

---

# Scaling the Indian Buffet Process via Submodular Maximization

---

Colorado Reed  
Zoubin Ghahramani

CR478@CAM.AC.UK  
ZOUBIN@ENG.CAM.AC.UK

Engineering Department, Cambridge University, Cambridge UK

## Abstract

Inference for latent feature models is inherently difficult as the inference space grows exponentially with the size of the input data and number of latent features. In this work, we use Kurihara & Welling (2008)’s maximization-expectation framework to perform approximate MAP inference for linear-Gaussian latent feature models with an Indian Buffet Process (IBP) prior. This formulation yields a submodular function of the features that corresponds to a lower bound on the model evidence. By adding a constant to this function, we obtain a nonnegative submodular function that can be maximized via a greedy algorithm that obtains at least a  $\frac{1}{3}$ -approximation to the optimal solution. Our inference method scales linearly with the size of the input data, and we show the efficacy of our method on the largest datasets currently analyzed using an IBP model.

## 1. Introduction

Nonparametric latent feature models experienced a surge of interest in the machine learning community following Griffiths & Ghahramani (2006)’s formulation of the Indian Buffet Process (IBP)—a nonparametric prior for equivalence classes of sparse binary matrices. These binary matrices have a finite number of exchangeable rows and an unbounded number of columns, where a 1 in row  $n$  and column  $k$  indicates that observation  $n$  expresses latent feature  $k$ . For example, given an image dataset of human faces, each observation is an image, and the latent features might be “is smiling,” “is wearing glasses,” etc. More generally, feature models can be viewed as a generalization of unsupervised clustering, see Broderick et al. (2012).

The IBP prior is often used in sparse matrix factorization models where a data matrix of  $N$   $D$ -dimensional observations is expressed as a product of two matrices that factor over  $K$  latent factors plus a noise term:  $\mathbf{X} = \mathbf{Z}\mathbf{A} + \mathbf{E}$ . Formally, this model has a binary feature matrix  $\mathbf{Z} \in \{0, 1\}^{N \times K}$  that linearly combines a latent factor matrix  $\mathbf{A} \in \mathbb{R}^{K \times D}$  plus a noise matrix  $\mathbf{E} \in \mathbb{R}^{N \times D}$  to form the observed data  $\mathbf{X} \in \mathbb{R}^{N \times D}$ . Placing an IBP prior on  $\mathbf{Z}$  lets  $K$  be unbounded and allows the number of active features  $K_+$  (those with non-zero  $\mathbf{Z}$  column sums) to be learned from the data while remaining finite with probability one. The IBP inspired several infinite-limit versions of classic matrix factorization models, e.g. infinite independent component analysis models (Knowles & Ghahramani, 2007).

Inference with IBP models is challenging as its discrete state space has  $2^{NK_+}$  possible assignments. In turn, the IBP has found limited application to large data in comparison to the Chinese Restaurant Process, which assigns one feature to each observation. In this paper, we use Kurihara & Welling (2008)’s Maximization-Expectation (ME) framework to perform approximate MAP inference with IBP matrix factorization models, termed MEIBP inference. For nonnegative  $\mathbf{A}$ , we show that we can obtain approximate MAP solutions for  $\mathbf{Z}$  by maximizing  $N$  submodular cost functions. The submodularity property enables the use of a simple greedy algorithm that obtains at least a  $\frac{1}{3}$ -approximation to the optimal solution. While the worst-case complexity of MEIBP inference is comparable to sampling and variational approaches, in §5 we show that MEIBP inference often converges to better solutions than variational methods and similar solutions as the best sampling techniques but in a fraction of the time.

This paper is structured as follows: in §2 we present background material that sets the foundation for our presentation of MEIBP inference in §3 and the resulting submodular maximization problem that arises, then in §4 we discuss related work, and in §5 we compare the MEIBP with other IBP inference techniques using both synthetic and real-world datasets.

## 2. Background

### 2.1. The Indian Buffet Process

Griffiths & Ghahramani (2006) derived the IBP prior by placing independent beta priors on Bernoulli generated entries of an  $N \times K$  binary matrix  $\mathbf{Z}$ , marginalizing over the beta priors, and letting  $K$  go to infinity. In this infinite limit, however,  $P(\mathbf{Z})$  is zero for any particular  $\mathbf{Z}$ . Griffiths & Ghahramani (2006) therefore take the limit of an equivalence classes of binary matrices,  $[\mathbf{Z}]$ , defined by the “left-order form” (*lof*) ordering of the columns and show that  $P([\mathbf{Z}]_{lof})$  has a non-zero probability as  $K$  goes to infinity.

The *lof* ordering arranges the columns of  $\mathbf{Z}$  such that the binary values of the columns are non-increasing, where the first row is the most significant bit. Ding et al. (2010) examine different “shifted” equivalence classes formed by shifting all-zero columns to the right of non-zero columns while maintaining the non-zero column ordering. Given  $K_+$  non-zero columns, the IBP prior for the shifted equivalence classes is

$$P([\mathbf{Z}]|\alpha) = \frac{\alpha^{K_+}}{K_+!} e^{-\alpha H_N} \prod_{k=1}^{K_+} \frac{(N - m_k)!(m_k - 1)!}{N!} \quad (1)$$

where  $\alpha$  is a hyperparameter,  $H_N$  is the  $N^{th}$  harmonic number, and  $m_k = \sum_{n=1}^N z_{nk}$ . The supplementary material has a derivation of  $P([\mathbf{Z}]|\alpha)$  as well as a comparison to the *lof* equivalence classes. The derivations in §3 can be applied using either equivalence class. However, the shifted equivalence classes simplify the mathematics and produce the same results in practice.

