
Scalable Discrete Sampling as a Multi-Armed Bandit Problem

Yutian Chen

Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK

YUTIAN.CHEN@ENG.CAM.AC.UK

Zoubin Ghahramani

Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK
Alan Turing Institute, 96 Euston Road, London NW1 2DB, UK

ZOUBIN@ENG.CAM.AC.UK

Abstract

Drawing a sample from a discrete distribution is one of the building components for Monte Carlo methods. Like other sampling algorithms, discrete sampling suffers from the high computational burden in large-scale inference problems. We study the problem of sampling a discrete random variable with a high degree of dependency that is typical in large-scale Bayesian inference and graphical models, and propose an efficient approximate solution with a subsampling approach. We make a novel connection between the discrete sampling and Multi-Armed Bandits problems with a finite reward population and provide three algorithms with theoretical guarantees. Empirical evaluations show the robustness and efficiency of the approximate algorithms in both synthetic and real-world large-scale problems.

1. Introduction

Sampling a random variable from a discrete (conditional) distribution is one of the core operations in Monte Carlo methods. It is an ubiquitous and often necessary component for inference algorithms such as Gibbs sampling and particle filtering. Applying discrete sampling for large-scale problems has been a challenging task like other Monte Carlo algorithms due to the high computational burden. Various approaches have been proposed to address different dimensions of “large scales”. For example, distributed algorithms have been used to sample a model with a large number of discrete variables (Newman et al., 2009; Bratires et al., 2010; Wu et al., 2011), smart transition kernels were described for Markov chain Monte Carlo (MCMC) algorithms to sample efficiently a single variable

with a large or even infinite state space (Li et al., 2014; Kalli et al., 2011). This paper is focused on another dimension of the “large-scales” where the variable to sample has a large degree of statistical dependency.

Consider a random variable with a finite domain $X \in \mathcal{X}$ and a distribution in the following form

$$p(X = x) \propto \tilde{p}(X = x), \text{ with } \tilde{p}(X = x) = f_0(x) \prod_{n=1}^N f_n(x), \quad (1)$$

where f_n can be any function of x . Such distribution occurs frequently in machine learning problems. For example, in Bayesian inference for a model with parameter X and N i.i.d. observations $\mathcal{D} = \{\mathbf{y}_n\}_{n=1}^N$, the posterior distribution of X depends on all the observations when sufficient statistics is not available. The unnormalized posterior distribution can be written as $\tilde{p}(X|\mathcal{D}) = p(X) \prod_{i=1}^N p(\mathbf{y}_i|X)$. In undirected graphical model inference problems where a node X_i appears in N potential functions, the distribution of X_i depends on the value of all of the N functions. The unnormalized conditional distribution is $\tilde{p}(X_i|\mathbf{x}_{-i}) = \prod_{n=1}^N \phi_n(X_i, \mathbf{x}_{-i})$, where \mathbf{x}_{-i} denotes the value of all the other nodes in the graph and ϕ_n denotes a potential function that includes X_i in the scope. In this paper we study how to sample a discrete random variable X in a manner that is scalable in N .

A common approach to address the big data problem is divide-and-conquer that uses parallel or distributed computing resources to process data in parallel and then synchronize the results periodically or merely once in the end (Scott et al., 2013; Medlar et al., 2013; Xu et al., 2014).

An orthogonal approach has been studied for the Metropolis-Hastings (MH) algorithm in a general state space by running a sampler with subsampled data. This approach can be combined easily with the distributed computing idea for even better scalability (e.g. Ahn et al., 2015).

Maclaurin & Adams (2015) introduced an MH algorithm

in an augmented state space that could achieve higher efficiency than the standard MH by processing only a subset of active data every iteration while still preserving the correct stationary distribution. But the introduction of auxiliary variables might also slow down the overall mixing rate in the augmented space.

Approximate MH algorithms have been proposed in the subsampling approach with high scalability. The stochastic gradient Langevin dynamics (SGLD) (Welling & Teh, 2011) and its extensions (Ahn et al., 2012; Chen et al., 2014; Ding et al., 2014) introduced efficient proposal distributions based on subsampled data. Approximate algorithms induce bias in the stationary distribution of the Markov chain. But given a fixed amount of runtime they could reduce the expected error in the Monte Carlo estimate via a proper trade-off between variance and bias by mixing faster w.r.t. the runtime. This is particularly important for large-scale learning problems when the runtime is one of the limiting factors for generalization performance (Bottou & Bousquet, 2008). However, the stochastic gradient MCMC approach usually skips the rejection step in order to obtain sublinear time complexity and the induced bias is very hard to estimate or control.

Another line of research on approximate subsampled MH algorithms does not ignore the rejection step but controls the error with an approximate rejection step based on a subset of data (Korattikara et al., 2014; Bardenet et al., 2014). The bias can thus be better controlled (Mitrophanov, 2005; Pillai & Smith, 2014). That idea has also been extended to slice sampling (DuBois et al., 2014) and Gibbs for binary variables (Korattikara et al., 2014).

In this paper we follow the last line of research and propose a novel approximate sampling algorithm to improve the scalability of sampling *discrete* distributions. We first reformulate the problem in Eq. 1 as a Multi-Armed Bandit (MAB) problem with a finite reward population via the Gumbel-Max trick (Papandreou & Yuille, 2011), and then propose three algorithms with theoretical guarantees on the approximation error and an upper bound of $N|\mathcal{X}|$ on the sample size. This is to our knowledge the first attempt to address discrete sampling problem with a large number of dependencies and our work will likely contribute to a more complete library of scalable MCMC algorithms. Moreover, the racing algorithm in Sec. 3.3 provides a unified framework for subsampling-based discrete sampling, MH (Korattikara et al., 2014; Bardenet et al., 2014) and slice sampling (DuBois et al., 2014) algorithms as discussed in Sec. 4. We also show in the experiments that our algorithm can be combined straightforwardly with stochastic gradient MCMC to achieve both high efficiency and controlled bias. Lastly, the proposed algorithms also deserve their own interest for MAB problems under this particular setting.

We first review an alternative way of drawing discrete variables and build a connection with MABs in Sec. 2, then propose three algorithms in Sec. 3. We discuss related work in Sec. 4 and evaluate the proposed algorithms on both synthetic data and real-world problems of Bayesian inference and graphical model inference in Sec. 5. Particularly, we show how our proposed sampler can be combined conveniently as a building component with other subsampling sampler for a hierarchical Bayesian model. Sec. 6 concludes the paper with a discussion.

2. Approximate Discrete Sampling

2.1. Discrete Sampling as an Optimization Problem

The common procedure to sample X from a discrete domain $\mathcal{X} = \{1, 2, \dots, D\}$ is to first normalize $\tilde{p}(X)$ and compute the CDF $F(X = x) = \sum_{i=1}^x p(X = i)$. Then draw a uniform random variable $u \sim \text{Uniform}(0, 1]$, and find x that satisfies $F(x - 1) < u \leq F(x)$. This procedure requires computing the sum of all the unnormalized probabilities. For \tilde{p} in the form of Eq. 1 this is $\mathcal{O}(ND)$.

