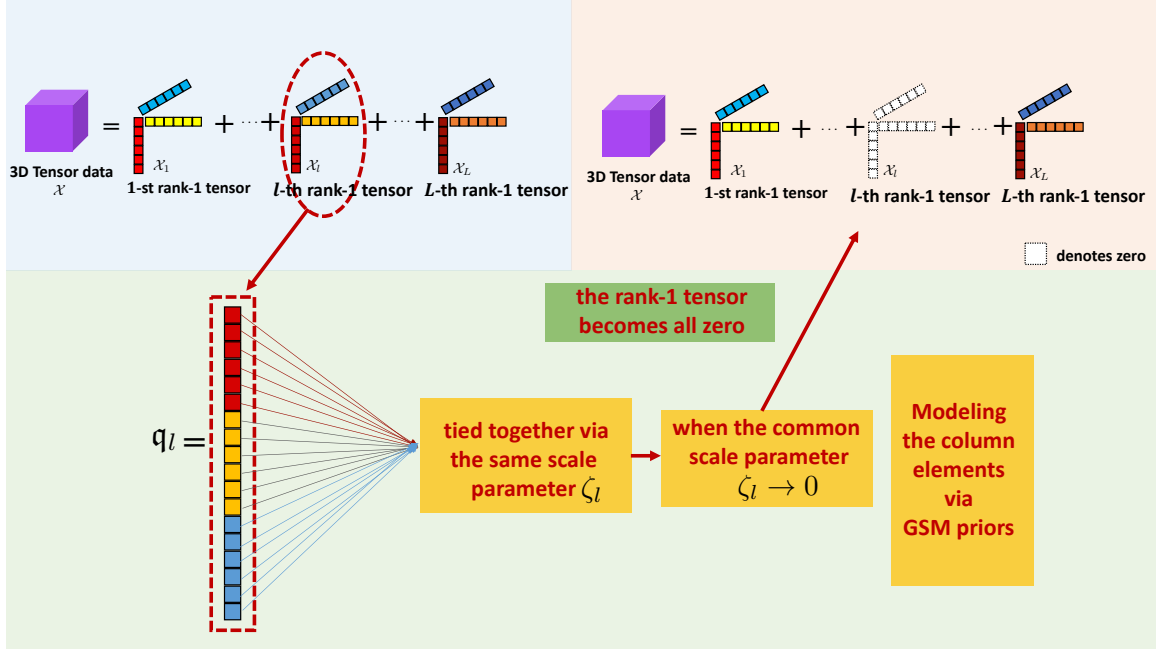


Fig. 12. Illustration of sparsity-aware modeling for rank-1 tensors using GSM priors.

By defining the term  $\mathbf{a}_r^{(1)} \circ \mathbf{a}_r^{(2)} \circ \dots \circ \mathbf{a}_r^{(P)}$  as a  $P$ -D rank-1 tensor, CPD essentially seeks for  $R$  rank-1 tensors/components from the observed dataset, each corresponding to one specific underlying source signal. Thus, the tensor rank  $R$  has a clear physical meaning, namely it corresponds to the number of the underlying source signals. Different from the matrix decomposition, where the rank-1 components are in general not unique, CPD for a  $P$ -D tensor ( $P > 2$ ) gives unique rank-1 components under mild conditions [42]. The uniqueness endows superior interpretability of the CPD model used in various unsupervised data analysis tasks.

2) *Low-Rank CPD and Sparsity-Aware Modeling*: In real-world data analysis, the number of underlying source signals is usually small. For instance, in brain-source imaging [44], [45], both the EEG and fMRI data analysis outcomes have shown that only a small fraction of source signals contribute to the observed brain activities. This suggests that the assumed CPD model should have a small tensor rank  $R$  to avoid data overfitting.

In the sequel, we show how the *low-rankness* is embedded into the CPD model through practicing the ideas reported in the previous two sections. First, we employ an over-parameterized model for CPD by assuming an upper-bound value  $L$  of tensor rank  $R$ , i.e.,  $L \gg R$ . The low-rankness implies that  $L - R$  rank-1 tensors should be zero, each specified by vectors  $\{\mathbf{a}_l^{(p)}\}_{p=1}^P, \forall l$ . In other words, let vector  $\mathbf{q}_l \triangleq [\mathbf{a}_l^{(1)}; \mathbf{a}_l^{(2)}; \dots; \mathbf{a}_l^{(P)}] \in \mathbb{R}^{\sum_{p=1}^P J_p}, \forall l$ . The low-rankness indicates that a number of vectors in the set

$\{\mathbf{q}_l\}_{l=1}^L$  are zero vectors. To model such sparsity, we adopt the following multivariate extension of GSM prior introduced in Section III-B, that is:

$$\begin{aligned} p(\mathbf{q}_l) &= \int \prod_{i=1}^{\sum_{p=1}^P J_p} \mathcal{N}([\mathbf{q}_l]_i; 0, \zeta_l) p(\zeta_l; \boldsymbol{\eta}_p) d\zeta_l, \\ &= \int \mathcal{N}(\mathbf{q}_l; 0, \zeta_l \mathbf{I}) p(\zeta_l; \boldsymbol{\eta}_p) d\zeta_l, \\ &= \int \prod_{p=1}^P \mathcal{N}(\mathbf{a}_l^{(p)}; 0, \zeta_l \mathbf{I}) p(\zeta_l; \boldsymbol{\eta}_p) d\zeta_l, \end{aligned} \quad (47)$$

where  $[\mathbf{q}_l]_i$  denotes the  $i$ -th element of vector  $\mathbf{q}_l$ . Since the elements in  $\mathbf{q}_l$  are assumed to be statistically independent, then according to the definition of a multivariate Gaussian distribution, we have the second and third lines of (47) showing the equivalent prior modeling on the concatenated vector  $\mathbf{q}_l$  and the associated set of vectors  $\{\mathbf{a}_l^{(p)}\}_{p=1}^P$ , respectively. The mixing distribution  $p(\zeta_l; \boldsymbol{\eta}_p)$  can be any one listed in Table I. Note that in (47), the elements in vector  $\mathbf{q}_l$  are tied together via a common hyper-parameter  $\zeta_l$ . Once the learning phase is over, if  $\zeta_l$  approaches zero, the elements in  $\mathbf{q}_l$  will shrink to zero simultaneously, thus nulling a rank-1 tensor, as illustrated in Fig. 12. Since the prior distribution given in (47) favors zero-valued rank-1 tensors, it promotes the low-rankness of the CPD model.

*Remark 2:* If the factor matrices are further constrained to be non-negative for enhanced interpretability in certain applications, simple modification, that is, multiplying a unit-step function  $\mathbf{U}(\mathbf{a}_l^{(p)} \geq 0)$  (which returns one when  $\mathbf{a}_l^{(p)} \geq 0$  or zero otherwise) to the prior derived in (47), can be made to embed both the non-negativeness and the low-rankness, see in-depth discussions in [11].

3) *Extensions to Other Tensor Decomposition Models:* Similar ideas have been applied to other tensor decomposition models including Tucker decomposition (TuckerD) [47] and tensor train decomposition (TTD) [48]–[50]. In these works, one first assumes an over-parametrized model by setting the model configuration parameters (e.g., multi-linear ranks in TuckerD and TT ranks in TTD) to be large numbers, and then impose GSM prior on the associated model parameters to control the model complexity, see detailed discussions in [47]–[50].

*Remark 3:* Some further suggestions are given on choosing appropriate tensor decomposition models for different data analysis tasks, see e.g., [51]. If the interpretability is crucial, one might try CPD (and its structured variants) first, due to its appealing uniqueness property. On the other hand, if the task is related to subspace learning, Tucker decomposition should be considered since its model parameters can be interpreted as the basis functions and the associated coefficients. For missing data imputation, TTD is a good choice as it disentangles different tensor dimensions. More concrete examples can be found in the recent overview paper [51].

## V. THE ART OF INFERENCE: EVIDENCE MAXIMIZATION AND VARIATIONAL APPROXIMATION

Having introduced sparsity-promoting priors, we are now at the stage of deriving the associated Bayesian inference algorithms that aim to learn both the posterior distributions of the unknown parameters/functions and the optimal configurations of the model hyper-parameters. In Section V-A, we will first show that the inference algorithms developed for our considered data analysis tools can be unified into a common evidence maximization framework. Then, for each data analysis tool, we will further show how to leverage recent advances in *variational approximation* and *non-convex optimization* to deal with specific problem structures for enhanced learning performance. Concretely, we introduce inference algorithms for GP in Section V-B, for tensor decompositions in Section V-C, and for Bayesian deep neural networks in Section V-D.

### A. Evidence Maximization Framework

Given a data analysis task and having selected the learning model,  $\mathcal{M}$ , that is the associated likelihood function  $p_{\mathcal{M}}(\mathcal{D}|\boldsymbol{\theta})$  and a sparsity-promoting prior  $p_{\mathcal{M}}(\boldsymbol{\theta}; \boldsymbol{\eta}_p)$ , the goal of Bayesian SAL is to infer the posterior distribution  $p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D}; \boldsymbol{\eta})$ , using the Bayes' theorem given in (1), and to compute the model hyper-parameters  $\boldsymbol{\eta}$  by maximizing the evidence  $p_{\mathcal{M}}(\mathcal{D}; \boldsymbol{\eta})$ .

We differentiate the following two cases of the evidence function. First, if the evidence  $p_{\mathcal{M}}(\mathcal{D}; \boldsymbol{\eta})$  defined in (5) can be derived analytically, such as (11) in the Bayesian linear regression example, the model hyper-parameters  $\boldsymbol{\eta}$  can be learnt via solving the evidence maximization problem, for which advanced non-convex optimization tools, e.g., [52]–[56], can be utilized to find high-quality solutions. In this case, since the prior, likelihood and evidence all have analytical expressions, applying the Bayes' theorem (1) yields a closed-form posterior distribution of the unknown parameters.

Unfortunately in most cases, the multiple integration required in computing the evidence (5) turns out to be analytically intractable. Inspired by the ideas of the *Minorize-Maximization* (also called Majorization-minimization (MM)) optimization framework [55], we can seek for a tractable lower bound (or a valid surrogate function in general) that minorizes the evidence function, and maximize the lower bound iteratively until convergence. It has been shown, see e.g., [57], [1], [19], that such an optimization process can push the evidence function to a stationary point. More concretely, the logarithm of the evidence function is lower bounded as follows:

$$\log p_{\mathcal{M}}(\mathcal{D}; \boldsymbol{\eta}) \geq \mathcal{L}(q(\boldsymbol{\theta}); \boldsymbol{\eta}), \quad (48)$$

where the lower bound

$$\mathcal{L}(q(\boldsymbol{\theta}); \boldsymbol{\eta}) \triangleq \int q(\boldsymbol{\theta}) \log \frac{p(\mathcal{D}, \boldsymbol{\theta}; \boldsymbol{\eta})}{q(\boldsymbol{\theta})} d\boldsymbol{\theta}, \quad (49)$$

is called *evidence lower bound* (ELBO), and  $q(\boldsymbol{\theta})$  is known as the variational distribution. The tightness of the ELBO is determined by the closeness between the variational distribution  $q(\boldsymbol{\theta})$  and the posterior  $p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D};\boldsymbol{\eta})$ , measured by the Kullback-Leibler (KL) divergence,  $\text{KL}(q(\boldsymbol{\theta})||p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D};\boldsymbol{\eta}))$ . In other words, the ELBO becomes tight, i.e., the lower bound becomes equal to the evidence when  $\text{KL}(q(\boldsymbol{\theta})||p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D};\boldsymbol{\eta})) = 0$ , which holds true if and only if  $q(\boldsymbol{\theta}) = p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D};\boldsymbol{\eta})$ . This is easy to see if we expand (49) and reformulate it as

$$\log p_{\mathcal{M}}(\mathcal{D};\boldsymbol{\eta}) = \mathcal{L}(q(\boldsymbol{\theta});\boldsymbol{\eta}) + \text{KL}(q(\boldsymbol{\theta})||p_{\mathcal{M}}(\boldsymbol{\theta}|\mathcal{D};\boldsymbol{\eta})). \quad (50)$$

Since the KL divergence is nonnegative, the equality in (48) holds if and only if it is equal to zero.

