

Bayesian Inference with Posterior Regularization and Applications to Infinite Latent SVMs

Jun Zhu

Ning Chen

Department of Computer Science and Technology

State Key Laboratory of Intelligent Technology and Systems

Tsinghua National Laboratory for Information Science and Technology

Tsinghua University

Beijing, 100084 China

DCSZJ@MAIL.TSINGHUA.EDU.CN

NINGCHEN@MAIL.TSINGHUA.EDU.CN

Eric P. Xing

School of Computer Science

Carnegie Mellon University

Pittsburgh, PA 15213, USA

EPXING@CS.CMU.EDU

Editor: Tony Jebara

Abstract

Existing Bayesian models, especially nonparametric Bayesian methods, rely on specially conceived priors to incorporate domain knowledge for discovering improved latent representations. While priors affect posterior distributions through Bayes' rule, imposing posterior regularization is arguably more direct and in some cases more natural and general. In this paper, we present *regularized Bayesian inference* (RegBayes), a novel computational framework that performs posterior inference with a regularization term on the desired post-data posterior distribution under an information theoretical formulation. RegBayes is more flexible than the procedure that elicits expert knowledge via priors, and it covers both directed Bayesian networks and undirected Markov networks. When the regularization is induced from a linear operator on the posterior distributions, such as the expectation operator, we present a general convex-analysis theorem to characterize the solution of RegBayes. Furthermore, we present two concrete examples of RegBayes, *infinite latent support vector machines* (iLSVM) and *multi-task infinite latent support vector machines* (MT-iLSVM), which explore the large-margin idea in combination with a nonparametric Bayesian model for discovering predictive latent features for classification and multi-task learning, respectively. We present efficient inference methods and report empirical studies on several benchmark data sets, which appear to demonstrate the merits inherited from both large-margin learning and Bayesian nonparametrics. Such results contribute to push forward the interface between these two important subfields, which have been largely treated as isolated in the community.

Keywords: Bayesian inference, posterior regularization, Bayesian nonparametrics, large-margin learning, classification, multi-task learning

1. Introduction

Over the past decade, nonparametric Bayesian models have gained remarkable popularity in machine learning and other fields, partly owing to their desirable utility as a “nonparamet-

ric” prior distribution for a wide variety of probabilistic models, thereby turning the largely heuristic model selection practice, such as determining the unknown number of components in a mixture model (Antoniak, 1974) or the unknown dimensionality of latent features in a factor analysis model (Griffiths and Ghahramani, 2005), as a Bayesian inference problem in an unbounded model space. Popular examples include Gaussian process (GP) (Rasmussen and Ghahramani, 2002), Dirichlet process (DP) (Ferguson, 1973; Antoniak, 1974), and Beta process (BP) (Thibaux and Jordan, 2007). DP is often described with a Chinese restaurant process (CRP) metaphor, and similarly BP is often described with an Indian buffet process (IBP) metaphor (Griffiths and Ghahramani, 2005). Such nonparametric Bayesian approaches allow the model complexity to grow as more data are observed, which is a key factor differing them from other traditional “parametric” Bayesian models.

One recent development in practicing Bayesian nonparametrics is to relax some unrealistic assumptions on data, such as homogeneity and exchangeability. For example, to handle heterogenous observations, predictor-dependent processes (MacEachern, 1999; Williamson et al., 2010) have been proposed; and to relax the exchangeability assumption, stochastic processes with various correlation structures, such as hierarchical structures (Teh et al., 2006), temporal or spatial dependencies (Beal et al., 2002; Blei and Frazier, 2010), and stochastic ordering dependencies (Hoff, 2003; Dunson and Peddada, 2007), have been introduced. A common principle shared by these approaches is that they rely on defining, or in some unusual cases learning (Welling et al., 2012) a nonparametric Bayesian prior¹ encoding some special structures, which *indirectly*² influences the posterior distribution of interest through an interplay with a likelihood model according to the Bayes’ rule (also known as Bayes’ theorem). In this paper, we explore a different principle known as *posterior regularization*, which offers an additional and arguably richer and more flexible set of means to augment a posterior distribution under rich side information, such as predictive margin, structural bias, etc., which can be harder, if possible, to be captured by a Bayesian prior.

Let Θ denote model parameters and H denote hidden variables. Then given a set of observed data \mathcal{D} , posterior regularization (Ganchev et al., 2010) is generally defined as solving a regularized maximum likelihood estimation (MLE) problem:

$$\text{Posterior Regularization : } \max_{\Theta} \mathcal{L}(\Theta; \mathcal{D}) + \Omega(p(H|\mathcal{D}, \Theta)), \quad (1)$$

where $\mathcal{L}(\Theta; \mathcal{D})$ is the marginal likelihood of \mathcal{D} , and $\Omega(\cdot)$ is a regularization function of the model posterior over latent variables. Note that here we view posterior as a generic post-data distribution on hidden variables in the sense of Ghosh and Ramamoorthi (2003, pp.15), not necessarily corresponding to a Bayesian posterior that must be induced by the Bayes’ rule. The regularizer can be defined as a KL-divergence between a desired distribution with certain properties over latent variables and the model posterior in question, or other constraints on the model posterior, such as those used in generalized expectation (Mann and McCallum, 2010) or constraint-driven semi-supervised learning (Chang

1. Although likelihood is another dimension that can incorporate domain knowledge, existing work on Bayesian nonparametrics has been mainly focusing on the priors. Following this convention, this paper assumes that a common likelihood model (e.g., Gaussian likelihood for continuous data) is given.
 2. A hard constraint on the prior (e.g., a truncated Gaussian) can directly affect the support of the posterior. RegBayes covers this as a special case as shown in Remark 7.

et al., 2007). An EM-type procedure can be applied to solve (1) approximately, and obtain an augmented MLE of the hidden variable model: $p(H|\mathcal{D}, \Theta_{\text{MLE}})$. When a distribution over the model parameter is of interest, going beyond the classical Bayesian theory, recent attempts toward learning a regularized posterior distribution of model parameters (and latent variables as well if present) include the “learning from measurements” (Liang et al., 2009), maximum entropy discrimination (MED) (Jaakkola et al., 1999; Zhu and Xing, 2009) and maximum entropy discrimination latent Dirichlet allocation (MedLDA) (Zhu et al., 2009). All these methods are parametric in that they give rise to distributions over a fixed and finite-dimensional parameter space. To the best of our knowledge, very few attempts have been made to impose posterior regularization in a nonparametric setting where model complexity depends on data, such as the case for nonparametric Bayesian latent variable models. A general formalism for (parametric and nonparametric) Bayesian inference with posterior regularization seems to be not yet available or apparent. In this paper, we present such a formalism, which we call *regularized Bayesian inference*, or RegBayes, built on the convex duality theory over distribution function spaces; and we apply this formalism to learn regularized posteriors under the Indian buffet process (IBP), conjoining two powerful machine learning paradigms, nonparametric Bayesian inference and SVM-style max-margin constrained optimization.

Unlike the regularized MLE formulation in (1), under the traditional formulation of Bayesian inference one is not directly optimizing an objective with respect to the posterior. To enable a regularized optimization formulation of RegBayes, we begin with a variational reformulation of the Bayes’ theorem, and define $\mathcal{L}(q(\mathbf{M}|\mathcal{D}))$ as the KL-divergence between a desired post-data posterior $q(\mathbf{M}|\mathcal{D})$ over model \mathbf{M} and the standard Bayesian posterior $p(\mathbf{M}|\mathcal{D})$ (see Section 3.1 for a recapitulation of the connection between KL-minimization and Bayes’ theorem). RegBayes solves the following optimization problem:

$$\mathbf{RegBayes} : \inf_{q(\mathbf{M}|\mathcal{D}) \in \mathcal{P}_{\text{prob}}} \mathcal{L}(q(\mathbf{M}|\mathcal{D})) + \Omega(q(\mathbf{M}|\mathcal{D})), \quad (2)$$

where the regularization $\Omega(\cdot)$ is a function of the post-data posterior $q(\mathbf{M}|\mathcal{D})$, and $\mathcal{P}_{\text{prob}}$ is the feasible space of well-defined distributions. By appropriately defining the model and its prior distribution, RegBayes can be instantiated to perform either parametric and nonparametric regularized Bayesian inference.

One particularly interesting way to derive the posterior regularization is to impose posterior constraints. Let $\boldsymbol{\xi}$ denote slack variables and $\mathcal{P}_{\text{post}}(\boldsymbol{\xi})$ denote the general soft posterior constraints (see Section 3.2 for a formal description), then, we can express the regularization term variationally:

$$\Omega(q(\mathbf{M}|\mathcal{D})) = \inf_{\boldsymbol{\xi}} U(\boldsymbol{\xi}), \quad \text{s.t.}: q(\mathbf{M}|\mathcal{D}) \in \mathcal{P}_{\text{post}}(\boldsymbol{\xi}), \quad (3)$$

where $U(\boldsymbol{\xi})$ is normally defined as a convex penalty function. The RegBayes formalism defined in (2) applies to a wide spectrum of models, including directed graphical models (i.e., Bayesian networks) and undirected Markov networks. For undirected models, when performing Bayesian inference the resulting posterior takes the form of a hybrid chain graphical model (Frydenberg, 1990; Murray and Ghahramani, 2004; Qi et al., 2005; Welling and Parise, 2006), which is usually much more challenging to regularize than for Bayesian

inference with directed graphical models. When the regularization term is convex and induced from a linear operator (e.g., expectation) of the posterior distributions, RegBayes can be solved with convex analysis theory.

By allowing direct regularization over posterior distributions, RegBayes provides a significant source of extra flexibility for post-data posterior inference, which applies to both parametric and nonparametric Bayesian learning (see the remarks after the main Theorem 6). In this paper, we focus on applying this technique to the later case, and illustrate how to use RegBayes to facilitate integration of Bayesian nonparametrics and large-margin learning, which have complementary advantages but have been largely treated as two disjoint subfields. Previously, it has been shown that, the core ideas of support vector machines (Vapnik, 1995) and maximum entropy discrimination (Jaakkola et al., 1999), as well as their structured extensions to the max-margin Markov networks (Taskar et al., 2003) and maximum entropy discrimination Markov networks (Zhu and Xing, 2009), have led to successful outcomes in many scenarios. But a large-margin model rarely has the flexibility of nonparametric Bayesian models to automatically handle model complexity from data, especially when latent variables are present (Jebara, 2001; Zhu et al., 2009). In this paper, we intend to bridge this gap using the RegBayes principle.

Specifically, we develop the *infinite latent support vector machines* (iLSVM) and *multi-task infinite latent support vector machines* (MT-iLSVM), which explore the discriminative large-margin idea to learn infinite latent feature models for classification and multi-task learning (Argyriou et al., 2007; Bakker and Heskes, 2003), respectively. We show that both models can be readily instantiated from the RegBayes master equation (2) by defining appropriate posterior regularization using the large-margin principle, and by employing an appropriate prior. For iLSVM, we use the IBP prior to allow the model to have an unbounded number of latent features *a priori*. For MT-iLSVM, we use a similar IBP prior to infer a latent projection matrix to capture the correlations among multiple predictive tasks while avoiding pre-specifying the dimensionality of the projection matrix. The regularized inference problems can be efficiently solved with an iterative procedure, which leverages existing high-performance convex optimization techniques.

The rest of the paper is organized as follows. Section 2 discusses related work. Section 3 presents regularized Bayesian inference (RegBayes), together with the convex duality results that will be needed in latter sections. Section 4 concretizes the ideas of RegBayes and presents two infinite latent feature models with large-margin constraints for both classification and multi-task learning. Section 5 presents some preliminary experimental results. Finally, Section 6 concludes and discusses future research directions.

2. Related Work

Expectation regularization or expectation constraints have been considered to regularize model parameter estimation in the context of semi-supervised learning or learning with weakly labeled data. Mann and McCallum (2010) summarized the recent developments of the generalized expectation (GE) criteria for training a discriminative probabilistic model with unlabeled data, e.g., maximum entropy models or conditional random fields (Lafferty et al., 2001). By providing appropriate side information, such as labeled features or estimates of label distributions, a GE-based penalty function is defined to regularize the model

distribution, e.g., the distribution of class labels. One commonly used GE function is the KL-divergence between empirical expectation and model expectation of some feature functions if the expectations are normalized or the general Bregman divergence for unnormalized expectations. Although the GE criteria can be used alone as a scoring function to estimate the unknown parameters of a discriminative model, it is more usually used as a regularization term to an estimation method, such as maximum (conditional) likelihood estimation. Bellare et al. (2009) presented a different formulation of using expectation constraints in semi-supervised learning by introducing an auxiliary distribution to GE, together with an alternating projection algorithm, which can be more efficient. Liang et al. (2009) proposed to use the general notion of “measurements” to encapsulate the variety of weakly labeled data for learning exponential family models. The measurements can be labels, partial labels or other constraints on model predictions. Under the EM framework, posterior constraints were used in Graca et al. (2007) to modify the E-step of an EM algorithm to project model posterior distributions onto the subspace of distributions that satisfy a set of auxiliary constraints.

Dudík et al. (2007) studied the generalized maximum entropy principle with a rich form of expectation constraints using convex duality theory, where the standard moment matching constraints of maximum entropy are relaxed to inequality constraints. But their analysis was restricted to KL-divergence minimization (maximum entropy is a special case) and the finite dimensional space of observations. Later on, Altun and Smola (2006) presented a more general duality theory for a family of divergence functions on Banach spaces. We have drawn inspiration from both papers to develop the regularized Bayesian inference framework using convex duality theory.

When using large-margin posterior regularization, RegBayes generalizes the previous work on maximum entropy discrimination (Jaakkola et al., 1999; Zhu and Xing, 2009). The present paper provides a full extension of our preliminary work on max-margin nonparametric Bayesian models (Zhu et al., 2011b,a). For example, the infinite SVM (iSVM) (Zhu et al., 2011b) is a latent class model, where each data example is assigned to a single mixture component (i.e., an 1-dimensional space), and both iLSVM and MT-iLSVM extend the ideas to infinite latent feature models. For multi-task learning, nonparametric Bayesian models have been developed by Xue et al. (2007) and Rai and Daume III (2010) for learning features shared by multiple tasks. However, these methods are based on standard Bayesian inference without a posterior regularization using, for example, the large-margin constraints. Finally, MT-iLSVM can be also regarded as a nonparametric Bayesian formulation of the popular multi-task learning methods (Ando and Zhang, 2005; Jebara, 2011).

3. Regularized Bayesian Inference

We begin by laying out a general formulation of regularized Bayesian inference, using an optimization framework built on convex duality theory.

3.1 Variational formulation of Bayes’ theorem

We first derive an optimization-theoretic reformulation of the Bayes’ theorem. Let \mathcal{M} denote the space of feasible models, and $\mathbf{M} \in \mathcal{M}$ represents an atom in this space. We assume that \mathcal{M} is a complete separable metric space endowed with its Borel σ -algebra $\mathcal{B}(\mathcal{M})$. Let

Π be a distribution (i.e., a probability measure) on the measurable space $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$. We assume that Π is absolutely continuous with respect to some background measure μ , so that there exists a density π such that $d\Pi = \pi d\mu$. Let $\mathcal{D} = \{\mathbf{x}_n\}_{n=1}^N$ be a collection of observed data, which we assume to be i.i.d. given a model. Let $P(\cdot|\mathbf{M})$ be the likelihood distribution, which is assumed to be dominated by a σ -finite measure λ for all \mathbf{M} with positive density, so that there exists a density $p(\cdot|\mathbf{M})$ such that $dP(\cdot|\mathbf{M}) = p(\cdot|\mathbf{M})d\lambda$. Then, the Bayes' conditionalization rule gives a posterior distribution with the density (Ghosh and Ramamoorthi, 2003, Chap.1.3):

$$p(\mathbf{M}|\mathcal{D}) = \frac{\pi(\mathbf{M})p(\mathcal{D}|\mathbf{M})}{p(\mathcal{D})} = \frac{\pi(\mathbf{M}) \prod_{n=1}^N p(\mathbf{x}_n|\mathbf{M})}{p(\mathbf{x}_1, \dots, \mathbf{x}_N)},$$

a density over \mathbf{M} with respect to the base measure μ , where $p(\mathcal{D})$ is the marginal likelihood of the observed data.

For reasons to be clear shortly, we now introduce a variational formulation of the Bayes' theorem. Let Q are an arbitrary distribution on the measurable space $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$. We assume that Q is absolutely continuous with respect to Π and denote by q its density with respect to the background measure μ .³ It can be shown that the posterior distribution of \mathbf{M} due to the Bayes' theorem is equivalent to the optimum solution of the following convex optimization problem:

$$\begin{aligned} \inf_{q(\mathbf{M})} \text{KL}(q(\mathbf{M})\|\pi(\mathbf{M})) - \int_{\mathcal{M}} \log p(\mathcal{D}|\mathbf{M})q(\mathbf{M})d\mu(\mathbf{M}) \quad (4) \\ \text{s.t. : } q(\mathbf{M}) \in \mathcal{P}_{\text{prob}}, \end{aligned}$$

where $\text{KL}(q(\mathbf{M})\|\pi(\mathbf{M})) = \int_{\mathcal{M}} q(\mathbf{M}) \log(q(\mathbf{M})/\pi(\mathbf{M}))d\mu(\mathbf{M})$ is the Kullback-Leibler (KL) divergence from $q(\cdot)$ to $\pi(\cdot)$, and $\mathcal{P}_{\text{prob}}$ represents the feasible space of all density functions over \mathbf{M} with respect to the measure μ . The proof is straightforward by noticing that the objective will become $\text{KL}(q(\mathbf{M})\|p(\mathbf{M}|\mathcal{D}))$ by adding the constant $\log p(\mathcal{D})$. It is noteworthy that $q(\mathbf{M})$ here represents the density of a general post-data posterior distribution in the sense of Ghosh and Ramamoorthi (2003, pp.15), not necessarily corresponding to a Bayesian posterior that is induced by the Bayes' rule. As we shall see soon later, when we introduce additional constraints, the post-data posterior $q(\mathbf{M})$ is different from the Bayesian posterior $p(\mathbf{M}|\mathcal{D})$, and moreover, it could even not be obtainable from any Bayesian conditionalization in a different model. In the sequel, in order to distinguish $q(\cdot)$ from the Bayesian posterior, we will call it post-data distribution in short or post-data posterior distribution in full.⁴ For notation simplicity, we have omitted the condition \mathcal{D} in the post-data posterior distribution $q(\mathbf{M})$.

Remark 1 *The optimization formulation in (4) implies that Bayes' rule is an information projection procedure that projects a prior density to a post-data posterior by taking account of the observed data. In general, Bayes's rule is a special case of the principle of minimum information (Williams, 1980).*

3. This assumption is necessary to make the KL-divergence between the two distributions Q and Π well-defined. This assumption (or constraint) will be implicitly included in $\mathcal{P}_{\text{prob}}$ for clarity.

4. Rigorously, $q(\cdot)$ is the density of the post-data posterior distribution $Q(\cdot)$. We simply call q a distribution if no confusion arises.

3.2 Regularized Bayesian Inference with Expectation Constraints

In the variational formulation of Bayes’ rule in (4), the constraints on $q(\mathbf{M})$ ensure that q is well-normalized and the objective is well-defined, i.e., $q(\mathbf{M}) \in \mathcal{P}_{\text{prob}}$, which do not capture any domain knowledge or structures of the model or data. Some previous efforts have been devoted to eliciting domain knowledge by constraining the prior or the base measure μ (Robert, 1995; Garthwaite et al., 2005). As we shall see, such constraints without considering data are special cases of RegBayes to be presented.

