
Scaling-up Split-Merge MCMC with Locality Sensitive Sampling (LSS)*

Chen Luo¹ Anshumali Shrivastava¹

Abstract

Split-Merge MCMC (Monte Carlo Markov Chain) is one of the essential and popular variants of MCMC for problems when an MCMC state consists of an unknown number of components or clusters. It is well known that state-of-the-art methods for split-merge MCMC do not scale well. Strategies for rapid mixing requires smart and informative proposals to reduce the rejection rate. However, all known smart proposals involve cost at least linear in the size of the data $\geq O(N)$, to suggest informative transitions. Thus, the cost of each iteration is prohibitive for massive scale datasets. It is further known that uninformative but computationally efficient proposals, such as random split-merge, leads to extremely slow convergence. This tradeoff between mixing time and per update cost seems hard to get around. In this paper, we get around this tradeoff by utilizing simple similarity information, such as cosine similarity, between the entity vectors to design a proposal distribution. Such information is readily available in almost all applications. We show that the recent use of locality sensitive hashing for efficient adaptive sampling can be leveraged to obtain a computationally efficient pseudo-marginal MCMC. The new split-merge MCMC has constant time update, just like random split-merge, and at the same time the proposal is informative and needs significantly fewer iterations than random split-merge. Overall, we obtain a sweet tradeoff between convergence and per update cost. As a direct consequence, our proposal, named LSHSM, is around 10x faster than the state-of-the-art sampling methods on both synthetic datasets and two large real datasets KDDCUP and PubMed with several millions of entities and thousands of cluster centers.

1. Introduction

Bayesian mixture models are of great interest due to their flexibility in fitting a countably infinite number of components which can grow with the data (Medvedovic et al., 2004). The growth of model complexity with the data is also in agreement with modern progress in machine learning over massive datasets. However, the appealing properties of Bayesian modeling come with hard computational problems. Even with simple mixture models, the mathematical problems associated with training and inference are intractable. As a result, recent research focuses on developing tractable computational techniques. In particular, the use of Markov chain Monte Carlo (MCMC) methods, to sample from the posterior distribution (Andrieu et al., 2003; Nasrabadi, 2007; Wang & Blei, 2012) is widely prevalent. The practical utility of these methods is illustrated in several applications including haplotype reconstruction (Eronen et al., 2003), nucleotide substitutions (Huelsenbeck & Ronquist, 2001), gene expression (Sharma & Adlakha, 2015), etc.

Metropolis-Hastings (MH) (Andrieu et al., 2003) is a favorite class of MCMC methods, which includes several state-of-the-art algorithms that have proven useful in practice. MH is associated with a transition kernel which provides a proposal step. This step is followed by appropriate stochastic acceptance process that ensures detailed balance. A notable example of MH is the Split-Merge MCMC algorithm (Jain & Neal, 2004; Wang & Russell, 2015) which is particularly useful for problems where an MCMC state can be thought of as consisting of a number of components (or clusters). Here as the name suggests, the proposal step comprises of either a split or a merge. A split move partitions an existing mixture component (or cluster) into two, while a merge move combines two mixture components into one.

In the seminal work of (Jain & Neal, 2004), split-merge MCMC procedure was proposed. To illustrate the process, the authors first introduce a random split-merge MCMC, where the split and the merge decision were taken uniformly at random. However, it was also pointed out, in the same paper, that due to the random nature of the proposal it was unlikely to lead to a new state x' with higher likelihood $\mathcal{L}(x')$ leading to low acceptance. To mitigate the slow progress, the authors then propose the restricted Gibbs split-merge (RGSM). In RGSM instead of a random proposal the idea

*Code is available on: https://github.com/rackingroll/mcmc_lsh

¹Department of Computer Science, Rice University, Houston, TX, USA. Correspondence to: Chen Luo <cl67@rice.edu>.

was to use restricted Gibbs sampling to generate proposals with a higher likelihood of acceptance. Thus, a less number of MCMC iterations were sufficient for convergence due to fewer rejections. However, the cost of restricted Gibbs is itself prohibitive. As a result, even though the iterations are less, each iteration is costly making the overall algorithm slow, especially for large datasets. Our experiments confirm this slow convergence of RGSM.

An essential and surprising observation about space asymmetry with smart proposals in split-merge MCMC was made in (Wang & Russell, 2015). The authors show the necessity to mix smart and dumb (random) proposals for faster progress. They proposed a Smart-Dumb/Dumb-Smart Algorithm (SDDS) as an alternative to RGSM. Instead of relying on Gibbs sampling, the SDDS algorithm instead uses the likelihood of the model itself as a guiding strategy for smart proposals. In other words, the SDDS method evaluates a large number of possible proposals x' based on the likelihood of each x' and choose the best ones. This strategy, as expected, ensures a higher chance of improving the state x with every proposal. However, from a computational perspective, it is not difficult to see that smart proposal x' obtained after evaluation of a large number of proposal states, based on the likelihood, is equivalent to evaluating all these states for acceptance/rejection as part of MH. As a result, the reduction in the number of iteration is not helpful in obtaining an efficient algorithm. Our experiments show that SDDS also has poor convergence.