An alternative procedure is to first draw D i.i.d. samples from the standard Gumbel distribution¹ $\varepsilon_i \sim \text{Gumbel}(0, 1)$, and then solve the following optimization problem:

$$x = \operatorname{argmax}_{i \in \mathcal{X}} \log \tilde{p}(i) + \varepsilon_i. \quad (2)$$

It is shown in Kuzmin & Warmuth (2005) that x follows the distribution $p(X)$. With this method after drawing random variables that do not depend on \tilde{p} , we turn a random sampling problem to an optimization problem. While the computational complexity is the same to draw an *exact* sample, an *approximate* algorithm may potentially save computations by avoiding computing accurate values of $\tilde{p}(X = x)$ when x is considered unlikely to be the maximum as discussed next.

2.2. Approximate Discrete Sampling as a Multi-Armed Bandits Problem

In a Multi-Armed Bandit (MAB) problem, the i 'th bandit is a slot machine with an arm, which when pulled generates an i.i.d. reward l_i from a distribution associated with that arm with an unknown mean μ_i . The optimal arm identification problem for MABs (Bechhofer, 1958; Paulson, 1964) in the fixed confidence setting is to find the arm with the highest mean reward with a confidence $1 - \delta$ using as few pulls as possible.

Under the assumption of Eq. 1, the solution in Eq. 2 can be

¹The Gumbel distribution is used to model the maximum extreme value distribution. If a random variable $Z \sim \text{Exp}(1)$, then $-\log(Z) \sim \text{Gumbel}(0, 1)$. ε can be easily drawn as $-\log(-\log(u))$ with $u \sim \text{U}[0, 1]$.

expressed as

$$\begin{aligned}
 x &= \operatorname{argmax}_{i \in \mathcal{X}} \sum_{n=1}^N \log f_n(i) + \log f_0(i) + \varepsilon_i \\
 &= \operatorname{argmax}_{i \in \mathcal{X}} \sum_{n=1}^N \underbrace{\left(\log f_n(i) + \frac{1}{N} (\log f_0(i) + \varepsilon_i) \right)}_{\stackrel{\text{def}}{=} l_{i,n}} \\
 &= \operatorname{argmax}_{i \in \mathcal{X}} \frac{1}{N} \sum_{n=1}^N l_{i,n} = \operatorname{argmax}_{i \in \mathcal{X}} \mathbb{E}_{l_i \sim \text{Uniform}(\mathcal{L}_i)} [l_i] \\
 &\stackrel{\text{def}}{=} \operatorname{argmax}_{i \in \mathcal{X}} \mu_i
 \end{aligned} \tag{3}$$

where $\mathcal{L}_i \stackrel{\text{def}}{=} \{l_{i,1}, l_{i,2}, \dots, l_{i,N}\}$. After drawing D Gumbel variables ε_i , we turn the discrete sampling problem into the optimal arm identification problem in MABs where the reward l_i is uniformly sampled from a finite population \mathcal{L}_i . An approximate algorithm that solves the problem with a fixed confidence may avoid drawing all the rewards from an obviously sub-optimal arm and save computations. We show the induced bias in the sample distribution as follows with the proof in Appx. A.1.

Proposition 1. *If an algorithm solves (2) exactly with a probability at least $1 - \delta$ for any value of ε , the total variation between the sample distribution \hat{p} and the true distribution is bounded by*

$$\|\hat{p}(X) - p(X)\|_{\text{TV}} \leq \delta \tag{4}$$

When applied in the MCMC framework as a transition kernel, we can apply immediately the theories in Mitrophanov (2005); Pillai & Smith (2014) to show that the approximate Markov chain satisfies uniform ergodicity under regular conditions and the analysis of convergence rate are readily available under various assumptions. So the discrete sampling problem of this paper reduces to finding a good MAB algorithm for Eq. 2 in our problem setting.

3. Algorithms for MABs with a Finite Population and Fixed Confidence

The key difference of our problem from the regular MABs is that our rewards are generated from a finite population while regular MABs assume i.i.d. rewards. Because one can obtain the exact mean by sampling all the N values $l_{i,n}$ for arm i *without replacement*, a good algorithm should pull no more than N times for each arm regardless of the mean gap between arms. We introduce three algorithms in this section whose sample complexity is upper bounded by $O(ND)$ in the worst case and can be very efficient when the mean gap is large.

3.1. Notations

The iteration of an algorithm is indexed by t . We denote the entire index set with $[N] = \{1, 2, \dots, N\}$, the sampled set of reward indices up to t 'th iteration from arm i with $\mathcal{N}_i^{(t)} \subseteq [N]$, and the corresponding number of sampled rewards with $T_i^{(t)}$. We define the estimated mean for i 'th arm with $\hat{\mu}_i^{(t)} \stackrel{\text{def}}{=} \frac{1}{|\mathcal{N}_i^{(t)}|} \sum_{n \in \mathcal{N}_i^{(t)}} l_{i,n}$, the natural variance (biased) estimate with $(\hat{\sigma}_i^{(t)})^2 \stackrel{\text{def}}{=} \frac{1}{|\mathcal{N}_i^{(t)}|} \sum_{n \in \mathcal{N}_i^{(t)}} (l_{i,n} - \hat{\mu}_i^{(t)})^2$, the variance estimate of the mean gap between two arm with $(\hat{\sigma}_{i,j}^{(t)})^2 \stackrel{\text{def}}{=} \frac{1}{|\mathcal{N}_i^{(t)}|} \sum_{n \in \mathcal{N}_i^{(t)}} ((l_{i,n} - l_{j,n}) - (\hat{\mu}_i^{(t)} - \hat{\mu}_j^{(t)}))^2$ (defined only when $\mathcal{N}_i^{(t)} = \mathcal{N}_j^{(t)}$), the bound of the reward value $C_i \stackrel{\text{def}}{=} \max_{n,n'} \{l_{i,n} - l_{i,n'}\}$. The subscripts and superscripts may be dropped for notational simplicity when the meaning is clear from the context.

3.2. Adapted lil'UCB

We first study one of the state-of-the-art algorithms for fixed-confidence optimal arm identification problem and adjust it for the finite population setting. The lil'UCB algorithm (Jamieson et al., 2014) maintains an upper confidence bound (UCB) of μ_i that is inspired by the law of the iterated logarithm (LIL) for every arm. At each iteration, it draws a single sample from the arm with the highest bound and updates it. The algorithm terminates when some arm is sampled much more often than all the other arms. We refer readers to Fig. 1 of Jamieson et al. (2014) for details. The time complexity for t iterations is $\mathcal{O}(\log(D)t)$. It was shown in Jamieson et al. (2014) that lil'UCB achieved the optimal sample complexity up to constants.