Specifically, the optimization-based formulation of Bayes’ rule makes it straightforward to generalize Bayesian inference to a richer type of posterior inference, by replacing the standard normality constraint on q with a wide spectrum of knowledge-driven and/or data-driven constraints or regularization. To contrast, we will refer to problem (4) as “unconstrained” or “unregularized”. Formally, we define *regularized Bayesian inference* (RegBayes) as a generalized posterior inference procedure that solves a constrained optimization problem due to such additional regularization imposed on q :

$$\begin{aligned} \inf_{q(\mathbf{M}), \boldsymbol{\xi}} \text{KL}(q(\mathbf{M})\|\pi(\mathbf{M})) - \int_{\mathcal{M}} \log p(\mathcal{D}|\mathbf{M})q(\mathbf{M})d\mu(\mathbf{M}) + U(\boldsymbol{\xi}) \quad (5) \\ \text{s.t. : } q(\mathbf{M}) \in \mathcal{P}_{\text{post}}(\boldsymbol{\xi}), \end{aligned}$$

where $\mathcal{P}_{\text{post}}(\boldsymbol{\xi})$ is a subspace of distributions that satisfy a set of additional constraints besides the standard normality constraint of a probability distribution. Using the variational formulation in (3), problem (5) can be rewritten in the form of the master equation (2), of which the objective is: $\mathcal{L}(q(\mathbf{M})) = \text{KL}(q(\mathbf{M})\|\pi(\mathbf{M})) - \int_{\mathcal{M}} \log p(\mathcal{D}|\mathbf{M})q(\mathbf{M})d\mu(\mathbf{M}) = \text{KL}(q(\mathbf{M})\|p(\mathbf{M}, \mathcal{D}))$ and the posterior regularization is $\Omega(q(\mathbf{M})) = \inf_{\boldsymbol{\xi}} U(\boldsymbol{\xi})$, s.t.: $q(\mathbf{M}|\mathcal{D}) \in \mathcal{P}_{\text{post}}(\boldsymbol{\xi})$. Note that when \mathcal{D} is given, the distribution $p(\mathbf{M}, \mathcal{D})$ is unnormalized for \mathbf{M} ; and we have abused the KL notation for unnormalized distributions in $\text{KL}(q(\mathbf{M})\|p(\mathbf{M}, \mathcal{D}))$, but with the same formula.

Obviously this formulation enables different types of constraints to be employed in practice. In this paper, we focus on the *expectation constraints*, of which each one is a function of $q(\mathbf{M})$ through an expectation operator. For instance, let $\boldsymbol{\psi} = (\psi_1, \dots, \psi_T)$ be a vector of feature functions, each of which is $\psi_t(\mathbf{M}; \mathcal{D})$ defined on \mathbf{M} and possibly data dependent. Then a subspace of feasible post-data distributions can be defined in the following form:

$$\mathcal{P}_{\text{post}}(\boldsymbol{\xi}) \stackrel{\text{def}}{=} \left\{ q(\mathbf{M}) \mid \forall t = 1, \dots, T, h(Eq(\psi_t; \mathcal{D})) \leq \xi_t \right\}, \quad (6)$$

where E is the expectation operator that maps $q(\mathbf{M})$ to a point in the space \mathbb{R}^T , and for each feature function ψ_t : $Eq(\psi_t; \mathcal{D}) \stackrel{\text{def}}{=} \mathbb{E}_{q(\mathbf{M})}[\psi_t(\mathbf{M}; \mathcal{D})]$. The function h can be of any form in theory, though a simple h function will make the optimization problem easy to solve. The auxiliary parameters $\boldsymbol{\xi}$ are usually nonnegative and interpreted as slack variables. The constraints with non-trivial $\boldsymbol{\xi}$ are soft constraints as illustrated in Figure 1(b). But we emphasize that by defining U as an indicator function, the formulation (5) covers the case where hard constraints are imposed. For instance, if we define

$$U(\boldsymbol{\xi}) = \sum_{t=1}^T \mathbb{I}(\xi_t = \gamma_t) = \mathbb{I}(\boldsymbol{\xi} = \boldsymbol{\gamma}),$$




Figure 1: Illustration for the (a) hard and (b) soft constraints in the simple setting which has only three possible models. For hard constraints, we have only one feasible subspace. In contrast, we have many (normally infinite for continuous ξ) feasible subspaces for soft constraints and each of them is associated with a different complexity or penalty, measured by the U function.

where $\mathbb{I}(c)$ is an indicator function that equals to 0 if the condition c is satisfied; otherwise ∞ , then all the expectation constraints in (6) are hard constraints. As illustrated in Figure 1(a), hard constraints define one single feasible subspace (assuming to be non-empty). In general, we assume that $U(\xi)$ is a convex function, which represents a penalty on the size of the feasible subspaces, as illustrated in Figure 1(b). A larger subspace typically leads to models with a higher complexity. In the classification models to be presented, U corresponds to a surrogate loss, e.g., hinge loss of a prediction rule, as we shall see.

Similarly, the formulation of RegBayes with expectation constraints in (6) can be equivalently written in an “unconstrained” form by using the rule in (3). Specifically, let $g(Eq(\psi; \mathcal{D})) \stackrel{\text{def}}{=} \inf_{\xi} U(\xi)$, s.t. : $h(Eq(\psi_t; \mathcal{D})) \leq \xi_t, \forall t$, we have the equivalent optimization problem:

$$\inf_{q(\mathbf{M}) \in \mathcal{P}_{\text{prob}}} \text{KL}(q(\mathbf{M}) \parallel \pi(\mathbf{M})) - \int_{\mathcal{M}} \log p(\mathcal{D} | \mathbf{M}) q(\mathbf{M}) d\mu(\mathbf{M}) + g(Eq(\psi; \mathcal{D})), \quad (7)$$

where $Eq(\psi; \mathcal{D})$ is a point in \mathbb{R}^T and the t -th coordinate is $Eq(\psi_t; \mathcal{D})$, a function of $q(\mathbf{M})$ as defined before. We assume that the real-valued function $g : \mathbb{R}^T \rightarrow \mathbb{R}$ is convex and lower semi-continuous. For each U , we can induce a g function by taking the infimum of $U(\xi)$ over ξ with the posterior constraints; vice versa. If we use hard constraints, similar as in regularized maximum entropy density estimation (Altun and Smola, 2006; Dudík et al., 2007), we have

$$g(Eq) = \sum_{t=1}^T \mathbb{I}(h(Eq(\psi_t; \mathcal{D})) \leq \gamma_t).$$

For the regularization function g , as well as U , we can have many choices, besides the above mentioned indicator function. For example, if the feature function ψ_t is an

indicator function and we could obtain ‘prior’ expectations $\mathbb{E}_{\tilde{p}}[\psi_t]$ from domain/expert knowledge about \mathbf{M} . If we further normalize the empirical expectations of T functions and denote the discrete distribution by $\tilde{p}(\mathbf{M})$, one natural regularization function would be the KL-divergence between prior expectations and the expectations computed from the normalized model posterior $q(\mathbf{M})$, i.e., $g(Eq) = \sum_t s(\mathbb{E}_{\tilde{p}}[\psi_t], Eq(\psi_t)) = \text{KL}(\tilde{p}(\mathbf{M})||q(\mathbf{M}))$, where $s(x, y) = x \log(x/y)$ for $x, y \in (0, 1)$. The general Bregman divergence can be used for unnormalized expectations. This kind of regularization function has been used in Mann and McCallum (2010) for label regularization, in the context of semi-supervised learning. Other choices of the regularization function include the ℓ_2^2 penalty or indicator function with equality constraints; see Table 1 in Dudík et al. (2007) for a summary.

Remark 2 *So far, we have focused on RegBayes in the context of full Bayesian inference. Indeed, RegBayes can be generalized to apply to empirical Bayesian inference, where some model parameters need to be estimated. More generally, RegBayes applies to both directed Bayesian networks (of which the hierarchical Bayesian models we have discussed are an example) and undirected Markov random fields. But for undirected models, a RegBayes treatment will have to deal with a chain graph resultant from Bayesian inference, which is more challenging due to existence of normalization factors. We will discuss some details and examples in Appendix A.*

3.3 Optimization with Convex Duality Theory

Depending on several factors, including the size of the model space, the data likelihood model, the prior distribution, and the regularization function, a RegBayes problem in general can be highly non-trivial to solve, either in the constrained or unconstrained form, as can be seen from several concrete examples of RegBayes models we will present in the next section and in the Appendix B. In this section, we present a representation theorem to characterize the solution the convex RegBayes problem (7) with expectation regularization. These theoretical results will be used later in developing concrete RegBayes models.

To make the subsequent statements general, we consider the following problem:

$$\inf_{x \in \mathcal{X}} f(x) + g(Ax),$$

where $f : \mathcal{X} \rightarrow \mathbb{R}$ is a convex function; $A : \mathcal{X} \rightarrow \mathcal{B}$ is a bounded linear operator; and $g : \mathcal{B} \rightarrow \mathbb{R}$ is also convex. Below we introduce some tools in convex analysis theory to study this problem. We begin by formulating the primal-dual space relationships of convex optimization problems in the general settings, where we assume both \mathcal{X} and \mathcal{B} are Banach spaces.⁵ An important result we build on is the Fenchel duality theorem.

Definition 3 (Convex Conjugate) *Let \mathcal{X} be a Banach space and \mathcal{X}^* be its dual space. The convex conjugate or the Legendre-Fenchel transformation of a function $f : \mathcal{X} \rightarrow [-\infty, +\infty]$ is $f^* : \mathcal{X}^* \rightarrow [-\infty, +\infty]$, where*

$$f^*(x^*) = \sup_{x \in \mathcal{X}} \{\langle x^*, x \rangle - f(x)\}.$$

5. A Banach space is a vector space with a metric that allows the computation of vector length and distance between vectors. Moreover, a Cauchy sequence of vectors always converges to a well defined limit in the space.

Theorem 4 (Fenchel Duality (Borwein and Zhu, 2005)) *Let \mathcal{X} and \mathcal{B} be Banach spaces, $f : \mathcal{X} \rightarrow \mathbb{R} \cup \{+\infty\}$ and $g : \mathcal{B} \rightarrow \mathbb{R} \cup \{+\infty\}$ be convex functions and $A : \mathcal{X} \rightarrow \mathcal{B}$ be a bounded linear map. Define the primal and dual values t, d by the Fenchel problems*

$$t = \inf_{x \in \mathcal{X}} \{f(x) + g(Ax)\} \text{ and } d = \sup_{x^* \in \mathcal{B}^*} \{-f^*(A^*x^*) - g^*(-x^*)\}.$$

Then these values satisfy the weak duality inequality $t \geq d$. If f, g and A satisfy either

$$0 \in \text{core}(\text{dom}g - A\text{dom}f) \text{ and both } f \text{ and } g \text{ are lower semicontinuous (lsc),}$$

or

$$A\text{dom}f \cap \text{cont}g \neq \emptyset,$$

then $t = d$ and the supremum to the dual problem is attainable if finite.

Let \mathcal{S} be a subset of a Banach space \mathcal{B} . In the above theorem, we say s is in the *core* of \mathcal{S} , denoted by $s \in \text{core}(\mathcal{S})$, provided that $\cup_{\lambda > 0} \lambda(\mathcal{S} - s) = \mathcal{B}$.

The Fenchel duality theorem has been applied to solve divergence minimization problems for density estimation (Altun and Smola, 2006; Dudík et al., 2007). Let $\boldsymbol{\psi} \stackrel{\text{def}}{=} (\psi_1, \dots, \psi_T)$ be a vector of feature functions. Each feature function is a mapping, $\psi_t : \mathcal{M} \rightarrow \mathbb{R}$. Therefore, \mathcal{B} is the product space \mathbb{R}^T , a simple Banach space. Let \mathcal{X} be the Banach space of finite signed measures (with total variation as the norm) that are absolutely continuous with respect to the measure μ , and let A be the expectation operator of the feature functions with respect to the distribution q on \mathcal{M} , that is, $Aq \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{M} \sim q}[\boldsymbol{\psi}(\mathbf{M})]$, where $\boldsymbol{\psi}(\mathbf{M}) = (\psi_1(\mathbf{M}), \dots, \psi_T(\mathbf{M}))$. Let $\tilde{\boldsymbol{\psi}}$ be a reference point in \mathbb{R}^T . As for density estimation, we have some observations of \mathbf{M} here, and $\tilde{\boldsymbol{\psi}} = Ap_{\text{emp}}[\boldsymbol{\psi}(\mathbf{M})]$, where p_{emp} is the empirical distribution. Then, when the f function is a KL-divergence and the constraints are relaxed moment matching constraints, the following result can be proven.

Lemma 5 (KL-divergence with Constraints (Altun and Smola, 2006))

$$\begin{aligned} & \inf_q \left\{ \text{KL}(q||p) \text{ s.t. : } \|\mathbb{E}_q[\boldsymbol{\psi}] - \tilde{\boldsymbol{\psi}}\|_{\mathcal{B}} \leq \epsilon \text{ and } q \in \mathcal{P}_{\text{prob}} \right\} \\ & = \sup_{\boldsymbol{\phi}} \left\{ \langle \boldsymbol{\phi}, \tilde{\boldsymbol{\psi}} \rangle - \log \int_{\mathcal{M}} p(\mathbf{M}) \exp(\langle \boldsymbol{\phi}, \boldsymbol{\psi}(\mathbf{M}) \rangle) d\mu(\mathbf{M}) - \epsilon \|\boldsymbol{\phi}\|_{\mathcal{B}^*} \right\}, \end{aligned}$$

where the unique solution is given by $\hat{q}_{\hat{\boldsymbol{\phi}}}(\mathbf{M}) = p(\mathbf{M}) \exp(\langle \hat{\boldsymbol{\phi}}, \boldsymbol{\psi}(\mathbf{M}) \rangle - \Lambda_{\hat{\boldsymbol{\phi}}})$; $\hat{\boldsymbol{\phi}}$ is the solution of the dual problem; and $\Lambda_{\hat{\boldsymbol{\phi}}}$ is the log-partition function.

Note that for this lemma and the ones to be presented below to hold, the problems need to meet some regularity conditions (or constraint qualifications), such as those in Theorem 4. In practice it can be difficult to check whether the constraint qualifications hold. One solution is to solve the dual optimization problem and examine if the conditions hold depending on whether the solution diverge or not (Altun and Smola, 2006).

The problem in the above lemma is subject to hard constraints, therefore the corresponding g is the indicator function $\mathbb{I}(\|\mathbb{E}_q[\boldsymbol{\psi}] - \tilde{\boldsymbol{\psi}}\|_{\mathcal{B}} \leq \epsilon)$ when applying the Fenchel duality

theorem. Other examples of the posterior constraints can be found in Dudík et al. (2007); Mann and McCallum (2010); Ganchev et al. (2010), as we have discussed in Section 3.2. In this paper, we consider the general soft constraints as defined in the RegBayes problem (5). Furthermore, we do not assume the existence of a fully observed data set to compute the empirical expectation $\tilde{\phi}$. Specifically, following a similar line of reasoning as in Altun and Smola (2006), though this time with an un-normalized p in $\text{KL}(q||p)$, we have the following result. The detailed proof is deferred to Appendix C.1.

Theorem 6 (Representation theorem of RegBayes) *Let E be the expectation operator with feature functions $\psi(\mathbf{M}; \mathcal{D})$, and assume g is convex and lower semicontinuous (lsc). We have*

$$\begin{aligned} & \inf_{q(\mathbf{M})} \left\{ \text{KL}(q(\mathbf{M})||p(\mathbf{M}, \mathcal{D})) + g(Eq) \text{ s.t. } : q(\mathbf{M}) \in \mathcal{P}_{\text{prob}} \right\} \\ &= \sup_{\hat{\phi}} \left\{ -\log \int_{\mathcal{M}} p(\mathbf{M}, \mathcal{D}) \exp(\langle \hat{\phi}, \psi(\mathbf{M}; \mathcal{D}) \rangle) d\mu(\mathbf{M}) - g^*(-\hat{\phi}) \right\}, \end{aligned}$$

where the unique solution is given by $\hat{q}_{\hat{\phi}}(\mathbf{M}) = p(\mathbf{M}, \mathcal{D}) \exp(\langle \hat{\phi}, \psi(\mathbf{M}; \mathcal{D}) \rangle - \Lambda_{\hat{\phi}})$; and $\hat{\phi}$ is the solution of the dual problem; and $\Lambda_{\hat{\phi}}$ is the log-partition function.

From the optimum solution $\hat{q}_{\hat{\phi}}(\mathbf{M})$, we can see that the form of the RegBayes posterior is symbolically similar to that of the Bayesian posterior; but instead of multiplying the likelihood term with a prior distribution, RegBayes introduces an extra term, $\exp(\langle \hat{\phi}, \psi(\mathbf{M}; \mathcal{D}) \rangle - \Lambda_{\hat{\phi}})$, whose coefficients are derived from an constrained optimization problem resultant from the constraints on the posterior. We make the following remarks.

Remark 7 (Putting constraints on priors is a special case of RegBayes) *If both the feature function $\psi(\mathbf{M}; \mathcal{D})$ and $\hat{\phi}$ depend on the model \mathbf{M} only, this extra term contributes to define a new prior $\pi'(\mathbf{M}) \propto \pi(\mathbf{M}) \exp(\langle \hat{\phi}, \psi(\mathbf{M}; \mathcal{D}) \rangle - \Lambda_{\hat{\phi}})$. For example, if we constrain the model space to a subset $\mathcal{M}_0 \subset \mathcal{M}$ a priori, this constraint can be incorporated in RegBayes by defining the expectation constraint on \mathbf{M} only. Specifically, define the single feature function $\psi(\mathbf{M})$: $\psi(\mathbf{M}) = 0$ if $\mathbf{M} \in \mathcal{M}_0$, otherwise 1; and define the simple posterior regularization $g(Eq) = \mathbb{I}(\mathbb{E}_q[\psi(\mathbf{M})] = 0)$. Then, by Theorem 6,⁶ we have $\hat{\phi} = -\infty$ and $\hat{q}_{\hat{\phi}}(\mathbf{M}) \propto \pi'(\mathbf{M})p(\mathcal{D}|\mathbf{M})$, where $\pi'(\mathbf{M}) \propto \pi(\mathbf{M})\mathbb{I}(\mathbf{M} \in \mathcal{M}_0)$ is the constrained prior. Therefore, such a constraint lets RegBayes cover the widely used truncated priors, such as truncated Gaussian (Robert, 1995).*

Remark 8 (RegBayes is more flexible than Bayes' rule) *For the more general case where $\psi(\mathbf{M}; \mathcal{D})$ depends on both \mathbf{M} and \mathcal{D} , the term $p(\mathbf{M}, \mathcal{D}) \exp(\langle \hat{\phi}, \psi(\mathbf{M}; \mathcal{D}) \rangle)$ implicitly defines a joint distribution on $(\mathbf{M}, \mathcal{D})$ if it has a finite measure. In this case, RegBayes is doing implicit Bayesian conditionalization, that is, the posterior $\hat{q}_{\hat{\phi}}(\mathbf{M})$ can be obtained through Bayes' rule with some well-defined prior and likelihood. However, it could be that the integral of $p(\mathbf{M}, \mathcal{D}) \exp(\langle \hat{\phi}, \psi(\mathbf{M}; \mathcal{D}) \rangle)$ with respect to $(\mathbf{M}, \mathcal{D})$ is not finite because of the way $\hat{\phi}$ varies with \mathcal{D} ,⁷ in which case there is no implicit prior and likelihood that give back*

6. We also used the fact that if $f(x) = \mathbb{I}(x = c)$ is an indicator function, its conjugate is $f^*(\mu) = c \cdot \mu$.

7. Note: this does not affect the well-normalization of the posterior $\hat{q}_{\hat{\phi}}(\mathbf{M})$ because its integral is taken over \mathbf{M} only, with \mathcal{D} fixed.

$\hat{q}_{\hat{\phi}}(\mathbf{M})$ through Bayesian conditionalization. Therefore, RegBayes is more flexible than the standard Bayesian inference, where the prior and likelihood model are explicitly defined, but no additional constraints or regularization can be systematically incorporated. The recent work (Mei et al., 2014) presents an example. Specifically, we show that incorporating domain knowledge via posterior regularization can lead to a flexible framework that automatically learns the importance of each piece of knowledge, thereby allowing for a robust incorporation, which is important in the scenarios where noisy knowledge is collected from crowds. In contrast, eliciting expert knowledge via fitting some priors is generally hard, especially in high-dimensional spaces, as experts are normally good at perceiving low-dimensional and well-behaved distributions but can be very bad in perceiving high-dimensional or skewed distributions (Garthwaite et al., 2005).