Unfortunately, most MCMC methodologies ignore the tradeoff between the number of iteration and computations associated with each iteration. They instead only focus on reducing the number of rejections, which is often achieved by informative proposals with increased per iteration cost. In this paper, we are interested in efficient split-merge MCMC algorithm which leads to overall fast convergence. Thus, reducing both is the aim of this work.

Parallelization is Complementary: Due to the significance of the problem there are several works which try to scale up MCMC by using parallelism. Parallelism is often achieved by running parallel MCMC chains on subsets of data and later merging them (Chang & Fisher III, 2013). Since our proposal reduces the overall cost of split-merge MCMC algorithm in general, it will reduce the cost of each of the parallel chains thereby increasing the effectiveness of these parallelisms on MCMC. Thus, existing advances in parallelizing MCMC is complementary to our proposal.

Our Contributions: In this work, we leverage several complementary ideas to design a computationally efficient split-merge MCMC algorithm. We first leverage a simple observation that clusters with entities similar in their vector representation should be favored. We use standard notions of vector similarity such as cosine similarity. However, this

observation is not sufficient in itself, as designing proposals favoring similar entities in the same cluster requires computing all pairwise similarities, which is a prohibitive quadratic time operation.

We then leverage the recent advances in LSH sampler (Luo & Shrivastava, 2017; Spring & Shrivastava, 2017b; Charikar & Siminelakis, 2017b) that can perform similarity sampling in linear cost. We use this efficient LSH sampling to guide our proposal design. We further leverage the surprising observation made in (Wang & Russell, 2015) to merge dumb and smart proposal for fast progress. Finally, to reduce the likelihood computation time we use the same LSH sampler, for proposal design, to produce an unbiased estimator of likelihood leading to a valid pseudo-marginal split-merge MCMC. Such unbiased estimators are superior to favorite random sampling based estimation. We name our method LSHSM (LSH Split-Merge) MCMC.

Overall, LSHSM obtains a sweet tradeoff between the number of iteration and computational cost per iteration. As a result, we reduce the overall computational cost. On several simulations as well as two large public datasets, LSHSM significantly outperforms other state-of-the-art split-merge MCMC algorithms in convergence speed as measured on wall clock time on the same machine. LSHSM is around 10x faster than the second best baseline on real datasets without loss in accuracy.

2. Background

Our work requires bridging Locality Sensitive Sampling with split-merge MCMC algorithm. We briefly review the necessary background.

2.1. Locality Sensitive Hashing

Locality-Sensitive Hashing (LSH) is a popular technique for efficient approximate nearest-neighbor search. LSH is a family of functions, such that a function uniformly sampled from this hash family has the property that, under the hash mapping, similar points have a high probability of having the same hash value. More precisely, consider \mathcal{H} a family of hash functions mapping \mathbb{R}^D to a discrete set $[0, R - 1]$.

Definition 1 *Locality Sensitive Hashing (LSH) Family* A family \mathcal{H} is called (S_0, cS_0, u_1, u_2) -sensitive if for any two points $x, y \in \mathbb{R}^d$ and h chosen uniformly from \mathcal{H} satisfies the following:

- if $\text{Sim}(x, y) \geq S_0$ then $\Pr_{\mathcal{H}}(h(x) = h(y)) \geq u_1$
- if $\text{Sim}(x, y) \leq cS_0$ then $\Pr_{\mathcal{H}}(h(x) = h(y)) \leq u_2$

A collision occurs when the hash values for two data vectors are equal, meaning that $h(x) = h(y)$.

LSH is a very well studied topic in computer science theory and database literature. There are many well-known LSH families in the literature. Please refer (Gionis et al., 1999) for details. The most popular one is Signed Random Projections (Charikar, 2002).

Signed Random Projections(SRP) is an LSH for the cosine similarity measure, which originates from the concept of **randomized rounding (SRP)** (Goemans & Williamson, 1994; Charikar, 2002). Given a vector x , SRP utilizes a random w vector with each component generated from i.i.d. normal, *i.e.*, $w_i \sim N(0, 1)$, and only stores the sign of the projection. Formally SRP family is given by

$$h_w(x) = \text{sign}(w^T x). \quad (1)$$

It was shown in the seminal work (Goemans & Williamson, 1994) that collision under SRP satisfies the following equation:

$$\Pr_w(h_w(x) = h_w(y)) = 1 - \frac{\theta}{\pi}, \quad (2)$$

where $\theta = \cos^{-1}\left(\frac{x^T y}{\|x\|_2 \|y\|_2}\right)$. $\frac{x^T y}{\|x\|_2 \|y\|_2}$, is the cosine similarity.

2.1.1. LOCALITY SENSITIVE SAMPLING (LSS) AND UNBIASED ESTIMATORS

LSH was considered a black-box algorithm for similarity search, similarity estimation and dimensionality reduction. Recently, it was found that LSH can be used for something more subtle but useful. It is a data structure that can be used for efficient dynamically adaptive sampling. We first describe the sampling algorithm of (Spring & Shrivastava, 2017b) and later comment on its known properties crucial for our proposal.