Since the ELBO in (49) involves two arguments, namely,  $q(\boldsymbol{\theta})$  and  $\boldsymbol{\eta}$ , solving the maximization problem

$$\max_{q(\boldsymbol{\theta}),\boldsymbol{\eta}} \mathcal{L}(q(\boldsymbol{\theta});\boldsymbol{\eta}), \quad (51)$$

can provide both an estimate of the model hyper-parameters and the posterior distributions. These two terms can be optimized in an alternating fashion. Different strategies for optimizing  $q(\boldsymbol{\theta})$  and  $\boldsymbol{\eta}$  result in different inference algorithms. For example, the variational distribution  $q(\boldsymbol{\theta})$  can be optimized either via functional optimization [58], or via Monte Carlo method [59], while the hyper-parameters  $\boldsymbol{\eta}$  can be optimized via various non-convex optimization methods [52]–[54], [60], [61].

In the following subsections, we will introduce some inference algorithms designed specifically for the three popular data analysis tools introduced in Section IV that have been equipped with certain sparsity-promoting priors.

### B. Inference Algorithms for GP Regression

Let us start with the GP model for regression, because in this case the evidence function  $p_{\mathcal{M}}(\mathcal{D};\boldsymbol{\eta})$  can be derived analytically owing to the Gaussian prior and likelihood assumed throughout the modeling process. In this subsection, we introduce an effective inference algorithm for GP regression based on the linear multiple kernel in (37). In Section IV, we have already derived the logarithm of the evidence function in analytical form, as shown in (38). Therefore, we can optimize it directly to obtain an estimate of the model hyper-parameters  $\boldsymbol{\eta}$ .

Traditionally, one could estimate the weights of the subkernels,  $\alpha_i, i = 1, 2, \dots, Q$ , as well as the precision parameter,  $\beta$ , using an iterative algorithm similar to the one derived in [40]. In particular, one sequentially solves for  $\alpha_i, i = 1, 2, \dots, Q$ , from the equation  $\gamma(\alpha_i) = 0$  derived in (41) by fixing the rest of the weights to their latest estimate and then check its relevance with the data in each iteration. This iterative method works quite well for various different datasets. In the sequel, however, we introduce a

potentially more effective numerical method in terms of the sensitivity to an initial guess and the data fitting performance than the original one [31].

Next, we take the GridSM kernel in (45) as an example of the linear multiple kernel, and rewrite the evidence maximization problem as

$$\boldsymbol{\eta}^* = \arg \min_{\boldsymbol{\eta}} l(\boldsymbol{\eta}) \triangleq g(\boldsymbol{\eta}) - h(\boldsymbol{\eta}), \quad (52)$$

where  $\boldsymbol{\eta} = [\boldsymbol{\alpha}^T; \beta]^T$  with  $\boldsymbol{\alpha} \geq \mathbf{0}$  and  $\beta > 0$ ,  $g(\boldsymbol{\eta}) \triangleq \mathbf{y}^T \mathbf{C}^{-1}(\boldsymbol{\eta}) \mathbf{y}$ , and  $h(\boldsymbol{\eta}) \triangleq -\log \det(\mathbf{C}(\boldsymbol{\eta}))$ . Let us introduce a short notation  $\mathbf{C}(\boldsymbol{\eta}) \triangleq \sum_{i=1}^Q \alpha_i \mathbf{K}_i + \beta^{-1} \mathbf{I}$ , where  $\mathbf{K}_i$  represents the  $i$ -th sub-kernel matrix evaluated with the training inputs. It can be shown that  $g(\boldsymbol{\eta})$  and  $h(\boldsymbol{\eta}) : \Theta \rightarrow \mathbb{R}$  are both convex and differentiable functions with  $\Theta$  being a convex set. Therefore, the cost function in (52) is a difference of two convex functions with respect to  $\boldsymbol{\eta}$ , and the optimization problem belongs to the well known *difference-of-convex* program (DCP) [62]. Instead of adopting the classic iterative procedure proposed for the RVM [40], we take advantages of the DCP optimization structure. Such a favorable structure may help the speed-up of the convergence process, and avoiding to be trapped in a bad local minimum of the optimization problem, and, thus, to further improve the level of sparsity [7].

1) *Sequential Majorization-Minimization (MM) Algorithm*: The main idea is to solve  $\min_{\boldsymbol{\eta} \in \Theta} l(\boldsymbol{\eta})$  with  $\Theta \subseteq \mathbb{R}^{Q+1}$  through an iterative scheme, where in each iteration a so-called majorization function  $\bar{l}(\boldsymbol{\eta}, \boldsymbol{\eta}^k)$  of  $l(\boldsymbol{\eta})$  at  $\boldsymbol{\eta}^k \in \Theta$  is minimized, i.e.,

$$\boldsymbol{\eta}^{k+1} = \arg \min_{\boldsymbol{\eta} \in \Theta} \bar{l}(\boldsymbol{\eta}, \boldsymbol{\eta}^k), \quad (53)$$

where the majorization function  $\bar{l}(\cdot, \cdot) : \Theta \times \Theta \rightarrow \mathbb{R}$  satisfies  $\bar{l}(\boldsymbol{\eta}, \boldsymbol{\eta}) = l(\boldsymbol{\eta})$  for  $\boldsymbol{\eta} \in \Theta$  and  $l(\boldsymbol{\eta}) \leq \bar{l}(\boldsymbol{\eta}, \boldsymbol{\eta}')$  for  $\boldsymbol{\eta}, \boldsymbol{\eta}' \in \Theta$ . We adopt the so-called linear majorization. Concretely, we make the convex function  $h(\boldsymbol{\eta})$  affine by performing the first-order Taylor expansion and obtain:

$$\bar{l}(\boldsymbol{\eta}, \boldsymbol{\eta}^k) \triangleq g(\boldsymbol{\eta}) - h(\boldsymbol{\eta}^k) - \nabla_{\boldsymbol{\eta}}^T h(\boldsymbol{\eta}^k) (\boldsymbol{\eta} - \boldsymbol{\eta}^k). \quad (54)$$

In this way, minimizing the cost function in (53) becomes a convex optimization problem in each iteration. By fulfilling the regularity conditions, the MM method is guaranteed to converge to a stationary point [55]. Next, we show how (53) with the linear majorization in (54) can be solved.

Since  $g(\boldsymbol{\eta})$  is a matrix fractional function, in each iteration (53) actually solves a convex matrix fractional minimization problem [62], which is equivalent to a semi-definite programming (SDP) problem via the Schur complement. This problem can be further cast into a second-order cone program (SOCP) problem and can efficiently and reliably be solved using the off-the-shelf convex solvers, e.g., MOSEK [7].

Although the previous MM algorithm can often lead to rather good solution, it cannot ensure local minimal in all cases. Occasionally, we found that it provides less satisfactory results, and they can be significantly improved by using a novel non-linearly constrained *alternating direction method of multipliers* (ADMM) algorithm as proposed in [7]. In general, the ADMM algorithm takes the form of a decomposition coordination procedure, where the original large problem is decomposed into a number of small local subproblems that can be solved in a coordinated way [63].

For our problem, the idea is to reformulate the original problem by introducing an  $N \times N$  matrix  $\mathbf{S}$  and solve instead

$$\begin{aligned} & \arg \min_{\mathbf{S}, \boldsymbol{\alpha}} \mathbf{y}^T \mathbf{S} \mathbf{y} - \log \det(\mathbf{S}), \\ & \text{s.t. } \mathbf{S} \left( \sum_i^Q \alpha_i \mathbf{K}_i + \beta^{-1} \mathbf{I} \right) = \mathbf{I}, \quad \boldsymbol{\alpha} \geq \mathbf{0}. \end{aligned} \quad (55)$$

Then, an augmented Lagrangian function can be formulated and solved by the ADMM algorithm through iteratively updating the auxiliary matrix variable  $\mathbf{S}$ , the kernel hyper-parameters,  $\boldsymbol{\alpha}$ , and some associated dual variables. By introducing an auxiliary matrix variable  $\mathbf{S}$ , all ADMM subproblems become convex; in particular, the weight parameters,  $\boldsymbol{\alpha}$ , are derived in closed form. From the experimental evaluation results given in [7], this ADMM algorithm can potentially find a better local minimum with improved prediction accuracy compared with the MM algorithm, however, at the cost of increased computational time in practice.

In addition to the above mentioned MM- and ADMM algorithms, one could also resort to some other advanced optimization algorithms to solve the problem; for instance, the successive convex approximation (SCA) algorithms reviewed in [55] are of great potential.

The computational complexity for one iteration of the MM algorithm is  $\mathcal{O}(n^2 \cdot \max(n, \sum_{i=1}^Q r_i))$ , where  $r_i$  stands for the rank of  $\mathbf{K}_i$ . The MM algorithm benefits from the low-rank property of the GSM subkernels. Let the average rank of the GSM subkernels be, i.e.,  $\bar{r} = \frac{1}{Q} \sum_{i=1}^Q r_i \ll n$ ; if  $Q\bar{r} > n$ , then the overall complexity of the MM algorithm scales as  $\mathcal{O}(Q \cdot \bar{r} \cdot n^2)$ ; otherwise, it scales as  $\mathcal{O}(n^3)$ . Similar conclusions hold for the ADMM algorithm too. These results show that the complexity also relies on the preselected number of subkernels,  $Q$ , which is often set to a larger value than the one that it is actually required; however, how to set this parameter adaptively and economically for different datasets remains an open challenge.

### C. Inference Algorithms for Bayesian Tensor Decompositions

In this subsection, we introduce the inference algorithm design for Bayesian tensor decompositions. Our focus will be on presenting the key ideas for deriving inference algorithms for the Bayesian tensor

CPD model [10]–[15] via the Gaussian likelihood and the GSM prior (introduced in Section IV-C). For other tensor decomposition formats, e.g., the Bayesian tensor TuckerD [47] and TTD [50], since they share the same prior design principle as that of CPD, the associated inference algorithm follows a similar rationale.

In the Bayesian tensor CPD, the goal of inference is to estimate the posterior distributions of factor matrices  $\{\mathbf{A}^{(p)} \in \mathbb{R}^{J_p \times L}\}_{p=1}^P$  from possibly incomplete  $P$ -D tensor data observations  $\mathcal{Y}_\Omega \in \mathbb{R}^{J_1 \times \dots \times J_P}$ , where  $\mathcal{Y}_{j_1, \dots, j_P}$  is observed if the  $P$ -tuple indices  $(j_1, \dots, j_P)$  belongs to the set  $\Omega$ . The forward problem is commonly modeled as a Gaussian likelihood:

$$p(\mathcal{Y}_\Omega | \{\mathbf{A}^{(p)}\}_{p=1}^P; \beta) = \prod_{(j_1, \dots, j_P) \in \Omega} \mathcal{N}(\mathcal{Y}_{j_1, \dots, j_P}; [\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(P)}]_{j_1, \dots, j_P}, \beta^{-1}), \quad (56)$$

where  $\beta$  is the precision (the inverse of variance) of the Gaussian noise. Since it is unknown, a non-informative prior (e.g., Jeffery prior,  $p(\beta) \propto 1/\beta$ ) can be employed. To promote the low-rankness, for the  $l$ -th columns of all the factor matrices, a GSM sparsity-promoting prior with latent variance variable  $\zeta_l$  has been adopted, see detailed discussions in Section IV-C. Usually, the hyper-parameters  $\boldsymbol{\eta}$  in the adopted GSM priors are pre-selected to make the prior non-informative, and thus need no further optimization. The unknown parameters  $\boldsymbol{\theta}$  include the factor matrices  $\{\mathbf{A}^{(p)}\}_{p=1}^P$ , the latent variance variables  $\{\zeta_l\}_{l=1}^L$  of the GSM priors, and the noise precision  $\beta$ . Under the evidence maximization framework, the inference problem can be formulated as (51) with unknown parameters<sup>5</sup>

$$\boldsymbol{\theta} \triangleq \{\{\mathbf{A}^{(p)}\}_{p=1}^P, \{\zeta_l\}_{l=1}^L, \beta\}, \quad (57)$$

and the joint pdf

$$p(\mathcal{D}, \boldsymbol{\theta}) \triangleq p(\mathcal{Y}_\Omega, \{\{\mathbf{A}^{(p)}\}_{p=1}^P, \{\zeta_l\}_{l=1}^L, \beta\}), \quad (58)$$

which can be computed by the product of the likelihood and the priors.