### 2.2. Maximization-Expectation

Kurihara & Welling (2008) presented the ME algorithm: an inference algorithm that exchanges the expectation and maximization variables in the EM algorithm. Consider a general probabilistic model  $p(\mathbf{X}, \mathbf{Z}, \mathbf{A})$ , where  $\mathbf{X}$  are the observed random variables (RVs),  $\mathbf{Z}$  are the local latent RVs, and  $\mathbf{A}$  are the global latent RVs. RVs are qualified as “local” if there is one RV for each observation, and RVs are “global” if their multiplicity is constant or inferred from the data.

ME can be viewed as a special case of a Mean-Field Variational Bayes (MFVB) approximation to a posterior that cannot be computed analytically,  $p(\mathbf{Z}, \mathbf{A}|\mathbf{X})$ . MFVB operates by approximating the posterior distribution of a given probabilistic model by assuming independent variational distributions,  $p(\mathbf{Z}, \mathbf{A}|\mathbf{X}) \approx q(\mathbf{Z})q(\mathbf{A})$  (Attias, 2000; Ghahramani & Beal, 2001). The independence constraint lets us compute the variational distribution  $q$  that minimizes the KL diver-

gence between the variational distribution and true posterior. Without this constraint, the distribution that minimizes the KL-divergence is the true posterior, returning us to our original problem. In MFVB, we determine the variational distributions and their parameters using coordinate ascent optimization in which we iteratively update:

$$q(\mathbf{Z}) \propto \exp \mathbb{E}_{q(\mathbf{A})} [\ln p(\mathbf{X}, \mathbf{Z}, \mathbf{A})] \quad (2)$$

$$q(\mathbf{A}) \propto \exp \mathbb{E}_{q(\mathbf{Z})} [\ln p(\mathbf{X}, \mathbf{Z}, \mathbf{A})], \quad (3)$$

which commonly has closed-form solutions.

The EM algorithm can be viewed as a special case of MFVB that obtains MAP values of the global RVs by letting  $q(\mathbf{A}) = \delta(\mathbf{A} - \mathbf{A}^*)$ , where  $\delta(\cdot)$  is the delta function and  $\mathbf{A}^*$  is the MAP assignment. The ME algorithm instead maximizes the local RVs  $\mathbf{Z}$  and computes the expectation over the global RVs  $\mathbf{A}$ , which can be viewed as MFVB with  $q(\mathbf{Z}) = \delta(\mathbf{Z} - \mathbf{Z}^*)$ . In the limit of large  $N$ , the ME algorithm recovers a Bayesian information criterion regularization term (Kurihara & Welling, 2008). Also, maintaining a variational distribution over the global RVs retains the model selection ability of MFVB, while using point estimates of the local RVs allows the use of efficient data structures and optimization techniques. As we will show, the ME algorithm leads to a scalable submodular optimization problem for latent feature models.

### 2.3. Submodularity

Submodularity is a set function property that makes optimization of the function tractable or approximable. Given ground set  $V$  and set function  $f : 2^V \rightarrow \mathbb{R}$ ,  $f$  is *submodular* if for all  $A \subseteq B \subseteq V$  and  $e \in V \setminus B$ :

$$f(A \cup \{e\}) - f(A) \geq f(B \cup \{e\}) - f(B), \quad (4)$$

which expresses a “diminishing returns” property, where the incremental benefit of element  $e$  diminishes as we include it in larger solution sets. Submodularity is desirable in discrete optimization because submodular functions are discrete analogs of convex functions and can be globally minimized in polynomial time (Lovász, 1983). However, global submodular maximization is NP-hard, but submodularity often enables approximation bounds via greedy algorithms. In the next section, we show that determining a MAP estimate of  $\mathbf{Z}$  in the ME algorithm is a scalable submodular maximization problem.

## 3. Maximization-Expectation IBP

Here we present the ME algorithm for nonnegative linear-Gaussian IBP models and show that approxi-

mate MAP inference arises as a submodular maximization problem. Boldface variables are matrices with (row, column) subscripts; a dot indicates all elements of the dimension, and lowercase variables are scalars.

### 3.1. Nonnegative Linear-Gaussian IBP Model

We consider the following probabilistic model:

$$p(\mathbf{X}, \mathbf{Z}, \mathbf{A}|\boldsymbol{\theta}) = p(\mathbf{X}|\mathbf{Z}, \mathbf{A}, \sigma_{\mathbf{X}}^2)p(\mathbf{A}|\sigma_{\mathbf{A}}^2)p(\mathbf{Z}|\alpha) \quad (5)$$

$$p(\mathbf{X}|\mathbf{Z}, \mathbf{A}, \sigma_{\mathbf{A}}^2) = \prod_{n=1}^N \mathcal{N}(\mathbf{X}_n; \mathbf{Z}_n \mathbf{A}, \sigma_{\mathbf{A}}^2 I) \quad (6)$$

$$p(\mathbf{A}|0, \sigma_{\mathbf{A}}^2) = \prod_{k=1}^K \prod_{d=1}^D \mathcal{TN}(a_{kd}; 0, \sigma_{\mathbf{A}}^2) \quad (7)$$

with  $p([\mathbf{Z}|\alpha)$  specified in Eq. 1. This is a nonnegative linear-Gaussian IBP model, where the prior over the latent factors,  $p(\mathbf{A}|0, \sigma_{\mathbf{A}}^2)$ , is a zero-mean i.i.d. truncated Gaussian with nonnegative support, denoted  $\mathcal{TN}$ . As we show below, this nonnegative prior yields a submodular maximization problem when optimizing  $\mathbf{Z}$ . We use a truncated Gaussian as it is conjugate to the Gaussian likelihood, but other nonnegative priors (e.g. exponential) can be used. For brevity we assume the hyperparameters,  $\boldsymbol{\theta} = \{\alpha, \sigma_{\mathbf{A}}^2, \sigma_{\mathbf{X}}^2\}$ , are known and discuss  $\boldsymbol{\theta}$  inference in the supplementary material.