It is worth mentioning that although the above theorem provides a generic representation of the solution to RegBayes, in practice we usually need to make additional assumptions in order to make either the primal or dual problem tractable to solve. Since such assumptions could make the feasible space non-convex, additional cautions need to be paid. For instance, the mean-field assumptions will lead to a non-convex feasible space (Wainright and Jordan, 2008), and we can only apply the convex analysis theory to deal with convex sub-problems within an EM-type procedure. More concrete examples will be provided later along the developments of various models. We should also note that the modeling flexibility of RegBayes comes with risks. For example, it might lead to inconsistent posteriors (Barron et al., 1999; Choi and Ramamoorthi, 2008). This paper focuses on presenting several practical instances of RegBayes and we leave a systematic analysis of the Bayesian asymptotic properties (e.g., posterior consistency and convergence rates) for future work.

Now, we derive the conjugate functions of three examples which will be used shortly for developing the infinite latent SVM models we have intended. We defer the proof to Appendix C. Specifically, the first one is the conjugate of a simple function, which will be used in a binary latent SVM classification model.

Lemma 9 *Let $g_0 : \mathbb{R} \rightarrow \mathbb{R}$ be defined as $g_0(x) = C \max(0, x)$. Then, we have*

$$g_0^*(\mu) = \mathbb{I}(0 \leq \mu \leq C).$$

The second function is slightly more complex, which will be used for defining a multi-way latent SVM classifier. Specifically, we define the function $g_1 : \mathbb{R}^L \rightarrow \mathbb{R}$ as

$$g_1(\mathbf{x}) = C \max(\mathbf{x}), \tag{8}$$

where $\max(\mathbf{x}) \stackrel{\text{def}}{=} \max(x_1, \dots, x_L)$. Apparently, g_1 is convex because it is a point-wise maximum (Boyd and Vandenberghe, 2004) of the simple linear functions $\phi_i(\mathbf{x}) = x_i$. Then, we have the following results.

Lemma 10 *The convex conjugate of $g_1(\mathbf{x})$ as defined above is*

$$g_1^*(\boldsymbol{\mu}) = \mathbb{I}\left(\forall i, \mu_i \geq 0; \text{ and } \sum_i \mu_i = C\right).$$

Let $y \in \mathbb{R}$ and $\epsilon \in \mathbb{R}_+$ are fixed parameters. The last function that we are interested in is $g_2 : \mathbb{R} \rightarrow \mathbb{R}$, where

$$g_2(x; y, \epsilon) = C \max(0, |x - y| - \epsilon).$$

Finally, we have the following lemma, which will be used in developing large-margin regression models.

Lemma 11 *The convex conjugate of $g_2(x)$ as defined above is*

$$g_2^*(\mu; y, \epsilon) = \mu y + \epsilon |\mu| + \mathbb{I}(|\mu| \leq C).$$

4. Infinite Latent Support Vector Machines

Given the general theoretical framework of RegBayes introduced in Section 3, now we are ready to present its application to the development of two interesting nonparametric RegBayes models. In these two models we conjoin the ideas behind the nonparametric Bayesian infinite feature model known as the Indian buffet process (IBP), and the large margin classifier known as support vector machines (SVM) to build a new class of models for simultaneous single-task (or multi-task) classification and feature learning. A parametric Bayesian model is presented in Appendix B.

Specifically, to illustrate how to develop latent large-margin classifiers and automatically resolve the unknown dimensionality of latent features from data, we demonstrate how to choose/define the three key elements of RegBayes, that is, *prior distribution*, *likelihood model*, and *posterior regularization*. We first present the single-task classification model. The basic setup is that we project each data example $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^D$ to a latent feature vector \mathbf{z} . Here, we consider binary features. Real-valued features can be easily considered by elementwise multiplication of \mathbf{z} by a Gaussian vector (Griffiths and Ghahramani, 2005). Given a set of N data examples, let \mathbf{Z} be the matrix, of which each row is a binary vector \mathbf{z}_n associated with data sample n . Instead of pre-specifying a fixed dimension of \mathbf{z} , we resort to the nonparametric Bayesian methods and let \mathbf{z} have an infinite number of dimensions. To make the expected number of active latent features finite, we employ an IBP as prior for the binary feature matrix \mathbf{Z} , as reviewed below.

4.1 Indian Buffet Process

Indian buffet process (IBP) was proposed in Griffiths and Ghahramani (2005) and has been successfully applied in various fields, such as link prediction (Miller et al., 2009) and multi-task learning (Rai and Daume III, 2010). We will make use of its stick-breaking construction (Teh et al., 2007), which is good for developing efficient inference methods. Let $\pi_k \in (0, 1)$ be a parameter associated with each column of the binary matrix \mathbf{Z} . Given π_k , each z_{nk} in column k is sampled independently from Bernoulli(π_k). The parameter $\boldsymbol{\pi}$ are generated by a stick-breaking process

$$\pi_1 = \nu_1, \text{ and } \pi_k = \nu_k \pi_{k-1} = \prod_{i=1}^k \nu_i,$$

where $\nu_i \sim \text{Beta}(\alpha, 1)$. Since each ν_i is less than 1, this process generates a decreasing sequence of π_k . Specifically, given a finite data set, the probability of seeing feature k decreases exponentially with k .

IBP has several properties. For a finite number of rows, N , the prior of the IBP gives zero mass on matrices with an infinite number of ones, as the total number of columns with non-zero entries is $\text{Poisson}(\alpha H_N)$, where H_N is the N th harmonic number, $H_N = \sum_{j=1}^N \frac{1}{j}$. Thus, \mathbf{Z} has almost surely only a finite number of non-zero entries, though this number is unbounded. A second property of IBP is that the number of features possessed by each data point follows a $\text{Poisson}(\alpha)$ distribution. Therefore, the expected number of non-zero entries in \mathbf{Z} is $N\alpha$.

4.2 Infinite Latent Support Vector Machines

Consider a single-task, but multi-way classification, where each training data is provided with a categorical label $y \in \mathcal{Y} \stackrel{\text{def}}{=} \{1, \dots, L\}$. Suppose that the latent features \mathbf{z}_n for document n are given, then we can define the *latent discriminant function* as linear

$$f(y, \mathbf{x}_n, \mathbf{z}_n; \boldsymbol{\eta}) \stackrel{\text{def}}{=} \boldsymbol{\eta}^\top \mathbf{g}(y, \mathbf{x}_n, \mathbf{z}_n), \tag{9}$$

where $\mathbf{g}(y, \mathbf{x}_n, \mathbf{z}_n)$ is a vector stacking L subvectors of which the y th is \mathbf{z}_n^\top and all the others are zero;⁸ $\boldsymbol{\eta}$ is the corresponding infinite-dimensional vector of feature weights. Since we are doing Bayesian inference, we need to maintain the entire distribution profile of the latent feature matrix \mathbf{Z} . However, in order to make a prediction on the observed data \mathbf{x} , we need to remove the uncertainty of \mathbf{Z} . Here, we define the *effective discriminant function* as an expectation⁹ (i.e., a weighted average considering all possible values of \mathbf{Z}) of the latent discriminant function. To fully explore the flexibility offered by Bayesian inference, we also treat $\boldsymbol{\eta}$ as random and aim to infer its posterior distribution from given data. For the prior, we assume all the dimensions of $\boldsymbol{\eta}$ are independent and each dimension η_k follows the standard normal distribution. This is in fact a Gaussian process (GP) prior as $\boldsymbol{\eta}$ is infinite dimensional. More formally, the effective discriminant function $f : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$ is

$$\begin{aligned} f(y, \mathbf{x}_n; q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) &\stackrel{\text{def}}{=} \mathbb{E}_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})} [f(y, \mathbf{x}_n, \mathbf{z}_n; \boldsymbol{\eta})] \\ &= \mathbb{E}_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})} [\boldsymbol{\eta}^\top \mathbf{g}(y, \mathbf{x}_n, \mathbf{z}_n)], \end{aligned} \tag{10}$$

where $q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})$ is the post-data posterior distribution we want to infer. We have included \mathbf{W} as a place holder for any other variables we may define, e.g., the variables arising from a data likelihood model. Since we are taking the expectation, the variables which do not appear in the feature map \mathbf{g} (i.e., \mathbf{W}) will be marginalized out.

Before moving on, we should note that since we require q to be absolutely continuous with respect to the prior to make the KL-divergence term well defined in the `RegBayes`

8. We can consider the input features \mathbf{x}_n or its certain statistics in combination with the latent features \mathbf{z}_n to define a classifier boundary, by simply concatenating them in the subvectors.

9. Although other choices such as taking the mode are possible, our choice could lead to a computationally easy problem because expectation is a linear functional of the distribution under which the expectation is taken. Moreover, expectation can be more robust than taking the mode (Khan et al., 2010), and it has been widely used in previous work (Zhu et al., 2009, 2011b).

problem, $q(\mathbf{Z})$ will also put zero mass on \mathbf{Z} 's with an infinite number of non-zero entries, because of the properties of the IBP prior. The sparsity of \mathbf{Z} is essential to ensure that the dot-product in (9) and the expectation in (10) are well defined, i.e., with finite values.¹⁰ Moreover, in practice, to make the problem computationally feasible, we usually set a finite upper bound K to the number of possible features, where K is sufficiently large and known as the truncation level (See Section 4.4 and Appendix D.2 for details). As shown in Doshi-Velez (2009), the ℓ_1 -distance truncation error of marginal distributions decreases exponentially as K increases. For a finite truncation level, all the expectations are definitely finite.

Let \mathcal{I}_{tr} denote the set of training data. Then, with the above definitions, we define the $\mathcal{P}_{\text{post}}(\boldsymbol{\xi})$ in problem (5) using soft large-margin constraints as¹¹

$$\mathcal{P}_{\text{post}}^c(\boldsymbol{\xi}) \stackrel{\text{def}}{=} \left\{ q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) \left| \begin{array}{l} \forall n \in \mathcal{I}_{\text{tr}} : \Delta f(y, \mathbf{x}_n; q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) \geq \ell_n^\Delta(y) - \xi_n, \forall y \\ \xi_n \geq 0 \end{array} \right. \right\},$$

where $\Delta f(y, \mathbf{x}_n; q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) \stackrel{\text{def}}{=} f(y_n, \mathbf{x}_n; q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) - f(y, \mathbf{x}_n; q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}))$ is the margin favored by the true label y_n over an arbitrary label y and the superscript is used to distinguish from the posterior constraints for multi-task iLSVM to be presented. We define the penalty function for classification as

$$U^c(\boldsymbol{\xi}) \stackrel{\text{def}}{=} C \sum_{n \in \mathcal{I}_{\text{tr}}} \xi_n^\kappa,$$

where $\kappa \geq 1$. If κ is 1, minimizing $U^c(\boldsymbol{\xi})$ is equivalent to minimizing the hinge-loss (or ℓ_1 -loss) \mathcal{R}_h^c of the averaging prediction rule (13), where

$$\mathcal{R}_h^c(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) = C \sum_{n \in \mathcal{I}_{\text{tr}}} \max_y (\ell_n^\Delta(y) - \Delta f(y_n, \mathbf{x}_n; q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})));$$

if κ is 2, the surrogate loss is the squared ℓ_2 -loss. For clarity, we consider the hinge loss. The non-negative cost function $\ell_n^\Delta(y)$ (e.g., 0/1-cost) measures the cost of predicting \mathbf{x}_n to be y when its true label is y_n . \mathcal{I}_{tr} is the index set of training data.

Besides performing the prediction task, we may also be interested in explaining observed data \mathbf{x} using the latent factors \mathbf{Z} . This can be done by defining a likelihood model $p(\mathbf{x}|\mathbf{Z})$. Here, we define the most common linear-Gaussian likelihood model for real-valued data

$$p(\mathbf{x}_n|\mathbf{z}_n, \mathbf{W}, \sigma_{n0}^2) = \mathcal{N}(\mathbf{x}_n|\mathbf{W}\mathbf{z}_n^\top, \sigma_{n0}^2 I),$$

where \mathbf{W} is a $D \times \infty$ random loading matrix. We assume \mathbf{W} follows an independent Gaussian prior and each entry has the prior distribution $\pi(w_{dk}) = \mathcal{N}(w_{dk}|0, \sigma_0^2)$. The hyperparameters σ_0^2 and σ_{n0}^2 can be set a priori or estimated from observed data (See Appendix D.2 for details). Figure 2 (a) shows the graphical structure of iLSVM as defined above, where the plate means N replicates.

10. A more rigorous derivation of finiteness of these quantities is beyond the scope of this work and could require additional technical conditions (Orbanz, 2012). We refer the readers to Stummer and Vajda (2012) for a generic definition of Bregman divergence (or KL divergence in particular) on Banach spaces and in the case where the second measure is unnormalized.

11. Hard constraints for the separable cases are covered by simply setting $\boldsymbol{\xi} = 0$.

Training: Putting the above definitions together, we get the RegBayes problem for iLSVM in the following two equivalent forms

$$\begin{aligned} & \inf_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}), \boldsymbol{\xi}} \text{KL}(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) \| p(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}, \mathcal{D})) + U^c(\boldsymbol{\xi}) \\ & \text{s.t. : } q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) \in \mathcal{P}_{\text{post}}^c(\boldsymbol{\xi}) \end{aligned} \quad (11)$$

$$\iff \inf_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) \in \mathcal{P}_{\text{prob}}} \text{KL}(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) \| p(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}, \mathcal{D})) + \mathcal{R}_h^c(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})), \quad (12)$$

where $p(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}, \mathcal{D}) = \pi(\boldsymbol{\eta})\pi(\mathbf{Z})\pi(\mathbf{W}) \prod_{n=1}^N p(\mathbf{x}_n | \mathbf{z}_n, \mathbf{W}, \sigma_{n0}^2)$ is the joint distribution of the model; $\pi(\mathbf{Z})$ is an IBP prior; and $\pi(\boldsymbol{\eta})$ and $\pi(\mathbf{W})$ are Gaussian process priors with identity covariance functions.

Directly solving the iLSVM problems is not easy because either the posterior constraints or the non-smooth regularization function \mathcal{R} is hard to deal with. Thus, we resort to convex duality theory, which will be useful for developing approximate inference algorithms. We can either solve the constrained form (11) using Lagrangian duality theory (Ito and Kunisch, 2008) or solve the unconstrained form (12) using Fenchel duality theory. Here, we take the second approach. In this case, the linear operator is the expectation operator, denoted by $E : \mathcal{P}_{\text{prob}} \rightarrow \mathbb{R}^{|\mathcal{I}_{\text{tr}}| \times L}$ and the element of Eq evaluated at y for the n th example is

$$Eq(n, y) \stackrel{\text{def}}{=} \Delta f(y, \mathbf{x}_n; q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) = \mathbb{E}_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})} \left[\boldsymbol{\eta}^\top \Delta \mathbf{g}_n(y, \mathbf{Z}) \right],$$

where $\Delta \mathbf{g}_n(y, \mathbf{Z}) \stackrel{\text{def}}{=} \mathbf{g}(y_n, \mathbf{x}_n, \mathbf{z}) - \mathbf{g}(y, \mathbf{x}_n, \mathbf{z})$. Then, let $g_1 : \mathbb{R}^L \rightarrow \mathbb{R}$ be a function defined in the same form as in (8). We have

$$\mathcal{R}_h^c(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) = \sum_{n \in \mathcal{I}_{\text{tr}}} g_1(\ell_n^\Delta - Eq(n)),$$

where $Eq(n) \stackrel{\text{def}}{=} (Eq(n, 1), \dots, Eq(n, L))$ and $\ell_n^\Delta \stackrel{\text{def}}{=} (\ell_n^\Delta(1), \dots, \ell_n^\Delta(L))$ are the vectors of elements evaluated for n th data. By the Fenchel's duality theorem and the results in Lemma 10, we can derive the conjugate of the problem (12). The proof is deferred to Appendix C.4.

Lemma 12 (Conjugate of iLSVM) *For the iLSVM problem, we have that*

$$\begin{aligned} & \inf_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) \in \mathcal{P}_{\text{prob}}} \text{KL}(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) \| p(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}, \mathcal{D})) + \mathcal{R}_h^c(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) \\ & = \sup_{\boldsymbol{\omega}} -\log Z(\boldsymbol{\omega} | \mathcal{D}) + \sum_{n \in \mathcal{I}_{\text{tr}}} \sum_y \omega_n^y \ell_n^\Delta(y) - \sum_n g_1^*(\boldsymbol{\omega}_n), \end{aligned}$$

where $\boldsymbol{\omega}_n = (\omega_n^1, \dots, \omega_n^L)$ is the subvector associated with data n . Moreover, The optimum distribution is the posterior distribution

$$\hat{q}(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) = \frac{1}{Z(\hat{\boldsymbol{\omega}} | \mathcal{D})} p(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}, \mathcal{D}) \exp \left\{ \sum_{n \in \mathcal{I}_{\text{tr}}} \sum_y \hat{\omega}_n^y \boldsymbol{\eta}^\top \Delta \mathbf{g}_n(y, \mathbf{Z}) \right\},$$

where $Z(\hat{\boldsymbol{\omega}} | \mathcal{D})$ is the normalization factor and $\hat{\boldsymbol{\omega}}$ is the solution of the dual problem.




Figure 2: Graphical structures of (a) infinite latent SVM (iLSVM); and (b) multi-task infinite latent SVM (MT-iLSVM). For MT-iLSVM, the dashed nodes (i.e., ζ_m) illustrate the task relatedness but do not exist.

Testing: to make prediction on test examples, we put both training and test data together to do regularized Bayesian inference. For training data, we impose the above large-margin constraints because of the awareness of their true labels, while for test data, we do the inference without the large-margin constraints since we do not know their true labels. Therefore, the classifier $q(\eta)$ is learned from the training data only, while both training and testing data influence the posterior distributions of the likelihood model \mathbf{W} . After inference, we make the prediction via the rule

$$y^* \stackrel{\text{def}}{=} \underset{y}{\operatorname{argmax}} f(y, \mathbf{x}; q(\mathbf{Z}, \eta, \mathbf{W})). \quad (13)$$

Note that the ability to generalize to test data relies on the fact that all the data examples share η and the IBP prior. We can also cast the problem as a transductive inference problem by imposing additional large-margin constraints on test data (Joachims, 1999). However, the resulting problem will be generally harder to solve because it needs to resolve the unknown labels of testing examples. We also note that the testing is different from the standard inductive setting (Zhu et al., 2011b), where the latent features of a new data example can be approximately inferred given the training data. Our empirical study shows little difference on performance between our setting and the standard inductive setting.

4.3 Multi-Task Infinite Latent Support Vector Machines

Different from classification, which is typically formulated as a single learning task, multi-task learning aims to improve a set of related tasks through sharing statistical strength among these tasks, which are performed jointly. Many different approaches have been developed for multi-task learning; see Jebara (2011) for a review. In particular, learning a common latent representation shared by all the related tasks has proven to be an effective way to capture task relationships (Ando and Zhang, 2005; Argyriou et al., 2007; Rai and Daume III, 2010). Below, we present the multi-task infinite latent SVM (MT-iLSVM) for learning a common binary projection matrix \mathbf{Z} to capture the relationships among multiple

tasks. Similar as in iLSVM, we also put the IBP prior on \mathbf{Z} to allow it to have an unbounded number of columns.

Suppose we have M related tasks. Let $\mathcal{D}_m = \{(\mathbf{x}_{mn}, y_{mn})\}_{n \in \mathcal{I}_{\text{tr}}^m}$ be the training data for task m . We consider binary classification tasks, where $\mathcal{Y}_m = \{+1, -1\}$. Extension to multi-way classification or regression can be easily done. A naïve way to solve this learning problem with multiple tasks is to perform the multiple tasks independently. In order to make the multiple tasks coupled and share statistical strength, MT-iLSVM introduces a latent projection matrix \mathbf{Z} . If the latent matrix \mathbf{Z} is given, we define the *latent discriminant function* for task m as

$$f_m(\mathbf{x}_{mn}, \mathbf{Z}; \boldsymbol{\eta}_m) \stackrel{\text{def}}{=} (\mathbf{Z}\boldsymbol{\eta}_m)^\top \mathbf{x}_{mn} = \boldsymbol{\eta}_m^\top (\mathbf{Z}^\top \mathbf{x}_{mn}),$$

where \mathbf{x}_{mn} is one data example in \mathcal{D}_m and $\boldsymbol{\eta}_m$ is the vector of parameters for task m . The dimension of $\boldsymbol{\eta}_m$ is the number of columns of the latent projection matrix \mathbf{Z} , which is unbounded in the nonparametric setting. This definition provides two views of how the M tasks get related.