The algorithm uses two parameters - (K, L) . We construct L independent hash tables from the collection \mathcal{C} . Each hash table has a meta-hash function H that is formed by concatenating K random independent hash functions from some appropriate locality sensitive hash family \mathcal{H} . The candidate sampling algorithm works in two phases [See (Spring & Shrivastava, 2017b) for details]:

1. **Pre-processing Phase:** We construct L hash tables from the data by storing all elements $x \in \mathcal{C}$. We only store pointers to the vector in the hash tables because storing whole data vectors is very memory inefficient. This is one-time linear cost.
2. **Sampling Phase:** Given a query q , we collect one bucket from a randomly selected hash table and return a random element from the bucket. If the bucket is empty, we reselect a different hash table again. Keep track of the number of different tables probed.

It is not difficult to show that an item returned as a candidate from a (K, L) -parameterized LSH algorithm is sampled with probability exactly $1 - (1 - p^K)^L$, where p is the collision probability of LSH function. The LSH family defines the precise form of p used to build the hash tables. This sampling view of LSH was first utilized to perform adaptive sparsification of deep networks in near-constant time, leading to efficient backpropagation algorithm (Spring & Shrivastava, 2017a). A year later, (Spring & Shrivastava, 2017b) demonstrated the first theory of using these samples for unbiased estimation of partition functions in log-linear models. More specifically, the authors showed that since we know the precise probability of sampled elements $1 - (1 - p^K)^L$, we could design provably unbiased estimators using importance sampling type idea. This was the first demonstration that random sampling could be beaten with roughly the same computational cost as vanilla sampling. (Luo & Shrivastava, 2017) used the same approach for unbiased estimation of anomaly scoring function. (Chen et al., 2017) used the sampling in a very different context of connected component estimation for unique entity counts. (Charikar & Siminelakis, 2017a) showed improvements in sample complexity of kernel density estimation problems.

The most important observation made in (Spring & Shrivastava, 2017b), is that the expression $1 - (1 - p^K)^L$ is a monotonically increasing function of p , which in turn is a monotonic function of cosine similarity if we use SRP as hashing scheme. Thus, given a query q , points with higher cosine similarity with q is more likely to be sampled. Similarity, points dissimilar with q is less likely. It should be noted that querying cost only involved few (like 5-10) hash computations followed by a couple of memory lookups which is $O(1)$, very similar to random sampling. Capitalizing on this unique efficiency (Chen et al., 2018) proposed LSD (locality sensitive descent), which was the first gradient descent algorithm that can beat the popular SGD, and any of its variants, on running time breaking what they call *the chicken-and-egg loop* in adaptive sampling.

Our proposal will heavily rely on this unusual probability expression $1 - (1 - p^K)^L$ to design an informative and proposal distribution. We will, in addition to the probability expression, also utilize hashing to obtain an unbiased estimate of the likelihood leading to our pseudo-marginal MCMC algorithm.

2.2. Split-Merge MCMC

Split-Merge MCMC (Hughes et al., 2012) is useful for dealing with the tasks such as clustering or topic modeling where the number of clusters or components are not known in advance. Split-Merge MCMC is a Metropolis-Hastings algorithm with two main transitions: Split and Merge. During a split, a cluster is partitioned into two components. On

the contrary, a merge takes two components and makes them into one.

During the MCMC inference process, split and merge moves simultaneously change the number of entities and change the assignments of entities to different clusters. (Jain & Neal, 2004) proposes the first non-trivial Restricted Gibbs Split-Merge (RGSM) algorithm, which was later utilized for efficient topic modeling over large datasets in (Wang & Blei, 2012).

(Wang & Russell, 2015) presented a surprising argument about information asymmetry. It was shown that both informative split and merge leads to poor acceptance ratio. The author proposed a combination of the smart split with dumb (random) merge and dumb split with smart merge as a remedy. The algorithm was named as Smart-Dumb/Dumb-Smart Split Merge algorithm (SDDS), which was superior to RGSM. To obtain non-trivial smart split (or merge), the authors propose to evaluate a large number of dumb proposals based on the likelihood and select the best. This search process made the proposal very expensive. It is not difficult to see that finding a smart split is computationally not very different from running a chain with several sequences of dumb (random) splits.

3. LSHM: LSS based Split-Merge MCMC

3.1. Intuition

Utilizing Similarity Information: In this paper we make an argument that similarity information, such as cosine similarity, between different entities is almost always available. For example, in the clustering, we almost always have a vector representation of the data using which we compute the likelihood. Even in an application where we deal with complex entities such as trees, it is not uncommon to have approximate embeddings (Bengio et al., 2010).

It is natural to believe that similar entities, in terms of cosine similarity, of the underlying vector representation, are more likely to go to the same cluster than non-similar ones. Thus, designing proposals which favor similar entries in the same cluster and dissimilar entities in different clusters is more likely to lead to acceptance than random proposals.