Without imposing any constraint on the pdf  $q(\boldsymbol{\theta})$ , the optimal solution is just the posterior, i.e.,  $q^*(\boldsymbol{\theta}) = p_{\mathcal{M}}(\boldsymbol{\theta} | \mathcal{Y}_\Omega)$ , whose computation using the Bayes' theorem will, however, encounter the intractable multiple integration challenge. To get over this difficulty, modern approximate inference techniques propose to solve problem (51) by further constraining  $q(\boldsymbol{\theta})$  into a functional family  $\mathcal{F}$ , i.e.,  $q(\boldsymbol{\theta}) \in \mathcal{F}$ . It is hoped that the family  $\mathcal{F}$  is as flexible as possible to allow accurate posterior estimates and at the same time simple enough to enable tractable optimization algorithm designs.

Among all the functional families, the *mean-field* family is undoubtedly the most favorable one in recent Bayesian tensor research [10]–[15]. It assumes that the variational pdf  $q(\boldsymbol{\theta}) = \prod_{k=1}^K q(\boldsymbol{\theta}_k)$ , where

<sup>5</sup>The expression of the objective function in (51) is quite lengthy, see e.g., [14], and thus is not included here.

$\theta$  is partitioned into mutually disjoint non-empty subsets  $\theta_k$  (i.e.,  $\cup_{k=1}^K \theta_k = \theta$  and  $\cap_{k=1}^K \theta_k = \emptyset$ ). In the context of the Bayesian tensor CPD, the mean-field assumption states that

$$q(\theta) = \prod_{p=1}^P q(\mathbf{A}^{(p)}) q(\{\zeta_l\}_{l=1}^L) q(\beta). \quad (59)$$

The factorized structure in (59) inspires the idea of block minimization in the optimization theory. In particular, for the ELBO maximization problem (51), specified after fixing the variational pdfs  $\{q(\theta_j)\}_{j \neq k}$ , the resulting subproblem that optimizes  $q(\theta_k)$  has been shown to have the following optimal solution, see e.g., [1]:

$$q^*(\theta_k) = \frac{\exp\left(\mathbb{E}_{\prod_{j \neq k} q(\theta_j)} [\ln p(\mathcal{D}, \theta)]\right)}{\int \exp\left(\mathbb{E}_{\prod_{j \neq k} q(\theta_j)} [\ln p(\mathcal{D}, \theta)]\right) d\theta_k}, \quad (60)$$

where  $\mathbb{E}_{q(\cdot)}[\cdot]$  denotes the expectation with respect to the variational pdf  $q(\cdot)$ . The inference framework under the mean-field assumption is termed as *mean-field variational inference* (MF-VI).

Whether the integration in the denominator (60) has a closed-form is determined by the functional forms of the likelihood and the priors. In particular, if they are conjugate pairs within the exponential family of probability distributions, see, discussions in, e.g., [1], [19], the optimal variational pdf in (60) will accept a closed-form expression. Fortunately, for the Bayesian tensor CPD adopting the Gaussian likelihood and the GSM prior for the columns of the factor matrices, this condition is usually satisfied, which enables the derivation of closed-form updates in recent advances [10]–[15].

*Remark 4:* To facilitate the algorithm derivation, MF-VI imposes a factorization structure on  $q(\theta)$ , which implies the statistical independence of the variables  $\theta_k$  given the observed dataset  $\mathcal{D}$ . If this is not the case, the mean-field approximation will lead to mismatch when approaching the ground-truth posteriors. In general, the MF-VI tends to provide posterior approximations that are more compact compared to the true ones, which means that the posterior estimates are usually “over-confident” [19]. To achieve more accurate posterior estimation, there is a research trend to employ more advanced variational approximation techniques than the mean-field approximation. For example, recent tensor-aided Bayesian deep neural network research [48] utilizes the kernelized Stein discrepancy to derive the inference algorithm that can approximate the posterior better. The interested reader may refer to [58] for some recent advances in variational approximation methods.

Some computational and theoretical difficulties that are commonly encountered in Bayesian tensor decompositions are summarized as follows. First, due to the block coordinate descent nature of the MF-VI [58], it is crucial to choose informative initial values to avoid poor local minima. On the other hand, the associated computational complexity is cubic with respect to the tensor rank, see, e.g., [10], [11],



[14], [50], which is high if the initial tensor rank is set to a large value. Finally, it was found that the algorithm performance significantly degrades in some challenging regimes, e.g., low signal-to-noise-ratio (SNR) and/or high rank, see, e.g., [10], [11], [14], [50]. To overcome these difficulties, suggestions on the real algorithm implementations are provided in Section VI-C.

#### D. Inference Algorithms for Bayesian Deep Neural Networks

The step of inference (training) for Bayesian deep neural networks follows the same backpropagation-type of philosophy as that of training their deterministic counterparts. There are, however, two notable differences. First, the unknown (synaptic) parameters/weights are now described via parameterized distributions. Thus, the cost function to be optimized has to be expressed in terms of the hyper-parameters that define the respective distributions, instead of the weights/synapses. This involves the so-called *reparameterization step* and we will describe it in Section V-D. Second, the evidence function to be maximized is not of a tractable form and it has to be approximated by its ELBO (see the definition in (49)).

In this subsection, we will outline the basic steps that are followed for variational inference in the case of a Bayesian deep network that comprises layers with: (a) stochastic LWTA blocks; (b) stochastic synaptic weights of Gaussian form; and (c) a sparsity-inducing mechanism imposed over the network synapses that is driven via an IBP prior. We have already discussed this type of network in Section IV-A.3; see, also, Fig. 10. To facilitate understanding, we provide a graphical illustration of the considered stochastic LWTA block in Fig. 13.

Without harming generality, let us focus on a specific layer, say, the  $f + 1$  one. In order to slightly unclutter notation, assume that for this layer, the input dimension  $a^f = L$  and the number of nodes (LWTA units)  $a^{f+1} = K$ . Thus the corresponding input matrix to the layer becomes  $\mathbf{X} \in \mathbb{R}^{N \times L}$  with  $N$  samples, each comprising  $L$  features. Under the stochastic LWTA-based modeling rationale, nodes (neurons) are replaced by LWTA blocks, each containing a set of  $J$  *competing linear* units. Thus, the layer input is now presented to each different block and each unit therein, via different weights. Thus, the weights for this layer are now represented via a three-dimensional matrix  $\mathbf{W} \in \mathbb{R}^{L \times K \times J}$  (we again refrain our notation on the dependence on  $f$ , i.e. the layer index). Recall from Section IV-A3 that each layer is associated with a latent discrete random vector,  $\boldsymbol{\xi}_n \in \text{one\_hot}(J)^K$ , that encodes the outcome of the local competition among the units in all  $K$  LWTA blocks of a network layer, when the  $n$ -th input sample is presented.

Furthermore, recall that each link connecting an input dimension, of the  $n$ -th sample, e.g.,  $x_{ni}$ , to an LWTA block, e.g., the  $k$ -th one, is weighted by a utility binary random variable,  $z_{ik}$ . This is set equal

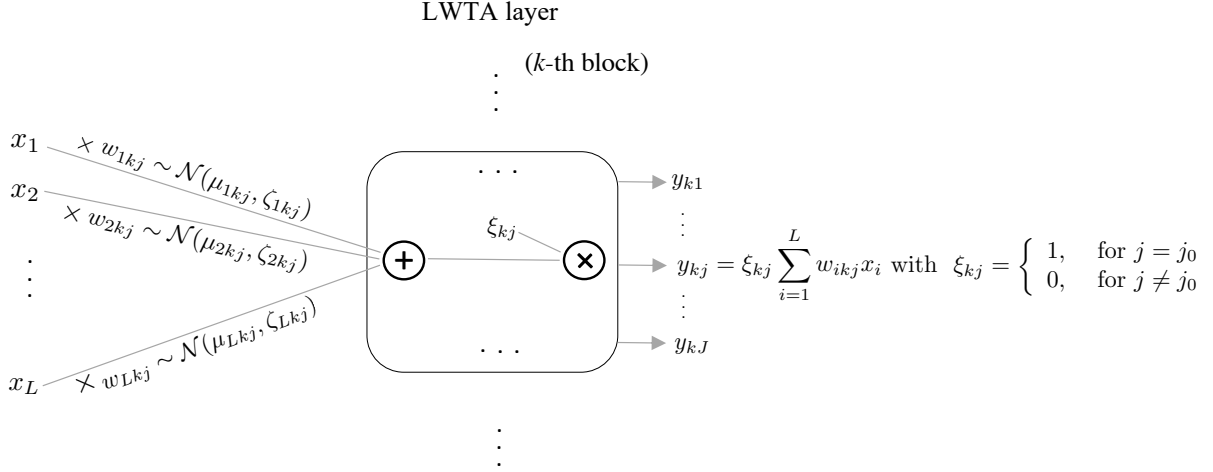


Fig. 13. A zoomed-in graphical illustration of the  $k$ -th block of a stochastic LWTA layer. Input  $\mathbf{x} = [x_1, x_2, \dots, x_L]$  is presented to each unit,  $j = 1, 2, \dots, J$ , in the block. Assume that the index of the winner unit is  $j = j_0$ . Then, the output of the block is a vector with a single non-zero value at index  $j_0$ .

to one, if the  $i$ -th dimension of the input is presented to the  $k$ -th LWTA block, otherwise  $z_{ik} = 0$ . We impose the sparsity-inducing IBP prior over these utility hidden variables.

We are now ready to write the output of a specific layer of the considered model, i.e.,  $\mathbf{y}_n \in \mathbb{R}^{K \cdot J}$ , as follows:

$$y_{nkj} = \xi_{nkj} \sum_{i=1}^L (w_{ikj} z_{ik}) x_{ni} \in \mathbb{R}, \quad (61)$$

where  $\mathbf{x}_n$  is the  $L$ -th dimensional input that coincides with the output of the previous layer. The involved random variables, whose *posterior* distributions are to be learnt during training, are: a) the synaptic weights,  $w_{ikj}$ ,  $i = 1, 2, \dots, L$ ,  $k = 1, 2, \dots, K$ ,  $j = 1, 2, \dots, J$ , for *all layers*, b) the utility variables,  $z_{ik}$ , for all layers and c) the indicator vectors,  $\xi_{nk}$ , for the  $n$ -th sample and the  $k$ -th LWTA, for all layers. The functional form of the respective distributions are:

■ Synaptic weights:

$$\text{Prior : } p(w_{ikj}) \sim \mathcal{N}(w_{ikj} | 0, 1), \quad \text{Posterior : } q(w_{ikj}) \sim \mathcal{N}(w_{ikj} | \mu_{ikj}, \zeta_{ikj}),$$

where the mean and variance,  $\mu_{ikj}$ ,  $\zeta_{ikj}$ , respectively, are learnt during training.

■ Utility binary random variables:

$$\text{Prior : } \text{Bernoulli}(z_{ik} | \pi_{ik}), \quad \text{Posterior : } q(z_{ik}) = \text{Bernoulli}(z_{ik} | \tilde{\pi}_{ik}),$$

where  $\pi_{ik}$  come from the IBP prior (Section III-C1) and  $\tilde{\pi}_{ik}$  are learnt during training. The use of the Bernoulli distribution is imposed by the binary nature of the variable.

■ Indicator random vectors,  $\xi_{nk}$ :

Prior :  $p(\xi_k) = \text{Categorical}(\xi_k | \frac{1}{J}, \dots, \frac{1}{J})$ , i.e., all linear units equiprobable.

Posterior :  $q(\xi_{nk}) = \text{Categorical}(\xi_{nk} | P_{nk1}, \dots, P_{nkJ})$ ,

where  $P_{nkj}$  is defined via the softmax operation, e.g., Eq. (36). The Categorical distribution is imposed because only one out of the  $J$  elements of  $\xi_{nk}$  is equal to 1 and the rest are zeros. The  $j$ -th element becomes 1 with probability  $P_{nkj}$ .

Note that besides the previous random variables, which are directly related to the DNN architecture, there is another set of hidden random variables, i.e., the  $u_j$ 's, which are used for generating the IBP prior. These have also to be considered as part of the palette of the involved random variables. As already said in Section III-C1, these follow the Beta distribution, with prior  $\text{Beta}(u_j | \alpha, 1)$  and posteriors  $\text{Beta}(u_j | a_j, b_j)$ , where  $a_j, b_j$  are learnt during training.