- (1) If we let $\varsigma_m = \mathbf{Z}\boldsymbol{\eta}_m$, then ς_m is the actual parameter of task m and all ς_m in different tasks are coupled by sharing the same latent matrix \mathbf{Z} ;
- (2) Another view is that each task m has its own parameters $\boldsymbol{\eta}_m$, but all the tasks share the same latent projection matrix \mathbf{Z} to extract latent features $\mathbf{Z}^\top \mathbf{x}_{mn}$, which is a projection of the input features \mathbf{x}_{mn} .

As such, our method can be viewed as a nonparametric Bayesian treatment of alternating structure optimization (ASO) (Ando and Zhang, 2005), which learns a single projection matrix with a pre-specified latent dimension. Moreover, different from Jebara (2011), which learns a binary vector with known dimensionality to select features or kernels on \mathbf{x} , we learn an unbounded projection matrix \mathbf{Z} using nonparametric Bayesian techniques.

As in iLSVM, we employ a Bayesian treatment of $\boldsymbol{\eta}_m$, and view it as random variables. We assume that $\boldsymbol{\eta}_m$ has a fully-factorized Gaussian prior, i.e., $\eta_{mk} \sim \mathcal{N}(0, 1)$. Then, we define the effective discriminant function for task m as the expectation

$$f_m(\mathbf{x}; q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) \stackrel{\text{def}}{=} \mathbb{E}_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})} [f_m(\mathbf{x}, \mathbf{Z}; \boldsymbol{\eta}_m)] = \mathbb{E}_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})} [\mathbf{Z}\boldsymbol{\eta}_m]^\top \mathbf{x},$$

where \mathbf{W} is a place holder for the variables that possibly arise from other parts of the model. As in iLSVM, since we are taking expectation, the variables which do not appear in the feature map (i.e., \mathbf{W}) will be marginalized out. Then, the prediction rule for task m is naturally $y_m^* \stackrel{\text{def}}{=} \text{sign} f_m(\mathbf{x})$. Similarly, we perform regularized Bayesian inference by defining:

$$U^{MT}(\boldsymbol{\xi}) \stackrel{\text{def}}{=} C \sum_{m, n \in \mathcal{I}_{\text{tr}}^m} \xi_{mn}$$

and imposing the following constraints:

$$\mathcal{P}_{\text{post}}^{MT}(\boldsymbol{\xi}) \stackrel{\text{def}}{=} \left\{ q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) \left| \forall m, \forall n \in \mathcal{I}_{\text{tr}}^m : y_{mn} \mathbb{E}_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})} [\mathbf{Z}\boldsymbol{\eta}_m]^\top \mathbf{x}_{mn} \geq 1 - \xi_{mn} \right. \right\}. \quad (14)$$

Finally, as in iLSVM we may also be interested in explaining observed data \mathbf{x} . Therefore, we relate \mathbf{Z} to the observed data \mathbf{x} by defining a likelihood model:

$$p(\mathbf{x}_{mn}|\mathbf{w}_{mn}, \mathbf{Z}, \lambda_{mn}^2) = \mathcal{N}(\mathbf{x}_{mn}|\mathbf{Z}\mathbf{w}_{mn}, \lambda_{mn}^2 I),$$

where \mathbf{w}_{mn} is a vector. We assume that \mathbf{W} (the collection of \mathbf{w}_{mn}) has an independent prior $\pi(\mathbf{W}) = \prod_{mn} \mathcal{N}(\mathbf{w}_{mn}|0, \sigma_{m0}^2 I)$. Fig. 2 (b) illustrates the graphical structure of MT-iLSVM.

For training, we can derive the similar convex conjugate as in the case of iLSVM. Similar as in iLSVM, minimizing $U^{MT}(\boldsymbol{\xi})$ is equivalent to minimizing the hinge-loss \mathcal{R}_h^{MT} of the multiple binary prediction rules, where

$$\mathcal{R}_h^{MT}(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) = C \sum_{m,n \in \mathcal{I}_{\text{tr}}^m} \max\left(0, 1 - y_{mn} \mathbb{E}_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})}[\mathbf{Z}\boldsymbol{\eta}_m]^\top \mathbf{x}_{mn}\right).$$

Thus, the RegBayes problem of MT-iLSVM can be equivalently written as

$$\inf_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})} \text{KL}(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})\|p(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}, \mathcal{D})) + \mathcal{R}_h^{MT}(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})). \quad (15)$$

Then, by the Fenchel's duality theorem and Lemma 9, we can derive the conjugate of MT-iLSVM. The proof is deferred to Appendix C.5.

Lemma 13 (Conjugate of MT-iLSVM) *For the MT-iLSVM problem, we have that*

$$\begin{aligned} & \inf_{q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) \in \mathcal{P}_{\text{prob}}} \text{KL}(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})\|p(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}, \mathcal{D})) + \mathcal{R}_h^{MT}(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) \\ &= \sup_{\boldsymbol{\omega}} -\log Z'(\boldsymbol{\omega}|\mathcal{D}) + \sum_{m,n} \omega_{mn} - \sum_{m,n} g_0^*(\omega_{mn}). \end{aligned}$$

Moreover, The optimum distribution is the posterior distribution

$$\hat{q}(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) = \frac{1}{Z'(\hat{\boldsymbol{\omega}}|\mathcal{D})} p(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}, \mathcal{D}) \exp\left\{\sum_{m,n} y_{mn} \hat{\omega}_{mn} (\mathbf{Z}\boldsymbol{\eta}_m)^\top \mathbf{x}_{mn}\right\},$$

where $Z'(\hat{\boldsymbol{\omega}}|\mathcal{D})$ is the normalization factor and $\hat{\boldsymbol{\omega}}$ is the solution of the dual problem.

For testing, we use the same strategy as in iLSVM to do Bayesian inference on both training and test data. The difference is that training data are subject to large-margin constraints, while test data are not. Similarly, the hyper-parameters σ_{m0}^2 and λ_{mn}^2 can be set a priori or estimated from data (See Appendix D.1 for details).

4.4 Inference with Truncated Mean-Field Constraints

Now we discuss how to perform regularized Bayesian inference with the large-margin constraints for both iLSVM and MT-iLSVM. From the primal-dual formulations, it is obvious that there are basically two methods to perform the regularized Bayesian inference. One is to directly solve the primal problem for the posterior distribution $q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})$, and the other is to first solve the dual problem for the optimum $\hat{\boldsymbol{\omega}}$ and then infer the posterior distribution. However, both the primal and dual problems are intractable for iLSVM and

Algorithm 1 Inference Algorithm for Infinite Latent SVMs

- 1: **Input:** corpus \mathcal{D} and constants (α, C) .
 - 2: **Output:** posterior distribution $q(\boldsymbol{\nu}, \mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})$.
 - 3: **repeat**
 - 4: infer $q(\boldsymbol{\nu}), q(\mathbf{W})$ and $q(\mathbf{Z})$ with $q(\boldsymbol{\eta})$ and $\boldsymbol{\omega}$ given;
 - 5: infer $q(\boldsymbol{\eta})$ and solve for $\boldsymbol{\omega}$ with $q(\mathbf{Z})$ given.
 - 6: **until** convergence
-

MT-iLSVM. The intrinsic hardness is due to the mutual dependency among the latent variables in the desired posterior distribution. Therefore, a natural approximation method is the mean field (Jordan et al., 1999), which breaks the mutual dependency by assuming that q is of some factorization form. This method approximates the original problems by imposing additional constraints. An alternative method is to apply approximate methods (e.g., MCMC sampling) to infer the true posterior distributions derived via convex conjugates as above, and iteratively estimate the dual parameters using approximate statistics (e.g., feature expectations estimated using samples) (Schofield, 2006). Below, we use MT-iLSVM as an example to illustrate the idea of the first strategy. A full discussion on the second strategy is beyond the scope of this paper. For iLSVM, the similar procedure applies and we defer its details to Appendix D.2.

To make the problem easier to solve, we use the stick-breaking representation of IBP, which includes the auxiliary variable $\boldsymbol{\nu}$, and infer the augmented posterior $q(\boldsymbol{\nu}, \mathbf{W}, \mathbf{Z}, \boldsymbol{\eta})$. The joint model distribution is now $q(\boldsymbol{\nu}, \mathbf{W}, \mathbf{Z}, \boldsymbol{\eta}, \mathcal{D})$. Furthermore, we impose the truncated mean-field constraint that

$$q(\boldsymbol{\nu}, \mathbf{W}, \mathbf{Z}, \boldsymbol{\eta}) = q(\boldsymbol{\eta}) \prod_{k=1}^K \left(q(\nu_k | \gamma_k) \prod_{d=1}^D q(z_{dk} | \psi_{dk}) \right) \prod_{mn} q(\mathbf{w}_{mn} | \Phi_{mn}, \sigma_{mn}^2 I), \quad (16)$$

where K is the truncation level, and we assume that

$$\begin{aligned} q(\nu_k | \gamma_k) &= \text{Beta}(\gamma_{k1}, \gamma_{k2}), \\ q(z_{dk} | \psi_{dk}) &= \text{Bernoulli}(\psi_{dk}), \\ q(\mathbf{w}_{mn} | \Phi_{mn}, \sigma_{mn}^2 I) &= \mathcal{N}(\mathbf{w}_{mn} | \Phi_{mn}, \sigma_{mn}^2 I). \end{aligned}$$

Then, we can use the duality theory¹² to solve the RegBayes problem by alternating between two substeps, as outlined in Algorithm 1 and detailed below.

Infer $q(\boldsymbol{\nu}), q(\mathbf{W})$ and $q(\mathbf{Z})$: Since $q(\boldsymbol{\nu})$ and $q(\mathbf{W})$ are not directly involved in the posterior constraints, we can solve for them by using standard Bayesian inference, i.e., minimizing a KL-divergence. Specifically, for $q(\mathbf{W})$, since the prior is also normal, we can easily derive the update rules for Φ_{mn} and σ_{mn}^2 . For $q(\boldsymbol{\nu})$, we have the same update rules as in Doshi-Velez (2009). We defer the details to Appendix D.1.

¹² Lagrangian duality (Ito and Kunisch, 2008) was used in Zhu et al. (2011a) to solve the constrained variational formulations, which is closely related to Fenchel duality (Magnanti, 1974) and leads to the same solutions for iLSVM and MT-iLSVM.

For $q(\mathbf{Z})$, it is directly involved in the posterior constraints. So, we need to solve it together with $q(\boldsymbol{\eta})$ using conjugate theory. However, this is intractable. Here, we adopt an alternating strategy that first infers $q(\mathbf{Z})$ with $q(\boldsymbol{\eta})$ and dual parameters $\boldsymbol{\omega}$ fixed, and then infers $q(\boldsymbol{\eta})$ and solves for $\boldsymbol{\omega}$. Specifically, since the large-margin constraints are linear of $q(\mathbf{Z})$, we can get the mean-field update equation as

$$\psi_{dk} = \frac{1}{1 + e^{-\vartheta_{dk}}},$$

where

$$\begin{aligned} \vartheta_{dk} = & \sum_{j=1}^k \mathbb{E}_q[\log v_j] - \mathcal{L}_k^\nu - \sum_{mn} \frac{1}{2\lambda_{mn}^2} \left((K\sigma_{mn}^2 + (\phi_{mn}^k)^2) \right. \\ & \left. - 2x_{mn}^d \phi_{mn}^k + 2 \sum_{j \neq k} \phi_{mn}^j \phi_{mn}^k \psi_{dj} \right) + \sum_{m,n \in \mathcal{I}_{\text{tr}}^m} y_{mn} \mathbb{E}_q[\eta_{mk}] x_{mn}^d, \end{aligned}$$

and \mathcal{L}_k^ν is a lower bound of $\mathbb{E}_q[\log(1 - \prod_{j=1}^k v_j)]$ (See Appendix D.1 for details). The last term of ϑ_{dk} is due to the large-margin posterior constraints as defined in (14). Therefore, from this equation we can see how the large-margin constraints regularize the procedure of inferring the latent matrix \mathbf{Z} .

Infer $q(\boldsymbol{\eta})$ and solve for $\boldsymbol{\omega}$: Now, we can apply the convex conjugate theory and show that the optimum posterior distribution of $\boldsymbol{\eta}$ is

$$q(\boldsymbol{\eta}) = \prod_m q(\boldsymbol{\eta}_m), \text{ where } q(\boldsymbol{\eta}_m) \propto \pi(\boldsymbol{\eta}_m) \exp \left\{ \boldsymbol{\eta}_m^\top \boldsymbol{\mu}_m \right\},$$

and $\boldsymbol{\mu}_m = \sum_{n \in \mathcal{I}_{\text{tr}}^m} y_{mn} \omega_{mn} (\boldsymbol{\psi}^\top \mathbf{x}_{mn})$. Here, we assume $\pi(\boldsymbol{\eta}_m)$ is standard normal. Then, we have $q(\boldsymbol{\eta}_m) = \mathcal{N}(\boldsymbol{\eta}_m | \boldsymbol{\mu}_m, I)$ and the optimum dual parameters can be obtained by solving the following M independent dual problems

$$\begin{aligned} \sup_{\boldsymbol{\omega}_m} \quad & -\frac{1}{2} \boldsymbol{\mu}_m^\top \boldsymbol{\mu}_m + \sum_{n \in \mathcal{I}_{\text{tr}}^m} \omega_{mn} \\ \forall n \in \mathcal{I}_{\text{tr}}^m, \text{ s.t. : } \quad & 0 \leq \omega_{mn} \leq C, \end{aligned}$$

where the constraints are from the conjugate function g_0^* in Lemma 13. These dual problems (or their primal forms) can be efficiently solved with a binary SVM solver, such as SVM-light or LibSVM.

5. Experiments

We present empirical results for both classification and multi-task learning. Our results appear to demonstrate the merits inherited from both Bayesian nonparametrics and large-margin learning.

5.1 Multi-way Classification

We evaluate the infinite latent SVM (iLSVM) for classification on the real TRECVID2003 and Flickr image data sets, which have been extensively evaluated in the context of learning

Model	TRECVID2003		Flickr	
	Accuracy	F1 score	Accuracy	F1 score
EFH+SVM	0.565 \pm 0.0	0.427 \pm 0.0	0.476 \pm 0.0	0.461 \pm 0.0
MMH	0.566 \pm 0.0	0.430 \pm 0.0	0.538 \pm 0.0	0.512 \pm 0.0
IBP+SVM	0.553 \pm 0.013	0.397 \pm 0.030	0.500 \pm 0.004	0.477 \pm 0.009
iLSVM	0.563 \pm 0.010	0.448 \pm 0.011	0.533 \pm 0.005	0.510 \pm 0.010

Table 1: Classification accuracy and F1 scores on the TRECVID2003 and Flickr image data sets (Note: MMH and EFH have zero std because of their deterministic initialization).




Figure 3: Accuracy and F1 score of MMH on the Flickr data set with different numbers of latent features.

finite latent feature models (Chen et al., 2012). TRECVID2003 consists of 1078 video keyframes that belong to 5 categories, including *Airplane scene*, *Basketball scene*, *Weather news*, *Baseball scene*, and *Hockey scene*. Each data example has two types of features, including a 1894-dimension binary vector of text features and a 165-dimension HSV color histogram. The Flickr image data set consists of 3411 natural scene images about 13 types of animals, including *squirrel*, *cow*, *cat*, *zebra*, *tiger*, *lion*, *elephant*, *whales*, *rabbit*, *snake*, *antlers*, *hawk and wolf*, downloaded from the Flickr website.¹³ Also, each example has two types of features, including 500-dimension SIFT bag-of-words and 634-dimension real-valued features (e.g., color histogram, edge direction histogram, and block-wise color moments). Here, we consider the real-valued features only by defining Gaussian likelihood distributions for \mathbf{x} ; and we define the discriminant function using latent features only as in (9). We follow the same training/testing splits as in Chen et al. (2012).

We compare iLSVM with the large-margin Harmonium (MMH) (Chen et al., 2012), which was shown to outperform many other latent feature models, and two decoupled approaches of *EFH+SVM* and *IBP+SVM*. EFH+SVM uses the exponential family Harmo-

13. The website is available at: <http://www.flickr.com/>.




Figure 4: (Up) the overall average values of the latent features with standard deviation over different classes; and (Bottom) the per-class average values of latent features learned by iLSVM on the TRECVID data set.




Figure 5: The overall average values of the latent features with standard deviation over different classes on the Flickr data set.

nium (EFH) (Welling et al., 2004) to discover latent features and then learns a multi-way SVM classifier. IBP+SVM is similar, but uses an IBP factor analysis model to discover latent features (Griffiths and Ghahramani, 2005). To initialize the learning algorithms for these models, we found that using the SVD factors of the input feature matrix as the initial weights for MMH and EFH can produce better results. Here, we also use the SVD factors as the initial mean of weights in the likelihood models for iLSVM. Both MMH and EFH+SVM are finite models and they need to pre-specify the dimensionality of latent features. We report their results on classification accuracy and F1 score (i.e., the average F1 score over all possible classes) (Zhu et al., 2011b) achieved with the best dimensionality in Table 1. Figure 3 illustrates the performance change of MMH when using different number of latent features, from which we can see that $K = 40$ produces the best performance and either increasing or decreasing K could make the performance worse. For




Figure 6: Six example features discovered iLSVM on the Flickr animal data set. For each feature, we show 5 top-ranked images.

iLSVM and IBP+SVM, we use the mean-field inference method and present the average performance with 5 randomly initialized runs (Please see Appendix D.2 for the algorithm and initialization details). We perform 5-fold cross-validation on training data to select hyperparameters, e.g., α and C (we use the same procedure for MT-iLSVM). We can see that iLSVM can achieve comparable performance with the nearly optimal MMH, without needing to pre-specify the latent feature dimension,¹⁴ and is much better than the decoupled approaches (i.e., IBP+SVM and EFH+SVM). For the two stage methods, we don't have a clear winner—IBP+SVM performs a bit worse than EFH+SVM on the TRECVID data set, while it outperforms EFH+SVM on the flickr data set. The reason for the difference may be due to the initialization or different properties of the data.

It is also interesting to examine the discovered latent features. Figure 4 shows the overall average values of latent features and the per-class average feature values of iLSVM in one run on the TRECVID data set. We can see that on average only about 45 features are active for the TRECVID data set. For the overall average, we also present the standard deviation over the 5 categories. A larger deviation means that the corresponding feature is more discriminative when predicting different categories. For example, feature 26 and feature 34 are generally less discriminative than many other features, such as feature 1

14. We set the truncation level to 300, which is sufficiently large.

and feature 30. Figure 5 shows the overall average feature values together with standard deviation on the Flickr data set. We omitted the per-class average because that figure is too crowded with 13 categories. We can see that as k increases, the probability that feature k is active decreases. The reason for the features with stable values (i.e., standard deviations are extremely small) is due to our initialization strategy (each feature has 0.5 probability to be active). Initializing ψ_{dk} as being exponentially decreasing (e.g., like the constructing process of $\boldsymbol{\pi}$) leads to a faster decay and many features will be inactive. To examine the semantics of each feature,¹⁵ Figure 6 presents some example features discovered on the Flickr animal data set. For each feature, we present 5 top-ranked images which have large values on this particular feature. We can see that most of the features are semantically interpretable. For instance, feature F1 is about squirrel; feature F2 is about ocean animal, which is whales in the Flickr data set; and feature F4 is about hawk. We can also see that some features are about different aspects of the same category. For example, feature F2 and feature F3 are both about whales, but with different background.