However, the problem is far from being solved. Any similarity based sampling requires computing all pairwise similarity as a prerequisite, which is a quadratic operation $O(n^2)$. Quadratic operations are prohibitive (near-infeasible) for large datasets. One critical observation is that with the modern view of LSH as samplers, described in section 2.1.1, we can get around this quadratic cost and design cheaper non-trivial proposals.

3.2. LSS based Proposal Design

This section discusses our informative proposal and how we compute the transition probabilities $q(x'|x)$, which is an important component of the acceptance ratio $\alpha(x'|x)$ (Jain & Neal, 2004). Here, x denote the state before split/merge, and x' denote the state after split/merge. Although the likelihood terms $\mathcal{L}(x)$ and $\mathcal{L}(x')$ in $\alpha(x'|x)$ can be approximated without changing the equilibrium distribution of MCMC (Section 3.3.1), $q(x'|x)$ still cannot be approximated. Thus, it is imperative that $q(x'|x)$ is easy to calculate as well as the proposed state x' is informative. Note that cheap approximation, like sampling, cannot be used for proposing x' as it will likely result in intractable (or expensive) expression of $q(x'|x)$. Thus, designing the right $q(x'|x)$ is the key to speed up computation. Following the intuition described in Section 3.1, we introduce our LSS based proposal design in the rest of this section.

We first create the hash tables data structure T for sampling (as described in section 2.1). We use Sign Random Projection as the LSH function, thus our notion of similarity is cosine and Eq. 2.1 gives the collision probability. We pay a one-time linear cost for this preprocessing. Note, we need significantly less K and L (both has value 10 in our experiments) compared to what is required for near-neighbor queries as we are only sampling. The sampling is informative (better than random) for any values of K and L

For our informative proposal, we will need capabilities to do both similarity sampling as well as dissimilarity sampling for merge and split respectively. The similarity sampling is the usual sampling algorithm discussed in section 2.1, which ensures that given a query u , points similar to u , in cosine similarity are more likely to be sampled. Analogously, we also need to sample points that are likely to be dissimilar. With cosine similarity, flipping the sign of the query, i.e., changing u to $-u$ will automatically do dissimilarity sampling.

Inspired from (Wang & Russell, 2015), we also leverage the information asymmetry and mix smart and dumb moves for better convergence. However, this time our proposals will be super efficient. At each iteration of MCMC, we start by choosing randomly between an LSH Smart-split/Dumb-merge or an LSH Smart-merge/Dumb-split operation. These two operations are defined below:

- **LSH Smart-split/Dumb-merge :** LSH based Split begins by randomly select an element u in the dataset. Then, we use LSS (Locality-sensitive Sampler) to sample points likely to be dissimilar to u . Thus, we query our data structure T with $-u$ as the query to get another element v which is likely far away from u . If u and v belong to the same cluster C , we split the cluster. During the split, we create two new clusters C_u and C_v .

We assign u to C_u and v to C_v . For every element in C , we randomly assign them to either C_u or C_v . Since we ensure that dissimilar points u and v are split, this is an informative or smart split. If we find u and v are already in a different cluster, we do a dumb merge.

The most important part is that we can precisely compute the probability of the proposed split move $q(x'|x)$ and the corresponding inverse move probability $q(x|x')$ as follow:

$$q(x'|x) = \frac{1}{n} (1 - (1 - Pr(-u, v)^K)^L) * \frac{|C_v \cap S^d|}{|S^d|} * \left(\frac{1}{2}\right)^{|C_u|+|C_v|-1} \quad (3)$$

$$q(x|x') = 1 \quad (4)$$

In the above, n is the number of data point. S^d (dissimilarity) is the set of data points that returned by querying in T using $-u$. C denotes the original component. C_u and C_v are the two new components after split with elements u and v in them. $|S^d|$ denotes the number of elements in S^d . K is the number of bits used for hashing, and L is the number of hash tables probed. $Pr(-u, v)$ is the collision probability between $-u$ and v .

- **LSH Smart-merge/Dumb-split:** LSH based Merge begins by randomly select an element u in the dataset. Then use LSS to sample from T to get another element v which is similar with u . Then, if the mixture component of u and v are different, then we do merge operation for the corresponding two mixture component. If u and v are in the same cluster, we do a dumb split randomly separating u and v . We provide the the probability of the merge move $q(x'|x)$ and the corresponding inverse probability $q(x|x')$:

$$q(x'|x) = \frac{1}{n} (1 - (1 - Pr(u, v)^K)^L) * \frac{|C_v \cap S^s|}{|S^s|} \quad (5)$$

$$q(x|x') = \left(\frac{1}{2}\right)^{|C_u|+|C_v|-1} \quad (6)$$

In the above, S^s (similarity) is the set of data points that returned by query in T using u . $|S^s|$ denotes the number of elements in S^s . All the other symbols have the same meaning as before. $Pr(u, v)$ is the collision probability between u and v .