To train the proposed model, we resort to the maximization of the ELBO. The trainable model parameters, in our case, are the set of all the weights' posterior means and variances, i.e.,  $\mu_{ikj}$  and  $\zeta_{ikj}$ , the synaptic utility indicator posterior probabilities,  $\tilde{\pi}_{ik}$ , and the stick-variable posterior parameters  $a_j$  and  $b_j$ , across all network blocks and layers. Let us refer to this set as  $\Theta$ . We shall assume that our task comprises  $C$  classes and the softmax nonlinearity is used in the output layer.

In the following, we denote  $\mathcal{D}$  as the input-output training dataset. In addition, let  $\mathbf{Z}$  be the set of the synaptic utility indicators across the network layers;  $\Xi$  be the set of winner unit indicators across all blocks of all layers;  $\mathbf{W}$  be the set of synapse weights across all layers; and  $\mathbf{U}$  be the set of the stick-variables of the sparsity-inducing priors imposed across the network layers. Employing the mean-field approximation on the joint posterior pdf, i.e., factorizing  $q(\mathbf{W}, \mathbf{Z}, \Xi, \mathbf{U})$ , it is readily shown, e.g., [2], that

$$\begin{aligned} \text{ELBO}(\Theta) = & -\mathbb{E}_q \left[ \sum_{n=1}^N \sum_{c=1}^C y_{nc} \ln \tilde{y}_{nc}(\mathbf{x}_n; \mathbf{W}, \mathbf{Z}, \Xi, \mathbf{U}) \right] + \underbrace{\mathbb{E}_q \ln \frac{p(\mathbf{Z}|\mathbf{U})}{q(\mathbf{Z})} + \mathbb{E}_q \ln \frac{p(\mathbf{U})}{q(\mathbf{U})}}_{\text{regularizing terms}} \\ & + \underbrace{\mathbb{E}_q \ln \frac{p(\Xi)}{q(\Xi|\mathbf{Z}, \mathbf{W})} + \mathbb{E}_q \ln \frac{p(\mathbf{W})}{q(\mathbf{W})}}_{\text{regularizing terms}}, \end{aligned} \quad (62)$$

where  $y_{nc}$  are the outputs in the training set and  $\tilde{y}_{nc}$ ,  $c = 1, 2, \dots, C$ , are the class probability outputs as estimated via the softmax nonlinearities that implement the output layer. Observe that the first term

on the right hand side is the expectation of the cross-entropy of the network. The only difference with the deterministic DNNs that use this function to optimize the network is that now the expectation over the posterior is involved. In practice, it turns out that drawing one sample from the involved distributions suffices to lead to good approximation, provided that this sample is written as a differentiable function of  $\Theta$  and some low-variance random variable, e.g., [64], [65]. The rest of the terms in Eq. (62) are Kullback-Leibler (KL) divergences that bias the posteriors to be as close as possible to the corresponding priors. In other words, they act as regularizers that bias the solution towards certain regions in the parameter space as dictated by the adopted priors. For example, in the last term, the posterior of the synaptic weights is biased towards the normal Gaussian and bears close similarities with the  $\ell_2$  regularization, when it is enforced on the synaptic weights.

Drawing samples to approximate the expectations in Eq. (62) is closely related to what we called before as reparameterization. Recall that in the framework of the backpropagation algorithm, during the forward pass, one needs specific values/samples of the involved random parameters in order to compute the outputs, given the input to the network. This is performed via sampling the respective distributions, based on the current estimates of the involved posteriors. However, in order to be able to optimize with respect to their defining hyper-parameters, the corresponding current estimates should be explicitly considered. Let us take the Gaussian synaptic weights as an example. Let  $\mathcal{N}(w_{ikj}|\mu_{ikj}, \zeta_{ikj})$  be the current estimate of some weight in some layer. Instead of sampling from this Gaussian, it is easy to see that it is equivalent to obtaining a corresponding sample of the weight, as

$$\tilde{w}_{ikj} = \mu_{ikj} + \zeta_{ikj}^{1/2} \cdot \epsilon, \quad \epsilon \sim \mathcal{N}(0, 1). \quad (63)$$

In this way, every link in the network is determined explicitly by the pair  $(\mu_{ijk}, \zeta_{ijk})$ , and the backpropagation optimizes with respect to the means and variances. Reparameterization of the rest of the involved random variables follows a similar rationale, yet the involved formulae are slightly more complex; however, they are still given in terms of analytic expressions. For example, for the utility variables,  $z_{ik}$ , reparameterization is achieved via the posteriors  $\tilde{\pi}_{ik}$  and the so-called Gumbel-Softmax relaxation, e.g., [66]. For the stick breaking variables, reparameterization is achieved via the so-called Kumaraswamy approximation [67]. Details and the exact formulae can be found in [2].

Once samples have been drawn for all the involved variables, the ELBO in (62) is expressed directly in terms of the drawn sample, i.e.,  $\tilde{\mathbf{W}}, \tilde{\mathbf{Z}}, \tilde{\mathbf{\Xi}}, \tilde{\mathbf{U}}$ , without any expectation being involved. The trainable parameters' set,  $\Psi$ , can be obtained by means of any off-the-shelf gradient-based optimizer, such as the Adam [52]. Note that, by adopting the reported reparameterizations, one yields low-variance stochastic gradients that ensure convergence.

Training a fully-Bayesian model slightly increases the required complexity, since more parameters are involved, e.g., instead of a single weight one has to train with respect to the respective mean value and variance, as well as the hidden utility variables. However, the training timing remains of the same order as that required by the deterministic versions.

Once training has been completed, during testing, given an input: a) One can use the mean values of the obtained posterior Gaussians,  $\mu_{ijk}$ , in place of the synaptic weights,  $w_{ijk}$ . Sampling from the distribution could also be another possibility. Usually, the mean values are used. b) One can employ a threshold value, e.g.,  $\tau$ , and remove all links where the corresponding posterior  $\tilde{\pi}$  is below this threshold. Sampling is also another alternative. c) One samples from the respective categorical distributions to determine which linear unit “fires” in each block. Selecting the one with the largest probability is another alternative.

*Remark 5:* At this point, we must stress that the learnt posterior variances over the network weights can be also used for reducing the floating-point bit precision required for storing them; this effectively results in memory footprint reduction. The main rationale behind this process consists of the fact that, the higher the posterior weight variance, the more their fluctuation under sampling at inference time. This implies that, eventually, some bits fluctuate too much under sampling, and therefore, storing and retaining their values does not contribute to inference accuracy. Thus, Bayesian methods offer this added luxury, to optimally control the required bit precision for individual nodes. For example, in [2], it is reported that for the case of LENET-300-100 trained on MNIST, the bit precision can be reduced from 23 bit mantissa to just 2 bits. For more details, see, e.g., [2].

*Remark 6:* In [2], it was shown that the resulting architectures are able to yield a significant reduction in their computational footprint, retaining, nevertheless, state-of-the-art performance; positively, the flexibility of the link-wise IBP allowed for a more potent sparsity-activation-aware blend based on *stochastic* LWTA, offering significant benefits compared to conventional non-linearities introduced by the activation functions such as sigmoid and ReLU. For example, in the case of CIFAR-10 VGG-like network, it turns out that only around 5% of the original nodes are only retained, without affecting the performance of the network, even though quantized arithmetic was also employed to reduce the required number of bits, as explained in the previous remark, see, e.g., [2].

## VI. APPLICATIONS IN SIGNAL PROCESSING AND MACHINE LEARNING

In this section, we showcase typical applications of the sparsity-promoting data analysis tools introduced in Section IV. More specifically, advanced time series prediction using Gaussian process models is considered in Section VI-A; adversarial learning using Bayesian deep neural networks is presented

in Section VI-B; and lastly social group clustering and image completion using unsupervised tensor decompositions are demonstrated in Section VI-C.

### A. Time Series Prediction via GPs

In the following, we present an important signal processing and machine learning application, namely the *time series prediction*, using non-parametric GPs. We will focus on the GP regression models with the family of sparse spectrum kernels introduced in Section IV-B. To demonstrate the advantages of the sparsity-promoting GP models over other competing counterparts, we selected a number of classic time series datasets such as *CO<sub>2</sub>*, *Electricity*, *Unemployment*<sup>6</sup> as well as a “fresh” real-world *5G wireless traffic* dataset in our tests. Data descriptions are given in Table II.

Table II: Descriptions of the selected datasets. The training data,  $\mathcal{D}$ , is used for optimizing the hyper-parameters of the learning model, while the test data,  $\mathcal{D}_*$ , is used for evaluating the prediction accuracy. The numbers given in the last two columns are the training sample size and test sample size, respectively.

Name	Data Description	Training $\mathcal{D}$	Test $\mathcal{D}_*$
ECG	Electrocardiography of an ordinary person measured over a period of time	680	20
CO <sub>2</sub>	Carbon dioxide concentration observed from 1958 to 2003	481	20
Electricity	Monthly average residential electricity usage in Iowa City from 1971 to 1979	86	20
Employment	Wisconsin employment status observed from January 1961 to October 1975	158	20
Hotel	Monthly hotel occupied rooms collected from 1963 to 1976	148	20
Passenger	Passenger miles flown domestic U.K. from July 1962 to May 1972	98	20
Clay	Monthly production of clay bricks from January 1956 to August 1995	450	20
Unemployment	Monthly U.S. female (16-19 years) unemployment figures from 1948 to 1981	380	20
5G wireless traffic	Downlink data usage in a small cell observed in four weeks of 2021	607	67

1) *Classic Datasets*: In the sequel, we compare the performance of the sparsity-promoting GP models using the original SM kernels [35], [36] and the modified GridSM kernel [7] with that of a classic deep learning based time series prediction model, namely the long-short-term-memory (LSTM) [68], as well as a canonical statistical model, namely the autoregressive integrated moving average (ARIMA) model<sup>7</sup> [69], from various different aspects. Furthermore, we compare GP models with recently proposed Transformer-based time series prediction model, called *Informer*, which successfully addressed the computation issues and some inherent limitations of the encoder-decoder architecture in the original Transformer model. For

<sup>6</sup>These datasets are available from the UCL repository

<sup>7</sup>ARMA model can be regarded as a special case of a GP model adopting a specific sparse kernel matrix.

Table III: Comparisons of various time series prediction models in terms of the prediction MSE. Herein, we let both the GSMGP and the SMGP employ  $Q = 500$  Gaussian mixture modes in their kernels, and the SSGP employs the same amount of trigonometric basis functions. The GSMGP samples  $Q$  normalized frequency parameters,  $\mu_i$ ,  $i = 1, 2, \dots, Q$ , uniformly from  $[0, 1/2)$ , while set the variance parameter to  $\sigma = 0.001$ . The LSTM model follows a standard setup with one hidden layer and the dimension of the hidden state is set to 30. The Informer model follows the default setup given in the original paper [70]. The ARIMA( $p, d, q$ ) model is a standard one with ( $p = 5, d = 1, q = 2$ ).

Name	GSMGP MSE	SSGP MSE	SMGP MSE	LSTM MSE	Informer MSE	ARIMA MSE
ECG	<b>1.3E-02</b>	1.6E-01	1.9E-02	2.1E-02	5.4E-02	1.8E-01
CO2	1.5E+00	2.0E+02	<b>1.1E+00</b>	2.1E+00	8.4E+01	4.9E+00
Electricity	<b>4.7E+03</b>	8.2E+03	7.5E+03	<b>4.7E+03</b>	8.3E+03	1.2E+04
Employment	1.1E+02	7.7E+01	<b>0.7E+02</b>	4.3E+02	2.0E+03	3.9E+02
Hotel	<b>8.9E+02</b>	1.9E+04	2.8E+03	7.8E+03	2.3E+04	1.7E+04
Passenger	1.9E+02	6.9E+02	1.6E+02	1.6E+02	<b>1.2E+02</b>	4.5E+03
Clay	1.9E+02	5.3E+02	3.3E+02	2.7E+02	<b>1.4E+02</b>	3.3E+02
Unemployment	3.6E+03	2.1E+04	1.4E+04	<b>3.5E+03</b>	3.8E+03	1.5E+04

brevity, we name the first sparse spectrum GP model (equivalent to using trigonometric basis functions) proposed in [35] as *SSGP*, the GP model using the original SM kernel [36] as *SMGP*, and the most recent one with the rectified GridSM kernel [7] as *GSMGP*. Their configurations can be found in detail in [7].