However, lil'UCB requires i.i.d. rewards for each arm i , that is, sampled with replacement from \mathcal{L}_i . Therefore, the total number of samples t is unbounded and could be $\gg ND$ when the means are close to each other. We adapt lil'UCB for our problem with the following modifications:

1. Samples $l_{i,n}$ *without replacement* for each arm but keep different arms independent.
2. When $T_i^{(t)} = N$ for some arm i , the estimate $\hat{\mu}_i^{(t)}$ becomes exact. So set its UCB to $\hat{\mu}_i^{(t)}$.
3. The algorithm terminates either with the original stopping criterion or when the arm with the highest upper bound has an exact mean estimate, whichever comes first.

The adapted algorithm satisfies all the theoretical guarantees in Thm. 2 of Jamieson et al. (2014) with additional properties as shown in the following proposition with proof in Appx. A.2.

Proposition 2. *Theorem 2 of Jamieson et al. (2014) holds for the adapted lil’UCB algorithm. Moreover $T_i^{(t)} \leq N, \forall i, t$. Therefore, when the algorithm terminates, $t = \sum_{i \in \mathcal{X}} T_i^{(t)} \leq ND$.*

Notice that Thm. 2 of Jamieson et al. (2014) shows that t scales roughly as $O(1/\Delta^2)$ with Δ being the mean gap and therefore $t \ll ND$ when the gap is large.

3.3. Racing Algorithm for a Finite Population

When rewards are sampled without replacement, the negative correlation between rewards would generally improve the convergence of $\hat{\mu}_i$. Unfortunately, the bound in lil’UCB ignores the negative correlation when $T_i^{(t)} < N$ even with the adaptations. We introduce a new family of racing algorithms (Maron & Moore, 1994) that takes advantage of the finite population setting as shown in Alg. 1. The choice of the uncertainty bound function G differentiates specific algorithms and two examples will be discussed in the following sections.

Alg. 1 maintains a set of candidate set \mathcal{D} initialized with all arms. At iteration t , a shared mini-batch of $m^{(t)}$ indices are drawn w/o replacement for all survived arms in \mathcal{D} . Then the uncertainty bound G is used to eliminate sub-optimal arms with a given confidence. The algorithm stops when only one arm remains. We require for $m^{(t)}$ that the total number of sampled indices $T^{(t^*)} = \sum_{t=1}^{t^*} m^{(t)}$ equals N at the last iteration t^* . Particularly, we take a doubling schedule $T^{(t)} = 2T^{(t-1)}$ (so $t^* = \lceil \log_2 \frac{N}{m^{(1)}} \rceil + 1$) and leave $m^{(1)}$ as a free parameter. We also require $G(\cdot, T, \cdot, \cdot) = 0$ whenever $T = N$ so that Alg. 1 always stops within t^* iterations. The computational complexity for t iterations is $\mathcal{O}(DT^{(t)})$ with the marginal estimate $\hat{\sigma}_i$ and $\mathcal{O}(D^2T^{(t)})$ with the pairwise estimate $\hat{\sigma}_{i,j}$. The former version is more efficient than the latter when D is large at the price of a looser bound.

Proposition 3. *If G satisfies*

$$\mathcal{E} \stackrel{\text{def}}{=} P(\exists t < t^*, \hat{\mu}^{(t)} - \mu > G(\delta, T^{(t)}, \hat{\sigma}^{(t)}, C)) \leq \delta, \quad (5)$$

for any $\delta \in (0, 1)$ with a probability at least $1 - \delta$, Alg. 1 returns the optimal arm with at most ND samples.

The proof is provided in Appx. A.3. Unlike adapted lil’UCB, Racing draws a shared set of sample indices among all the arms and could provide a tighter bound with pairwise variance estimates $\hat{\sigma}_{i,j}$ when there is positive correlation, a typical case in Bayesian inference problems.

3.3.1. RACING WITH SERFLING CONCENTRATION BOUNDS FOR G

Serfling (1974) studied the concentration inequalities of sampling without replacement and obtained an improved

Algorithm 1 Racing Algorithm with a Finite Reward Population

input Number of arms D , population size N , mini-batch sizes $\{m^{(t)}\}_{t=1}^{t^*}$, confidence level $1 - \delta$, uncertainty bound function $G(\delta, T, \hat{\sigma}, C)$, range of samples C_i (optional).

$t \leftarrow 0, T \leftarrow 0, \mathcal{D} \leftarrow \{1, 2, \dots, D\}, \mathcal{N} \leftarrow \emptyset$

while $|\mathcal{D}| > 1$ **do**

$t \leftarrow t + 1$

Sample w/o replacement $m^{(t)}$ indices $\mathcal{M} \subseteq [N] \setminus \mathcal{N}$, and set $\mathcal{N} \leftarrow \mathcal{N} \cup \mathcal{M}, T \leftarrow T + m^{(t)}$

Compute $l_{i,n}, \forall i \in \mathcal{D}, n \in \mathcal{M}$, and update $\hat{\mu}_i$ and $\hat{\sigma}_i$ (or $\hat{\sigma}_{i,j}$), $\forall i \in \mathcal{D}$.

Find the best arm $x \leftarrow \operatorname{argmax}_{i \in \mathcal{D}} \hat{\mu}_i$

Eliminate sub-optimal arms when the estimated gap is large $\mathcal{D} \leftarrow \mathcal{D} \setminus \{i : \hat{\mu}_x - \hat{\mu}_i > G(\frac{\delta}{D}, T, \hat{\sigma}_x, C_x) + G(\frac{\delta}{D}, T, \hat{\sigma}_i, C_i)\}$ (or $\mathcal{D} \leftarrow \mathcal{D} \setminus \{i : \hat{\mu}_x - \hat{\mu}_i > G(\frac{\delta}{D-1}, T, \hat{\sigma}_{x,i}, C_x + C_i)\}$)

end while

output \mathcal{D}

Hoeffding bound. Bardenet & Maillard (2013) extended the work and provided an empirical Bernstein-Serfling bound that was later used for the subsampling-based MH algorithm (Bardenet et al., 2014): for any $\delta \in (0, 1]$ and any $n \leq N$, with probability $1 - \delta$, it holds that

$$\hat{\mu}_n - \mu \leq \hat{\sigma}_n \sqrt{\frac{2\rho_n \log(5/\delta)}{n}} + \frac{\kappa C \log(5/\delta)}{n} \stackrel{\text{def}}{=} B_{\text{EBS}}(\delta, n, \hat{\sigma}_n, C) \quad (6)$$

where $\kappa = \frac{7}{3} + \frac{3}{\sqrt{2}}$, and $\rho_n = \begin{cases} 1 - \pi_{n-1} & \text{if } n \leq N/2 \\ (1 - \pi_n)(1 + \frac{1}{n}) & \text{if } n > N/2 \end{cases}$, with $\pi_n \stackrel{\text{def}}{=} \frac{n}{N}$. The extra term ρ_n that is missing in regular empirical Bernstein bounds reduces the bound significantly when n is close to N . We set $m^{(1)} = 2$ in Alg. 1 to provide a valid $\hat{\sigma}^{(t)}$ for any t and set the uncertain bound G with the empirical Bernstein-Serfling (EBS) bounds as

$$G_{\text{EBS}}(\delta, T, \hat{\sigma}, C) = B_{\text{EBS}}\left(\frac{\delta}{t^* - 1}, T, \hat{\sigma}, C\right) \quad (7)$$

It is trivial to prove that G_{EBS} satisfies the condition in Eq. 5 using a union bound over $t < t^*$.

