

Original citation:

Wade, Sara and Ghahramani, Zoubin. (2017) Bayesian cluster analysis : point estimation and credible balls. Bayesian Analysis.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/92153>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work of researchers of the University of Warwick available open access under the following conditions.

This article is made available under the Creative Commons Attribution 4.0 International license (CC BY 4.0) and may be reused according to the conditions of the license. For more details see: <http://creativecommons.org/licenses/by/4.0/>

A note on versions:

The version presented in WRAP is the published version, or, version of record, and may be cited as it appears here.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

Bayesian Cluster Analysis: Point Estimation and Credible Balls

Sara Wade* and Zoubin Ghahramani†

Abstract. Clustering is widely studied in statistics and machine learning, with applications in a variety of fields. As opposed to popular algorithms such as agglomerative hierarchical clustering or k-means which return a single clustering solution, Bayesian nonparametric models provide a posterior over the entire space of partitions, allowing one to assess statistical properties, such as uncertainty on the number of clusters. However, an important problem is how to summarize the posterior; the huge dimension of partition space and difficulties in visualizing it add to this problem. In a Bayesian analysis, the posterior of a real-valued parameter of interest is often summarized by reporting a point estimate such as the posterior mean along with 95% credible intervals to characterize uncertainty. In this paper, we extend these ideas to develop appropriate point estimates and credible sets to summarize the posterior of the clustering structure based on decision and information theoretic techniques.

Keywords: mixture model, random partition, variation of information, Binder’s loss.

1 Introduction

Clustering is widely studied in statistics and machine learning, with applications in a variety of fields. Numerous models and algorithms for clustering exist, and new studies which apply these methods to cluster new datasets or develop novel models or algorithms are constantly being produced. Classical algorithms such as agglomerative hierarchical clustering or the k-means algorithm (Hartigan and Wong (1979)) are popular but only explore a nested subset of partitions or require specifying the number of clusters apriori. Moreover, they are largely heuristic and not based on formal models, prohibiting the use of statistical tools, for example, in determining the number of clusters.

Model-based clustering methods utilize finite mixture models, where each mixture component corresponds to a cluster (Fraley and Raftery (2002)). Problems of determining the number of clusters and the component probability distribution can be dealt with through statistical model selection, for example, through various information criteria. The expectation-maximization (EM) algorithm is typically used for maximum likelihood estimation (MLE) of the mixture model parameters. Given the MLEs of the parameters, the posterior probability that a data point belongs to a class can be computed through Bayes rule. The cluster assignment of the data point corresponds to the class with maximal posterior probability, with the corresponding posterior probability reported as a

*University of Warwick, s.wade@warwick.ac.uk

†University of Cambridge, zoubin@eng.cam.ac.uk

measure of uncertainty. Importantly, however, this measure of uncertainty ignores uncertainty in the parameter estimates. As opposed to MLE, Bayesian mixture models incorporate prior information on the parameters and allow one to assess uncertainty in the clustering structure unconditional on the parameter estimates.

Bayesian nonparametric mixture models assume that the number of components is infinite. As opposed to finite mixture models, this not only avoids specification of the number of components but also allows the number of clusters present in the data to grow unboundedly as more data is collected. Bayesian nonparametric mixture models induce a random partition model (Quintana (2006)) of the data points into clusters, and the posterior of the random partition reflects our belief and uncertainty of the clustering structure given the data.

However, an important problem in Bayesian nonparametric cluster analysis is how to summarize this posterior; indeed, often the first question one asks is what is an appropriate point estimate of the clustering structure based on the posterior. Such a point estimate is useful for concisely representing the posterior and often needed in applications. Moreover, a characterization of the uncertainty around this point estimate would be desirable in many applications. Even in studies of Bayesian nonparametric models where the latent partition is used simply as a tool to construct flexible models, such as in mixture models for density estimation (Lo (1984)), it is important to understand the behavior of the latent partition to improve understanding of the model. To do so, the researcher needs to be equipped with appropriate summary tools for the posterior of the partition.