### 3.2. MEIBP Evidence

In the ME framework, we approximate the true posterior distribution via a MFVB assumption:

$$p(\mathbf{Z}, \mathbf{A}|\mathbf{X}, \boldsymbol{\theta}) \approx q(\mathbf{A})\delta(\mathbf{Z} - \mathbf{Z}^*). \quad (8)$$

That is, we maintain a variational distribution over the latent factors  $\mathbf{A}$  and optimize the latent features  $\mathbf{Z}$ . Given the MFVB constraint, we determine the variational distributions by minimizing the KL-divergence between the variational distributions and the true posterior, which is equivalent to maximizing a lower bound on the evidence (Attias, 2000):

$$\begin{aligned} \ln p(\mathbf{X}|\boldsymbol{\theta}) &= \mathbb{E}_q[\ln p(\mathbf{X}, \mathbf{A}, \mathbf{Z}|\boldsymbol{\theta})] + H[q] + D(q||p) \\ &\geq \mathbb{E}_q[\ln p(\mathbf{X}, \mathbf{A}, \mathbf{Z}|\boldsymbol{\theta})] + H[q] \equiv \mathcal{F} \end{aligned} \quad (9)$$

where  $H[q]$  is the entropy of  $q$  and  $D(q||p)$  represents the KL-divergence between the variational distribution and the true posterior. The evidence lower bound,  $\mathcal{F}$ , for the nonnegative linear-Gaussian IBP model is:

$$\begin{aligned} &\frac{1}{\sigma_{\mathbf{X}}^2} \sum_{n=1}^N \left[ -\frac{1}{2} \mathbf{Z}_n \boldsymbol{\Phi} \boldsymbol{\Phi}^T \mathbf{Z}_n^T + \mathbf{Z}_n \boldsymbol{\xi}_n^T \right] - \ln K_+ \\ &+ \sum_{k=1}^{K_+} \left[ \ln \frac{(N - m_k)(m_k - 1)!}{N!} + \eta_k \right] + \text{const} \end{aligned} \quad (10)$$

with

$$\xi_{nk} = \boldsymbol{\Phi}_k \cdot \mathbf{X}_n^T + \frac{1}{2} \sum_{d=1}^D [\mathbb{E}[a_{kd}]^2 - \mathbb{E}[a_{kd}^2]] \quad (11)$$

and

$$\eta_k = \sum_{d=1}^D \left[ -\frac{\ln \frac{\pi \sigma_{\mathbf{A}}^2}{2\alpha^{2/D}}}{2} - \frac{\mathbb{E}[a_{kd}^2]}{2\sigma_{\mathbf{A}}^2} + H(q(a_{kd})) \right] \quad (12)$$

where  $\boldsymbol{\Phi}_k = (\mathbb{E}[a_{k1}], \dots, \mathbb{E}[a_{kD}])$ , and all expectations are with respect to  $q(\mathbf{A})$ , which is defined in the next subsection. In §3.5 we show that maximizing this lower bound with respect to  $\mathbf{Z}$  can be formulated as a submodular maximization problem.

### 3.3. Variational Factor Updates

Maximizing Eq. 10 with respect to  $q(\mathbf{A})$  yields

$$q(\mathbf{A}) = \prod_{k=1}^K \prod_{d=1}^D \mathcal{TN}(a_{kd}; \tilde{\mu}_{kd}, \tilde{\sigma}_{kd}^2), \quad (13)$$

with parameter updates

$$\tilde{\mu}_{kd} = \rho_k \sum_{n=1}^N z_{nk} \left( x_{nd} - \sum_{k' \neq k} z_{nk'} \mathbb{E}[a_{k'd}] \right) \quad (14)$$

$$\tilde{\sigma}_{kd}^2 = \rho_k \sigma_{\mathbf{X}}^2, \quad (15)$$

where  $\rho_k = (m_k + \frac{\sigma_{\mathbf{X}}^2}{\sigma_{\mathbf{A}}^2})^{-1}$ . These updates take  $O(NK^2D)$ , and the relevant moments are:

$$\mathbb{E}[a_{kd}] = \tilde{\mu}_{kd} + \tilde{\sigma}_{kd} \frac{\sqrt{2/\pi}}{\text{erfcx}(\wp_{kd})} \quad (16)$$

$$\mathbb{E}[a_{kd}^2] = \tilde{\mu}_{kd}^2 + \tilde{\sigma}_{kd}^2 + \tilde{\sigma}_{kd} \tilde{\mu}_{kd} \frac{\sqrt{2/\pi}}{\text{erfcx}(\wp_{kd})} \quad (17)$$

with  $\wp_{kd} = -\frac{\tilde{\mu}_{kd}}{\tilde{\sigma}_{kd}\sqrt{2}}$  and  $\text{erfcx}(y) = e^{y^2}(1 - \text{erf}(y))$  representing the scaled complementary error function.

### 3.4. Evidence Lower Bound as $K \rightarrow \infty$

Here we show that the evidence lower bound [Eq. 10] is well-defined in the limit  $K \rightarrow \infty$ ; in fact, all instances of  $K$  are simply replaced by  $K_+$ . Therefore, similar to variational IBP methods, a user must specify a maximum model complexity  $K_+$ . A benefit over variational IBP methods, however, is that the  $q(\mathbf{Z})$  updates are not affected by inactive features—see §4.