3.3.2. RACING WITH A NORMAL ASSUMPTION FOR G

The concentration bounds often give a conservative strategy as they assume an arbitrary bounded reward distribution. When the number of drawn samples is large, the central limit theorem suggests that $\hat{\mu}^{(t)}$ follows approximately a Gaussian distribution. Korattikara et al. (2014) made such

an assumption and obtained a tighter bound. We first provide an immediate corollary of Prop. 2 in Appx. A of [Korattikara et al. \(2014\)](#).

Corollary 4. *Let $\hat{\mu}_{\text{unit}}^{(t)}$, $t = 1, 2, \dots, t^*$ be the estimated means using sampling without replacement from any finite population with mean μ and unit variance. The joint normal random variables $\tilde{\mu}^{(t)}$ that match the mean and covariance matrix with $\hat{\mu}_{\text{unit}}^{(t)}$ follow a Gaussian random walk process as*

$$p_{\mu}(\tilde{\mu}^{(t)} | \tilde{\mu}^{(1)}, \dots, \tilde{\mu}^{(t-1)}) = \mathcal{N}(m_t(\tilde{\mu}^{(t-1)}), S_t) \quad (8)$$

where $m_t = \mu + A_t(\tilde{\mu}_{t-1} - \mu)$, $S_t = \frac{B_t}{T^{(t)}} \left(1 - \frac{T^{(t)} - 1}{N-1}\right)$, $A_t = \frac{\pi_{t-1}(1-\pi_t)}{\pi_t(1-\pi_{t-1})}$, $B_t = \frac{\pi_t - \pi_{t-1}}{\pi_t(1-\pi_{t-1})}$ with π_t short for $\pi_{T^{(t)}}$.

Remark 5. *The marginal distribution $p(\tilde{\mu}^{(t)}) = \mathcal{N}\left(\mu, \frac{1}{T^{(t)}} \left(1 - \frac{T^{(t)} - 1}{N-1}\right)\right)$ where the variance approaches 0 when $T^{(t)} \rightarrow N$.*

Assumption 6. *When $T^{(t)} \gg 1, \forall t$, we assume $\hat{\sigma}^{(t)} \approx \sigma$ and the central limit theorem suggests that the joint distribution of $\hat{\mu}^{(t)}/\hat{\sigma}^{(t)}$ can be approximated by the joint distribution of $\tilde{\mu}^{(t)}$.*

With the normal assumption, we choose the uncertainty bound G in the following form

$$G_{\text{Normal}}(\delta, T, \hat{\sigma}) = \frac{\hat{\sigma}}{\sqrt{T}} \left(1 - \frac{T-1}{N-1}\right)^{1/2} B_{\text{Normal}} \quad (9)$$

Intuitively we use a constant confidence level, $\Phi(B_{\text{Normal}})$, for all marginal distributions of $\hat{\mu}^{(t)}$ over t where $\Phi(\cdot)$ is the CDF of the standard normal. To choose the constant B_{Normal} , we plug G_{Normal} into the condition for G in Eq. 5 and apply the normal distribution (8) to solve the univariate equation $\mathcal{E}(B) = \delta$. This way of computing G gives a tighter bound than applying the union bound across t as in the previous section because it takes into account the correlation of mean estimates across iterations. Appx. B provides a lookup table and a plot of $B_{\text{Normal}}(\delta) = \mathcal{E}^{-1}(\delta)$. Notice that B_{Normal} only needs to be computed once and we can obtain it for any δ by either interpolating the table or computing numerically with code to be shared (runtime < 1 second). For the parameter of the first mini-batch size $m^{(1)}$, a value of 50 performs robustly in all experiments.

We provide the sample complexity below with the proof in Appx. A.4. Particularly, $T^*(\Delta) \rightarrow DN$ as $\Delta \rightarrow 0$, and $T^*(\Delta) = Dm^{(1)}$ when $\Delta \geq 2B_{\text{Normal}}(\delta/D')\sqrt{(N/m^{(1)} - 1)/(N-1)}$.

Proposition 7. *Let x^* be the best arm and Δ be the minimal normalized gap of means from other arms, defined as $\min_{i \neq x^*} \frac{\mu_{x^*} - \mu_i}{\sigma_{x^*} + \sigma_i}$ when using marginal variance estimate $\hat{\sigma}_i$ and $\min_{i \neq x^*} \frac{\mu_{x^*} - \mu_i}{\sigma_{x^*, i}}$ when using pairwise variance estimate $\hat{\sigma}_{x, i}$. If Assump. 6 holds, with a probability at least*

$1 - \delta$ *Racing-Normal draws no more rewards than*

$$T^*(\Delta) = D \left[\frac{N}{(N-1) \frac{\Delta^2}{4B_{\text{Normal}}^2(\delta/D')} + 1} \right]_{m^{(1)}} \quad (10)$$

where $\lceil n \rceil_m \stackrel{\text{def}}{=} m2^{\lceil \log_2 n/m \rceil} \wedge N \geq n, \forall n \leq N$. $D' \stackrel{\text{def}}{=} D$ if using $\hat{\sigma}_i$ and is $D-1$ if using $\hat{\sigma}_{x, i}$.

3.4. Variance Reduction for Random Rewards with Control Variates

The difficulty of MABs depends heavily on the ratio of the mean gap to the reward noise, Δ . To improve the signal noise ratio, we exploit the control variates technique ([Wilson, 1984](#)) to reduce the reward variance. Consider a variable $h_{i, n}$ whose expectation $\mathbb{E}_{n \sim [N]}[h_{i, n}]$ can be computed efficiently. The residue reward $l_{i, n} - h_{i, n} + \mathbb{E}_n[h_{i, n}]$ has the same mean as $l_{i, n}$ and the variance is reduced if $h_{i, n} \approx l_{i, n}$. In the Bayesian inference experiment where the factor $f_n(X = i) = p(\mathbf{y}_n | X = i)$, we adopt a similar approach as [Wang et al. \(2013\)](#) and take the Taylor expansion of $l_{i, n}$ around a reference point $\hat{\mathbf{y}}$ as

$$l_{i, n} \approx l_i(\hat{\mathbf{y}}) + \mathbf{g}_i^T(\mathbf{y}_n - \hat{\mathbf{y}}) + \frac{1}{2}(\mathbf{y}_n - \hat{\mathbf{y}})^T H_i(\mathbf{y}_n - \hat{\mathbf{y}}) \stackrel{\text{def}}{=} h_{i, n} \quad (11)$$

where \mathbf{g}_i and H_i are the gradient and Hessian matrix of $\log p(\mathbf{y} | i)$ respectively evaluated at $\hat{\mathbf{y}}$. $\mathbb{E}[h_{i, n}]$ can be computed analytically with the first two moments of \mathbf{y}_n . A typical choice of $\hat{\mathbf{y}}$ is $\mathbb{E}[\mathbf{y}]$.