5.2 Multi-task Learning

Now, we evaluate the multi-task infinite latent SVM (MT-iLSVM) on several well-studied real data sets.

5.2.1 DESCRIPTION OF THE DATA

Scene and Yeast Data: These data sets are from the UCI repository, and each data example has multiple labels. As in Rai and Daume III (2010), we treat the multi-label classification as a multi-task learning problem, where each label assignment is treated as a binary classification task. The Yeast data set consists of 1500 training and 917 test examples, each having 103 features, and the number of labels (or tasks) per example is 14. The Scene data set consists 1211 training and 1196 test examples, each having 294 features, and the number of labels (or tasks) per example for this data set is 6.

School Data: This data set comes from the Inner London Education Authority and has been used to study the effectiveness of schools. It consists of examination records of 15,362 students from 139 secondary schools in years 1985, 1986 and 1987. The data set is publicly available and has been extensively evaluated in various multi-task learning methods (Bakker and Heskes, 2003; Bonilla et al., 2008; Zhang and Yeung, 2010), where each task is defined as predicting the exam scores of students belonging to a specific school based on four student-dependent features (year of the exam, gender, VR band and ethnic group) and four school-dependent features (percentage of students eligible for free school meals, percentage of students in VR band 1, school gender and school denomination). In order to compare with the above methods, we follow the same setup described by Argyriou et al. (2007) and Bakker and Heskes (2003) and similarly we create dummy variables for those features that are categorical forming a total of 19 student-dependent features and 8 school-dependent features. We use the same 10 random splits¹⁶ of the data, so that 75% of the examples from each school (task) belong to the training set and 25% to the test set.

15. The interpretation of latent features depends heavily on the input data.

16. The splits are available at: <http://ttic.uchicago.edu/~argyriou/code/index.html>.

Data set	Model	Acc	F1-Micro	F1-Macro
Yeast	YaXue	0.5106	0.3897	0.4022
	Piyushrai-1	0.5212	0.3631	0.3901
	Piyushrai-2	0.5424	0.3946	0.4112
	MT-IBP+SVM	0.5475 ± 0.005	0.3910 ± 0.006	0.4345 ± 0.007
	MT-iLSVM	0.5792 ± 0.003	0.4258 ± 0.005	0.4742 ± 0.008
Scene	YaXue	0.7765	0.2669	0.2816
	Piyushrai-1	0.7756	0.3153	0.3242
	Piyushrai-2	0.7911	0.3214	0.3226
	MT-IBP+SVM	0.8590 ± 0.002	0.4880 ± 0.012	0.5147 ± 0.018
	MT-iLSVM	0.8752 ± 0.004	0.5834 ± 0.026	0.6148 ± 0.020

Table 2: Multi-label classification performance on Scene and Yeast data sets.

On average, the training set includes about 80 students per school and the test set about 30 students per school.

5.2.2 RESULTS

Scene and Yeast Data: We compare with the closely related nonparametric Bayesian methods, including kernel stick-breaking (YaXue) (Xue et al., 2007) and the basic and augmented infinite predictor subspace models (i.e., Piyushrai-1 and Piyushrai-2) proposed by Rai and Daume III (2010). These nonparametric Bayesian models were shown to outperform the independent Bayesian logistic regression and a single-task pooling approach in previous work (Rai and Daume III, 2010). We also compare with a decoupled method *MT-IBP+SVM* that uses an IBP factor analysis model to find shared latent features among multiple tasks and then builds separate SVM classifiers for different tasks.¹⁷ For MT-iLSVM and MT-IBP+SVM, we use the mean-field inference method in Sec 4.4 and report the average performance with 5 randomly initialized runs (See Appendix D.1 for initialization details). For comparison with Rai and Daume III (2010) and Xue et al. (2007), we use the overall classification accuracy, F1-Macro and F1-Micro as performance measures. Table 2 shows the results. On both data sets, MT-iLSVM needs less than 50 latent features on average. We can see that the large-margin MT-iLSVM performs much better than other nonparametric Bayesian methods and MT-IBP+SVM, which separates the inference of latent features from learning the classifiers.

School Data: We use the percentage of explained variance (Bakker and Heskes, 2003) as the measure of the regression performance, which is defined as the total variance of the data minus the sum-squared error on the test set as a percentage of the total variance. Since we use the same settings, we can compare with the state-of-the-art results of

- (1) Bayesian multi-task learning (BMTL) (Bakker and Heskes, 2003);
- (2) Multi-task Gaussian processes (MTGP) (Bonilla et al., 2008);

¹⁷. This decoupled approach is in fact an one-iteration MT-iLSVM, where we first infer the shared latent matrix \mathbf{Z} and then learn an SVM classifier for each task.




Figure 7: Percentage of explained variance by various models on the School data set.

(3) Convex multi-task relationship learning (MTRL) (Zhang and Yeung, 2010);

and single-task learning (STL) as reported in Bonilla et al. (2008) and Zhang and Yeung (2010). For MT-iLSVM and MT-IBP+SVM, we also report the results achieved by using both the latent features (i.e., $\mathbf{Z}^\top \mathbf{x}$) and the original input features \mathbf{x} through vector concatenation, and we denote the corresponding methods by $MT-iLSVM^f$ and $MT-IBP+SVM^f$, respectively. On average the multi-task latent SVM (i.e., MT-iLSVM) needs about 50 latent features to get sufficiently good and robust performance. From the results in Figure 7, we can see that the MT-iLSVM achieves better results than the existing methods that have been tested in previous studies. Again, the joint MT-iLSVM performs much better than the decoupled method MT-IBP+SVM, which separates the latent feature inference from the training of large-margin classifiers. Finally, using both latent features and the original input features can boost the performance slightly for MT-iLSVM, while much more significantly for the decoupled MT-IBP+SVM.

5.3 Sensitivity Analysis

Figure 8 shows how the performance of MT-iLSVM changes against the hyper-parameter α and regularization constant C on the Yeast and School data sets. We can see that on the Yeast data set, MT-iLSVM is insensitive to both α and C . For the School data set, MT-iLSVM is very insensitive to α , and it is stable when C is set between 0.3 and 1.

Figure 9 shows how the training size affects the performance and running time of MT-iLSVM on the School data set. We use the first $b\%$ ($b = 50, 60, 70, 80, 90, 100$) of the training data in each of the 10 random splits as training set and use the corresponding test data as test set. We can see that as training size increases, the performance and running time generally increase; and MT-iLSVM achieves the state-of-art performance when using about 70% training data. From the running time, we can also see that MT-iLSVM is generally quite efficient by using mean-field inference.

Finally, we investigate how the performance of MT-iLSVM changes against the hyper-parameters σ_{m0}^2 and λ_{mn}^2 . We initially set $\sigma_{m0}^2 = 1$ and compute λ_{mn}^2 from observed data. If we further estimate them by maximizing the objective function, the performance does

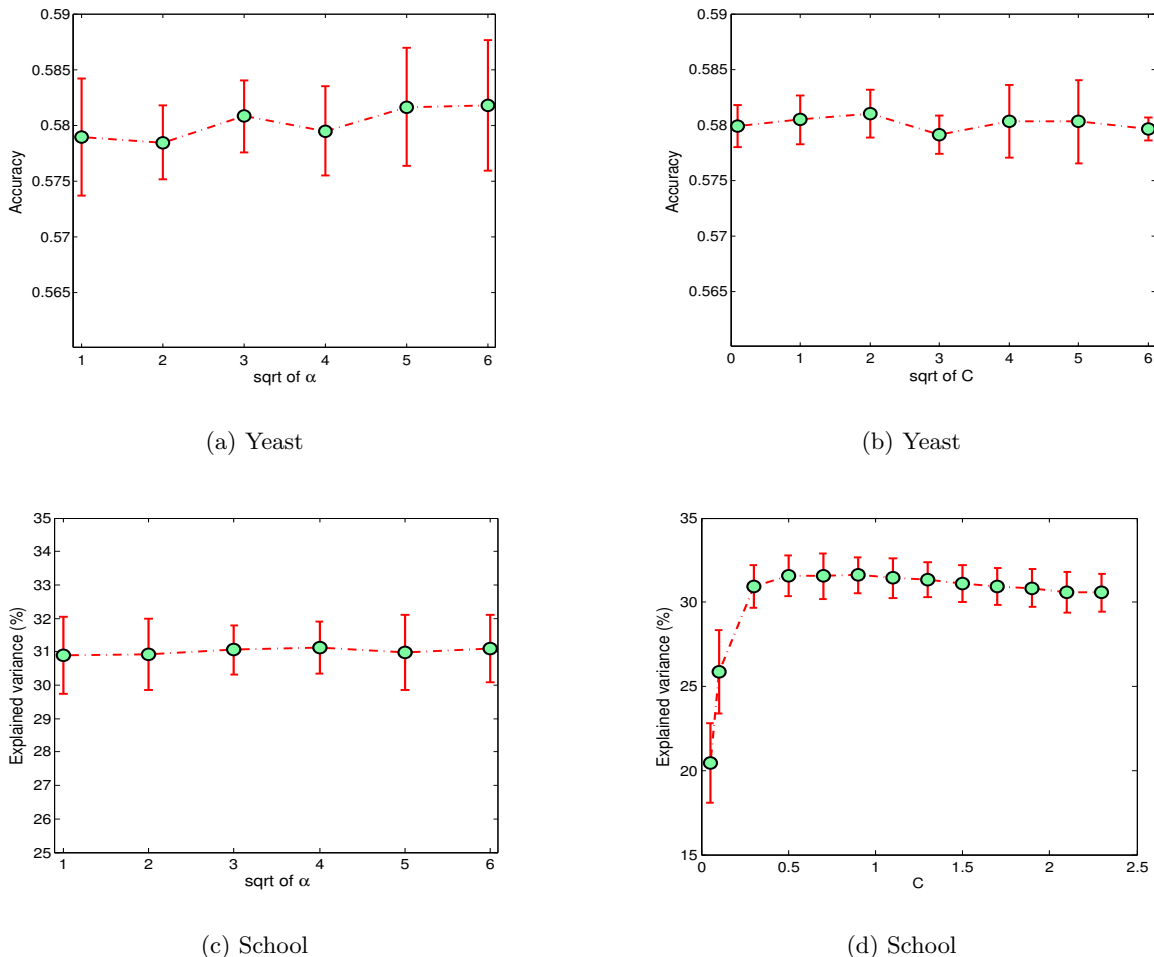


Figure 8: Sensitivity study of MT-iLSVM: (a) classification accuracy with different α on Yeast data; (b) classification accuracy with different C on Yeast data; (c) percentage of explained variance with different α on School data; and (d) percentage of explained variance with different C on School data.

not change much ($\pm 0.3\%$ for average explained variance on the School data set). We have similar observations for iLSVM.

6. Conclusions and Discussions

We present regularized Bayesian inference (RegBayes), a computational framework to perform post-data posterior inference with a rich set of regularization/constraints on the desired post-data posterior distributions. RegBayes is formulated as an information-theoretical optimization problem, and it is applicable to both directed and undirected graphical models. We present a general theorem to characterize the solution of RegBayes, when the posterior




Figure 9: Percentage of explained variance and running time by MT-iLSVM with various training sizes.

regularization is induced from a linear operator (e.g., expectation). Furthermore, we particularly concentrate on developing two large-margin nonparametric Bayesian models under the RegBayes framework to learn predictive latent features for classification and multi-task learning, by exploring the large-margin principle to define posterior constraints. Both models allow the latent dimension to be automatically resolved from the data. The empirical results on several real data sets appear to demonstrate that our methods inherit the merits from both Bayesian nonparametrics and large-margin learning.

RegBayes offers a flexible framework for considering posterior regularization in performing parametric or nonparametric Bayesian inference. For future work, we plan to study other posterior regularization beyond the large-margin constraints, such as posterior constraints defined on manifold structures (Huh and Fienberg, 2012) and those represented in the form of first-order logic, and investigate how posterior regularization can be used in other interesting nonparametric Bayesian models (Beal et al., 2002; Teh et al., 2006; Blei and Frazier, 2010) in different contexts, such as link prediction (Miller et al., 2009) for social network analysis and low-rank matrix factorization for collaborative prediction. Some of our preliminary results (Xu et al., 2012; Zhu, 2012; Mei et al., 2014) have shown great promise. It is interesting to investigate more carefully along this direction. Moreover, as we have stated, RegBayes can be developed for undirected MRFs. But the inference would be even harder. We plan to do a systematic investigation along this direction too. We have some preliminary results presented in Chen et al. (2014), but there is a lot of room to further improve. Finally, regularized Bayesian inference in general leads to a highly nontrivial inference problem. Although the general solution can be derived with convex analysis theory, it is normally intractable to infer them directly. Therefore, approximate inference techniques such as the truncated mean-field approximation have to be used. For the current truncated inference methods, one key limit is to pre-specify the truncation level. A too conservative truncation level could lead to a waste of computing resources. So, it is important to develop

inference algorithms that could adaptively determine the number of latent features, such as Monte Carlo methods. We have some preliminary progress along this direction as reported in Jiang et al. (2012) and Zhu et al. (2013). It is interesting to extend these techniques to deal with other challenging nonparametric Bayesian models.

Acknowledgements

We thank the anonymous reviewers and editors for many helpful comments to improve the manuscript. NC and JZ are supported by National Key Foundation R&D Projects (No.s 2013CB329403, 2012CB316301), National Natural Science Foundation of China (Nos. 61322308, 61332007, 61305066), Tsinghua University Initiative Scientific Research Program (No. 20121088071), a Microsoft Research Asia Research Fund (No. 20123000007), and the China Postdoctoral Science Foundation Grant (No. 2013T60117) to NC. EX is supported by AFOSR FA95501010247, ONR N000140910758, NSF Career DBI-0546594 and an Alfred P. Sloan Research Fellowship.

Appendix A. Generalization Beyond Bayesian Networks

Standard Bayesian inference and the proposed regularized Bayesian inference implicitly make the assumption that the model can be graphically drawn as a Bayesian network as illustrated in Figure 10(a).¹⁸ Here, we consider a more general formulation which could cover both directed and undirected latent variable models, such as the well-studied Boltzmann machines (Murray and Ghahramani, 2004; Welling et al., 2004), as well as the case where a model could have some unknown parameters (e.g., hyper-parameters) and need an estimation procedure, such as maximum likelihood estimation (MLE), besides posterior inference. The latter is also known as empirical Bayesian methods, which are frequently employed by practitioners.

Extension 1: Empirical Bayesian Inference with Unknown Parameters: As illustrated in Figure 10(b), in some cases we need to perform the empirical Bayesian inference in the presence of unknown parameters. For instance, in a linear-Gaussian Bayesian model, we may choose to estimate its covariance matrix using MLE; and in a latent Dirichlet allocation (LDA) (Blei et al., 2003) model, we may choose to estimate the unknown topical dictionary, although in principle we can treat these parameters as random variables and perform full Bayesian inference. In such cases, we need some mechanisms to estimate the unknown parameters when doing Bayesian inference. Let Θ be model parameters. We can formulate empirical Bayesian inference as solving¹⁹

$$\begin{aligned} \inf_{\Theta, q(\mathbf{M})} \text{KL}(q(\mathbf{M}) \parallel \pi(\mathbf{M})) - \int_{\mathcal{M}} \log p(\mathcal{D} | \mathbf{M}, \Theta) q(\mathbf{M}) d\mathbf{M} \\ \text{s.t. : } q(\mathbf{M}) \in \mathcal{P}_{\text{prob.}} \end{aligned} \quad (17)$$

18. The structure within \mathbf{M} can be arbitrary, either a directed, undirected or hybrid chain graph.

19. The objective can be derived using variational techniques. It is in fact a variational upper bound of the negative log-likelihood.




Figure 10: Illustration graphs for three different types of models that involve Bayesian inference: (a) a Bayesian generative model; (b) a Bayesian generative model with unknown parameters Θ ; and (c) a chain graph model.

Although the problem is convex over $q(\mathbf{M})$ for any fixed Θ , it is not jointly convex in general. A natural algorithm to solve this problem is the well-known EM procedure (Dempster et al., 1977), which converges to a local optimum. Specifically, we have the following result.

Lemma 14 *For problem (17), the optimum solution of $q(\mathbf{M})$ is equivalent to the posterior distribution by Bayes' theorem for any Θ ; and the optimum Θ^* is the MLE*

$$\Theta^* = \operatorname{argmax}_{\Theta} \log p(\mathcal{D}|\Theta).$$

Proof According to the variational formulation of Bayes' rule in (4), we get that the optimum solution is $q(\mathbf{M}) = p(\mathbf{M}|\mathcal{D}, \Theta)$ for any Θ . Substituting the optimum solution of q into the objective, we get the optimization problem of Θ . ■

Extension 2: Chain Graph: In the above cases, we have assumed that the observed data are generated by some model in a directed causal sense. This assumption holds in directed latent variable models. However, in many cases, we may choose alternative formulations to define the joint distribution of a model and the observed data. Figure 10(c) illustrates one such scenario, where the model \mathbf{M} consists of two subsets of random variables. One subset H is connected to the observed data via an undirected graph and the other subset Z is connected to the observed data and H using directed edges. This graph is known as a chain graph. Due to the Markov properties of chain graph (Frydenberg, 1990), we know that the joint distribution has the factorization form as

$$p(\mathbf{M}, \mathcal{D}) = p(Z)p(H, \mathcal{D}|Z), \quad (18)$$

where $p(H, \mathcal{D}|Z)$ is a Markov random field (MRF). One concrete example of such a hybrid chain model is the Bayesian Boltzman machines (Murray and Ghahramani, 2004), which

treat the parameters of a Boltzmann machine as random variables and perform Bayesian inference with MCMC sampling methods.

The insights that RegBayes covers undirected or chain graph latent variable models come from the observation that the objective $\mathcal{L}(q(\mathbf{M}))$ of problem (4) is in fact an KL-divergence, namely, we can show that

$$\mathcal{L}(q(\mathbf{M})) = \text{KL}(q(\mathbf{M})\|p(\mathbf{M}, \mathcal{D})), \tag{19}$$

where $p(\mathbf{M}, \mathcal{D})$ is the joint distribution. Note that when \mathcal{D} is given, the distribution $p(\mathbf{M}, \mathcal{D})$ is non-normalized for \mathbf{M} ; and we have abused the KL notation for non-normalized distributions in (19), but with the same formula. For directed Bayesian networks (Zhu et al., 2011a), we naturally have $p(\mathbf{M}, \mathcal{D}) = \pi(\mathbf{M})p(\mathcal{D}|\mathbf{M})$. For the undirected MRF models, we have $\mathbf{M} = \{Z, H\}$ and again we can define the joint distribution as in (18).

Putting the above two extensions of Bayesian inference together, the regularized Bayesian inference with estimating unknown model parameters can be generally formulated as

$$\begin{aligned} \inf_{\Theta, q(\mathbf{M}), \xi} \mathcal{L}(\Theta, q(\mathbf{M})) + U(\xi) \quad \text{or} \quad \inf_{\Theta, q(\mathbf{M})} \mathcal{L}(\Theta, q(\mathbf{M})) + g(Eq(\mathbf{M})) \tag{20} \\ \text{s.t. : } q(\mathbf{M}) \in \mathcal{P}_{\text{post}}(\xi) \quad \quad \quad \text{s.t. : } q(\mathbf{M}) \in \mathcal{P}_{\text{prob}}, \end{aligned}$$

where $\mathcal{L}(\Theta, q(\mathbf{M}))$ is the objective function of problem (17). These two formulations are equivalent. We will call the former a *constrained* formulation and call the latter an *unconstrained* formulation by ignoring the standard normalization constraints, which are easy to deal with.

Appendix B. MedLDA—A RegBayes Model with Finite Latent Features

This section presents a new interpretation of MedLDA (maximum entropy discrimination latent Dirichlet allocation) (Zhu et al., 2009) under the framework of regularized Bayesian inference. MedLDA is a max-margin supervised topic model, an extension of latent Dirichlet allocation (LDA) (Blei et al., 2003) for supervised learning tasks. In MedLDA, each data example is projected to a point in a finite dimensional latent space, of which each feature corresponds to a topic, i.e., a unigram distribution over the terms in a vocabulary. MedLDA represents each data as a probability distribution over the features, which results in a conservation constraint (i.e., the more a data expresses on one feature, the less it can express others) (Griffiths and Ghahramani, 2005). The infinite latent feature models discussed in Section 4 do not have such a constraint.