These probability expressions are obtained by combining the fact that $\frac{1}{n}$ is the probability of choosing u . $(1 - (1 - Pr(u, v)^K)^L)$ is the probability of having v in the buckets probed. $\frac{|C_v \cap S^s|}{|S^s|}$ is the probability of getting $v \in C_v$

by randomly sampling the bucket. $\left(\frac{1}{2}\right)^{|C_u|+|C_v|-1}$ is the probability of uniform spitting with u and v in different component. Merge probability, given two cluster is 1.

3.3. From MCMC to an Efficient Pseudo-Marginal MCMC

As we introduced before, our proposed LSHSM algorithm belongs to the general framework of metropolis-hastings algorithm (Andrieu et al., 2003). After each split/merge move, we need to calculate the acceptance rate $\alpha(x'|x)$ for this move which is given by:

$$\alpha(x'|x) = \min\left\{1, \frac{\mathcal{L}(x')q(x|x')}{\mathcal{L}(x)q(x'|x)}\right\}, \quad (7)$$

where x' is the proposed new state, x is the previous state, $q(x'|x)$ here is the designed proposal distribution, and it can be calculated as introduced in previous sections. $\mathcal{L}(x)$ is the likelihood value of the state x .

The likelihood of the data is generally in the form of:

$$\mathcal{L}(x) = \prod_D p_i(e_i), \quad (8)$$

where $p_i(e_i)$ is the probability of $e_i \in D$ in it's corresponding component C_i . D denotes the total dataset dataset. The corresponding log-likelihood is

$$\mathcal{L}(x) = \sum_D \log(p_i(e_i)), \quad (9)$$

Specifically, if we use the Gaussian mixture model, then

$$p_i(e) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(e_i - \mu_i)^2}{2\sigma_i^2}\right), \quad (10)$$

where μ_i and σ_i is the corresponding mean and variance of component C_i .

We can see from the above equation that, computing the likelihood of the data requires a complete pass over the entire dataset, which makes it computational expensive and does not scale well for large datasets.

3.3.1. SUPERIOR UNBIASED ESTIMATION OF LIKLIHOOD

Fortunately, replacing $\mathcal{L}(x)$ and $\mathcal{L}(x')$ in the computation of $\alpha(x'|x)$ with an unbiased estimator of \mathcal{L} suffices to guarantee the same equilibrium distribution. A method popularly knows as Pseudo-Marginal MCMC (Andrieu et al., 2009). Random sampling based unbiased estimator is a default choice to speed up the likelihood computation. It turns out that with existing LSH structure T we can get a better unbiased estimator.

Algorithm 1 LSHSM Algorithm**Input** : Dataset D , Parameter K, L **Output** : \overline{G}_T Pre-processing all the data point in D into LSH Data structure HT .**while** *Convergence* **do**

Choose a move type randomly: TYPE = (split, merge)

switch TYPE **do** **case** *Split* **do**

Do Smart-split/Dumb-merge operation as introduced in Section 3.2

end **case** *Merge* **do**

Do Smart-merge/Dumb-split operation as introduced in Section 3.2

end **end** Calculate the likelihood function $\mathcal{L}(x)$ as introduced in Section 3.3 Calculate the acceptance ratio α Apply the proposal with probability α **end**

To get a better unbiased estimation of likelihood, we can leverage insights from (Spring & Shrivastava, 2017b). Instead of calculating the likelihood over a random sample of the dataset, we can instead sample a small fraction of the data S use the LSS with cluster means as the query. Then use the small sampled data set S to approximate the likelihood function $\mathcal{L}(x) = \sum_S \log p_i(e_i)$.

To see why cluster means μ_i is better sampling, consider $p_i(e_i)$ in Equation 10. Note that the value of $p_i(e_i)$ is higher if e_i is closer to the cluster mean μ_i . Thus, if we query a small set of sample with mean μ_i as the query, LSS favors sampling heavier entries and therefore is a superior estimator of the sum than random sampling. Also, since we also know the probability of sampling we can design an unbiased sampler.

$$\begin{aligned} \mathcal{L}(x) &= \sum_S \frac{\log p_i(e_i)}{Pr_i} \\ &= \sum_S \log \left(\frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(e_i - \mu_i)^2}{2\sigma_i^2}\right) \right) - \log Pr_i, \end{aligned} \quad (11)$$

where S is the set obtained by querying each μ_i in the existing LSH tables, and Pr_i is the collision probability between the element e_i and the corresponding probability mean μ_i , as Eq. 2. The LSS estimators is unbiased and superior in variance than the plain random sampling estimators. Please refer to (Spring & Shrivastava, 2017b) for more details.

The vital point to note is that since estimators $\mathcal{L}(x)$ is unbiased we guarantee the desired equilibrium distribution of

*MCMC.***3.4. LSHSM Algorithm**

The overall procedure is then summarized as Algorithm 1. This algorithm implements the methods we have introduced in the above subsections.

4. Empirical Study

In this section, we demonstrate the advantage LSHSM by applying it to the Gaussian Mixture model inference and compare it with state-of-the-art sampling methods ¹.