Table III shows the obtained prediction accuracy of the various methods quantified in terms of the prediction *mean-squared-error* (MSE). It is readily observed that the sparsity-promoting GP models, and in particular SMGP and GSMGP, outperform all other competitors by far. The following facts need to be mentioned. For both the classic LSTM and Informer models to achieve good performance, the time series should, in general, be long so that the underlying pattern can be learnt during the training phase. The ARIMA model strongly relies on the optimal configuration of the parameters ( $p, d, q$ ), and it is incapable for long term prediction. In contrast, the sparsity-promoting GP models can automatically fit the underlying data pattern through solving the hyper-parameters from maximizing the evidence function. We have also shown in the supplement of [7] that the GSMGP is also superior to the GP models with elementary kernel (such as the SE kernel, rational quadratic kernel, etc.) as well as a hybrid of those.

Besides the improved prediction performance, the training time of the GSMGP outperforms the original

SMGP by far.<sup>8</sup> For the selected datasets, SMGP requires training time in the magnitude of  $10^3$  seconds, while the GSMGP only requires  $10^2$  seconds. By reducing the number of Gaussian modes,  $Q$ , the SMGP is able to reduce its training time albeit at the cost of sacrificing the fitting performance. On the other hand, the GSMGP improves its training time by fixing the frequency and variance parameters of the original SM kernel to known grids, so that the evidence maximization task enjoys the favorable difference-of-convex structure that can be efficiently handled by the MM algorithm introduced in Section V-B. In addition to the reduced training time, the overall optimization performance (including the convergence speed, chance of being trapped in a bad local minimum, etc.) and the sparsity level of the solution have been significantly improved. Detailed comparisons and pictorial illustrations can be found in [7]. In comparison, the LSTM and ARIMA models require the least training time in the magnitude of  $10^1$  seconds on average. However, due to the huge architecture adopted in the Informer model, the computational time is in the magnitude of  $10^2$  seconds on average. As it is readily observed from the results, the sparsity-promoting property helps reducing the computational time significantly; more importantly, the sparse solution identifies the most effective frequency components of the data and, thus, leads to good model interpretability.

2) *Real 5G Dataset*: Next, we focus on another favorable advantage of the GP models over their deep learning counterparts, namely the natural uncertainty region of a point prediction. We specifically select the real 5G wireless traffic dataset for visualization purposes due to the high demand of such a wireless application on a reasonable prediction uncertainty [25], [71]. The dataset was collected in a small cell of a southern city in China, and it contains the downlink data volume consumed by the mobile users located in the cell within each hour during a period of four weeks. Accurate prediction of the future downlink data consumption is vital to the operators for tuning the transmit power of the base station and switching on/off it automatically. In this application, uncertainty information is even more crucial because wrongly reducing the transmit power may largely influence the mobile users' surfing experience.

For this 5G wireless traffic prediction example, we constrained ourselves to a GSMGP with  $Q = 500$ , an SMGP with  $Q = 500$ , a standard GP with a hybrid of 3 elementary kernels (two periodic kernels plus an SE kernel) as was used in [25], and a classic LSTM deep learning model as described above. We used the data collected in the first 607 samples for training the models and used the last 67 samples to test their prediction performance. For comparing their prediction accuracy, we chose the *mean-absolute-percentage-error* (MAPE) measure, which is commonly used for evaluating the wireless traffic prediction

<sup>8</sup>Here, training time refers to the computational time required for training the learning models introduced in Section V.



error. The MAPE averaged over multiple test points is given by

$$e_{\text{MAPE}} = \frac{1}{n_*} \sum_{i=1}^{n_*} \left| \frac{y_i - \hat{y}_i}{y_i} \right| \times 100\%. \quad (64)$$

The prediction performance of the above learning models is shown in Fig. 14. It is readily seen that the GSMGP gives the best point prediction in terms of the MAPE. Moreover, as we mentioned before, the focus of this example is primarily on the uncertainty quantification. For GPs, the desired uncertainty region can be obtained naturally by computing the posterior variances associated with the test samples. In contrast, the classic LSTM model can only provide point predictions without any uncertainty quantification. A recent technique using the so-called deep ensembles [72] can be applied to quantify the predictive uncertainty of the LSTM model. The common characteristic of these techniques lies in that one has to train the models for multiple times, using different configurations (such as different initial guesses, step sizes, etc). However, such an approach increases substantially the computational load compared to the Bayesian approach, especially when complex (deterministic) learning models are involved. When comparing the uncertainty regions of the GP models, we can observe that the one using a mixture of elementary kernels tend to be conservative and show the largest uncertainty region. In contrast, both SMGP and GSMGP provide rather accurate point predictions as well as smaller uncertainty levels. It is noteworthy that SMGP presents less accurate point prediction (using its posterior mean) compared to that of GSMGP, but its uncertainty level is modestly larger than its counterpart. This suggests that, in this case, SMGP is less favorable because wrong decisions of switching on/off the BS are more likely made.

The above fitting results clearly demonstrate the advantages of the sparse spectrum kernel-based GP models; however, the obtained performance depends on the quality of the initialization. In particular, the method is sensitive to the initial guess of the SM kernel. According to our experience, a reliable initial guess can be obtained by fitting a periodogram (namely a nonparametric approximation of the true spectral density) in the frequency domain. We could also combine this strategy with the bootstrap technique to generate a number of candidate initial guesses for avoiding bad local minima. Codes for implementing the GSM kernel-based GP model are online available from [https://github.com/Paalis/MATLAB\\_GSM](https://github.com/Paalis/MATLAB_GSM).

### B. Adversarial Learning via Bayesian Deep Neural Networks

Despite the widespread success of DNNs, recent investigations have revealed their high susceptibility to *adversarial examples*; that is, cleverly crafted examples whose sole purpose is that of *fooling* a considered model into *misclassification*. Adversarial examples can be constructed using various approaches, e.g. FSGM [73] and CW [74]. A popular and powerful attack is the projected gradient descent (PGD) [75] attack. Under this scheme, the adversary is assumed to have access to the objective function of the

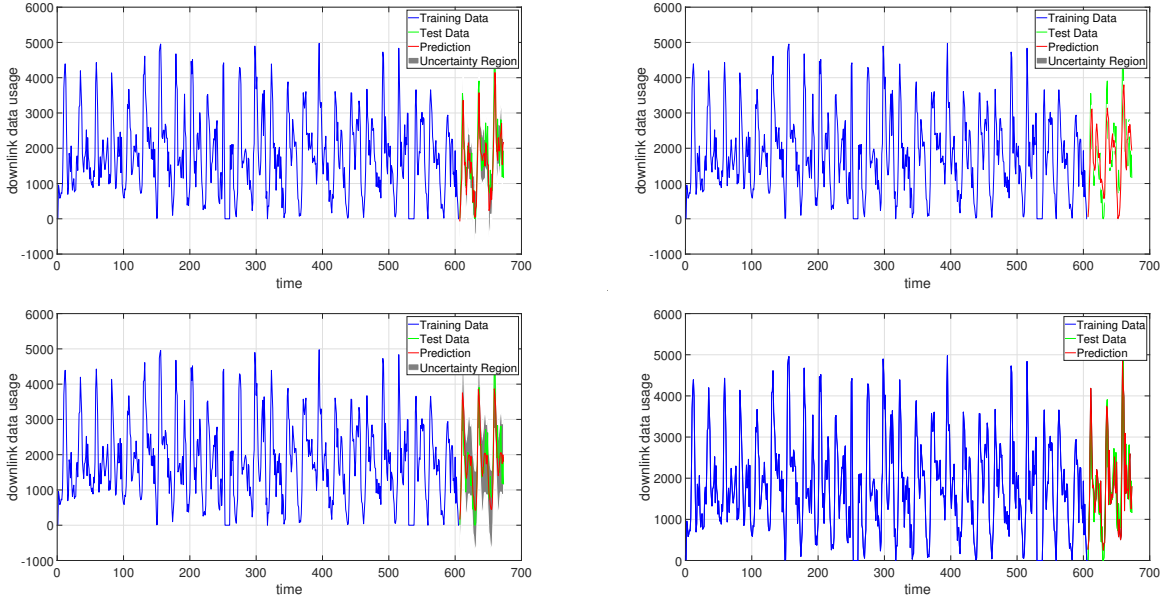


Fig. 14. Comparison of 5G wireless traffic prediction performance obtained from different models. Top left: GSMGP model with  $Q = 500$  fixed grids whose  $e_{\text{MAPE}} = 0.28$ ; Top right: SMGP model with  $Q = 500$  modes whose  $e_{\text{MAPE}} = 0.42$ ; Bottom left: a standard GP with a hybrid of 3 elementary kernels whose  $e_{\text{MAPE}} = 0.30$ ; Bottom right: LSTM model whose  $e_{\text{MAPE}} = 1.12$ . The gray shaded areas represent the uncertainty region (computed as the posterior variances) of the GP models.

target model,  $L(\mathbf{w}, \mathbf{x}, \mathbf{y})$ , where  $\mathbf{w}$  are the model trainable parameters,  $\mathbf{x}$  the input, and  $\mathbf{y}$  the predicted output variables. On this basis, the adversary performs an iterative computation; at each iteration,  $t$ , the adversary computes an (e.g.,  $\ell_\infty$ -bounded) adversarial perturbation of the training set examples  $\mathbf{x}$ , based on a multi-step PGD procedure that reads:

$$\mathbf{x}^{t+1} = \prod_{\mathbf{x} + \mathcal{S}} (\mathbf{x}^t + a \operatorname{sgn}(\nabla_{\mathbf{x}} L(\mathbf{w}, \mathbf{x}, \mathbf{y}))) \quad (65)$$

where  $\mathcal{S}$  is the set of allowed perturbations, that is the manipulative power of the adversary, e.g.  $\ell_\infty$ -ball around  $\mathbf{x}$ ;  $\operatorname{sgn}(\cdot)$  denotes the sign function that extracts the sign of a real number. In this context, even some minor, and many times imperceptible modifications, can successfully “*attack*” the model, resulting in severe performance degradation. This *frailness* of DNNs casts serious doubt about their deployment in *safety-critical applications*, e.g., autonomous driving [76].

Drawing upon this vulnerability, significant research effort has been recently devoted towards more reliable and robust DNNs. On this basis, several adversarial attacks and defenses have been proposed in the literature, e.g. adversarial training [75], [77], [78]. Among these, lies the *stochastic modeling* rationale; its main operating principle is founded upon the introduction of *stochasticity* in the considered

architecture, e.g., by randomizing the input data and/or the learning model itself [79]–[81]. Clearly, the Bayesian reasoning, which treats parameters as random entities instead of deterministic values, seeking to infer an appropriate generative process, seems to offer a natural stochastic defense framework towards more adversarially robust networks. It must be emphasized that the Bayesian techniques differ from the more standard randomized ones, which simply rely on the randomization of deterministic variables, in the context of the standard deterministic neural networks. Such techniques can be fairly easily handled and attacked. In contrast, in the Bayesian framework, the whole modeling and learning is built upon statistical arguments and the training involves learning of distributions.

Therefore, in the following, we focus on a recent application of the Bayesian rationale towards *adversarial robustness*. Specifically, we present the novel Bayesian deep network design paradigm proposed in [3] that yields state-of-the-art performance against powerful gradient-based adversarial attacks, e.g. PGD [75]. The key aspect of this method is its doubly stochastic nature stemming from two separate sampling processes relying on Bayesian arguments: a) the sparsity inducing non-parametric link-wise IBP prior introduced in Section IV-A3, and b) a stochastic adaptation of the biologically inspired and competition-based LWTA activation, as discussed in Sections IV-A3 and V-D.

We investigate the potency of LWTA-based networks against adversarial attacks under an Adversarial Training regime; we employ a PGD adversary [75]. To this end, we use the well-known WideResNet-34 [82] architecture, considering three different widen factors: 1, 5, and 10; note from the definition of the WideResNet-34, the larger the widen factor the larger the network. We focus on the CIFAR-10 dataset and adopt experimental settings similar to [83]. We use a batch size of 128 and an initial learning rate of 0.1; we halve the learning rate at every epoch after the 75-th epoch. We use a single sample for prediction. All experiments were performed using a single NVIDIA Quadro P6000.