Inference in Bayesian nonparametric partition models usually relies on Markov chain Monte Carlo (MCMC) techniques, which produce a large number of partitions that represent approximate samples from the posterior. Due to the huge dimension of the partition space and the fact that many of these partitions are quite similar differing only in a few data points, the posterior is typically spread out across a large number of partitions. Clearly, describing all the unique partitions sampled would be unfeasible, further emphasizing the need for appropriate summary tools to communicate our findings.

In a typical Bayesian analysis, the posterior of a univariate parameter of interest is often summarized by reporting a point estimate such as the posterior mean, median, or mode, along with a 95% credible interval to characterize uncertainty. In this paper, we aim to extend these ideas to develop summary tools for the posterior on partitions. In particular, we seek to answer the two questions: 1) What is an appropriate point estimate of the partition based on the posterior? 2) Can we construct a 95% credible region around this point estimate to characterize our uncertainty?

We first focus on the problem of finding an appropriate point estimate. A simple solution is to use the posterior mode. If the marginal likelihood of the data given the partition, that is with all mixture component parameters integrated out, and the prior of the partition are available in closed form, the posterior mode can be estimated based on the MCMC output by the sampled partition which maximizes the non-normalized posterior. In practice, a closed form for the marginal likelihood or prior is often unavailable, specifically, if conjugate priors for the component specific parameters do not exist

or are not utilized or hyperpriors are assigned to any hyperparameters. More generally, the posterior mode can be found by reporting the partition visited most frequently in the sampler. Yet this approach can be problematic, as producing reliable frequency counts is intractable due to the huge dimension of the partition space. In fact, in many examples, the MCMC chain does not visit a partition more than once. To overcome this, alternative search techniques have been developed to locate the posterior mode (Heller and Ghahramani (2005), Heard et al. (2006), Dahl (2009), Raykov et al. (2014)). However, it is well-known that the mode can be unrepresentative of the center of a distribution.

Alternative methods have been proposed based on the posterior similarity matrix. For a sample size of N , the elements of this N by N matrix represent the probability that two data points are in the same cluster, which can be estimated by the proportion of MCMC samples that cluster the two data points together. Then, classical hierarchical or partitioning algorithms are applied based on the similarity matrix (Medvedovic and Sivaganesan (2002), Medvedovic et al. (2004), Rasmussen et al. (2009), Molitor et al. (2010)). These methods have the disadvantage of being ad-hoc.

A more elegant solution is based on decision theory. In this case, one defines a loss function over clusterings. The optimal point estimate is that which minimizes the posterior expectation of the loss function. For example, for a real-valued parameter θ , the optimal point estimate is the posterior mean under the squared error loss $L_2(\theta, \hat{\theta}) = (\theta - \hat{\theta})^2$ and the posterior median under the absolute error loss $L_1(\theta, \hat{\theta}) = |\theta - \hat{\theta}|$.

The question to answer then becomes what is an appropriate loss function on the space of clusterings. The 0-1 loss function, a simple choice which leads to the posterior mode as the point estimate, is not ideal as it does not take into account the similarity between two clusterings. More general loss functions were developed by Binder (1978), and the so-called Binder’s loss, which measures the disagreements in all possible pairs of observations between the true and estimated clusterings, was studied in a Bayesian nonparametric setting by Lau and Green (2007). Alternative loss functions considered in Bayesian nonparametrics can be found in Quintana and Iglesias (2003) and Fritsch and Ickstadt (2009).

In this paper, we propose to use the variation of information developed by Meilă (2007) as a loss function in a Bayesian nonparametric setting. Both the variation of information and Binder’s loss possess the desirable properties of being metrics on the space of partitions and being *aligned* with the lattice of partitions. We provide a detailed comparison of these two metrics and discuss the advantages of the variation of information over Binder’s loss as a loss function in Bayesian cluster analysis. Additionally, we propose a novel algorithm to locate the optimal partition, taking advantage of the fact that both metrics are aligned on the space of partitions.