4.0.1. GAUSSIAN MIXTURE MODEL

We briefly review the Gaussian Mixture Model. A Gaussian mixture density is a weighted sum of component densities. For a M -class clustering task, we could have a set of GMMs associated with each cluster. For a D -dimensional feature vector denoted as \vec{x} , the mixture density is defined as $p(\vec{x}) = \sum_{i=1}^M w_i p_i(\vec{x})$, where $w_i, i = 1, \dots, M$ are the mixture weights which satisfy the constraint that $\sum_i^M w_i = 1$ and $w_i \geq 0$. The mixture density is a weighted linear combination of component M uni-model Gaussian density functions $p_i(\vec{x}), i = 1, \dots, M$. The Gaussian mixture density is parameterized by the mixture weights, mean vectors and covariance vectors from all components densities.

For a GMM-based clustering task, the goal of the model training is to estimate the parameters of the GMM so that the Gaussian mixture density can best match the distribution of the training feature vectors. Estimating the parameters of the GMM using the expectation-maximization (EM) algorithm (Nasrabadi, 2007) is popular. However, in most of the real world applications, the number of clusters M is not known, which is required by the EM algorithm. On the other hand, Split-Merge based MCMC algorithms are used for inference when M is unknown, which is also the focus of this paper. We therefore only compare our proposal LSHSM and other state-of-the-art split-merge algorithms on GMM clustering which does not require the prior knowledge of the number of clusters.

4.1. Competing Algorithms

We compare following three split-merge MCMC sampling algorithm on GMM with an unknown number of clusters:

- **RGSM** Restricted Gibbs split-merge MCMC algorithm (Jain & Neal, 2004) is considered as one of the state-of-the-art sampling algorithm.

¹Code is available on https://github.com/rackingroll/mcmc_lsh

Table 1. The statistics for the two real world dataset

Dataset	Samples	Dim	True #Clusters
KDDCUP	145751	74	2000
PubMed	8200000	141043	10000

- **SDDS** Smart-Dumb/Dumb-Smart Split Merge algorithm (Wang & Russell, 2015). SDDS combines “smart” split/merge move that proposes plausible splits of heterogeneous clusters with a “dumb” merge move that proposes merging random pairs of clusters.
- **LSHSM** This is the proposed method in this paper. In the LSHSM method, we use fixed $K = 10$ and $L = 10$ for all the dataset. We fix the hashing scheme to be signed random projection.

4.2. Dataset

We evaluate the effectiveness of our algorithm on both synthetic datasets and two large real-world datasets.

4.2.1. SYNTHETIC DATA SET

Synthetic data is a standard way of testing GMM models (Nasrabadi, 2007). So, in this paper, we first use synthetic datasets as a sanity check to evaluate the performance of different methods. The process of generating the synthetic dataset is as follow: Randomly generate k different Gaussian distributions (with different corresponding mean and variance). We fix the $k = 10$ in our experiment. Then based on the randomly generated Gaussian distributions, we generate a set of data points for each Gaussian distribution. Here we fix the dimensionality of each data point to 25.

In this experiment, we generate three synthetic dataset with different size (e.g. 100, 1000, 10000). We name the three synthetic dataset as **S1**, **S2**, **S3**.

4.2.2. REAL WORD DATA SET

We also evaluate the performance of all the methods on two real word datasets:

KDDCUP: This dataset was used in the KDD Cup 2004 data mining competition. It contains 145751 data point. The dimensionality of the dataset is 74. We have 2000 ground truth cluster labels for this dataset.²

PubMed The PubMed abstraction dataset contains 8200000 abstractions that extracted from the PubMed³. All the documents represented as the bag-of-words representation. In the data set, we have 141043, different words. This data set is ideal for document clustering or topic modeling. The

dataset is available from the UCI machine learning dataset Repository.⁴ The statistics of the two real-world datasets is shown in Table 1.

4.3. Result and Analysis

4.3.1. SPEED COMPARISON

We first plot the evolution of likelihood both as a function of iterations as well as the time of all the three competing methods. The evolution of likelihood with iterations on the synthetic dataset and two real-world data is shown in Fig. 1 and Fig. 3 respectively. Fig. 2 and Fig. 3 plots the same evolution but with time instead of iteration.

We can see a consistent trend in the evolution of likelihood, which holds true for both simulated as well as real datasets. First of all, RGSM consistently performs poorly and requires both more iterations as well as time. The need of combining smart and dumb moves for faster convergence made in (Wang & Russell, 2015), seems necessary. RGSM does not use it and hence leads to poor, even iteration wise, convergence.

SDDS seems to do quite well, compared to our proposed LSHSM when we look at iteration wise convergence. However, when we look at the time, the picture is completely changed. LSHSM is significantly faster than SDDS, even if the convergence is slower iteration wise. This is not surprising because the per-iteration cost of LSHSM is orders of magnitude less than SDDS. SDDS hides the computations inside the iteration by evaluating every possible state in each iteration, based on likelihood, is equivalent to several random iterations combined. Such costly evaluation per iteration can give a false impression of less iteration.

It is clear from the plots that merely comparing iterations and acceptance ratio can give a false impression of superiority. Time wise comparison is a legitimate comparison of overall computational efficiency. Clearly, LSHSM outperforms the other baselines by a large margin.