For evaluating the robustness of this structure, we initially consider the conventional PGD attack with 20 steps, step size 0.007 and  $\epsilon = 8/255$ , which are the two parameters required by the PGD. In Table IV, we compare the robustness of LWTA-based WideResNet networks against the baseline results of [83]. As we observe, the Stochastic LWTA-based networks yield significant improvements in robustness under a traditional PGD attack; they retain extremely high natural accuracy (up to  $\approx 13\%$  better), while exhibiting a staggering, up to  $\approx 32.6\%$ , difference in robust accuracy compared to the *exact same architectures* employing the conventional ReLU-based nonlinearities and trained in the *same fashion*. Natural accuracy refers to the performance based on non-adversarial examples, while robust accuracy refers to the case where the network is tested against adversarial examples.

Further, to ensure that this approach does not cause the well-known obfuscated gradient problem [87], stronger parameter-free attacks were adopted using the newly introduced AutoAttack (AA) framework

Table IV: Natural and Robust accuracy under a conventional PGD attack with 20 steps and 0.007 step-size using WideResNet-34 models with different widen factors. We use the same PGD-based Adversarial Training scheme for all models [75].

Adversarial Training-PGD				
Widen Factor	Natural Accuracy (%)		Robust Accuracy (%)	
	Baseline	Stochastic LWTA	Baseline	Stochastic LWTA
1	74.04	<b>87.0</b>	49.24	<b>81.87</b>
5	83.95	<b>91.88</b>	54.36	<b>83.4</b>
10	85.41	<b>92.26</b>	55.78	<b>84.3</b>

Table V: Robust Accuracy (%) comparison under the AutoAttack framework. † denotes models that are trained with additional unlabeled data. The AutoAttack performance corresponds to the final robust accuracy after employing all the attacks in AA. Results directly from the AA leaderboard.

Method	AutoAttack
HE [84]	53.74
WAR [83]	54.73
Pre-training [85]†	54.92
[86]†	65.88
WAR [83]†	61.84
Ours (Stochastic-LWTA/PGD/WideResNet-34-1)	<b>74.71</b>
Ours (Stochastic-LWTA/PGD/WideResNet-34-5)	<b>81.22</b>
Ours (Stochastic-LWTA/PGD/WideResNet-34-10)	<b>82.60</b>

[88]. AA comprises an ensemble of four powerful white-box and black-box attacks, e.g., the commonly employed APGD attack; this is a step-free variant of the standard PGD attack [75], which avoids the complexity and ambiguity of step-size selection. In addition, for the entailed  $L_\infty$  attack, the common  $\epsilon = 8/255$  value was used. Thus, in Table V, we compare the LWTA-based networks to several recent state-of-the-art approaches evaluated on AA<sup>9</sup>. The reported accuracies correspond to the final reported robust accuracy of the methods after sequentially performing all the considered AA attacks. Once again, we observe that these stochastic and sparse networks yield state-of-the-art (SOTA) robustness against

<sup>9</sup><https://github.com/fra31/auto-attack>

all SOTA methods, with an improvement of  $\approx 16.72\%$ , even when compared with methods that employ substantial data augmentation to increase robustness, e.g. [86]. More results are reported in [3]. All results vouch for the potency of Stochastic LWTA networks in adversarial settings.

Finally, since the newly proposed networks consist of stochastic components, i.e., the competitive random sampling procedure to determine the winner in each LWTA block, the output of the classifier might change at each iteration; this obstructs the attacker from altering the final decision. To counter such randomness in the involved computations, in [88] the APGD attack is combined with an averaging procedure of 20 computations of the gradient at the same point. This technique is known as Expectation over Transformation (EoT) [87]. Thus, AA was used jointly with EoT for further performance evaluation of the LWTA-based networks. The corresponding results are presented in Table VI. As we observe, all of the considered networks retain state-of-the-art robustness against the powerful AA & EoT attacks. This conspicuously supports the usefulness of the stochastic LWTA activations towards adversarial robustness. Further explanations on why this performance is obtained are provided in, e.g., [2].

Table VI: Robustness against AA combined with 20 iterations of EoT. APGD-DLR corresponds to the APGD attack, using a different loss, i.e., the Difference of Logits Ratio [88].

Widen Factor	Natural Accuracy	APGD	APGD-DLR
1	87.00	79.67	76.15
5	91.88	81.67	77.65
10	92.26	82.55	79.00

### C. Unsupervised Learning via Bayesian Tensor Decompositions

In this subsection, we present some recent advances of Bayesian tensor decompositions in two unsupervised learning applications: *social group clustering* and *image completion*. The first application adopts Bayesian tensor CPD [11], [14], while the second one employs Bayesian tensor TTD [48]–[50]. Exploiting the GSM-based sparsity-promoting prior as introduced in Section IV-C and the effective MF-VI inference as introduced in Section V-C, the resulting algorithms offer nice features bypassing the need of hyper-parameters tuning and in dealing with overfitting.

1) *Bayesian Tensor CPD for Social Group Clustering*: In contrast to matrix decompositions, which are not unique in general (unless certain constraints are imposed), tensor CPD is provably unique under mild conditions [42]. This appealing property has made CPD an important tool for extracting the underlying signals/patterns from the observed data. The interpretability of CPD model can be further enhanced by

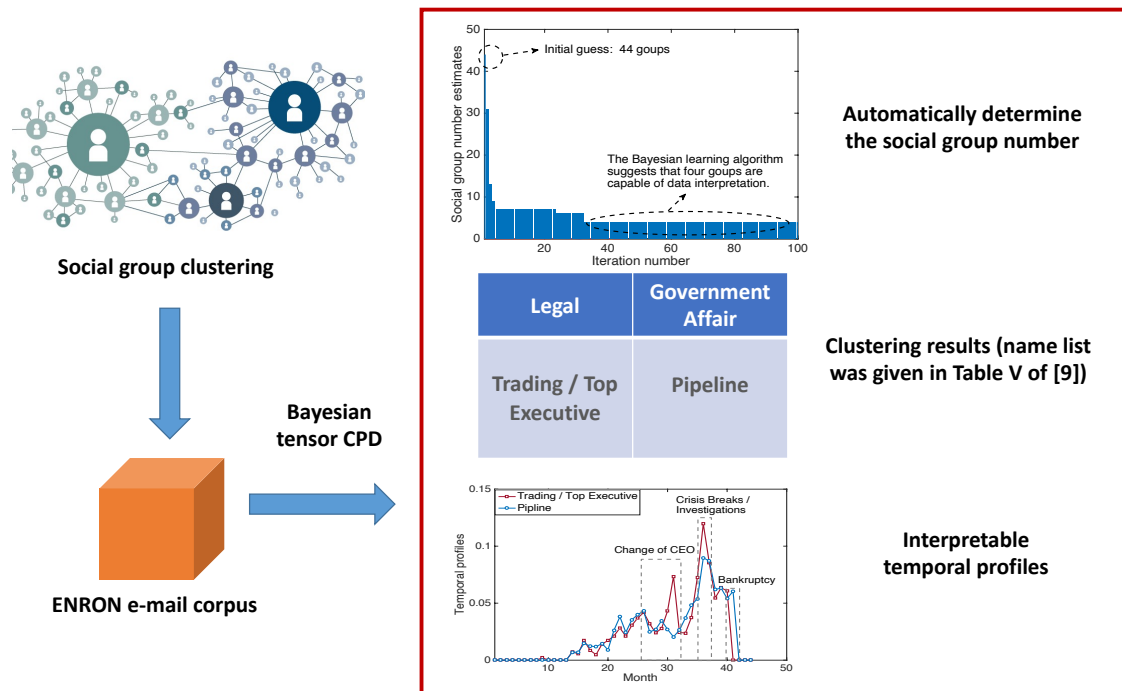


Fig. 15. Bayesian tensor CPD for social group clustering.

incorporating some side information, e.g., non-negativeness [11], into model learning. Here, using the *ENRON E-mail corpus* dataset (a 3-D tensor with the size  $184 \times 184 \times 44$ ), we demonstrate how the Bayesian tensor CPD (with non-negative factor matrices) [11] can be used to simultaneously determine the number of social groups, cluster people into different groups, and extract interpretable temporal profiles of different social groups.

The considered dataset records the number of E-mail exchanges between 184 people within 44 months. In particular, each entry is the number of E-mails exchanged between two people within a certain month. The physical meaning of three tensor dimensions are: the people who sent E-mails, the people who received E-mails, and the months, respectively. After applying Bayesian tensor CPD, the automatically determined tensor rank can be interpreted as the number of underlying social groups. For the first two factor matrices, the physical meaning for each element is that it quantifies the “score” that a particular person belongs to a particular E-mail sending and receiving group, respectively. For the third factor matrix, each column corresponds to the temporal profile of the associated social group, see discussion in [11].

Typically, we set the initial number of the social groups (i.e., tensor rank) large, (e.g., the minimal dimension of the tensor data, 44, in the above example), and then run the Bayesian learning algorithm

[11] to automatically determine the number of social groups that best interprets the data. As seen in Fig. 15, the estimated number of social groups gradually reduces to the value 4, indicating four underlying social groups. This is consistent with the results published in [43], [89], which are obtained via trial-and-error experiments. The clustering results can be read from the first factor matrix, which is of size  $184 \times 4$ . Specifically, for each column, only a few elements have non-zero values, and they can be used to identify the significance of the corresponding people in this social group. After sorting the scores of each column in the first factor matrix, the people with top 10 scores in each social group are shown in Table V of [11]. The clustering results are well interpretable as illustrated in Fig. 15. For example, the people in the first group work either in legal department or as lawyers, thus are clustered together. Moreover, interesting temporal patterns can be observed from the third factor matrix. It is clear that when the company has important events such as the change of CEO, crisis breaks and bankruptcy, distinct peaks appear.<sup>10</sup> The E-mail data analysis results have showcased the appealing advantage of Bayesian SAL in the context of tensor CPD, that is, the automatic determination of the social group number. This is important, since it leads to interpretable results on the group member and the temporal profiles can be naturally obtained.

2) *Bayesian Tensor TTD for Image Completion*: Color images are naturally 3-D tensor (with two spatial dimensions and one RGB dimension). To fully exploit the inherent structures of images, recent image completion works [90]–[92] usually fold an image into a higher dimensional tensor (e.g., 9-D tensor), and then apply tensor decompositions to recover the missing pixels, among which TTD is one of the most important tools due to its excellent performance. The folding operation is called tensor augmentation, see details in, e.g., [90]–[92]. For a  $P$ -D tensor, TTD has  $P - 1$  hyper-parameters (called TT ranks). Manually tuning different combinations of these hyper-parameters for overfitting avoidance is time-consuming. To facilitate this process, recent advances, see e.g., [48]–[50], first assume large values for TT-ranks, employ the sparsity-promoting GSM prior, and use variational inference for effective Bayesian SAL. The resulting algorithms can automatically learn the most suitable TT-ranks to match the underlying data pattern [48]–[50].

As an illustration, we consider the image completion of 5 images. Each image is with size  $256 \times 256 \times 3$ , and 80 percent of its pixels are randomly removed. After tensor augmentation, they are folded into a 9-D tensor with size  $16 \times 4 \times 4 \times 4 \times 4 \times 4 \times 4 \times 4 \times 3$ . We assume the initial TT-ranks is set as large as 60, and apply the Bayesian TTD algorithm [50] to complete the missing pixels. In comparison, we present the image completion results from other recent TTD algorithms: TTC-TV [90], TMAC-TT [91], and STTO [92], with the suggested hyper-parameter settings in their papers. The widely-used metrics, e.g., the peak

<sup>10</sup>This information can be acquired via checking the E-mail content and related research works, e.g., [43], [89].

signal-to-noise ratio (PSNR), were reported in Table II of [50], from which it can be concluded that the Bayesian TTD algorithm achieves the best overall performance. Particularly, in most cases, the Bayesian TTD algorithm recovers images with 1-5 dB higher PSNR than other algorithms. This is visually evident in the recovered images shown in Fig. 16. In this example, the Bayesian SAL-based tensor TTD [50] gets rid of the costly hyper-parameter tuning process for balancing the trade-off between data fitting and the noise overfitting; it directly learns these hyper-parameters from observations and shows excellent image restoration performance.