We take this limit by breaking the evidence into components  $1, \dots, K_+$  and  $K_+ + 1, \dots, K$  and note that

when  $m_k = 0$ :  $\tilde{\mu}_{kd} = 0$ ,  $\tilde{\sigma}_{kd}^2 = \sigma_{\mathbf{A}}^2$ , and  $H(a_{kd}) = \frac{1}{2} \ln \frac{\pi e \sigma_{\mathbf{A}}^2}{2}$ . After some algebra, the evidence becomes:

$$\psi_{K_+} + \frac{1}{2} \sum_{k=K_++1}^K \sum_{d=1}^D \left[ -\ln \frac{\pi \sigma_{\mathbf{A}}^2}{2} - \frac{\mathbb{E}[a_{kd}^2]}{\sigma_{\mathbf{A}}^2} + \ln \frac{\pi e \sigma_{\mathbf{A}}^2}{2} \right] \quad (18)$$

where  $\psi_{K_+}$  is Eq. 10 but with  $K_+$  replacing all  $K$ . From Eq. 17, we see that  $\mathbb{E}[a_{kd}^2] = \sigma_{\mathbf{A}}^2$  when  $m_k = 0$ , which causes all terms to cancel in Eq. 18 except  $\psi_{K_+}$ .

The evidence lower bound remains well-defined because both the likelihood and IBP prior terms do not depend on inactive features, so for inactive features the KL-divergence between the posterior and variational distributions is simply the KL-divergence between  $p(\mathbf{A})$  and  $q(\mathbf{A})$ . For inactive features,  $p(\mathbf{A}) = q(\mathbf{A})$ , and as a result, the KL-divergence is zero.

### 3.5. Z Objective Function

Given  $q(\mathbf{A})$ , we compute MAP estimates of  $\mathbf{Z}$  by maximizing the evidence [Eq. 10] for each  $n \in \{1, \dots, N\}$  while holding constant all  $n' \in \{1, \dots, N\} \setminus n$ . Decomposing Eq. 10 into terms that depend on  $\mathbf{Z}_n$  and those that do not yields (see the supplementary material):

$$\mathcal{F}(\mathbf{Z}_n) = -\frac{1}{2\sigma_{\mathbf{X}}^2} \mathbf{Z}_n \cdot \mathbf{\Phi} \mathbf{\Phi}^T \mathbf{Z}_n^T + \mathbf{Z}_n \cdot \boldsymbol{\omega}_n^T + \text{const} - \ln \left( K_{+\setminus n} + \sum_{k=1}^{K_+} \left[ \mathbf{1}_{\{m_{k\setminus n}=0\}} z_{nk} \right] \right)! \quad (19)$$

$$\mathbf{\Phi}_{k\cdot} = \left( \mathbb{E}[a_{k1}], \dots, \mathbb{E}[a_{kD}] \right)$$

$$\omega_{nk} = \frac{1}{\sigma_{\mathbf{X}}^2} \left( \mathbf{\Phi}_{k\cdot} \cdot \mathbf{X}_n^T + \frac{1}{2} \sum_{d=1}^D [\mathbb{E}[a_{kd}]^2 - \mathbb{E}[a_{kd}^2]] \right) + \nu(z_{nk} = 1) - \nu(z_{nk} = 0) + \mathbf{1}_{\{m_{k\setminus n} > 0\}} \eta_k,$$

which is a quadratic pseudo-Boolean function plus a term that penalizes  $K_+$ , where  $\mathbf{1}_{\{\cdot\}}$  is the indicator function, a “ $\setminus n$ ” subscript indicates the given variable is determined after removing the  $n^{\text{th}}$  row from  $\mathbf{Z}$ , and

$$\nu(z_{nk}) = \begin{cases} 0, & \text{if } m_{k\setminus n} = 0 \text{ and } z_{nk} = 0 \\ \ln(N - m_{k\setminus n} - z_{nk})! \\ \quad + \ln(m_{k\setminus n} + z_{nk} - 1)!, & \text{otherwise} \end{cases}$$

We can prove  $\mathcal{F}(\mathbf{Z}_n)$  is submodular given the following two well-known propositions, see Fujishige (2005):

**Proposition 1.** *Nonnegative linear combinations of submodular functions are submodular.*

**Proposition 2.** *A quadratic pseudo-Boolean function with quadratic weight matrix  $\mathbf{W}$  is submodular if and only if  $W_{ij} \leq 0$  for all  $i, j$ .*

Via Proposition 2, we see that  $-\frac{1}{2\sigma_{\mathbf{X}}^2} \mathbf{Z}_n \cdot \mathbf{\Phi} \mathbf{\Phi}^T \mathbf{Z}_n^T + \mathbf{Z}_n \cdot \boldsymbol{\omega}_n^T$  is submodular when  $\mathbf{\Phi}$  is nonnegative. From Proposition 1, Eq. 19 is submodular if and only if

$$\mathcal{G}(\mathbf{Z}_n) = -\ln \left( K_{+\setminus n} + \sum_{k=1}^{K_+} \left[ \mathbf{1}_{\{m_{k\setminus n}=0\}} z_{nk} \right] \right)! \quad (20)$$

is submodular. We prove this property by rephrasing  $\mathcal{G}(\mathbf{Z}_n)$  as a set function and using the definition of submodularity given by Eq. 4. Let  $A_n \subseteq B_n \subseteq V$  where  $V = \{1, \dots, K_+\}$  and  $A_n, B_n \in 2^V$  with  $\mathcal{G}(A_n) = -\ln(K_{+\setminus n} + K_{A_n})!$ . Here we let  $K_{A_n} = \sum_{k=1}^{K_+} \mathbf{1}_{\{m_{k\setminus n}=0\}} \mathbf{1}_{\{k \in A_n\}}$  where  $k \in A_n$  indicates  $z_{nk} = 1$ .  $\mathcal{G}$  is submodular if for all  $e \in V \setminus B_n$ :

$$\begin{aligned} \mathcal{G}(A_n \cup \{e\}) - \mathcal{G}(A_n) &\geq \mathcal{G}(B_n \cup \{e\}) - \mathcal{G}(B_n) \\ \ln \frac{(K_{+\setminus n} + K_{A_n})!}{(K_{+\setminus n} + K_{A_n \cup \{e\}})!} &\geq \ln \frac{(K_{+\setminus n} + K_{B_n})!}{(K_{+\setminus n} + K_{B_n \cup \{e\}})!} \end{aligned} \quad (21)$$