The control variate method is mostly useful for Racing-Normal. For algorithms depending on a reward bound C in order to get a tight bound for $l_{i, n} - h_{i, n}$ it requires a more restrictive condition for C as in [Bardenet et al. \(2015\)](#) and we might end up with an even more conservative strategy in general cases.

4. Related Work

The Gumbel-Max trick has been exploited in [Kuzmin & Warmuth \(2005\)](#); [Papandreou & Yuille \(2011\)](#); [Maddison et al. \(2014\)](#) for different problems. The closest work is [Maddison et al. \(2014\)](#) where this trick is extended to draw continuous random variables with a Gumbel process, reminiscent to adaptive rejection sampling.

Our work is closely related to the optimal arm identification problem for MABs with a fixed confidence. This is, to our knowledge, the first work to consider MABs with a finite population. The proposed algorithms tailored under this setting could be of interest beyond the discrete sampling problem. The normal assumption in Sec. 3.3.2 is similar to UCB-Normal in [Auer et al. \(2002\)](#) but the latter assumes a normal distribution for *individual* rewards and will perform poorly when it does not hold.

The bounds in Sec. 3.3 are based on subsampling-based MH algorithms in [Bardenet et al. \(2014\)](#); [Korattikara et al. \(2014\)](#). The proposed algorithm extends those ideas from MH to discrete sampling. In fact, let x and x' be the current and proposed value in an MH iteration, Racing-EBS and Racing-Normal reduce to the algorithms in [Bardenet et al. \(2014\)](#) and [Korattikara et al. \(2014\)](#) respectively if we set

$$\begin{aligned} \mathcal{X} &= \{x, x'\}, & f_0(1) &= u p(x)q(x'|x), \\ f_0(2) &= p(x')q(x|x'), & f_n(x) &= p(y_n|x) \end{aligned} \quad (12)$$

where $p(x)$ is the prior distribution, $u \sim \text{Uniform}[0, 1]$ and $q(\cdot|\cdot)$ is the proposal distribution. The difference with [Bardenet et al. \(2014\)](#) is that we distribute the error δ evenly across t in Eq. 7 while [Bardenet et al. \(2014\)](#) set $\delta_t = (p - 1)/(p(T^{(t)})^p)\delta$ with p a free parameter. The differences with [Korattikara et al. \(2014\)](#) are that we take a doubling schedule for $m^{(t)}$ and replace the t-test with the normal assumption. We find that our algorithms are more efficient and robust than both original algorithms in practice. Moreover, the binary Gibbs sampling in Appx. F of [Korattikara et al. \(2014\)](#) is also a special case of Racing-Normal with $D = 2$. Therefore, Alg. 1 provides a unifying approach to a family of subsampling-based samplers.

The variance reduction technique is similar to the proxies in [Bardenet et al. \(2015\)](#), but the control variate here is a function in the data space while the proxy in the latter is a function in the parameter space. We do not assume the posterior distribution is approximate Gaussian and our algorithm works with multi-modal distributions.

It is important not to confuse the focus of our algorithm for the big N problem in Eq. 1 with other algorithms that address sampling for a large state space (big D) or similarly a high-dimensional vector of discrete variables (exponentially large D). The combination of these two approaches for problems with both big N and big D is possible but beyond the scope of this paper.

5. Experiments

Since this is the first work to discuss efficient discrete sampling for problem (1), we compare the adapted lil'UCB, Racing-EBS, Racing-Normal with the exact sampler only. We report the result of Racing-Normal in real data experiments only as the speed gains of the other two are marginal.

5.1. Synthetic Data

We construct a distribution with $D = 10$ by sampling $N = 10^5$ rewards of $l_{i,n}$ for each state from one of the three distributions $\mathcal{N}(0, 1)$, $\text{Uniform}[0, 1]$, $\text{LogNormal}(0, 2)$. We normalized $l_{i,n}$ to have a fixed distribution $p(X)$ in Fig. 1(a) and a reward variance σ^2 that controls the difficulty. The normal distribution is the ideal setting for

Racing-Normal, and the uniform distribution is desirable for adapted lil'UCB and Racing-EBS as the reward bound is close to σ . The LogNormal distribution, whose ex. kurtosis ≈ 4000 , is difficult for all due to the heavy tail. We use a tight bound $C = \max\{l_{i,n} - l_{i,n'}\}$ for Racing-EBS. We set the scale parameter of adapted lil'UCB with $C/2$ and other parameters with the heuristic setting in [Jamieson et al. \(2014\)](#). Racing uses the pairwise variance estimate.

Fig. 1(b)-(d) show the empirical error of best arm identification by drawing 10^4 samples of X for each setting and vary the target error bound $\delta \in [10^{-3}, 0.1]$. The bound appears very loose for lil'UCB and Racing-EBS but is sharp for Racing-Normal when the noise is large (1(b)) and $\delta \ll 1$. This is consistent with the direct comparison of uncertainty bounds in Fig. 1(e). Consequently, given the same error tolerance δ Racing-Normal requires much fewer rewards than the other conservative strategies in all the settings except when $\sigma = 10^{-5}$ and $l_{i,n} \sim \text{Uniform}[0, 1]$, as shown in Fig. 1(f)-(h). We verify the observations with more experiments in Appx. C.1 with $D \in \{2, 100\}$ and marginal estimate $\hat{\sigma}_i$.

Surprisingly, Racing-Normal performs robustly regardless of reward distributions with the first mini-batch size $m^{(1)} = 50$ while it was shown in [Bardenet et al. \(2014\)](#) that the algorithm with the same normal assumption in [Korattikara et al. \(2014\)](#) failed with LogNormal even when $m^{(1)} = 500$. The dramatic improvement in robustness is mainly due to our doubling scheme where central limit theorem applies quickly with $m^{(t)}$ increasing exponentially. We do not claim that the single trick will solve the problem completely because there still exist cases in theory with extremely heavy-tailed reward distributions where our normal assumption does not hold and the algorithm will fail to meet the confidence level. In practice, we do not observe that pathological case in any of the experiments.