Some suggestions are provided on the real implementations of Bayesian tensor decomposition algorithms. a) Initialization: To assist the algorithm to avoid being trapped in a poor local minima, the initial factor matrix is usually set equal to the singular value decomposition approximation of the matrix, which is obtained by unfolding the tensor data along a specific dimension, see e.g., [10], [11], [14], [50]. b) On-the-fly pruning: To accelerate the learning process while not affecting the convergence, in each iteration, if some of the columns in the factor matrices are found to be indistinguishable from an all-zeros column vector, they can be safely pruned, see, e.g., the discussion in [10], [11]. c) Robustness against strong noise: When the corrupting noise sources are of large power, it was shown in [10] that slowing the noise precision learning can increase the robustness of the algorithm. Demo codes of Bayesian tensor CPD and TTD algorithms are online available from <https://github.com/leicheng-tensor?tab=repositories>.

## VII. CONCLUDING REMARKS AND FUTURE RESEARCH

In this article, we have presented an overview of some state-of-the-art sparsity-promoting priors for both Bayesian linear and non-linear modeling, as well as parametric and non-parametric models. In particular, these priors are incorporated into three advanced data analysis tasks, namely, the GP models, Bayesian deep neural networks, and tensor decomposition models, that can be applied to a wide spectrum of signal processing and machine learning applications. Commonly used inference algorithms for estimating the associated hyper-parameters and the (approximate) posterior distributions have also been discussed.

To demonstrate the effectiveness of the considered advanced sparsity-promoting models, we have carefully selected four important use cases, namely the time series prediction via the GP regression, adversarial learning via Bayesian deep neural networks, social group clustering and image completion using tensor decompositions. The reported results indicate that: a) Sparsity-promoting priors are able to adapt themselves to the given data and enable automatic model structure selection; b) The resulting sparse solution can better reveal the underlying (physical) characteristics of a target system/signal with only a few effective components; c) Sparsity-promoting priors, acting as the counterparts of the regularizers in optimization-based methods, can effectively help to avoid data overfitting, especially when the data size





Fig. 16. Experimental results for visual comparison on image completion with 80% missing data. Ground-truth images are in the top row. The second row includes the images with missing values. The third to the bottom rows include results from Bayesian TTD, TTC-TV, TMAC-TT, and STTO.

is relatively small; and (4) Sparsity-promoting priors lead to natural and more reasonable uncertainty quantification which is hard to obtain via traditional deep learning models.

Despite the rapid development of Bayesian learning and the enumerated advantages of sparsity-promoting models, still, such models are confronted with some challenges. Some open research directions are summarized as follows.

- *Quality of the posterior/predictive distribution.* As we mentioned before, a unique feature of Bayesian learning models lies in its posterior distribution that can be used to generate a point prediction and meanwhile provide an uncertainty quantification of the point prediction. Various recent works [23], [93], [94] indicate that the quality of the posterior distribution that is derived via the Bayesian deep neural networks and GP models can be significantly improved by using *cold tempering*:

$$p(\boldsymbol{\theta}|\mathbf{X}, \mathbf{y}) \propto (p(\mathbf{y}|\mathbf{X}, \boldsymbol{\theta}) \cdot p(\boldsymbol{\theta}))^{1/T}, \quad T < 1. \quad (66)$$

Two conjectures lie in the misspecification of the learning model and careless adoption of an inadequate, unintentionally informative prior [93]. Deeper analysis of such behavior is highly demanded. It is of great value to verify either analytically or experimentally if adopting the sparsity-promoting priors can help to avoid the use of cold tempering. It is also interesting to investigate the generalization property of the sparsity-promoting Bayesian learning models.

Another path to improve the quality of posterior/predictive distribution is through designing more effective inference methods. There is a recent trend to integrate the strengths of the variational inference [58] and Monte Carlo sampling [95] in a principled fashion, see, e.g., [96], in order to achieve the best trade-off between the inference accuracy and the computational efficiency. Given multiple inference results, Bayesian deep ensembles, e.g., [23], were also proposed for improved posterior/predictive distribution. It will be interesting to investigate how these recent general-purpose advances can be tailored to the sparsity-aware Bayesian modeling introduced in this article.

- *More emerging applications in complex systems.* We have witnessed various applications of Bayesian learning models, and they will surely continue to play important role in *large and complex systems*, such as 6G wireless communication systems [25] and autonomous systems [76], that are constantly facing rapid changing environments and critical decision making. Sparsity-promoting models are flexible enough to adapt themselves (for instance by nulling irrelevant basis kernels in the GP models) to changing data profile and provide rather reliable uncertainty quantification with small computational expense.
- *More and tighter interactions of the three data analysis tools.* Each of the three data analysis tool (introduced in this article) has already tapped into the design of other tools, see, e.g., DNN and GP

[16], GP and tensor [97], DNN and tensor [98], to achieve performance enhancement by borrowing the strengths of other tools. However, many of these works are not under the framework of Bayesian SAL, and thus do not possess the associated comparative advantages. It is promising to investigate how to combine the strengths of the three popular models, especially under the Bayesian SAL umbrella, to tackle challenging tasks such as non-linear regression for multi-dimensional and even heterogeneous data with deep kernels.

- *Sparsity-awareness in emerging learning paradigms.* Recently, we have witnessed various new paradigms, including, for instance, federated learning, life-long learning, meta learning, etc. We strongly believe that by further encoding sparsity-awareness through Bayesian sparse learning strategies, these emerging learning paradigms can further improve their learning efficiency over the learning models that were introduced in this article.

## REFERENCES

- [1] S. Theodoridis, *Machine learning: A Bayesian and optimization perspective, 2nd edition.* Academic press, 2020.
- [2] K. Panousis, S. Chatzis, and S. Theodoridis, “Nonparametric Bayesian deep networks with local competition,” in *Proc. International Conference on Machine Learning (ICML)*, 2019, pp. 4980–4988.
- [3] K. Panousis, S. Chatzis, A. Alexos, and S. Theodoridis, “Local competition and stochasticity for adversarial robustness in deep learning,” in *Proc. International Conference on Artificial Intelligence and Statistics (AISTAT)*, vol. 130, 2021, pp. 3862–3870.
- [4] K. Panousis, S. Chatzis, and S. Theodoridis, “Stochastic local winner-takes-all networks enable profound adversarial robustness,” in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [5] C. Louizos, K. Ullrich, and M. Welling, “Bayesian compression for deep learning,” in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 3288–3298.
- [6] S. Ghosh, J. Yao, and F. Doshi-Velez, “Model selection in Bayesian neural networks via horseshoe priors,” *Journal of Machine Learning Research*, vol. 20, no. 182, pp. 1–46, 2019.
- [7] F. Yin, L. Pan, T. Chen, S. Theodoridis, and Z.-Q. Luo, “Linear multiple low-rank kernel based stationary Gaussian processes regression for time series,” *IEEE Transactions on Signal Processing*, vol. 68, pp. 5260–5275, 2020.
- [8] T. Paananen, J. Piironen, M. R. Andersen, and A. Vehtari, “Variable selection for Gaussian processes via sensitivity analysis of the posterior predictive distribution,” in *Proc. International Conference on Artificial Intelligence and Statistics (AISTAT)*, 2019, pp. 1743–1752.
- [9] H. Kim and Y. W. Teh, “Scaling up the automatic statistician: Scalable structure discovery using Gaussian processes,” in *Proc. International Conference on Artificial Intelligence and Statistics (AISTAT)*, vol. 84, Lanzarote, Spain, 2018, pp. 575–584.
- [10] L. Cheng, Z. Chen, Q. Shi, Y.-C. Wu, and S. Theodoridis, “Towards flexible sparsity-aware modeling: Automatic tensor rank learning using the generalized hyperbolic prior,” *IEEE Transactions on Signal Processing*, vol. 1, no. 1, pp. 1–16, 2022.
- [11] L. Cheng, X. Tong, S. Wang, Y.-C. Wu, and H. V. Poor, “Learning nonnegative factors from tensor data: Probabilistic modeling and inference algorithm,” *IEEE Transactions on Signal Processing*, vol. 68, pp. 1792–1806, 2020.

- [12] Y. Zhou and Y.-M. Cheung, “Bayesian low-tubal-rank robust tensor factorization with multi-rank determination,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 62–76, 2019.
- [13] L. Cheng, Y.-C. Wu, and H. V. Poor, “Probabilistic tensor canonical polyadic decomposition with orthogonal factors,” *IEEE Transactions on Signal Processing*, vol. 65, no. 3, pp. 663–676, 2016.
- [14] Q. Zhao, L. Zhang, and A. Cichocki, “Bayesian CP factorization of incomplete tensors with automatic rank determination,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 9, pp. 1751–1763, 2015.
- [15] Z. Zhang and C. Hawkins, “Variational Bayesian inference for robust streaming tensor factorization and completion,” in *Proc. IEEE International Conference on Data Mining (ICDM)*, Singapore, 2018, pp. 1446–1451.
- [16] Y. Dai, T. Zhang, Z. Lin, F. Yin, S. Theodoridis, and S. Cui, “An interpretable and sample efficient deep kernel for Gaussian process,” in *Proc. International Conference on Uncertainty in Artificial Intelligence (UAI)*, 2020, pp. 759–768.
- [17] L. Cheng, Y.-C. Wu, and H. V. Poor, “Scaling probabilistic tensor canonical polyadic decomposition to massive data,” *IEEE Transactions on Signal Processing*, vol. 66, no. 21, pp. 5534–5548, 2018.
- [18] L. Cheng and Q. Shi, “Towards overfitting avoidance: Tuning-free tensor-aided multi-user channel estimation for 3D massive MIMO communications,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 3, pp. 832–846, 2021.
- [19] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [20] S. F. Gull, “Bayesian inductive inference and maximum entropy,” in *Maximum-entropy and Bayesian methods in science and engineering*. Springer, 1988, pp. 53–74.
- [21] G. Schwarz, “Estimating the dimension of a model,” *The annals of statistics*, pp. 461–464, 1978.
- [22] D. J. MacKay, “Bayesian interpolation,” *Neural computation*, vol. 4, no. 3, pp. 415–447, 1992.
- [23] A. G. Wilson and P. Izmailov, “Bayesian deep learning and a probabilistic perspective of generalization,” in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [24] C. E. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [25] Y. Xu, F. Yin, W. Xu, J. Lin, and S. Cui, “Wireless traffic prediction with scalable Gaussian process: Framework, algorithms, and verification,” *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 6, pp. 1291–1306, June 2019.
- [26] D. J. C. MacKay, “Gaussian processes – a replacement for supervised neural networks?” 1997. [Online]. Available: <http://www.inference.org.uk/mackay/gp.pdf>
- [27] J. Lee, J. Sohl-dickstein, J. Pennington, R. Novak, S. Schoenholz, and Y. Bahri, “Deep neural networks as Gaussian processes,” in *Proc. of International Conference on Learning Representations (ICLR)*, Vancouver, BC, Canada, 2018.
- [28] Y. C. Eldar and G. Kutyniok, *Compressed sensing: theory and applications*. Cambridge University Press, 2012.
- [29] M. Elad, *Sparse and redundant representations: from theory to applications in signal and image processing*. Springer Science Business Media, 2010.
- [30] D. F. Andrews and C. L. Mallows, “Scale mixtures of normal distributions,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 36, no. 1, pp. 99–102, 1974.
- [31] D. P. Wipf and B. D. Rao, “Sparse Bayesian learning for basis selection,” *IEEE Transactions on Signal processing*, vol. 52, no. 8, pp. 2153–2164, 2004.
- [32] D. J. MacKay, “Probable networks and plausible predictions—a review of practical bayesian methods for supervised neural networks,” *Network: computation in neural systems*, vol. 6, no. 3, pp. 469–505, 1995.
- [33] J. Bridle, “Training stochastic model recognition algorithms as networks can lead to maximum mutual information estimation of parameters,” *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, vol. 2, 1989.
- [34] D. Duvenaud, J. R. Lloyd, R. Grosse, J. B. Tenenbaum, and Z. Ghahramani, “Structure discovery in nonparametric