Eq. 21 has two cases: (1)  $m_{e\setminus n} > 0$  so  $K_{B_n \cup \{e\}} = K_{B_n}$  and  $K_{A_n \cup \{e\}} = K_{A_n}$ , yielding  $0 \geq 0$  for Eq. 21, which is true for all  $e \in V \setminus B_n$  and  $A_n \subseteq B_n$ , (2)  $m_{e\setminus n} = 0$  so  $K_{B_n \cup \{e\}} = K_{B_n} + 1$  and  $K_{A_n \cup \{e\}} = K_{A_n} + 1$ . After some algebra this yields  $K_{B_n \cup \{e\}} \geq K_{A_n \cup \{e\}}$  for Eq. 21, which is again true for all  $e \in V \setminus B_n$  and  $A_n \subseteq B_n$ . As a result, both components of Eq. 19 are submodular, and by Proposition 1, adding these terms yields a submodular function.

### 3.6. Z Optimization

Eq. 19 is an unconstrained nonmonotone submodular function. Feige et al. (2011) prove that an approximability guarantee is NP-hard for this class of functions. However, Feige et al. (2011) also show that a local-search (ls) algorithm obtains a constant-factor approximation to the optimal solution, provided the submodular objective function is nonnegative. For a submodular function  $\mathcal{F} : 2^V \rightarrow \mathbb{R}$  with ground set  $V = \{1, \dots, K_+\}$  and solution set  $A \subseteq V$ , the ls-algorithm operates as follows:

1. *initialize*: let  $A = \{\arg \max_{w \in V} \mathcal{F}(\{w\})\}$
2. *grow*: while there is an element  $w \in V \setminus A$  s.t.  $\mathcal{F}(A \cup \{w\}) > (1 + \frac{\epsilon}{|V|^2}) \mathcal{F}(A)$ : let  $A := A \cup \{w\}$
3. *prune*: if there is an element  $w \in A$  s.t.  $\mathcal{F}(A \setminus \{w\}) > (1 + \frac{\epsilon}{|V|^2}) \mathcal{F}(A)$ : let  $A := A \setminus \{w\}$ , goto 2.
4. *return*: maximum of  $\mathcal{F}(A)$  and  $\mathcal{F}(V \setminus A)$ .

The ls-algorithm obtains a solution that is greater than  $\frac{1}{3}(1 + \frac{\epsilon}{|V|^2}) \text{OPT}$ —where  $\epsilon$  is a parameter and OPT is

the maximum value of  $\mathcal{F}$ . The ls-algorithm performs  $O(|V|^3 \log |V|)$  function calls in the grow/prune steps.

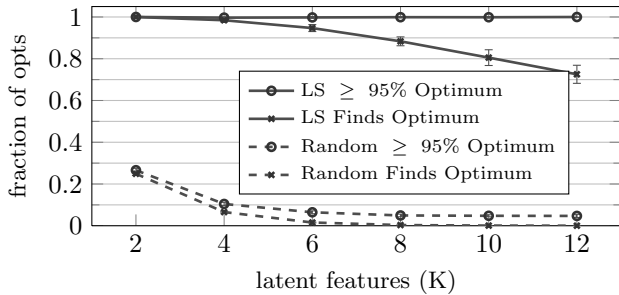


Figure 1. Fraction of ls-algorithm and random solutions that obtain [within 95% of] the true optimum using data generated from the nonnegative linear-Gaussian model with  $N = 500, D = 50, \sigma_X = 1.0$ .

Since Eq. 19 is strictly negative, we use its *normalized cost function* to interpret the ls-approximability guarantee:  $\mathcal{F}(\mathbf{Z}_{n\cdot}) - \mathcal{F}_{n0}$ , where  $\mathcal{F}_{n0}$  is the minimum value of  $\mathcal{F}(\mathbf{Z}_{n\cdot})$ . Using the normalized cost function, we obtain the following optimality guarantee:

$$\mathcal{F}(\mathbf{Z}_{n\cdot}^{\text{ls}}) \geq \mathcal{F}_{n0} + \frac{1}{3} (\mathcal{F}(\mathbf{Z}_{n\cdot}^*) - \mathcal{F}_{n0}) \quad (22)$$

where the superscript “ls” denotes the solution from the greedy ls-algorithm and an asterisk denotes the set that obtains the true maximum. This inequality states that the ls-algorithm solution is guaranteed to perform better as the disparity between the maximum and minimum solution increases. However, we emphasize that this inequality does not provide an optimality guarantee for the global MAP solution.

We studied the empirical performance of the ls-algorithm by generating high noise data ( $\sigma_X = 1$ ) from the nonnegative linear-Gaussian model with  $N = 500, D = 50$  and compared the ls-algorithm with the brute-force optimal solution as  $K$  varied from 2 to 12, performing  $1000K$  total optimizations for each of ten randomly generated datasets. Furthermore, we compared the ls-algorithm with randomly sampled  $\mathbf{Z}_{n\cdot}$  solutions to demonstrate that the optimization space was not skewed to favor solutions near the optimal value.

Figure 1 shows the fraction of solutions that obtain the true optimum as well as the fraction of solutions that were greater than 95% of  $\mathcal{F}(\mathbf{Z}_{n\cdot}^*) - \mathcal{F}_{n0}$ , where the error bars indicate the combined standard deviation over the  $10 \times 1000K$  optimizations. The ls-algorithm found the optimal solution roughly 70% of the time for  $K = 12$  and obtained within 95% of the optimal solution over 99.9% of the time for all  $K$ —meaning we could empirically replace the  $\frac{1}{3}$  in Eq. 22 with  $\frac{19}{20}$ . The

random sampling comparison indicated that the optimization space did not favor nearly-optimal solutions: its convergence to 5% for *within-95% optimal solutions* was characteristic of a uniform solution space.