5.2. Bayesian ARCH Model Selection

We evaluate Racing-Normal in a Bayesian model selection problem for the auto-regressive conditional heteroskedasticity (ARCH) models. The discrete sampler is integrated in the Markov chain as a building component to sample the hierarchical model. Specifically, we consider a mixture of ARCHs for the return r_t of stock price series with student-t innovations, each component with a different order q :

$$\begin{aligned} r_t &= \sigma_t z_t, & z_t &\stackrel{iid}{\sim} t_\nu(0, 1), & \sigma_t^2 &= \alpha_0 + \sum_{i=1}^q \alpha_i r_{t-i}^2, \\ q &\sim \text{Discrete}(\boldsymbol{\pi}), & \alpha_i, \nu &\stackrel{iid}{\sim} \text{Gamma}(1, 1) \end{aligned}$$

where $\boldsymbol{\pi} = \{\pi_q : q \in \mathbb{Q}\}$ is the prior distribution of a candidate model in the set \mathbb{Q} . The random variables to infer include the discrete model choice q and continuous param-

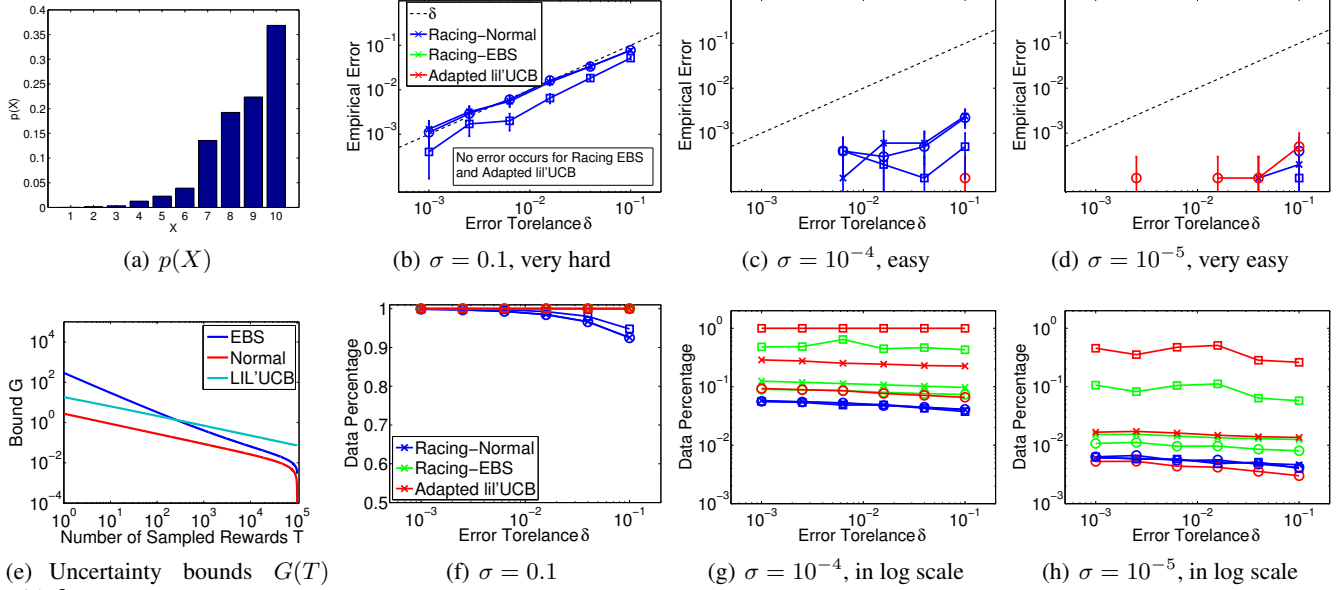


Figure 1. Synthetic data. ((b),(c),(d)) Estimated error with 95% confidence interval. **Plots not shown if no error occurred.** ((f),(g),(h)) proportion of sampled rewards. $l_{i,n}$ is sampled from Normal (\times), Uniform (\circ) and LogNormal (\square) distributions. Plots of Racing-Normal overlap in ((f),(g),(h)).

eters $\{\alpha_i\}_{i=0}^q, \nu$. We adopt the augmented MCMC algorithm in Carlin & Chib (1995) to avoid transdimensional moves. We apply subsampling-based scalable algorithms to sample all variables with subsampled observations $\{r_t\}$: Racing-Normal Gibbs for q , stochastic gradient Langevin dynamics (SGLD) (Welling & Teh, 2011) corrected with Racing-Normal MH (Sec. 4) for α_i and ν . We use adjusted priors $\tilde{\pi}_q$ as suggested by Carlin & Chib (1995) for sufficient mixing between all models and tune them with adaptive MCMC. The adjusted posterior $\tilde{p}(q|\mathbf{r}) \propto \tilde{\pi}_q p(\mathbf{r}|q)$ is then close to uniform and the value $\pi_q/\tilde{\pi}_q$ provides an estimate to the real unnormalized posterior $p(q|\mathbf{r})$. Control variates are also applied to reduce variance. Details of the sampling algorithm are provided in Appx. C.2.

We apply the model on the 5-minute Shanghai stock exchange composite index of one year consisting of about 13,000 data points (Fig. 2(a)). $\mathbb{Q} = \{5, 10, 15, 20, 25, 30\}$. We set $m^{(1)} = 50$ and $\delta = 0.05$. The control variate method reduces the reward variance by 2~3 orders of magnitude. Fig. 2(b) shows the estimated log-posterior of q by normalizing $\pi_q/\tilde{\pi}_q$ in the adaptive MCMC as a function of the number of likelihood evaluations (proportional to runtime). The subsampling-based sampler (Sub) converges about three times faster. We then fix $\tilde{\pi}_q$ for a fixed stationary distribution and run MCMC for 10^5 iterations to compare Sub with the exact sampler. The empirical error rates for Racing-Normal Gibbs and MH are about 4×10^{-4} and 2×10^{-3} respectively. Fig. 2(c) shows estimated adjusted posterior with 5 runs, and (d) compares the auto-correlation of sample q . Sub obtains over twice the effective sample

size without noticeable bias after the burn-in period.