From the experiment on two large datasets we can see that: on the KDD CUP data set, LSHSM can converge in less than 400 seconds, while RGSM and SDDS need nearly one hour (3600 seconds) to converge. On the PubMed dataset, the LSHSM method can converge in less than one hour, while SDDS need more than 10 hours to converge, and RGSM requires more than 20 hours to converge. This demonstrates that our proposed LSHSM algorithm can be at least 10 times faster than the state of the art algorithms in the large dataset.

²<https://cs.joensuu.fi/sipu/datasets/>

³www.pubmed.gov

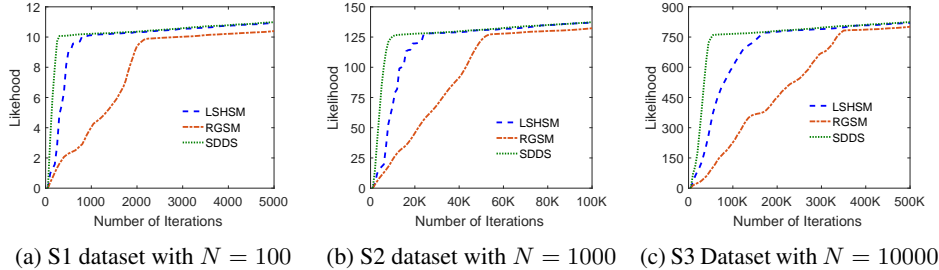


Figure 1. Comparing the evolution of likelihood during each iteration on different synthetic dataset.




Figure 2. The time wise comparison of the likelihood for difference methods on the Synthetic Dataset. LSHSM outperforms the other baselines by a large margin. It is also clear that requiring less iteration does not mean faster convergence.

Table 2. Clustering Accuracy for Different Methods

Methods	Metric	S1	S2	S3	KDD	Pub
RGSM	NMI	0.96	0.93	0.88	0.74	0.63
	Acc	0.95	0.92	0.87	0.68	0.62
SDDS	NMI	0.97	0.96	0.95	0.86	0.80
	Acc	0.98	0.97	0.94	0.85	0.77
LSHSM	NMI	0.96	0.95	0.96	0.84	0.78
	Acc	0.97	0.93	0.95	0.83	0.76

4.3.2. CLUSTERING ACCURACY COMPARISON

To evaluate the clustering performance of different algorithms, we use two widely used measures (Accuracy and NMI (Nasrabadi, 2007)). We briefly review the definition of these two measures below:

Normalized Mutual Information (NMI) (Nasrabadi, 2007): is widely used for measuring the performance of clustering algorithms. It can be calculated as follow:

$$NMI(C, C') = \frac{I(C; C')}{\sqrt{H(C)H(C')}},$$

where $H(C)$ and $H(C')$ are the marginal entropies, $I(C; C')$ is the mutual information between C' and C .

Accuracy: The accuracy measure, which is calculated as the percentage of target objects going to the correct cluster,

is defined as follow:

$$Accuracy = \frac{\sum_{i=1}^k a_i}{n},$$

where a_i is the number of data objects clustered to its corresponding true cluster, k is the number of cluster and n is the number of data objects in the dataset.

Table 2 shows the clustering accuracy of different competing methods. We can see that the LSHSM and SDDS are much more accurate than RGSM. This observation is in agreement with the likelihood plots. On the other hand, the accuracy difference between LSHSM and SDDS is negligible.

From the experiments, we can conclude that the proposed LSHSM method converged much faster than the state-of-the-art algorithm (10x faster on the large dataset), while achieving the same accuracy

5. Conclusion

The Split-Merge MCMC (Monte Carlo Markov Chain) is one of the essential and popular variants of MCMC for problems with an unknown number of components. It is a well known that the inference process of Split-Merge MCMC is computational expensive which is not applicable for the large-scale dataset. Existing approaches that try to speed up the split-merge MCMC are stuck in a computational chicken-and-egg loop problem.

In this paper, we proposed LSHSM, accelerating Split

⁴<https://archive.ics.uci.edu/ml/index.php>

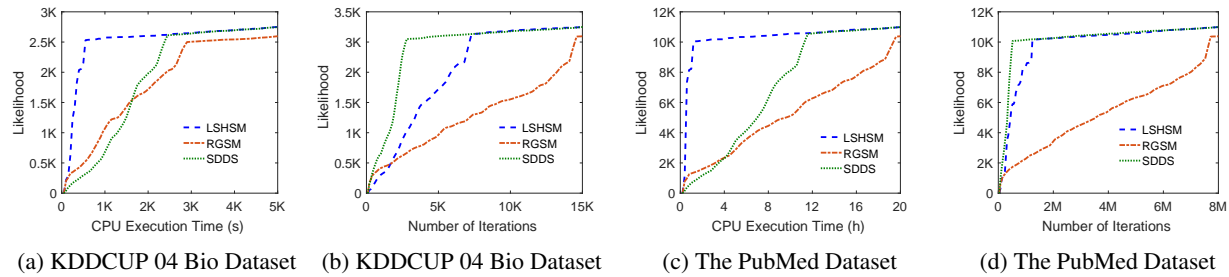


Figure 3. The time and iteration wise comparison of the likelihood for different methods on the two real datasets. It is obvious that our proposed LSHSM algorithm can be at least 10 times faster than the state of the art algorithms in the real large dataset.