Without loss of generality, we consider the MedLDA regression model as an example (classification model is similar), whose graphical structure is shown in Figure 11. We assume that all data examples have the same length V for notation simplicity. Each document is associated with a response variable Y , which is observed in the training phase but unobserved in testing. We will use y to denote an instance value of Y . Let K be the number of topics or the dimensionality of the latent topic space. MedLDA builds an LDA model to describe the observed words. The generating process of LDA is that each document n has a mixing proportion $\theta_n \sim \text{Dirichlet}(\alpha)$; each word w_{nm} is associated with a topic $z_{nm} \sim \theta_n$, which indexes the topic that generates the word, i.e., $w_{nm} \sim \beta_{z_{nm}}$. Define $\bar{Z}_n = \frac{1}{V} \sum_{m=1}^V Z_{nm}$ as the average topic assignment for document n . Let $\Theta = \{\alpha, \beta, \delta^2\}$ denote the unknown




Figure 11: Graphical structure of MedLDA.

model parameters and $\mathcal{D} = \{y_n, w_{nm}\}$ be the training set. MedLDA was defined as solving a regularized MLE problem with expectation constraints

$$\begin{aligned} \inf_{\Theta, \xi, \xi^*} \quad & -\log p(\{y_n, w_{nm}\}|\Theta) + C \sum_{n=1}^N (\xi_n + \xi_n^*) \\ \text{s.t. } \forall n : \quad & \begin{cases} y_n - \mathbb{E}_p[\boldsymbol{\eta}^\top \bar{Z}_n] \leq \epsilon + \xi_n \\ -y_n + \mathbb{E}_p[\boldsymbol{\eta}^\top \bar{Z}_n] \leq \epsilon + \xi_n^* \\ \xi_n, \xi_n^* \geq 0 \end{cases} \end{aligned}$$

The posterior constraints are imposed following the large-margin principle and they correspond to a quality measure of the prediction results on training data. In fact, it is easy to show that minimizing $U(\xi, \xi^*) = C \sum_{n=1}^N (\xi_n + \xi_n^*)$ under the above constraints is equivalent to minimizing an ϵ -insensitive loss (Smola and Schölkopf, 2003)

$$\mathcal{R}_\epsilon \left(p(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\mathcal{D}, \Theta) \right) = C \sum_{n=1}^N \max \left(0, \left| y_n - \mathbb{E}_p \left[\boldsymbol{\eta}^\top \bar{Z}_n \right] \right| - \epsilon \right).$$

of the expected linear prediction rule $\hat{y}_n = \mathbb{E}_p[\boldsymbol{\eta}^\top \bar{Z}_n]$.

To practically learn an MedLDA model, since the above problem is intractable, variational methods were used by introducing an auxiliary distribution $q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta)$ to approximate the true posterior $p(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\mathcal{D}, \Theta)$,²⁰ replacing the negative data likelihood with its upper bound $\mathcal{L}(q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta))$, and replacing p by q in the constraints. The variational MedLDA regression model is

$$\begin{aligned} \inf_{q, \Theta, \xi, \xi^*} \quad & \mathcal{L}(q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta)) + C \sum_{n=1}^N (\xi_n + \xi_n^*) \\ \text{s.t. } \forall n : \quad & \begin{cases} y_n - \mathbb{E}_q[\boldsymbol{\eta}^\top \bar{Z}_n] \leq \epsilon + \xi_n \\ -y_n + \mathbb{E}_q[\boldsymbol{\eta}^\top \bar{Z}_n] \leq \epsilon + \xi_n^* \\ \xi_n, \xi_n^* \geq 0 \end{cases} \end{aligned}$$

where $\mathcal{L}(q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta)) = -\mathbb{E}_q[\log p(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}, \mathcal{D}|\Theta)] - \mathcal{H}(q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta))$ is a variational upper-bound of the negative data log-likelihood. The upper bound is tight if no

20. We have explicitly written the condition on model parameters.

restricting constraints are made on the variational distribution q . In practice, additional assumptions (e.g., mean-field) can be made on q to derive a practical approximate algorithm.

Based on the previous discussions on the extensions of RegBayes and the duality in Lemma 14, we can reformulate the MedLDA regression model as an example of RegBayes. Specifically, for the MedLDA regression model, we have $\mathbf{M} = \{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}$. According to (19), we can easily show that

$$\begin{aligned} \mathcal{L}\left(q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta)\right) &= \text{KL}\left(q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta)\|p(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}, \{w_{nm}, y_n\}|\Theta)\right) \\ &= \mathcal{L}_B\left(\Theta, q(\mathbf{M}|\Theta)\right). \end{aligned}$$

Then, the MedLDA problem is a RegBayes model in (20) with

$$\mathcal{P}_{\text{post}}^{\text{MedLDA}}(\Theta, \boldsymbol{\xi}, \boldsymbol{\xi}^*) \stackrel{\text{def}}{=} \left\{ q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta) \left| \begin{array}{l} \forall n: y_n - \mathbb{E}_q[\boldsymbol{\eta}^\top \bar{Z}_n] \leq \epsilon + \xi_n \\ -y_n + \mathbb{E}_q[\boldsymbol{\eta}^\top \bar{Z}_n] \leq \epsilon + \xi_n^* \\ \xi_n, \xi_n^* \geq 0 \end{array} \right. \right\}. \quad (21)$$

For the MedLDA problem, we can use Lagrangian methods to solve the constrained formulation. Alternatively, we can use the convex duality theorem to solve an equivalent unconstrained form. For the variational MedLDA, the ϵ -insensitive loss is $\mathcal{R}_\epsilon(q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta))$. Its conjugate can be derived using the results of Lemma 11. Specifically, we have the following result, whose proof is deferred to Appendix C.6.

Lemma 15 (Conjugate of MedLDA) *For the variational MedLDA problem, we have*

$$\begin{aligned} &\inf_{\Theta, q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta) \in \mathcal{P}_{\text{prob}}} \mathcal{L}(q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta), \Theta) + \mathcal{R}_\epsilon(q(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta)) \\ &= \sup_{\boldsymbol{\omega}} -\log Z'(\boldsymbol{\omega}, \Theta^*) - \sum_n g_2^*(\boldsymbol{\omega}_n; -y_n + \epsilon, y_n + \epsilon), \end{aligned}$$

where $\boldsymbol{\omega}_n = (\omega_n, \omega'_n)$. Moreover, The optimum distribution is the posterior distribution

$$\hat{q}(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}|\Theta^*) = \frac{1}{Z'(\hat{\boldsymbol{\omega}}, \Theta^*|\mathcal{D})} p(\{\boldsymbol{\theta}_n, z_{nm}, \boldsymbol{\eta}\}, \mathcal{D}|\Theta^*) \exp\left\{\sum_n (\hat{\omega}_n - \hat{\omega}'_n) \boldsymbol{\eta}^\top \bar{z}_n\right\}, \quad (22)$$

where $Z'(\hat{\boldsymbol{\omega}}, \Theta|\mathcal{D})$ is the normalization factor and the optimum parameters are

$$\Theta^* = \underset{\Theta}{\operatorname{argmax}} \log p(\mathcal{D}|\Theta).$$

Note that although in general, either the primal or the dual problem is hard to solve exactly, the above conjugate results are still useful when developing approximate inference algorithms. For instance, we can impose additional mean-field assumptions on q in the primal formulation and iteratively solve for each factor; and in this process convex conjugates are useful to deal with the large-margin constraints (Zhu et al., 2009). Alternatively, we can apply approximate methods (e.g., MCMC sampling) to infer q based on its solution in (22), and iteratively solves for the dual parameters $\boldsymbol{\omega}$ using approximate statistics (Schofield, 2006). We will discuss more on this when presenting the inference algorithms for iLSVM and MT-iLSVM.

In the above discussions, we have treated the topics β as fixed unknown parameters. A fully Bayesian formulation would treat β as random variables, e.g., with a Dirichlet prior (Blei et al., 2003; Griffiths and Steyvers, 2004). Under the RegBayes interpretation, we can easily do such an extension of MedLDA, simply by moving β from Θ to \mathbf{M} .

Appendix C. Proof of the Theorems and Lemmas

This section provides the proof of the theorems and lemmas.

C.1 Proof of Theorem 6

Proof The adjoint of the linear operator E is given by $\langle Ex, \phi \rangle = \langle E^* \phi, x \rangle$. In this theorem, E is the expectation with respect to q . Thus, we have

$$\begin{aligned} \langle Eq, \phi \rangle &= \left\langle \int q(\mathbf{M}) \psi(\mathbf{M}, \mathcal{D}) d\mu(\mathbf{M}), \phi \right\rangle \\ &= \int q(\mathbf{M}) \langle \psi(\mathbf{M}, \mathcal{D}), \phi \rangle d\mu(\mathbf{M}) \\ &= (E^* \phi)(q), \end{aligned}$$

where $E^* \phi = \langle \phi, \psi(\cdot) \rangle$.

By definition, we have $\text{KL}(q(\mathbf{M}) \| p(\mathbf{M}, \mathcal{D})) = \text{KL}(q(\mathbf{M}) \| p(\mathbf{M} | \mathcal{D})) + c$, where c is a constant, $c = -\log p(\mathcal{D})$. Let $f(q(\mathbf{M}))$ denote the KL-divergence $\text{KL}(q(\mathbf{M}) \| p(\mathbf{M} | \mathcal{D}))$. The following proof is similar to the proof of the Fenchel duality theorem (Borwein and Zhu, 2005). Let t and d denote the primal value and the dual value, respectively. By Lemma 4.3.1 (Borwein and Zhu, 2005), under appropriate regularity conditions, there is a $\hat{\phi}$ such that

$$t \leq \left[f(q) - \langle \hat{\phi}, Eq \rangle \right] + \left[g(\phi) + \langle \hat{\phi}, \phi \rangle \right] + c.$$

For any μ , setting $\phi = Eq + \mu$ in the above inequality, we have

$$\begin{aligned} t &\leq f(q) + g(Eq + \mu) + \langle \hat{\phi}, \mu \rangle + c \\ &= \left\{ f(q) - \langle E^* \hat{\phi}, q \rangle \right\} + \left\{ g(Eq + \mu) - \langle -\hat{\phi}, Eq + \mu \rangle \right\} + c. \end{aligned}$$

Taking the infimum over all points μ , we have

$$t \leq \left\{ f(q) - \langle E^* \hat{\phi}, q \rangle \right\} - g^*(-\hat{\phi}) + c.$$

Then, taking the infimum over all points $q \in \mathcal{P}_{\text{prob}}$, we have

$$\begin{aligned} t &\leq \inf_{q \in \mathcal{P}_{\text{prob}}} \left\{ f(q) - \langle E^* \hat{\phi}, q \rangle \right\} - g^*(-\hat{\phi}) + c \\ &= -f^*(E^* \hat{\phi}) - g^*(-\hat{\phi}) + c \\ &\leq d, \end{aligned} \tag{23}$$

where

$$f^*(E^*\phi) = \log \int p(\mathbf{M}|\mathcal{D}) \exp(\langle \phi, \psi(\mathbf{M}, \mathcal{D}) \rangle) d\mu(\mathbf{M})$$

is the convex conjugate of the KL-divergence.

Since $d \leq t$ due to the Fenchel weak duality theorem (Borwein and Zhu, 2005) (Theorem 4.4.2), we have the strong duality that $t = d$, and $\hat{\phi}$ attains the supremum in the dual problem. During the deviation of the infimum in (23), we get the optimum solution of q :

$$\begin{aligned} \hat{q}_{\hat{\phi}}(\mathbf{M}) &\propto p(\mathbf{M}|\mathcal{D}) \exp(\langle \hat{\phi}, \psi(\mathbf{M}; \mathcal{D}) \rangle) \\ &= p(\mathbf{M}, \mathcal{D}) \exp(\langle \hat{\phi}, \psi(\mathbf{M}; \mathcal{D}) \rangle - \Lambda_{\hat{\phi}}). \end{aligned}$$

Absorbing the constant c into f^* , we get the dual objective of Theorem 6. ■

C.2 Proof of Lemma 9

Proof By definition, $g_0^*(\mu) = \sup_{x \in \mathbb{R}}(x\mu - C \max(0, x))$. We consider two cases. First, if $\mu < 0$, we have

$$g_0^*(\mu) \geq \sup_{x < 0} (x\mu - C \max(0, x)) = \sup_{x < 0} x\mu = \infty.$$

Therefore, we have $g_0^*(\mu) = \infty$ if $\mu < 0$. Second, if $\mu \geq 0$, we have

$$g_0^*(\mu) = \sup_{x \geq 0} (x\mu - Cx) = \mathbb{I}(\mu \leq C).$$

Putting the above results together, we prove the claim. ■

C.3 Proof of Lemma 10

Proof The proof has a similar structure as the proof of Lemma 9. By definition, we have

$$g_1^*(\boldsymbol{\mu}) = \sup_{\mathbf{x}} \left\{ \boldsymbol{\mu}^\top \mathbf{x} - g_1(\mathbf{x}) \right\} = \sup_{\mathbf{x}} \left\{ \sum_j \mu_j x_j - \max(x_1, \dots, x_L) \right\}.$$

We first show that $\forall i, \mu_i \geq 0$ in order to have finite g_1^* values. Suppose that $\exists j, \mu_j < 0$. Then, we define

$$\mathcal{G}_j = \{\mathbf{x} \in \mathbb{R}^L : x_j < 0\}, \text{ and } \mathcal{G}_j^o = \{\mathbf{x} \in \mathcal{G}_j : x_i = 0, \text{ if } i \neq j\}.$$

Since $\mathcal{G}_j^o \subset \mathcal{G}_j \subset \mathbb{R}^L$, we have

$$g_1^*(\boldsymbol{\mu}) \geq \sup_{\mathbf{x} \in \mathcal{G}_j} \left\{ \boldsymbol{\mu}^\top \mathbf{x} - g_1(\mathbf{x}) \right\} \geq \sup_{\mathbf{x} \in \mathcal{G}_j^o} \left\{ \boldsymbol{\mu}^\top \mathbf{x} - g_1(\mathbf{x}) \right\} = \sup_{x_j \in \mathbb{R}_-} \{x_j \mu_j - 0\} = \infty.$$

Therefore, $g_1^*(\boldsymbol{\mu}) = \infty$ if $\exists j, \mu_j < 0$.

Now, we consider the second case, where $\forall i, \mu_i \geq 0$. We can easily show that

$$\forall \mathbf{x} \in \mathbb{R}^L, \quad \boldsymbol{\mu}^\top \mathbf{x} - g_1(\mathbf{x}) \leq \sum_i \mu_i \max(\mathbf{x}) - g_1(\mathbf{x}).$$

Therefore

$$g_1^*(\boldsymbol{\mu}) \leq \sup_{\mathbf{x} \in \mathbb{R}^L} \left\{ \left(\sum_i \mu_i - C \right) \max(\mathbf{x}) \right\} = \mathbb{I} \left(\sum_i \mu_i = C \right).$$

Moreover, let $\mathcal{G}^1 = \{\mathbf{x} \in \mathbb{R}^L : \mathbf{x} = x\mathbf{e}, x \in \mathbb{R}\}$, where \mathbf{e} is a vector with every element being 1. Then, we have

$$g_1^*(\boldsymbol{\mu}) \geq \sup_{\mathbf{x} \in \mathcal{G}^1} \left\{ \boldsymbol{\mu}^\top \mathbf{x} - g_1(\mathbf{x}) \right\} = \sup_{x \in \mathbb{R}} \left\{ \left(\sum_i \mu_i - C \right) x \right\} = \mathbb{I} \left(\sum_i \mu_i = C \right).$$

Putting the above results together proves the claim. \blacksquare

C.4 Proof of Lemma 11

Proof By definition, the conjugate is

$$\begin{aligned} g_2^*(\mu) &= \sup_{x \in \mathbb{R}} \{ \mu x - C \max(0, |x - y| - \epsilon) \}. \\ &= - \inf_{x \in \mathbb{R}} \{ -\mu x + C \max(0, |x - y| - \epsilon) \}. \\ &= - \inf_{x \in \mathbb{R}; t \geq 0; t \geq |x - y| - \epsilon} \{ -\mu x + Ct \} \\ &= - \sup_{\alpha, \beta \geq 0} \left\{ \inf_{x, t \in \mathbb{R}} \{ -\mu x + Ct - \alpha(t - |x - y| + \epsilon) - \beta t \} \right\} \\ &= - \sup_{\alpha, \beta \geq 0} \left\{ \inf_{x \in \mathbb{R}} \{ -\mu x + \alpha|x - y| \} + \inf_{t \in \mathbb{R}} \{ Ct - \alpha t - \beta t \} - \alpha\epsilon \right\} \end{aligned}$$

For the second infimum, it is easy to show that

$$\inf_{t \in \mathbb{R}} \{ Ct - \alpha t - \beta t \} = -\mathbb{I}(\alpha + \beta = C).$$

For the first infimum, we can show that

$$\inf_{x \in \mathbb{R}} \{ -\mu x + \alpha|x - y| \} = -\mu y + \inf_{x' \in \mathbb{R}} \{ -\mu x' + \alpha|x'| \} = -\mu y - \mathbb{I}(|\mu| \leq \alpha).$$

Thus, we have

$$\begin{aligned} g_2^*(\mu) &= - \sup_{\alpha, \beta \geq 0} \left\{ -\mu y - \alpha\epsilon - \mathbb{I}(|\mu| \leq \alpha) - \mathbb{I}(\alpha + \beta = C) \right\} \\ &= -(-\mu y - \epsilon|\mu| - \mathbb{I}(|\mu| \leq C)) \\ &= \mu y + \epsilon|\mu| + \mathbb{I}(|\mu| \leq C), \end{aligned}$$

where the second equality holds by setting $\alpha = |\mu|$, under the condition that ϵ is positive; the condition $|\mu| \leq C$ is induced from the conditions $\alpha + \beta = C$ and $\beta \geq 0$. \blacksquare

C.5 Proof of Lemma 12

Proof By definition, we have $g(Eq) \stackrel{\text{def}}{=} \mathcal{R}_h^c(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W})) = \sum_n g_1(\ell_n^\Delta - Eq(n))$. Let $\boldsymbol{\mu}_n = Eq(n)$. We have the conjugate

$$\begin{aligned} g^*(\boldsymbol{\omega}) &= \sup_{\boldsymbol{\mu}} \left\{ \boldsymbol{\omega}^\top \boldsymbol{\mu} - \sum_n g_1(\ell_n^\Delta - \boldsymbol{\mu}_n) \right\} \\ &= \sum_n \sup_{\boldsymbol{\mu}_n} \left\{ \boldsymbol{\omega}_n^\top \boldsymbol{\mu}_n - g_1(\ell_n^\Delta - \boldsymbol{\mu}_n) \right\} \\ &= \sum_n \sup_{\boldsymbol{\nu}_n} \left\{ \boldsymbol{\omega}_n^\top (\ell_n^\Delta - \boldsymbol{\nu}_n) - g_1(\boldsymbol{\nu}_n) \right\} \\ &= \sum_n \left(\boldsymbol{\omega}_n^\top \ell_n^\Delta + g_1^*(-\boldsymbol{\omega}_n) \right). \end{aligned}$$

Thus,

$$g^*(-\boldsymbol{\omega}) = \sum_n \left(-\boldsymbol{\omega}_n^\top \ell_n^\Delta + g_1^*(\boldsymbol{\omega}_n) \right).$$