By precomputing  $\Phi\Phi^T$  and maintaining an auxiliary vector of  $K_+$  weights, we can evaluate Eq. 19 in constant time when adding/removing elements to the solution set. In turn, the ls-algorithm optimizes  $\mathcal{F}(\mathbf{Z}_{n\cdot})$  in  $K_+^2 D + O(K_+^3 \log K_+)$  operations. The  $O(K_+^3 \log K_+)$  component arises from the add/removal operations, but as we show in Figure 2, it is a loose upper bound that scales sub-quadratically in practice.




Figure 2. Number of  $O(1)$  updates per ls-optimization using data generated from the nonnegative linear-Gaussian model with  $N = 1000, D = 1000, \sigma_X = 1.0$ .

## 4. Related Work

Several proposals have been made for efficient inference with latent feature models. Table 1 summarizes the per-iteration complexity of the methods discussed below. In the next section we compare these methods on two synthetic and three real-world datasets.

Doshi-Velez et al. (2009) formulated a coordinate ascent variational inference technique for IBP models (VIBP). This method used the “stick breaking” formulation of the IBP, which maintained coupled beta-distributed priors on the entries of  $\mathbf{Z}$ —marginalizing these priors does not allow closed-form MFVB updates. Unlike MEIBP inference, maintaining the beta priors has the undesirable consequence that inactive features contribute to the evidence lower bound and must be ignored when updating the variational distributions. This was not a problem for Doshi-Velez et al. (2009)’s finite variational IBP, which computes variational distributions for a linear-Gaussian likelihood with a parametric beta-Bernoulli prior on the latent features. The inference complexity for both methods is  $O(NK_+^2 D)$ , which is dominated by updating  $q(\mathbf{Z})$ .

Ding et al. (2010) used mixed expectation-propagation style updates with MFVB inference in order to perform variational inference for a nonnegative linear-Gaussian IBP model (INMF). The expectation-propagation style updates are more complicated than MFVB updates and have per-iteration complexity  $O(N(K^3D + KD^2))$ . Ding et al. (2010) motivated this framework by stating that the evidence lower bound of a linear-Gaussian likelihood with a truncated Gaussian prior on the latent factors is negative infinity. This is only true if the variational distribution is a Gaussian, however the free-form variational distribution for their model is a truncated Gaussian, which has a well-defined evidence lower bound.

Doshi-Velez & Ghahramani (2009) presented a linear-time “accelerated” Gibbs sampler for conjugate IBP models that effectively marginalized over the latent factors (AIBP). The per-iteration complexity was  $O(N(K^2 + KD))$ . This is comparable to the uncollapsed IBP sampler (UGibbs) that has per-iteration complexity  $O(NDK^2)$  but does not marginalize over the latent factors, and as a result, takes longer to mix. In terms of both complexity and empirical performance, the accelerated Gibbs sampler is the most scalable sampling-based IBP inference technique currently available. One constraint of the accelerated IBP is that the latent factor distribution must be conjugate to the likelihood, which for instance, does not allow nonnegative priors on the latent factors.

Rai & Daume III (2011) introduced a beam-search heuristic for locating approximate MAP solutions to linear-Gaussian IBP models (BS-IBP). This heuristic sequentially adds a single data point to the model and determines the latent feature assignments by scoring all  $2^{K+}$  latent feature combinations. The scoring heuristic uses an estimate of the joint probability,  $P(\mathbf{X}, \mathbf{Z})$  to score assignments, which evaluates the collapsed likelihood  $P(\mathbf{X}|\mathbf{Z})$  for all  $2^{K+}$  possible assignments: an expensive  $N^3(K_+ + D)$  operation, yielding a per-iteration complexity of  $O(N^3(K_+ + D)2^{K+})$ .

## 5. Experiments

We evaluated the inference quality and efficiency of MEIBP inference on two synthetic and three real-world datasets. We used the runtime and predictive likelihood of held-out observations as our performance criteria and compared MEIBP inference with the methods listed in Table 1 (the finite and infinite VIBP are differentiated with an “f-” and “i-” prefix). We used a truncated Gaussian prior on the latent factors for UGibbs and INMF, and Gaussian priors for the AIBP and variational methods. In our evalua-

Table 1. Worst-case per-iteration complexity given a linear-Gaussian likelihood model for  $N$   $D$ -dimensional observations and  $K_+$  active latent features.

Algorithm	Iteration Complexity
MEIBP	$O(N(K_+^2D + K_+^3 \ln K_+))$
VIBP (Doshi-Velez et al., 2009)	$O(NK_+^2D)$
AIBP (Doshi-Velez & Ghahramani, 2009)	$O(N(K_+^2 + K_+D))$
UGibbs (Doshi-Velez & Ghahramani, 2009)	$O(NK_+^2D)$
BS-IBP (Rai & Daume III, 2011)	$O(N^3(K_+ + D)2^{K_+})$
INMF (Ding et al., 2010)	$O(N(K_+^3D + K_+D^2))$

tions, we also included Schmidt et al. (2009)’s iterated conditional modes algorithm, which computes a MAP estimate of a parametric nonnegative matrix factorization model:  $\mathbf{X} = \mathbf{B}\mathbf{A} + \mathbf{E}$ , where  $\mathbf{B}$  and  $\mathbf{A}$  have exponential priors and  $\mathbf{E}$  is zero-mean Gaussian noise. We abbreviate this model “BNMF”; it has a per-iteration complexity of  $O(N(K_+^2 + K_+D))$ .