- regression through compositional kernel search,” in *Proc. International Conference on Machine Learning (ICML)*, Atlanta, USA, 2013, pp. 1166–1174.
- [35] M. Lázaro-Gredilla, J. Quiñonero Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal, “Sparse spectrum Gaussian process regression,” *Journal of Machine Learning Research*, vol. 11, pp. 1865–1881, August 2010.
- [36] A. Wilson and R. P. Adams, “Gaussian process kernels for pattern discovery and extrapolation,” in *Proc. International Conference on Machine Learning (ICML)*, Atlanta, USA, 2013, pp. 1067–1075.
- [37] J. Shawe-Taylor and N. Cristianini, *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- [38] A. Rahimi and B. Recht, “Random features for large-scale kernel machines,” in *Proc. International Conference on Neural Information Processing Systems*, Vancouver, British Columbia, Canada, 2007, pp. 1177–1184.
- [39] M. W. Seeger, “Bayesian inference and optimal design for the sparse linear model,” *Journal of Machine Learning Research*, vol. 9, pp. 759–813, June 2008.
- [40] M. E. Tipping and A. Faul, “Fast marginal likelihood maximisation for sparse Bayesian models,” in *Proc. International Conference on Artificial Intelligence and Statistics (AISTAT)*, Key West, Florida, USA, 2003, pp. 3–6.
- [41] N. I. Achieser, *Theory of Approximation*. Dover Publications, Inc., New York, 1992.
- [42] N. D. Sidiropoulos, L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos, “Tensor decomposition for signal processing and machine learning,” *IEEE Transactions on Signal Processing*, vol. 65, no. 13, pp. 3551–3582, 2017.
- [43] E. E. Papalexakis, N. D. Sidiropoulos, and R. Bro, “From k-means to higher-way co-clustering: Multilinear decomposition with sparse latent factors,” *IEEE Transactions on Signal Processing*, vol. 61, no. 2, pp. 493–506, 2013.
- [44] H. Becker, L. Albera, P. Comon, R. Gribonval, F. Wendling, and I. Merlet, “Brain-source imaging: From sparse to tensor models,” *IEEE Signal Processing Magazine*, vol. 32, no. 6, pp. 100–112, 2015.
- [45] C. Chatzichristos, E. Kofidis, M. Morante, and S. Theodoridis, “Blind fMRI source unmixing via higher-order tensor decompositions,” *Journal of neuroscience methods*, vol. 315, pp. 17–47, 2019.
- [46] H. Chen, F. Ahmad, S. Vorobyov, and F. Porikli, “Tensor decompositions in wireless communications and MIMO radar,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 3, pp. 438–453, 2021.
- [47] Q. Zhao, L. Zhang, and A. Cichocki, “Bayesian sparse tucker models for dimension reduction and tensor completion,” *ArXiv*, 2015. [Online]. Available: <https://arxiv.org/abs/1505.02343>
- [48] C. Hawkins and Z. Zhang, “Bayesian tensorized neural networks with automatic rank selection,” *Neurocomputing*, vol. 453, pp. 172–180, Sep. 2021.
- [49] L. Xu, L. Cheng, N. Wong, and Y.-C. Wu, “Probabilistic tensor train decomposition with automatic rank determination from noisy data,” in *2021 IEEE Statistical Signal Processing Workshop (SSP)*. IEEE, 2021, pp. 461–465.
- [50] —, “Overfitting avoidance in tensor train factorization and completion: Prior analysis and inference,” in *2021 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2021, pp. 1439–1444.
- [51] Y. Panagakis, J. Kossaiifi, G. G. Chrysos, J. Oldfield, M. A. Nicolaou, A. Anandkumar, and S. Zafeiriou, “Tensor methods in computer vision and deep learning,” *Proceedings of the IEEE*, vol. 109, no. 5, pp. 863–890, 2021.
- [52] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [53] M. Razaviyayn, M. Hong, and Z.-Q. Luo, “A unified convergence analysis of block successive minimization methods for nonsmooth optimization,” *SIAM Journal on Optimization*, vol. 23, no. 2, pp. 1126–1153, 2013.
- [54] M. Hong, Z.-Q. Luo, and M. Razaviyayn, “Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems,” *SIAM Journal on Optimization*, vol. 26, no. 1, pp. 337–364, 2016.
- [55] Y. Sun, P. Babu, and D. P. Palomar, “Majorization-minimization algorithms in signal processing, communications, and machine learning,” *IEEE Transactions on Signal Processing*, vol. 65, no. 3, pp. 794–816, 2016.

- [56] T.-H. Chang, M. Hong, H.-T. Wai, X. Zhang, and S. Lu, “Distributed learning in the nonconvex world: From batch data to streaming and beyond,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 26–38, 2020.
- [57] G. Parisi and R. Shankar, *Statistical Field Theory*. Westview Press, 1988.
- [58] C. Zhang, J. Bütetpage, H. Kjellström, and S. Mandt, “Advances in variational inference,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 8, pp. 2008–2026, 2018.
- [59] R. Zhang, C. Li, J. Zhang, C. Chen, and A. G. Wilson, “Cyclical stochastic gradient MCMC for Bayesian deep learning,” in *Proc. International Conference on Learning Representations (ICLR)*, 2020.
- [60] L. Bottou, “Large-scale machine learning with stochastic gradient descent,” in *Proc. International Conference on Computational Statistics (COMPSTAT)*. Paris, France: Springer, 2010, pp. 177–186.
- [61] G. Lan, *First-order and Stochastic Optimization Methods for Machine Learning*. Springer, 2020.
- [62] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [63] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, January 2011.
- [64] O. Shayer, D. Levi, and E. Fetaya, “Learning discrete weights using the local reparameterization trick,” in *Proc. International Conference on Learning Representations (ICLR)*, Vancouver, BC, Canada, 2018.
- [65] D. P. Kingma and M. Welling, “Auto-encoding variational bayes.” In *International Conference on Learning Representations*, 2014.
- [66] C. J. Maddison, A. Mnih, and Y. W. Teh, “The concrete distribution: A continuous relaxation of discrete random variables.” In *International Conference on Learning Representations*, 2017.
- [67] P. Kumaraswamy, “A generalized probability density function for double-bounded random processes,” *Journal of hydrology*, vol. 46, no. 1-2, pp. 79–88, 1980.
- [68] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [69] P. J. Brockwell and R. A. Davis, *Time Series: Theory and Methods*. Berlin, Heidelberg: Springer-Verlag, 1986.
- [70] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, “Informer: Beyond efficient transformer for long sequence time-series forecasting,” in *Proc. AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, May 2021, pp. 11 106–11 115.
- [71] Y. Xu, F. Yin, W. Xu, C.-H. Lee, J. Lin, and S. Cui, “Scalable learning paradigms for data-driven wireless communication,” *IEEE Communications Magazine*, vol. 58, no. 10, pp. 81–87, 2020.
- [72] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” in *Proc. International Conference on Neural Information Processing Systems (NIPS)*, 2017, pp. 6405–6416.
- [73] I. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” in *Proc. International Conference on Learning Representation (ICLR)*, San Diego, CA, USA, 2015.
- [74] N. Carlini and D. A. Wagner, “Towards evaluating the robustness of neural networks,” in *Proc. IEEE Symposium on Security and Privacy*. San Jose, CA, USA: IEEE Computer Society, May 2017, pp. 39–57.
- [75] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” *arXiv preprint arXiv:1706.06083*, 2017.
- [76] R. McAllister, Y. Gal, A. Kendall, M. van der Wilk, A. Shah, R. Cipolla, and A. Weller, “Concrete problems for autonomous vehicle safety: Advantages of Bayesian deep learning,” in *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, 2017, pp. 4745–4753.



- [77] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel, “Ensemble adversarial training: Attacks and defenses,” in *Proc. International Conference on Learning Representations (ICLR)*, Vancouver, BC, Canada, 2018.
- [78] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb, “Learning from simulated and unsupervised images through adversarial training,” in *Proc. IEEE conference on computer vision and pattern recognition (CVPR)*, Honolulu, HI, USA, 2017, pp. 2107–2116.
- [79] C. Xie, J. Wang, Z. Zhang, Z. Ren, and A. Yuille, “Mitigating adversarial effects through randomization,” in *Proc. International Conference on Learning Representations (ICLR)*, Vancouver, BC, Canada, 2018.
- [80] A. Prakash, N. Moran, S. Garber, A. DiLillo, and J. Storer, “Deflecting adversarial attacks with pixel deflection,” in *Proc. IEEE conference on computer vision and pattern recognition (CVPR)*, Salt Lake City, UT, USA, 2018, pp. 8571–8580.
- [81] G. S. Dhillon, K. Azizzadenesheli, Z. C. Lipton, J. Bernstein, J. Kossaifi, A. Khanna, and A. Anandkumar, “Stochastic activation pruning for robust adversarial defense,” in *Proc. International Conference on Learning Representations (ICLR)*, Vancouver, BC, Canada, 2018.
- [82] S. Zagoruyko and N. Komodakis, “Wide residual networks,” in *Proc. BMVC*, 2016.
- [83] B. Wu, J. Chen, D. Cai, X. He, and Q. Gu, “Do wider neural networks really help adversarial robustness?” in *In Proc. NIPS*, 2021.
- [84] T. Pang, X. Yang, Y. Dong, K. Xu, J. Zhu, and H. Su, “Boosting adversarial training with hypersphere embedding,” in *Proc. NIPS*, 2020.
- [85] D. Hendrycks, K. Lee, and M. Mazeika, “Using pre-training can improve model robustness and uncertainty,” *In Proc. ICML*, 2019.
- [86] S. Gowal, C. Qin, J. Uesato, T. Mann, and P. Kohli, “Uncovering the limits of adversarial training against norm-bounded adversarial examples,” *arXiv preprint arXiv:2010.03593*, 2021.
- [87] A. Athalye, N. Carlini, and D. A. Wagner, “Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples,” in *Proc. ICML*, 2018.
- [88] F. Croce and M. Hein, “Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks,” in *Proc. ICML*, 2020.
- [89] B. W. Bader, T. G. Kolda, and R. A. Harshman, “Temporal analysis of social networks using three-way decom.” Sandia National Lab.(SNL-NM), Albuquerque, NM (United States); Sandia, Tech. Rep., 2006.
- [90] C.-Y. Ko, K. Batselier, L. Daniel, W. Yu, and N. Wong, “Fast and accurate tensor completion with total variation regularized tensor trains,” *IEEE Transactions on Image Processing*, vol. 29, pp. 6918–6931, 2020.
- [91] J. A. Bengua, H. N. Phien, H. D. Tuan, and M. N. Do, “Efficient tensor completion for color image and video recovery: Low-rank tensor train,” *IEEE Transactions on Image Processing*, vol. 26, no. 5, pp. 2466–2479, 2017.
- [92] L. Yuan, Q. Zhao, and J. Cao, “High-order tensor completion for data recovery via sparse tensor-train optimization,” in *Proc. IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2018, pp. 1258–1262.
- [93] F. Wenzel, K. Roth, B. S. Veeling, J. Swiatkowski, L. Tran, S. Mandt, J. Snoek, T. Salimans, R. Jenatton, and S. Nowozin, “How good is the Bayes Posterior in deep neural networks really?” in *Proc. International Conference on Machine Learning (ICML)*, vol. 119, 2020, pp. 10 248–10 259.
- [94] B. Adlam, J. Snoek, and S. L. Smith, “Cold posteriors and aleatoric uncertainty,” *CoRR*, vol. abs/2008.00029, 2020. [Online]. Available: <https://arxiv.org/abs/2008.00029>
- [95] E. Angelino, M. J. Johnson, and R. P. Adams, “Patterns of scalable Bayesian inference,” *Foundations and Trends® in Machine Learning*, vol. 9, no. 2-3, pp. 119–247, 2016.

- [96] J. Han, F. Ding, X. Liu, L. Torresani, J. Peng, and Q. Liu, “Stein variational inference for discrete distributions,” in *Proc. International Conference on Artificial Intelligence and Statistics (AISTAT)*. PMLR, 2020, pp. 4563–4572.
- [97] P. Izmailov, A. Novikov, and D. Kropotov, “Scalable Gaussian processes with billions of inducing inputs via tensor train decomposition,” in *Proc. International Conference on Artificial Intelligence and Statistics (AISTAT)*. Playa Blanca, Lanzarote, Canary Islands, Spain: PMLR, 2018, pp. 726–735.
- [98] A. Tjandra, S. Sakti, and S. Nakamura, “Compressing recurrent neural network with tensor train,” in *Proc. International Joint Conference on Neural Networks (IJCNN)*. Anchorage, AK, USA: IEEE, 2017, pp. 4451–4458.