Merge MCMC via probabilistic hashing. The new split-merge MCMC has constant time update, and at the same time the proposal is informative and needs significantly fewer iterations than random split-merge. Overall, we obtain a sweet tradeoff between convergence and per update cost. Experiments with Gaussian Mixture Model on both synthetic dataset and two real-world datasets demonstrate much faster convergence and better scaling to large datasets.

6. Acknowledgments

This work was supported by National Science Foundation IIS-1652131, RI-1718478, and a GPU grant from NVIDIA.

References

- Andrieu, Christophe, De Freitas, Nando, Doucet, Arnaud, and Jordan, Michael I. An introduction to mcmc for machine learning. *Machine learning*, 50(1-2):5–43, 2003.
- Andrieu, Christophe, Roberts, Gareth O, et al. The pseudo-marginal approach for efficient monte carlo computations. *The Annals of Statistics*, 37(2):697–725, 2009.
- Bengio, Samy, Weston, Jason, and Grangier, David. Label embedding trees for large multi-class tasks. In *Advances in Neural Information Processing Systems*, pp. 163–171, 2010.
- Chang, Jason and Fisher III, John W. Parallel sampling of dp mixture models using sub-cluster splits. In *Advances in Neural Information Processing Systems*, pp. 620–628, 2013.
- Charikar, Moses and Siminelakis, Paris. Hashing-based-estimators for kernel density in high dimensions. FOCS, 2017a.
- Charikar, Moses and Siminelakis, Paris. Hashing-based-estimators for kernel density in high dimensions. FOCS, 2017b.
- Charikar, Moses S. Similarity estimation techniques from rounding algorithms. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pp. 380–388. ACM, 2002.
- Chen, Beidi, Shrivastava, Anshumali, and Steorts, Rebecca C. Unique entity estimation with application to the syrian conflict. *arXiv preprint arXiv:1710.02690*, 2017.
- Chen, Beidi, Xu, Yingchen, and Shrivastava, Anshumali. Lsh-sampling breaks the computational chicken-and-egg loop in adaptive stochastic gradient estimation. 2018.
- Eronen, Lauri, Geerts, Floris, and Toivonen, Hannu. A markov chain approach to reconstruction of long haplotypes. In *Biocomputing 2004*, pp. 104–115. World Scientific, 2003.
- Gionis, Aristides, Indyk, Piotr, Motwani, Rajeev, et al. Similarity search in high dimensions via hashing. In *VLDB*, volume 99, pp. 518–529, 1999.
- Goemans, Michel X and Williamson, David P. .879-approximation algorithms for max cut and max 2sat. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pp. 422–431. ACM, 1994.
- Huelsenbeck, John P and Ronquist, Fredrik. Mrbayes: Bayesian inference of phylogenetic trees. *Bioinformatics*, 17(8):754–755, 2001.
- Hughes, Michael C, Fox, Emily, and Sudderth, Erik B. Effective split-merge monte carlo methods for nonparametric models of sequential data. In *Advances in neural information processing systems*, pp. 1295–1303, 2012.
- Jain, Sonia and Neal, Radford M. A split-merge markov chain monte carlo procedure for the dirichlet process mixture model. *Journal of Computational and Graphical Statistics*, 13(1):158–182, 2004.
- Luo, Chen and Shrivastava, Anshumali. Arrays of (locality-sensitive) count estimators (ace): High-speed anomaly detection via cache lookups. *arXiv preprint arXiv:1706.06664*, 2017.

- Medvedovic, Mario, Yeung, Ka Yee, and Bumgarner, Roger Eugene. Bayesian mixture model based clustering of replicated microarray data. *Bioinformatics*, 20(8): 1222–1232, 2004.
- Nasrabadi, Nasser M. Pattern recognition and machine learning. *Journal of electronic imaging*, 16(4):049901, 2007.
- Sharma, Amit and Adlakha, Neeru. A computational model to study the concentrations of dna, mrna and proteins in a growing cell. *Journal of Medical Imaging and Health Informatics*, 5(5):945–950, 2015.
- Spring, Ryan and Shrivastava, Anshumali. Scalable and sustainable deep learning via randomized hashing. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 445–454. ACM, 2017a.
- Spring, Ryan and Shrivastava, Anshumali. A new unbiased and efficient class of lsh-based samplers and estimators for partition function computation in log-linear models. *arXiv preprint arXiv:1703.05160*, 2017b.
- Wang, Chong and Blei, David M. A split-merge mcmc algorithm for the hierarchical dirichlet process. *arXiv preprint arXiv:1201.1657*, 2012.
- Wang, Wei and Russell, Stuart J. A smart-dumb/dumb-smart algorithm for efficient split-merge mcmc. In *UAI*, pp. 902–911, 2015.