The VIBP and MEIBP inference methods specify a maximum  $K$  value, while the sampling methods are unbounded. Therefore, we also included truncated versions of the sampling methods (indicated by a “t-” prefix) for a fairer comparison. We centered all input data to have a 0-mean for the models with 0-mean Gaussian priors and a 0-minimum for nonnegative models, and all inferred matrices were initialized randomly from their respective priors. Following Doshi-Velez & Ghahramani (2009), we fixed the hyperparameters  $\sigma_{\mathbf{X}}$  and  $\sigma_{\mathbf{A}}$  to  $\frac{3}{4}\sigma$ , where  $\sigma$  was the standard deviation across all dimensions of the data, and set  $\alpha = 3$ . We ran each algorithm until the multiplicative difference of the average training log-likelihood differed by less than  $10^{-4}$  between blocks of five iterations with a maximum runtime of 36 hours. Our experiments used MATLAB implementations of the algorithms, as provided by the respective authors, on 3.20 GHz processors.

**Synthetic Data** We created high-noise synthetic datasets in the following way: (1) sample  $z_{n,k} \sim \text{Bernoulli}(p = 0.4)$ , (2) generate  $\mathbf{A}$  with  $K$  random, potentially overlapping binary factors, (3) let  $\mathbf{X} = \mathbf{Z}\mathbf{A} + \mathbf{E}$ , where  $\mathbf{E} \sim \mathcal{N}(0, 1)$ . We evaluated the predictive likelihood on 20% of the dimensions from the last half of the data (see supplementary information).

Figure 3 shows the evolution of the test log-likelihood over time for a small dataset with  $N = 500, D = 500, K = 20$  and a large dataset with  $N = 10^5, D = 10^3, K = 50$ . All models were initialized randomly




Figure 3. Evolution of test log-likelihood over time for a synthetic dataset; **Left**: dataset with  $N = 500, D = 500, K = 20$  (NB: x-axis is log-scale) **Right**: dataset with  $N = 10^5, D = 10^3, K = 50$  (NB: x-axis is linear-scale).

with the true number of latent features, and the error regions display the standard deviation over five random restarts. The BS-IBP and INMF methods were removed from our experiments following the synthetic dataset tests as both methods took at least an order of magnitude longer than the other methods: in 36 hours, the BS-IBP did not complete a single iteration on the small dataset, and the INMF did not complete a single iteration on the large dataset.

MEIBP converged quickest among the IBP models for both the small and large dataset, while the parametric BNMF model converged much faster than all IBP models. However, the IBP models captured the sparsity of the latent features and the MEIBP and UGibbs eventually outperformed the BNMF on the small dataset, while only the MEIBP outperformed the BNMF on the large dataset. The VIBP methods converged quicker than the sampling counterparts but had trouble escaping local optima. The uncollapsed samplers eventually performed as well as the MEIBP on the small dataset but did not mix to a well-performing distribution for the large dataset.

**Real Data** Table 2 summarizes the real-world datasets used in our experiments. Piano and Yale-BC are dense real-valued datasets, whereas the Flickr dataset is a sparse binary dataset (0.81% filled). For the Piano and Flickr datasets, we evaluated the predictive likelihood on a held-out portion of 25% of the dimensions from the last half of the datasets. The Yale-BC dataset had roughly sixty-four facial images of thirty-eight subjects, and we removed the bottom half of five images from each subject for testing.

Figure 4 shows the test log-likelihood and convergence time for all inference methods applied to the real-world datasets, averaged over five random restarts. All inference methods were initialized with  $K = \{10, 25, 50\}$  as indicated by the size of the marker (the smallest marker shows the  $K = 10$  results). The sampling

Table 2. Summary of real-world datasets.

Dataset	Size ( $N \times D$ )	Details
Piano (Poliner & Ellis, 2006)	$16000 \times 161$	DFT of piano recordings
Flickr (Kollar & Roy, 2009)	$25000 \times 1500$	binary image-tag indicators
Yale-BC (Lee et al., 2005)	$2414 \times 32256$	face images with various lightings

methods (AIBP, UGibbs) also include a large faded marker that shows the results for unbounded  $K$ .

The Piano results were similar to the small synthetic dataset. The BNMF converged much faster than the IBP models, and the MEIBP performed best among the IBP models in terms of runtime and test log-likelihood—it converged to a similar solution as the AIBP in one-third the time. Though UGibbs has the best per-iteration complexity, it got stuck in poor local optima when randomly initialized. The VIBP methods and MEIBP expressed large uncertainty about the latent factors early on and overcame these poor local optima. By using hard latent feature assignments, the MEIBP took larger steps in the inference space than the VIBP methods, which was beneficial for this dataset, and achieved similar results to the AIBP.

MEIBP inference performed comparable to the best IBP sampling technique for the sparse binary Flickr dataset and converged over an order of magnitude faster. Surprisingly, the dense BNMF inference performed very well on this dataset even though the dataset was sparse and binary. The BNMF converged slower than the MEIBP because it inferred a sparse matrix from a dense prior, which took over four times as many iterations to converge compared to the dense datasets. While the t-AIBP converged to a better solution than the MEIBP, it took over an order-of-magnitude longer to surpass the MEIBP’s performance. As we demonstrate with the Flickr, ini-




Figure 4. Inference results on real-world datasets. The size of the marker indicates the  $K$  value for  $K = \{10, 25, 50\}$ , with larger markers indicating a larger  $K$ . AIBP and UGibbs also include a larger faded marker that shows the inference results for unbounded  $K$ . The Flickr plot also shows the result of initializing the AIBP using the MEIBP result. The error-bars indicate the standard deviation of convergence time and test log-likelihood over five random restarts.

tializing the AIBP with the MEIBP outcome obtained a similar solution in a fraction of the time (indicated as “meibp+aibp” on the figure).