Using the results of Theorem 6 proves the claim. \blacksquare

C.6 Proof of Lemma 13

Proof Similar structure as the proof of Lemma 12. In this case, the linear expectation operator is $E : \mathcal{P}_{\text{prob}} \rightarrow \mathbb{R}^{\sum_m |\mathcal{I}_{\text{tr}}^m|}$ and the element of Eq evaluated at the n th example for task m is

$$Eq(n, m) \stackrel{\text{def}}{=} y_{mn} \mathbb{E}_{q(\mathbf{Z}, \boldsymbol{\eta})} [\mathbf{Z} \boldsymbol{\eta}_m]^\top \mathbf{x}_{mn} = \mathbb{E}_{q(\mathbf{Z}, \boldsymbol{\eta})} \left[y_{mn} (\mathbf{Z} \boldsymbol{\eta}_m)^\top \mathbf{x}_{mn} \right].$$

Then, let $g_0 : \mathbb{R} \rightarrow \mathbb{R}$ be a function defined in Lemma 9. We have

$$g(Eq) \stackrel{\text{def}}{=} \mathcal{R}_h^{MT} \left(q(\mathbf{Z}, \boldsymbol{\eta}, \mathbf{W}) \right) = \sum_{m, n \in \mathcal{I}_{\text{tr}}^m} g_0 \left(1 - Eq(n, m) \right).$$

Let $\boldsymbol{\mu} = Eq$. By definition, the conjugate is

$$\begin{aligned} g^*(\boldsymbol{\omega}) &= \sup_{\boldsymbol{\mu}} \left\{ \boldsymbol{\omega}^\top \boldsymbol{\mu} - \sum_{m, n \in \mathcal{I}_{\text{tr}}^m} g_0(1 - \mu_{mn}) \right\} \\ &= \sum_{m, n \in \mathcal{I}_{\text{tr}}^m} \sup_{\mu_{mn}} \left\{ \omega_{mn} \mu_{mn} - g_0(1 - \mu_{mn}) \right\} \\ &= \sum_{m, n \in \mathcal{I}_{\text{tr}}^m} \sup_{\nu_{mn}} \left\{ \omega_{mn} (1 - \nu_{mn}) - g_0(\nu_{mn}) \right\} \\ &= \sum_{m, n \in \mathcal{I}_{\text{tr}}^m} \left(\omega_{mn} + g_0^*(-\omega_{mn}) \right). \end{aligned}$$

Thus,

$$g^*(-\boldsymbol{\omega}) = \sum_{m, n \in \mathcal{I}_{\text{tr}}^m} \left(-\omega_{mn} + g_0^*(\omega_{mn}) \right).$$

By the results in Theorem 6 and Lemma 9, we can derive the conjugate of the problem (15). ■

C.7 Proof of Lemma 15

Proof Similar structure as the proof of Lemma 12. In this case, the linear expectation operator is $E : \mathcal{P}_{\text{prob}} \rightarrow \mathbb{R}^N$ and the elements of Eq evaluated at the n th example is

$$\mu_n = \mathbb{E}_{q(\{\theta_n, z_{nm}, \boldsymbol{\eta}\}|\Theta)} \left[\boldsymbol{\eta}^\top \bar{z}_n \right].$$

Then, using the g_2 function defined in Lemma 11, we have

$$g(Eq) \stackrel{\text{def}}{=} \mathcal{R}_\epsilon(q(\{\theta_n, z_{nm}, \boldsymbol{\eta}\}|\Theta)) = \sum_n g_2(\mu_n; y_n, \epsilon).$$

Therefore $g^*(\boldsymbol{\omega}) = \sum_n g_2^*(\omega_n; y_n, \epsilon)$ and $g^*(-\boldsymbol{\omega}) = \sum_n g_2^*(-\omega_n; y_n, \epsilon)$. By the results in Theorem 6 and Lemma 9, we can derive the conjugate and the optimum solution of \hat{q} . The optimum solution of Θ is due to Lemma 14. Note that the constraints are not directly dependent on Θ . ■

Appendix D. Inference Algorithms for Infinite Latent SVMs

D.1 Inference for MT-iLSVM

In this section, we provide the derivation of the inference algorithm for MT-iLSVM, which is outlined in Algorithm 2 and detailed below.

For MT-iLSVM, the model \mathbf{M} consists of all the latent variables $(\boldsymbol{\nu}, \mathbf{W}, \mathbf{Z}, \boldsymbol{\eta})$. Let $L_{mn}(q) \stackrel{\text{def}}{=} \mathbb{E}_q[\log p(\mathbf{x}_{mn}|\mathbf{Z}, \mathbf{w}_{mn}, \lambda_{mn}^2)]$ be the expected data likelihood. Then, under the truncated mean-field assumption (16), we have

$$L_{mn}(q) = -\frac{\mathbf{x}_{mn}^\top \mathbf{x}_{mn} - 2\mathbf{x}_{mn}^\top \mathbb{E}_q[\mathbf{Z}\mathbf{w}_{mn}] + \mathbb{E}_q[\mathbf{w}_{mn}^\top \mathbf{U}\mathbf{w}_{mn}]}{2\lambda_{mn}^2} - \frac{D \log(2\pi\lambda_{mn}^2)}{2},$$

where $\mathbf{x}_{mn}^\top \mathbb{E}_q[\mathbf{Z}\mathbf{w}_{mn}] = \sum_k \mathbf{x}_{mn}^\top \boldsymbol{\psi}_{.k}$; $\boldsymbol{\psi}_{.k} \stackrel{\text{def}}{=} (\psi_{1k} \cdots \psi_{Dk})^\top$ is the k th column of $\boldsymbol{\psi} = \mathbb{E}_q[\mathbf{Z}]$;

$$\mathbb{E}_q[\mathbf{w}_{mn}^\top \mathbf{U}\mathbf{w}_{mn}] = 2 \sum_{j < k} \phi_{mn}^j \phi_{mn}^k \mathbf{U}_{jk} + \sum_k \mathbf{U}_{kk} \left(K\sigma_{mn}^2 + \Phi_{mn}^\top \Phi_{mn} \right);$$

and $\mathbf{U} \stackrel{\text{def}}{=} \mathbb{E}_q[\mathbf{Z}^\top \mathbf{Z}]$ is a $K \times K$ matrix, whose element is

$$\mathbf{U}_{ij} = \begin{cases} \sum_d \psi_{di}, & \text{if } i = j \\ \sum_d \psi_{di} \psi_{dj}, & \text{otherwise.} \end{cases}$$

For the KL-divergence term, we have $\text{KL}(q(\mathbf{M})\|\pi(\mathbf{M})) = \text{KL}(q(\boldsymbol{\nu})\|\pi(\boldsymbol{\nu})) + \text{KL}(q(\mathbf{W})\|\pi(\mathbf{W})) + \mathbb{E}_{q(\boldsymbol{\nu})}[\text{KL}(q(\mathbf{Z})\|\pi(\mathbf{Z}|\boldsymbol{\nu}))] + \text{KL}(q(\boldsymbol{\eta})\|\pi(\boldsymbol{\eta}))$, where the individual terms are

$$\begin{aligned} \text{KL}(q(\boldsymbol{\nu})\|\pi(\boldsymbol{\nu})) &= \sum_{k=1}^K \left((\gamma_{k1} - \alpha)(\varphi(\gamma_{k1}) - \varphi(\gamma_{k1} + \gamma_{k2})) - \log \frac{\Gamma(\gamma_{k1})\Gamma(\gamma_{k2})}{\Gamma(\gamma_{k1} + \gamma_{k2})} \right. \\ &\quad \left. + (\gamma_{k2} - 1)(\varphi(\gamma_{k2}) - \varphi(\gamma_{k1} + \gamma_{k2})) \right) - K \log \alpha, \\ \mathbb{E}_{q(\boldsymbol{\nu})}[\text{KL}(q(\mathbf{Z})\|\pi(\mathbf{Z}|\boldsymbol{\nu}))] &= \sum_{dk} \left(-\psi_{dk} \sum_{j=1}^k \mathbb{E}_q[\log \nu_j] - (1 - \psi_{dk}) \mathbb{E}_q[\log(1 - \prod_{j=1}^k \nu_j)] \right. \\ &\quad \left. + \psi_{dk} \log \psi_{dk} + (1 - \psi_{dk}) \log(1 - \psi_{dk}) \right) \\ \text{KL}(q(\mathbf{W})\|\pi(\mathbf{W})) &= \sum_{mn} \left(\frac{K\sigma_{mn}^2 + \Phi_{mn}^\top \Phi_{mn}}{2\sigma_{m0}^2} - \frac{K(1 + \log \frac{\sigma_{mn}^2}{\sigma_{m0}^2})}{2} \right). \end{aligned}$$

where $\varphi(\cdot)$ is the digamma function and $\mathbb{E}_q[\log \nu_j] = \varphi(\gamma_{j1}) - \varphi(\gamma_{j1} + \gamma_{j2})$. For $\text{KL}(q(\boldsymbol{\eta})\|\pi(\boldsymbol{\eta}))$, we do not need to write it explicitly, as we shall see. Finally, the effective discriminant function is

$$f_m(\mathbf{x}_{mn}; q(\mathbf{Z}, \boldsymbol{\eta})) = \mathbb{E}_q[\boldsymbol{\eta}_m]^\top \boldsymbol{\psi}^\top \mathbf{x}_{mn} = \sum_{k=1}^K \mathbb{E}_q[\eta_{mk}] \boldsymbol{\psi}_{\cdot k}^\top \mathbf{x}_{mn}.$$

All the above terms can be easily computed, except the term $\mathbb{E}_q[\log(1 - \prod_{j=1}^k \nu_j)]$. Here, we adopt the multivariate lower bound (Doshi-Velez, 2009)

$$\begin{aligned} \mathbb{E}_q \left[\log \left(1 - \prod_{j=1}^k \nu_j \right) \right] &\geq \sum_{m=1}^k q_{km} \varphi(\gamma_{m2}) + \sum_{m=1}^{k-1} \left(\sum_{n=m+1}^k q_{kn} \right) \varphi(\gamma_{m1}) \\ &\quad - \sum_{m=1}^k \left(\sum_{n=m}^k q_{kn} \right) \varphi(\gamma_{m1} + \gamma_{m2}) + \mathcal{H}(q_{k\cdot}), \end{aligned}$$

where the variational parameters $q_{k\cdot} = (q_{k1} \cdots q_{kk})^\top$ belong to the k -simplex, and $\mathcal{H}(q_{k\cdot})$ is the entropy of $q_{k\cdot}$. The tightest lower bound is achieved by setting $q_{k\cdot}$ to be the optimum value

$$q_{km} = \frac{1}{Z_k} \exp \left(\varphi(\gamma_{m2}) + \sum_{n=1}^{m-1} \varphi(\gamma_{n1}) - \sum_{n=1}^m \varphi(\gamma_{n1} + \gamma_{n2}) \right), \quad (24)$$

where Z_k is a normalization factor to make $q_{k\cdot}$ be a distribution. We denote the tightest lower bound by \mathcal{L}_k^ν . Replacing the term $\mathbb{E}_q[\log(1 - \prod_{j=1}^k \nu_j)]$ with its lower bound \mathcal{L}_k^ν , we can have an upper bound of $\text{KL}(q(\mathbf{M})\|\pi(\mathbf{M}))$ and we denote this upper bound by $\mathcal{L}(q)$.

With the above terms and the upper bound $\mathcal{L}(q)$, we can implement the general procedure outlined in Algorithm 1 to solve the MT-iLSVM problem. Specifically, the inference procedure iteratively solves the following steps, as summarized in Algorithm 2:

Algorithm 2 Inference Algorithm of MT-iLSVM

- 1: **Input:** data $\mathcal{D} = \{(\mathbf{x}_{mn}, y_{mn})\}_{m,n \in \mathcal{I}_{\text{tr}}^m} \cup \{\mathbf{x}_{mn}\}_{m,n \in \mathcal{I}_{\text{tst}}^m}$, constants α and C
 - 2: **Output:** distributions $q(\boldsymbol{\nu})$, $q(\mathbf{Z})$, $q(\mathbf{W})$, $q(\boldsymbol{\eta})$ and hyper-parameters σ_{m0}^2 and λ_{mn}^2
 - 3: Initialize $\gamma_{k1} = \alpha$, $\gamma_{k2} = 1$, $\psi_{dk} = 0.5 + \epsilon$, where $\epsilon \sim \mathcal{N}(0, 0.001)$, $\Phi_{mn} = 0$, $\sigma_{mn}^2 = \sigma_{m0}^2 = 1$, $\boldsymbol{\mu}_m = 0$, λ_{mn}^2 is computed from \mathcal{D} .
 - 4: **repeat**
 - 5: **repeat**
 - 6: update $(\gamma_{k1}, \gamma_{k2})$ using (26), $\forall 1 \leq k \leq K$;
 - 7: update ϕ_{mn}^k and σ_{mn}^2 using (25), $\forall m, \forall n, \forall 1 \leq k \leq K$;
 - 8: update ψ_{dk} using (27), $\forall 1 \leq d \leq D, \forall 1 \leq k \leq K$;
 - 9: **until** relative change of L is less than τ (e.g., $1e^{-3}$) or iteration number is T (e.g., 10)
 - 10: **for** $m = 1$ **to** M **do**
 - 11: solve the dual problem (28) using a binary SVM learner.
 - 12: **end for**
 - 13: update the hyper-parameters σ_{m0}^2 using (29) and λ_{mn}^2 using (30). (*Optional*)
 - 14: **until** relative change of L is less than τ' (e.g., $1e^{-4}$) or iteration number is T' (e.g., 20)
-

Infer $q(\boldsymbol{\nu})$, $q(\mathbf{Z})$ and $q(\mathbf{W})$: For $q(\mathbf{W})$, since both the prior $\pi(\mathbf{W})$ and $q(\mathbf{W})$ are Gaussian, we can easily derive the update rules, similar as in Gaussian mixture models

$$\begin{aligned} \phi_{mn}^k &= \frac{\sum_d x_{mn}^d \psi_{dk} - \sum_{j \neq k} \phi_{mn}^j \mathbf{U}_{kj}}{\lambda_{mn}^2} \left(\frac{1}{\sigma_{m0}^2} + \frac{\sum_d \psi_{dk}}{\lambda_{mn}^2} \right)^{-1} \\ \sigma_{mn}^2 &= \left(\frac{1}{\sigma_{m0}^2} + \frac{1}{K} \sum_k \frac{\mathbf{U}_{kk}}{\lambda_{mn}^2} \right)^{-1} \end{aligned} \quad (25)$$

For $q(\boldsymbol{\nu})$, we have the update rules similar as in (Doshi-Velez, 2009), that is,

$$\begin{aligned} \gamma_{k1} &= \alpha + \sum_{m=k}^K \sum_{d=1}^D \psi_{dm} + \sum_{m=k+1}^K \left(D - \sum_{d=1}^D \psi_{dm} \right) \left(\sum_{i=k+1}^m q_{mi} \right) \\ \gamma_{k2} &= 1 + \sum_{m=k}^K \left(D - \sum_{d=1}^D \psi_{dm} \right) q_{mk}. \end{aligned} \quad (26)$$

For $q(\mathbf{Z})$, we have the mean-field update equation as

$$\psi_{dk} = \frac{1}{1 + e^{-\vartheta_{dk}}}, \quad (27)$$

where

$$\begin{aligned} \vartheta_{dk} &= \sum_{j=1}^k \mathbb{E}_q[\log v_j] - \mathcal{L}_k^\nu - \sum_{mn} \frac{1}{2\lambda_{mn}^2} \left((K\sigma_{mn}^2 + (\phi_{mn}^k)^2) \right. \\ &\quad \left. - 2x_{mn}^d \phi_{mn}^k + 2 \sum_{j \neq k} \phi_{mn}^j \phi_{mn}^k \psi_{dj} \right) + \sum_{m,n \in \mathcal{I}_{\text{tr}}^m} y_{mn} \mathbb{E}_q[\eta_{mk}] x_{mn}^d. \end{aligned}$$

Infer $q(\boldsymbol{\eta})$ and solve for $\boldsymbol{\omega}$: By the convex duality theory, we have the solution

$$\begin{aligned} q(\boldsymbol{\eta}) &\propto \pi(\boldsymbol{\eta}) \exp \left\{ \sum_{m,n \in \mathcal{I}_{\text{tr}}^m} y_{mn} \omega_{mn} \boldsymbol{\eta}_m^\top \boldsymbol{\psi}^\top \mathbf{x}_{mn} \right\} \\ &= \prod_{m=1}^M \pi(\boldsymbol{\eta}_m) \exp \left\{ \boldsymbol{\eta}_m^\top \left(\sum_{n \in \mathcal{I}_{\text{tr}}^m} y_{mn} \omega_{mn} \boldsymbol{\psi}^\top \mathbf{x}_{mn} \right) \right\}. \end{aligned}$$

Therefore, we can see that although we did not assume $q(\boldsymbol{\eta})$ is factorized, we can get the induced factorization form $q(\boldsymbol{\eta}) = \prod_m q(\boldsymbol{\eta}_m)$, where

$$q(\boldsymbol{\eta}_m) \propto \pi(\boldsymbol{\eta}_m) \exp \left\{ \boldsymbol{\eta}_m^\top \left(\sum_{n \in \mathcal{I}_{\text{tr}}^m} y_{mn} \omega_{mn} \boldsymbol{\psi}^\top \mathbf{x}_{mn} \right) \right\}.$$

Here, we assume $\pi(\boldsymbol{\eta}_m)$ is standard normal. Then, we have $q(\boldsymbol{\eta}_m) = \mathcal{N}(\boldsymbol{\eta}_m | \boldsymbol{\mu}_m, I)$, where

$$\boldsymbol{\mu}_m = \sum_{n \in \mathcal{I}_{\text{tr}}^m} y_{mn} \omega_{mn} \boldsymbol{\psi}^\top \mathbf{x}_{mn}.$$

The optimum dual parameters can be obtained by solving the following M independent dual problems

$$\sup_{\boldsymbol{\omega}_m} -\frac{1}{2} \boldsymbol{\mu}_m^\top \boldsymbol{\mu}_m + \sum_{n \in \mathcal{I}_{\text{tr}}^m} \omega_{mn} \quad \text{s.t.} : 0 \leq \omega_{mn} \leq C, \forall n \in \mathcal{I}_{\text{tr}}^m, \quad (28)$$

which (and its primal form) can be efficiently solved with a binary SVM solver, such as SVM-light.

As we have stated, the hyperparameters σ_0^2 and λ_{mn}^2 can be set a priori or estimated from the data. The empirical estimation can be easily done with closed form solutions by optimizing the RegBayes objective with all the variational terms fixed. For MT-iLSVM, we have

$$\sigma_{m0}^2 = \frac{\sum_{n=1}^{N_m} (K \sigma_{mn}^2 + \Phi_{mn}^\top \Phi_{mn})}{K N_m} \quad (29)$$

$$\lambda_{mn}^2 = \frac{\mathbf{x}_{mn}^\top \mathbf{x}_{mn} - 2 \mathbf{x}_{mn}^\top \mathbb{E}_q[\mathbf{Z} \mathbf{w}_{mn}] + \mathbb{E}_q[\mathbf{w}_{mn}^\top \mathbf{U} \mathbf{w}_{mn}]}{D}. \quad (30)$$

D.2 Inference for Infinite Latent SVM

In this section, we develop the inference algorithm for iLSVM based on the stick-breaking construction of the IBP prior. The algorithm is outlined in Algorithm 3.