5.3. Author Coreference

We then study the performance in a large-scale graphical model inference problem. The author coreference problem for a database of scientific paper citations is to cluster the mentions of authors into real persons. Singh et al. (2012) addressed this problem with a conditional random field model with pairwise factors. The joint and conditional distributions are respectively

$$p_{\theta}(\mathbf{y}|\mathbf{x}) \propto \exp \left(\sum_{y_i=y_j, i \neq j, \forall i,j} f_{\theta}(x_i, x_j) \right),$$

$$p_{\theta}(Y_i = y_i | \mathbf{y}_{-i}, \mathbf{x}) \propto \exp \left(\sum_{y_j \in C_y = \{j: y_j = y, j \neq i\}} f_{\theta}(x_i, x_j) \right)$$

where $\mathbf{x} = \{x_i\}_{i=1}^N$ is the set of observed author mentions and $y_i \in \mathbb{N}^+$ is the cluster index for i 'th mention. The factor $f_{\theta}(x_i, x_j)$ measures the similarity between two mentions based on author names, coauthors, paper title, etc, parameterized by θ . In the conditional distribution, y_i can take a value of any non-empty cluster or another empty cluster index. When a cluster C_y contains a lot of mentions, a typical case for common author names, the number of factors to be evaluated $N_y = |C_y|$ will be large. We consider the MAP inference problem with fixed θ using annealed Gibbs sampling (Finkel et al., 2005). We apply Racing-Normal to sample Y_i by subsampling C_y for each

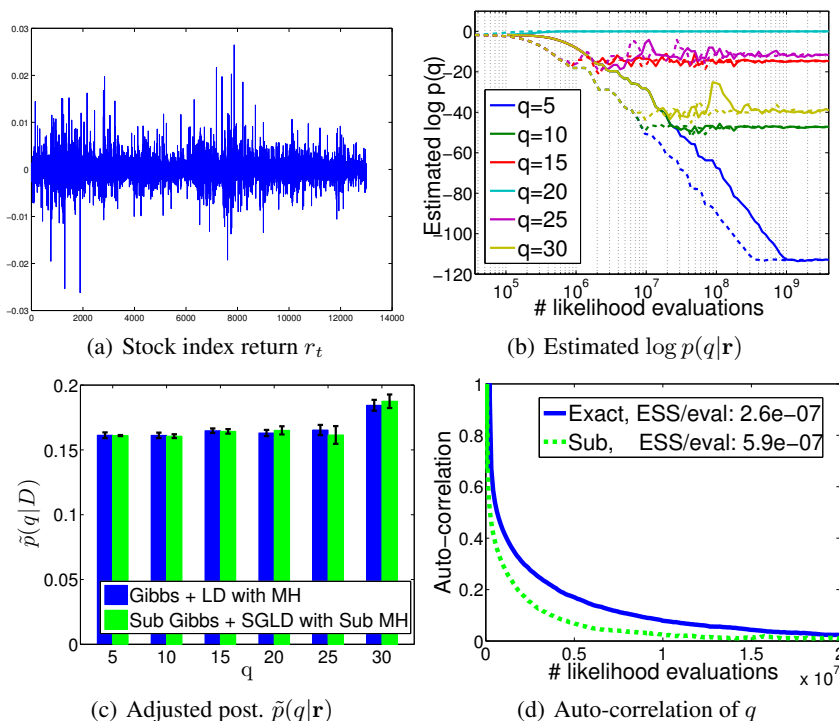


Figure 2. Bayesian ARCH Model Selection. Solid: exact, dashed: approximate using Sub Gibbs + SGLD with Sub MH.

candidate value y . An important difference of this problem from Eq. 1 is that $N_y \neq N_{y'}, \forall y \neq y'$ and N_y has a heavy tail distribution. We let the mini-batch size depend on N_y with details provided in Appx. C.3.

We run the experiment on the union of an unlabeled DBLP dataset of BibTex entries with about 5M authors and a Rexa corpus of about 11K author mentions with 3160 entries labeled. We monitor the clustering performance on the labeled subset with the B^3 F-1 score (Bagga & Baldwin, 1998). We use $\delta = 0.05$ and the empirical error rate is about 0.046. The number of candidate values D varies in $2 \sim 215$ and N_y varies in $1 \sim 1829$ upon convergence. Fig. 3(a) shows the F-1 score as a function of the number of factor evaluations with 7 random runs for each algorithm. Sub Gibbs converges about three times faster than exact Gibbs. Fig. 3(b) shows F-1 as a function of iterations that renders almost identical behavior for both algorithms, which suggests negligible bias in Sub Gibbs. The relative number of the evaluated factors of sub to exact Gibbs indicates about a 5-time speed up near convergence. The initial speed up is small because every cluster is initialized with a single mention, i.e. $N_y = 1$.

6. Discussion

We consider the discrete sampling problem with a high degree of dependency and proposed three approximate al-

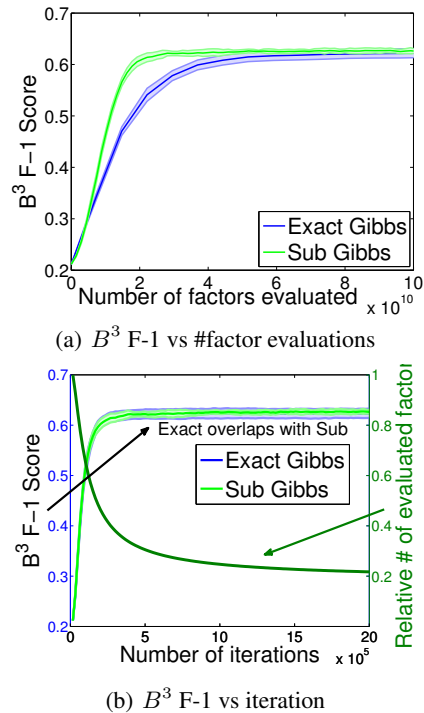


Figure 3. Author Coreference. Bigger B^3 F-1 score is better.

gorithms under the framework of MABs with theoretical guarantees. The Racing algorithm provides a unifying approaches to various subsampling-based Monte Carlo algorithms and also improves the robustness of the original MH algorithm in Korattikara et al. (2014). This is also the first work to discuss MABs under the setting of a finite reward population.

Empirical evaluations show that Racing-Normal achieves a robust and the highest speed-up among all competitors. Whilst adaptive lil’UCB shows inferior empirical performance to Racing-Normal, it has a better sample complexity w.r.t. the number of arms D . It will be a future direction to combine the bound of Racing-Normal with other MAB algorithms including lil’UCB for a better scalability in D . Another important problem is on how to relax the assumptions for Racing-Normal without sacrificing the performance.

It would also be an interesting direction to extend our work to draw continuous random variables efficiently with the Gumbel process (Maddison et al., 2014). In continuous state space, there are infinitely many “arms” and a naive application of our algorithm will lead to infinitely large error bound. This problem can be alleviated with algorithms for contextual MAB problems.

Acknowledgements

We thank Matt Hoffman for helpful discussions on the connection of our work to the MAB problems. We also thank all the reviewers for their constructive comments. We acknowledge funding from the Alan Turing Institute, Google, Microsoft Research and EPSRC Grant EP/I036575/1.