The MEIBP converged faster than the other IBP methods for the Yale-BC dataset but to a lower test likelihood. The UGibbs and BNMF also experienced difficulty for this dataset, where BNMF converged to a test log-likelihood around  $-3.6 \times 10^6$  (not visible in the figure). These linear-Gaussian models with nonnegative priors performed worse than the models with Gaussian priors because the dataset contained many images with dark shadows covering part of the face. The nonnegative priors appeared to struggle with reconstructing these shadows, because unlike the Gaussian priors, they could not infer negative-valued “shadow” factors that obscured part of the image.

In the above experiments, the MEIBP consistently exhibited a sudden convergence whereby it obtained a local optima and the ls-algorithm did not change any  $\mathbf{Z}$  assignments. This is a characteristic of using hard assignments with a greedy algorithm: at a certain point, changing any latent feature assignments decreased the objective function. This abrupt convergence, in combination with the speed of the ls-algorithm, helped the MEIBP consistently converge faster than other IBP methods. Furthermore, the submodular maximization algorithm converged to local optima that were comparable or better than the sampling or variational results, though at the cost of only obtaining a MAP solution. Like the variational methods, it maintained a distribution over  $\mathbf{A}$  that prevented it from getting stuck in local optima early on, and like the sampling methods, the MEIBP used hard  $\mathbf{Z}$  assignments to take larger steps in the inference space and obtain better optima.

## 6. Summary and Future Work

We presented a new inference technique for IBP models that used Kurihara & Welling (2008)’s ME framework to perform approximate MAP inference via submodular maximization. Our key insight was to exploit the submodularity inherent in the evidence lower bound formulated in §3, which arose from the quadratic pseudo-Boolean component of the linear-Gaussian model. MEIBP inference converged faster than competing IBP methods and obtained comparable solutions on various datasets.

There are many discrete Bayesian nonparametric priors, such as the Dirichlet process, and an interesting area for future research will be to generalize our results in order to phrase inference with these priors as submodular optimization problems. Furthermore, we used a simple local-search algorithm to obtain a  $\frac{1}{3}$ -approximation bound, but concurrently with this work, Buchbinder et al. (2012) proposed a simpler stochastic algorithm for unconstrained submodular maximization that obtains an expected  $\frac{1}{2}$ -approximation bound. Using this algorithm, MEIBP inference has an improved worst case complexity of  $O(NK_+^2D)$ . We will investigate this algorithm in an extended technical version of this paper.

**Code:** A MATLAB implementation of MEIBP is available at <https://github.com/cjrd/MEIBP>.

**Acknowledgements:** CR was supported by the Winston Churchill Foundation of the United States, and ZG was supported by EPSRC grant EP/I036575/1 and grants from Google and Microsoft. We thank the anonymous reviewers for their helpful comments.



## References

- Attias, H. A variational bayesian framework for graphical models. *Advances in Neural Information Processing Systems*, 12:209–215, 2000.
- Broderick, T., Jordan, M. I., and Pitman, J. Clusters and features from combinatorial stochastic processes. *arXiv:1206.5862*, 2012.
- Buchbinder, N., Feldman, M., Naor, J., and Schwartz, R. A tight linear time (1/2)-approximation for unconstrained submodular maximization. In *53rd Annual Symposium on Foundations of Computer Science*, pp. 649–658. IEEE, 2012.
- Ding, N., Qi, Y.A., Xiang, R., Molloy, I., and Li, N. Nonparametric Bayesian matrix factorization by Power-EP. In *14th Int'l Conf. on AISTATS*, volume 9, pp. 169–176, 2010.
- Doshi-Velez, F. and Ghahramani, Z. Accelerated sampling for the Indian buffet process. In *Proceedings of the 26th Annual Int'l Conference on Machine Learning*, pp. 273–280, 2009.
- Doshi-Velez, F., Miller, K. T., Van Gael, J., and Teh, Y. W. Variational inference for the Indian buffet process. In *13th Int'l Conf. on AISTATS*, pp. 137–144, 2009.
- Feige, U., Mirrokni, V.S., and Vondrak, J. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.
- Fujishige, S. *Submodular functions and optimization*, volume 58. Elsevier Science Limited, 2005.
- Ghahramani, Z. and Beal, M.J. Propagation algorithms for variational bayesian learning. In *Advances in Neural Information Processing Systems*, volume 13, 2001.
- Griffiths, T. L. and Ghahramani, Z. Infinite latent feature models and the Indian buffet process. In *Advances in Neural Information Processing Systems*, volume 18, 2006.
- Knowles, D. and Ghahramani, Z. Infinite sparse factor analysis and infinite independent components analysis. *Independent Component Analysis and Signal Separation*, pp. 381–388, 2007.
- Kollar, T. and Roy, N. Utilizing object-object and object-scene context when planning to find things. In *IEEE International Conference on Robotics and Automation*, pp. 2168–2173, 2009.
- Kurihara, K. and Welling, M. Bayesian k-means as a maximization-expectation algorithm. *Neural Computation*, 21(4):1145–1172, 2008.
- Lee, K.C., Ho, J., and Kriegman, D. Acquiring linear subspaces for face recognition under variable lighting. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 27(5):684–698, 2005.
- Lovász, L. Submodular functions and convexity. *Mathematical programming: the state of the art*, pp. 235–257, 1983.
- Poliner, G.E. and Ellis, D.P.W. A discriminative model for polyphonic piano transcription. *EURASIP Journal on Advances in Signal Processing*, 2006.
- Rai, P. and Daume III, H. Beam search based map estimates for the Indian buffet process. In *Proceedings of the 28th Annual Int'l Conference on Machine Learning*, 2011.
- Schmidt, M. N., Winther, O., and Hansen, L. K. Bayesian non-negative matrix factorization. In *Int'l Conference on Independent Component Analysis and Signal Separation*, volume 5441 of *Lecture Notes in Computer Science (LNCS)*, pp. 540–547. Springer, 2009.