Similar as in the inference for MT-iLSVM, we make the additional constraint about the feasible distribution

$$q(\boldsymbol{\nu}, \mathbf{W}, \mathbf{Z}, \boldsymbol{\eta}) = q(\boldsymbol{\eta}) q(\mathbf{W} | \Phi, \Sigma) \prod_n \left(\prod_{k=1}^K q(z_{nk} | \psi_{nk}) \right) \prod_{k=1}^K q(\nu_k | \gamma_k),$$

where $q(\mathbf{W}|\Phi, \Sigma) = \prod_k \mathcal{N}(\mathbf{W}_{.k}|\Phi_{.k}, \sigma_k^2 I)$; $q(z_{nk}|\phi_{nk}) = \text{Bernoulli}(\phi_{nk})$; and $q(\nu_k|\gamma_k) = \text{Beta}(\gamma_{k1}, \gamma_{k2})$; and K is the truncation level. Then, we solve the unconstrained problem using convex duality with dual parameters being ω . Let $L_n(q) \stackrel{\text{def}}{=} \mathbb{E}_q[\log p(\mathbf{x}_n|\mathbf{z}_n, \mathbf{W})]$. We have

$$L_n(q) = -\frac{\mathbf{x}_n^\top \mathbf{x}_n - 2\mathbf{x}_n^\top \Phi \mathbb{E}_q[\mathbf{z}_n]^\top + \mathbb{E}_q[\mathbf{z}_n \mathbf{A} \mathbf{z}_n^\top]}{2\sigma_{n0}^2} - \frac{D \log(2\pi\sigma_{n0}^2)}{2},$$

where $\mathbf{A} \stackrel{\text{def}}{=} \mathbb{E}_q[\mathbf{W}^\top \mathbf{W}]$ is a $K \times K$ matrix; $\mathbf{x}_n^\top \Phi \mathbb{E}_q[\mathbf{z}_n]^\top = 2 \sum_k \psi_{nk}(\mathbf{x}_n^\top \Phi_{.k})$; and

$$\mathbb{E}_q[\mathbf{z}_n \mathbf{A} \mathbf{z}_n^\top] = 2 \sum_{j < k} \psi_{nj} \psi_{nk} \mathbf{A}_{jk} + \sum_k \psi_{nk} (D\sigma_k^2 + \mathbf{A}_{kk}).$$

The effective discriminant function is $f(y, \mathbf{x}_n) = \sum_k \mathbb{E}_q[\eta_y^k] \psi_{nk}$. Again, for computational tractability, we need the lower bound \mathcal{L}_k^ν of the term $\mathbb{E}_q[\log(1 - \prod_{j=1}^k v_j)]$. Using this lower bound, we can get an upper bound of the KL-divergence term. Then, the inference procedure iteratively solves the following steps:

Infer $q(\nu)$, $q(\mathbf{Z})$ and $q(\mathbf{W})$: For $q(\mathbf{W})$, we have the update rules

$$\begin{aligned} \Phi_{.k} &= \sum_n \frac{\psi_{nk}}{\sigma_{n0}^2} \left(\mathbf{x}_n - \sum_{j \neq k} \psi_{nj} \Phi_{.j} \right) \left(1 + \sum_n \frac{\psi_{nk}}{\sigma_{n0}^2} \right)^{-1} \\ \sigma_k^2 &= \left(1 + \sum_n \frac{\psi_{nk}}{\sigma_{n0}^2} \right)^{-1}. \end{aligned} \quad (31)$$

For $q(\nu)$, we have the update rules similar as in (Doshi-Velez, 2009), that is,

$$\begin{aligned} \gamma_{k1} &= \alpha + \sum_{m=k}^K \sum_{n=1}^N \psi_{nm} + \sum_{m=k+1}^K \left(N - \sum_{n=1}^N \psi_{nm} \right) \left(\sum_{i=k+1}^m q_{mi} \right) \\ \gamma_{k2} &= 1 + \sum_{m=k}^K \left(N - \sum_{n=1}^N \psi_{nm} \right) q_{mk}, \end{aligned} \quad (32)$$

where $q_{.k}$ is computed in the same way as in (24). For $q(\mathbf{Z})$, the mean-field update equation for ψ is

$$\psi_{nk} = \frac{1}{1 + e^{-\vartheta_{nk}}}, \quad (33)$$

where

$$\begin{aligned} \vartheta_{nk} &= \sum_{j=1}^k \mathbb{E}_q[\log v_j] - \mathcal{L}_k^\nu(q) - \frac{1}{2\sigma_{n0}^2} \left(D\sigma_k^2 + \Phi_{.k}^\top \Phi_{.k} \right) \\ &\quad + \frac{1}{\sigma_{n0}^2} \Phi_{.k}^\top \left(\mathbf{x}_n - \sum_{j \neq k} \psi_{nj} \Phi_{.j} \right) + \sum_y \omega_n^y \mathbb{E}_q[\eta_{y_n}^k - \eta_{y_n}^k]. \end{aligned}$$

Algorithm 3 Inference Algorithm of iLSVM

- 1: **Input:** data $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n \in \mathcal{I}_{\text{tr}}} \cup \{\mathbf{x}_n\}_{n \in \mathcal{I}_{\text{tst}}}$, constants α and C
 - 2: **Output:** distributions $q(\boldsymbol{\nu})$, $q(\mathbf{Z})$, $q(\mathbf{W})$, $q(\boldsymbol{\eta})$ and hyper-parameters σ_0^2 and σ_{n0}^2
 - 3: Initialize $\gamma_{k1} = \alpha$, $\gamma_{k2} = 1$, $\psi_{nk} = 0.5 + \epsilon$, where $\epsilon \sim \mathcal{N}(0, 0.001)$, $\Phi_{.k} = 0$, $\sigma_k^2 = \sigma_0^2 = 1$, $\boldsymbol{\mu} = 0$, σ_{n0}^2 is computed from \mathcal{D} .
 - 4: **repeat**
 - 5: **repeat**
 - 6: update $(\gamma_{k1}, \gamma_{k2})$ using (32), $\forall 1 \leq k \leq K$;
 - 7: update $\Phi_{.k}$ and σ_k^2 using (31), $\forall 1 \leq k \leq K$;
 - 8: update ψ_{nk} using (33), $\forall n \in \mathcal{I}_{\text{tr}}, \forall 1 \leq k \leq K$;
 - 9: update ψ_{nk} using (33), but ϑ_{nk} doesn't have the last term, $\forall n \in \mathcal{I}_{\text{tst}}, \forall 1 \leq k \leq K$;
 - 10: **until** relative change of L is less than τ (e.g., $1e^{-3}$) or iteration number is T (e.g., 10)
 - 11: solve the dual problem (34) (or its primal form) using a multi-class SVM learner.
 - 12: update the hyper-parameters σ_0^2 using (35) and σ_{n0}^2 using (36). (*Optional*)
 - 13: **until** relative change of L is less than τ' (e.g., $1e^{-4}$) or iteration number is T' (e.g., 20)
-

For testing data, ϑ_{nk} does not have the last term because of the absence of large-margin constraints.

Infer $q(\boldsymbol{\eta})$ and solve for $\boldsymbol{\omega}$: By the convex duality theory, we have

$$q(\boldsymbol{\eta}) \propto \pi(\boldsymbol{\eta}) \exp \left\{ \boldsymbol{\eta}^\top \left(\sum_{n \in \mathcal{I}_{\text{tr}}} \sum_y \omega_n^y \mathbb{E}_q[\mathbf{g}(y_n, \mathbf{x}_n, \mathbf{z}_n) - \mathbf{g}(y, \mathbf{x}_n, \mathbf{z}_n)] \right) \right\}.$$

For the standard normal prior $\pi(\boldsymbol{\eta})$, we have that $q(\boldsymbol{\eta})$ is also normal, with mean

$$\boldsymbol{\mu} = \sum_{n \in \mathcal{I}_{\text{tr}}} \sum_y \omega_n^y \mathbb{E}_q[\mathbf{g}(y_n, \mathbf{x}_n, \mathbf{z}_n) - \mathbf{g}(y, \mathbf{x}_n, \mathbf{z}_n)]$$

and identity covariance matrix. The dual problem is

$$\sup_{\boldsymbol{\omega}} -\frac{1}{2} \boldsymbol{\mu}^\top \boldsymbol{\mu} + \sum_{n \in \mathcal{I}_{\text{tr}}} \sum_y \omega_n^y \quad \text{s.t.} : \omega_n^y \geq 0, \sum_y \omega_n^y = C, \forall n \in \mathcal{I}_{\text{tr}}, \quad (34)$$

which (and its primal form) can be efficiently solved with a multi-class SVM solver.

Similar as in MT-iLSVM, the hyperparameters σ_0^2 and σ_{n0}^2 can be set a priori or estimated from the data. The empirical estimation can be easily done with closed form solutions. For iLSVM, we have

$$\sigma_0^2 = \frac{\sum_{k=1}^K (D\sigma_k^2 + \Phi_{.k}^\top \Phi_{.k})}{KD} \quad (35)$$

$$\sigma_{n0}^2 = \frac{\mathbf{x}_n^\top \mathbf{x}_n - 2\mathbf{x}_n^\top \Phi \mathbb{E}_p[\mathbf{z}_n]^\top + \mathbb{E}_q[\mathbf{z}_n \mathbf{A} \mathbf{z}_n^\top]}{D}. \quad (36)$$

References

- Yasemin Altun and Alex Smola. Unifying divergence minimization and statistical inference via convex duality. In *International Conference on Learning Theory*, pages 139–153, 2006.
- Rie Kubota Ando and Tong Zhang. A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of Machine Learning Research*, (6):1817–1853, 2005.
- Charles E. Antoniak. Mixture of Dirichlet process with applications to Bayesian nonparametric problems. *Annals of Statistics*, (273):1152–1174, 1974.
- Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Multi-task feature learning. In *Advances in Neural Information Processing Systems*, pages 41–48, 2007.
- Bart Bakker and Tom Heskes. Task clustering and gating for Bayesian multitask learning. *Journal of Machine Learning Research*, (4):83–99, 2003.
- Andrew Barron, Mark Schervish, and Larry Wasserman. The consistency of posterior distributions in nonparametric problems. *The Annals of Statistics*, 27(2):536–561, 1999.
- Matthew J. Beal, Zoubin Ghahramani, and Carl E. Rasmussen. The infinite hidden Markov model. In *Advances in Neural Information Processing Systems*, pages 577–584, 2002.
- Kedar Bellare, Gregory Druck, and Andrew McCallum. Alternating projections for learning with expectation constraints. In *Uncertainty in Artificial Intelligence*, 2009.
- David Blei and Peter Frazier. Distance dependent Chinese restaurant process. In *International Conference on Machine Learning*, pages 87–94, 2010.
- David Blei, Andrew Ng, and Michael I. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, (3):993–1022, 2003.
- Edwin Bonilla, Kian Ming Chai, and Christopher Williams. Multi-task Gaussian process prediction. In *Advances in Neural Information Processing Systems*, pages 153–160, 2008.
- Jonathan Borwein and Qiji Zhu. *Techniques of Variational Analysis: An Introduction*. Springer, New York, NY, 2005.
- Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- Ming-Wei Chang, Lev Ratinov, and Dan Roth. Guiding semi-supervision with constraint-driven learning. In *the 45th Annual Meeting of the Association of Computational Linguistics*, pages 280–287, 2007.
- Ning Chen, Jun Zhu, Fuchun Sun, and Eric P. Xing. Large-margin predictive latent subspace learning for multiview data analysis. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, 34(12):2365–2378, 2012.

- Ning Chen, Jun Zhu, Fuchun Sun, and Bo Zhang. Learning harmonium models with infinite latent features. *IEEE Transactions on Neural Networks and Learning Systems*, 25(3): 520–532, 2014.
- Taeryon Choi and R. V. Ramamoorthi. Remarks on consistency of posterior distributions. *IMS Collections 3. Pushing the limits of contemporary statistics: Contributions in honor of Jayanta K. Ghosh*. eds. S. Ghosal and B. Clarke, pages 170–186, 2008.
- Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B*, 39(1):1–38, 1977.
- Finale Doshi-Velez. The Indian buffet process: Scalable inference and extensions. Master’s thesis, The University of Cambridge, Aug 2009.
- Miroslav Dudík, Steven J. Phillips, and Robert E. Schapire. Maximum entropy density estimation with generalized regularization and an application to species distribution modeling. *Journal of Machine Learning Research*, (8):1217–1260, 2007.
- David Dunson and Shyamal Peddada. Bayesian nonparametric inferences on stochastic ordering. *ISDS Discussion Paper*, 2, 2007.
- Thomas Ferguson. A Bayesian analysis of some nonparametric problems. *Annals of Statistics*, 1(2):209–230, 1973.
- Morten Frydenberg. The chain graph Markov property. *Scandinavian Journal of Statistics*, 17:333–353, 1990.
- Kuzman Ganchev, João Graca, Jennifer Gillenwater, and Ben Taskar. Posterior regularization for structured latent variable models. *Journal of Machine Learning Research*, (11):2001–2094, 2010.
- Paul Garthwaite, Joseph Kadane, and Anthony O’Hagan. Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association*, 100(470):680–700, 2005.
- Jayanta K. Ghosh and R.V. Ramamoorthi. *Bayesian Nonparametrics*. Springer, New York, NY, 2003.
- João Graca, Kuzman Ganchev, and Ben Taskar. Expectation maximization and posterior constraints. In *Advances in Neural Information Processing Systems*, pages 569–576, 2007.
- Thomas Griffiths and Zoubin Ghahramani. Infinite latent feature models and the Indian buffet process. Technical report, University College London, GCNU TR2005-001, 2005.
- Thomas L. Griffiths and Mark Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101:5228–5235, 2004.
- Peter D. Hoff. Bayesian methods for partial stochastic orderings. *Biometrika*, 90:303–317, 2003.

- Seungil Huh and Stephen Fienberg. Discriminative topic modeling based on manifold learning. *ACM Transactions on Knowledge Discovery from Data*, 5(4), 2012.
- Kazufumi Ito and Karl Kunisch. *Lagrange Multiplier Approach to Variational Problems and Applications*. Advances in Design and Control, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2008.
- Tommi Jaakkola, Meila Meila, and Tony Jebara. Maximum entropy discrimination. In *Advances in Neural Information Processing Systems*, pages 470–476, 1999.
- Tony Jebara. *Discriminative, Generative and Imitative Learning*. PhD thesis, Media Laboratory, MIT, Dec 2001.
- Tony Jebara. Multitask sparsity via maximum entropy discrimination. *Journal of Machine Learning Research*, (12):75–110, 2011.
- Qixia Jiang, Jun Zhu, Maosong Sun, and Eric P. Xing. Monte Carlo methods for maximum margin supervised topic models. In *Advances in Neural Information Processing Systems*, pages 1592–1600, 2012.
- Thorsten Joachims. Transductive inference for text classification using support vector machines. In *International Conference on Machine Learning*, 1999.
- Michael I. Jordan, Zoubin Ghahramani, Tommi Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37:183–233, 1999.
- Mohammad E. Khan, Guillaume Bouchard, Benjamin Marlin, and Kevin Murphy. Variational bounds for mixed-data factor analysis. In *Advances in Neural Information Processing Systems*, pages 1108–1116, 2010.
- John Lafferty, Andrew McCallum, and Fernando Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *International Conference on Machine Learning*, pages 282–289, 2001.
- Percy Liang, Michael I. Jordan, and Dan Klein. Learning from measurements in exponential families. In *International Conference on Machine Learning*, pages 641–648, 2009.
- Steven N. MacEachern. Dependent nonparametric process. In *the Section on Bayesian Statistical Science of ASA*, 1999.
- Thomas L. Magnanti. Fenchel and Lagrange duality are equivalent. *Mathematical Programming*, (7):253–258, 1974.
- Gideon Mann and Andrew McCallum. Generalized expectation criteria for semi-supervised learning with weakly labeled data. *Journal of Machine Learning Research*, (11):955–984, 2010.
- Shike Mei, Jun Zhu, and Xiaojin Zhu. Robust RegBayes: Selectively incorporating first-order logic domain knowledge into bayesian models. In *International Conference on Machine Learning*, pages 253–261, 2014.

- Kurt Miller, Thomas Griffiths, and Michael I. Jordan. Nonparametric latent feature models for link prediction. In *Advances in Neural Information Processing Systems*, pages 1276–1284, 2009.
- Iain Murray and Zoubin Ghahramani. Bayesian learning in undirected graphical models: Approximate MCMC algorithms. In *Uncertainty in Artificial Intelligence*, 2004.
- Peter Orbanz. Nonparametric priors on complete separable metric spaces. *Working paper*, 2012.
- Yuan (Alan) Qi, Martin Szummer, and Thomas P. Minka. Bayesian conditional random fields. In *International Conference on Artificial Intelligence and Statistics*, pages 269–276, 2005.
- Piyush Rai and Hal Daume III. Infinite predictor subspace models for multitask learning. In *International Conference on Artificial Intelligence and Statistics*, pages 613–620, 2010.
- Charles E. Rasmussen and Zoubin Ghahramani. Infinite mixtures of Gaussian process experts. In *Advances in Neural Information Processing Systems*, pages 881–888, 2002.
- Christian P. Robert. Simulation of truncated normal variables. *Statistics and Computing*, 5(2):121–125, 1995.
- Edward Schofield. *Fitting maximum-entropy models on large sample spaces*. PhD thesis, PhD thesis, Department of Computing, Imperial College London, 2006.
- Wolfgang Stummer and Igor Vajda. On bregman distances and divergences of probability measures. *IEEE Trans. on Information Theory*, 58(3):1277–1288, 2012.
- Ben Taskar, Carlos Guestrin, and Daphne Koller. Max-margin Markov networks. In *Advances in Neural Information Processing Systems*, 2003.
- Yee Whye Teh, Dilan Görür, and Zoubin Ghahramani. Stick-breaking construction of the Indian buffet process. In *International Conference on Artificial Intelligence and Statistics*, pages 556–563, 2007.
- Yee Whye Teh, Michael I. Jordan, Matthew J. Beal, and David Blei. Hierarchical Dirichlet process. *Journal of the American Statistical Association*, 101(476):1566–1581, 2006.
- Romain Thibaux and Michael I. Jordan. Hierarchical Beta processes and the Indian buffet process. In *International Conference on Artificial Intelligence and Statistics*, pages 564–571, 2007.
- Vladimir Vapnik. *The Nature of Statistical Learning Theory*. Springer, New York, 1995.
- Martin Wainright and Michael I. Jordan. Graphical models, exponential family, and variational methods. *Foundations and Trends in Machine Learning*, 1(1):1–305, 2008.
- Max Welling and Sridevi Parise. Bayesian random fields: The Bethe-Laplace approximation. In *Uncertainty in Artificial Intelligence*, 2006.

- Max Welling, Ian Porteous, and Kenichi Kurihara. Exchangeable inconsistent priors for Bayesian posterior inference. In *Workshop on Information Theory and Applications*, 2012.
- Max Welling, M. Rosen-Zvi, and G. Hinton. Exponential family harmoniums with an application to information retrieval. In *Advances in Neural Information Processing Systems*, pages 1481–1488, 2004.
- Peter M. Williams. Bayesian conditionalisation and the principle of minimum information. *The British Journal for the Philosophy of Science*, 31(2), 1980.
- Sinead Williamson, Peter Orbanz, and Zoubin Ghahramani. Dependent Indian buffet processes. In *International Conference on Artificial Intelligence and Statistics*, pages 924–931, 2010.
- Minjie Xu, Jun Zhu, and Bo Zhang. Bayesian nonparametric max-margin matrix factorization for collaborative prediction. In *Advances in Neural Information Processing Systems*, pages 64–72, 2012.
- Ya Xue, David Dunson, and Lawrence Carin. The matrix stick-breaking process for flexible multi-task learning. In *International Conference on Machine Learning*, pages 1063–1070, 2007.
- Yu Zhang and Dit-Yan Yeung. A convex formulation for learning task relationships in multi-task learning. In *Uncertainty in Artificial Intelligence*, 2010.
- Jun Zhu. Max-margin nonparametric latent feature relational models for link prediction. In *International Conference on Machine Learning*, pages 719–726, 2012.
- Jun Zhu, Amir Ahmed, and Eric P. Xing. MedLDA: Maximum margin supervised topic models for regression and classification. In *International Conference on Machine Learning*, pages 1257–1264, 2009.
- Jun Zhu, Ning Chen, Hugh Perkins, and Bo Zhang. Gibbs max-margin topic models with fast inference algorithms. In *International Conference on Machine Learning*, pages 124–132, 2013.
- Jun Zhu, Ning Chen, and Eric P. Xing. Infinite latent SVM for classification and multi-task learning. In *Advances in Neural Information Processing Systems*, pages 1620–1628, 2011a.
- Jun Zhu, Ning Chen, and Eric P. Xing. Infinite SVM: a Dirichlet process mixture of large-margin kernel machines. In *International Conference on Machine Learning*, pages 617–624, 2011b.
- Jun Zhu and Eric P. Xing. Maximum entropy discrimination Markov networks. *Journal of Machine Learning Research*, (10):2531–2569, 2009.