References

- Ahn, Sungjin, Korattikara, Anoop, and Welling, Max. Bayesian posterior sampling via stochastic gradient fisher scoring. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pp. 1591–1598, 2012.
- Ahn, Sungjin, Korattikara, Anoop, Liu, Nathan, Rajan, Suju, and Welling, Max. Large-scale distributed bayesian matrix factorization using stochastic gradient mcmc. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 9–18. ACM, 2015.
- Auer, Peter, Cesa-Bianchi, Nicolo, and Fischer, Paul. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Bagga, Amit and Baldwin, Breck. Algorithms for scoring coreference chains. In *LREC workshop on linguistics coreference*, volume 1, pp. 563–566, 1998.
- Bardenet, Rémi and Maillard, Odalric-Ambrym. Concentration inequalities for sampling without replacement. *arXiv preprint arXiv:1309.4029*, 2013.
- Bardenet, Rémi, Doucet, Arnaud, and Holmes, Chris. Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach. In *Proceedings of The 31st International Conference on Machine Learning*, pp. 405–413, 2014.
- Bardenet, Rémi, Doucet, Arnaud, and Holmes, Chris. On markov chain monte carlo methods for tall data. *arXiv preprint arXiv:1505.02827*, 2015.
- Bechhofer, Robert E. A sequential multiple-decision procedure for selecting the best one of several normal populations with a common unknown variance, and its use with various experimental designs. *Biometrics*, 14(3): 408–429, 1958.
- Bottou, L. and Bousquet, O. The tradeoffs of large scale learning. In *NIPS*, volume 20, pp. 161–168, 2008.
- Bratires, S., van Gael, J., Vlachos, A., and Ghahramani, Z. Scaling the iHMM: Parallelization versus hadoop. In *Computer and Information Technology (CIT), 2010 IEEE 10th International Conference on*, pp. 1235–1240, June 2010.
- Carlin, B.P. and Chib, S. Bayesian model choice via Markov chain Monte Carlo. *Journal of the Royal Statistical Society, Series B*, 57:473–484, 1995.
- Chen, Tianqi, Fox, Emily, and Guestrin, Carlos. Stochastic gradient hamiltonian monte carlo. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pp. 1683–1691, 2014.
- Ding, Nan, Fang, Youhan, Babbush, Ryan, Chen, Changyou, Skeel, Robert D, and Neven, Hartmut. Bayesian sampling using stochastic gradient thermostats. In *Advances in neural information processing systems*, pp. 3203–3211, 2014.
- DuBois, Christopher, Korattikara, Anoop, Welling, Max, and Smyth, Padhraic. Approximate slice sampling for Bayesian posterior inference. In *Proceedings of AIS-TATS*, pp. 185–193, 2014.
- Finkel, Jenny Rose, Grenager, Trond, and Manning, Christopher. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd ACL*, pp. 363–370, 2005.
- Hoeffding, Wassily. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- Jamieson, Kevin, Malloy, Matthew, Nowak, Robert, and Bubeck, Sébastien. lil’UCB: An optimal exploration algorithm for multi-armed bandits. In *Proceedings of The 27th COLT*, pp. 423–439, 2014.
- Kalli, Maria, Griffin, Jim E, and Walker, Stephen G. Slice sampling mixture models. *Statistics and computing*, 21(1):93–105, 2011.
- Korattikara, Anoop, Chen, Yutian, and Welling, Max. Austerity in MCMC land: Cutting the Metropolis-Hastings budget. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pp. 181–189, 2014.
- Kuzmin, Dima and Warmuth, Manfred K. Optimum follow the leader algorithm. In *Proceedings of the 18th annual conference on Learning Theory*, pp. 684–686. Springer-Verlag, 2005.
- Li, Aaron, Ahmed, Amr, Ravi, Sujith, and Smola, Alex. Reducing the sampling complexity of topic models. In *Proceedings of the ACM Conference on Knowledge Discovery and Data Mining (KDD)*, 2014.
- Maclaurin, Dougal and Adams, Ryan Prescott. Firefly monte carlo: Exact mcmc with subsets of data. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

- Maddison, Chris J, Tarlow, Daniel, and Minka, Tom. A* sampling. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., and Weinberger, K.Q. (eds.), *NIPS*, pp. 3086–3094, 2014.
- Maron, Oded and Moore, Andrew W. Hoeffding races: Accelerating model selection search for classification and function approximation. In *Advances in Neural Information Processing Systems 6*, pp. 59–66. Morgan-Kaufmann, 1994.
- Medlar, Alan, Głowacka, Dorota, Stanescu, Horia, Bryson, Kevin, and Kleta, Robert. Swiftlink: parallel mcmc linkage analysis using multicore cpu and gpu. *Bioinformatics*, 29(4):413–419, 2013.
- Mitrophanov, A Yu. Sensitivity and convergence of uniformly ergodic markov chains. *Journal of Applied Probability*, pp. 1003–1014, 2005.
- Newman, David, Asuncion, Arthur, Smyth, Padhraic, and Welling, Max. Distributed algorithms for topic models. *The Journal of Machine Learning Research*, 10:1801–1828, 2009.
- Papandreou, G. and Yuille, A. Perturb-and-MAP random fields: Using discrete optimization to learn and sample from energy models. In *Proceedings of ICCV*, pp. 193–200, Barcelona, Spain, November 2011.
- Paulson, Edward. A sequential procedure for selecting the population with the largest mean from k normal populations. *The Annals of Mathematical Statistics*, pp. 174–180, 1964.
- Pillai, Natesh S and Smith, Aaron. Ergodicity of approximate MCMC chains with applications to large data sets. *arXiv preprint arXiv:1405.0182*, 2014.
- Scott, Steven L, Blocker, Alexander W, Bonassi, Fernando V, Chipman, H, George, E, and McCulloch, R. Bayes and big data: The consensus monte carlo algorithm. *EFaBBayes 250 conference*, 16, 2013.
- Serfling, R. J. Probability inequalities for the sum in sampling without replacement. *Ann. Statist.*, 2(1):39–48, 01 1974.
- Singh, Sameer, Wick, Michael, and McCallum, Andrew. Monte Carlo MCMC: efficient inference by approximate sampling. In *Proceedings of EMNLP-CoNLL 2012*, pp. 1104–1113, 2012.
- Wang, Chong, Chen, Xi, Smola, Alex J, and Xing, Eric P. Variance reduction for stochastic gradient optimization. In *Advances in Neural Information Processing Systems*, pp. 181–189, 2013.
- Welling, Max and Teh, Yee W. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of ICML 2011*, pp. 681–688, 2011.
- Wilson, James R. Variance reduction techniques for digital simulation. *American Journal of Mathematical and Management Sciences*, 4(3-4):277–312, 1984.
- Wu, Yao, Yan, Qiang, Bickson, Danny, Low, Yucheng, and Yang, Qing. Efficient multicore collaborative filtering. In *ACM KDD CUP workshop*, 2011.
- Xu, M., Teh, Y. W., Zhu, J., and Zhang, B. Distributed context-aware bayesian posterior sampling via expectation propagation. In *Advances in Neural Information Processing Systems*, 2014.