Next, to address the problem of characterizing uncertainty around the point estimate, we propose to construct a credible ball around the point estimate. As both Binder’s loss and the variation of information are metrics on the partition space, we can easily construct such a ball. Interestingly, the two metrics can produce very different credible balls, and we discuss this in detail. In existing literature, quantifications of uncertainty include reporting a heat map of the estimated posterior similarity matrix.

However, there is no precise quantification of how much uncertainty is represented by the posterior similarity matrix, and in a comparison with the 95% credible balls, we find that the uncertainty is under-represented by the posterior similarity matrix. Finally, we provide an algorithm to construct the credible ball and discuss ways to depict or report it.

The paper is organized as follows. Section 2 provides a review of Bayesian nonparametric clustering and existing point estimates of the clustering structure from a decision theoretic approach. In Section 3, we give a detailed comparison of two loss functions, Binder’s loss and the variation of information, pointing out advantages of the latter. The optimal point estimate under the variation of information is derived in Section 4 and a novel algorithm to locate the optimal partition is proposed. In Section 5, we construct a credible ball around the point estimate to characterize posterior uncertainty and discuss how to compute and depict it. Finally, simulated and real examples are provided in Section 6.

2 Review

This section provides a review of Bayesian nonparametric clustering models and existing point estimates of the clustering in literature.

2.1 Bayesian nonparametric clustering

Mixture models are one of the most popular modeling tools in Bayesian nonparametrics. The data is assumed conditionally i.i.d. with density

$$f(y|P) = \int K(y|\theta)dP(\theta),$$

where $K(y|\theta)$ is a specified parametric density on the sample space with mixing parameter $\theta \in \Theta$ and P is a probability measure on Θ . In a Bayesian setting, the model is completed with a prior on the unknown parameter, which in this case, is the unknown mixing measure. In the most general setting, this parameter P can be any probability measure on Θ , requiring a nonparametric prior. Typically the nonparametric prior has discrete realizations almost surely (a.s.) with

$$P = \sum_{j=1}^{\infty} w_j \delta_{\theta_j} \text{ a.s.},$$

where it is often assumed that the weights (w_j) and atoms (θ_j) are independent and the θ_j are i.i.d. from some base measure P_0 . Thus, the density is modeled with a countably infinite mixture model

$$f(y|P) = \sum_{j=1}^{\infty} w_j K(y|\theta_j).$$

Since P is discrete a.s., this model induces a latent partitioning \mathbf{c} of the data where two data points belong to the same cluster if they are generated from the same mixture

component. The partition can be represented by $\mathbf{c} = (C_1, \dots, C_{k_N})$, where C_j contains the indices of data points in the j^{th} cluster and k_N is the number of clusters in the sample of size N . Alternatively, the partition can be represented by $\mathbf{c} = (c_1, \dots, c_N)$, where $c_n = j$ if the n^{th} data point is in the j^{th} cluster.

A key difference with finite mixture models is that the number of mixture components is infinite; this allows the data to determine the number of clusters k_N present in the data, which can grow unboundedly with the data. Letting $\mathbf{y}_j = \{y_n\}_{n \in C_j}$, the marginal likelihood for the data $y_{1:N}$ given the partition is

$$f(y_{1:N}|\mathbf{c}) = \prod_{j=1}^{k_N} m(\mathbf{y}_j) = \prod_{j=1}^{k_N} \int \prod_{n \in C_j} K(y_n|\theta) dP_0(\theta).$$

The posterior of the partition, which reflects our beliefs and uncertainty in the clustering given the data, is simply proportional to the prior times the marginal likelihood

$$p(\mathbf{c}|y_{1:N}) \propto p(\mathbf{c}) \prod_{j=1}^{k_N} m(\mathbf{y}_j), \quad (1)$$

where the prior of the partition is obtained from the selected prior on the mixing measure. For example, a Dirichlet process prior (Ferguson (1973)) for P with mass parameter α corresponds to

$$p(\mathbf{c}) = \frac{\Gamma(\alpha)}{\Gamma(\alpha + N)} \alpha^{k_N} \prod_{j=1}^{k_N} \Gamma(n_j),$$

where $n_j = |C_j|$ is the number of data points in cluster j . Various other priors developed in Bayesian nonparametric literature can be considered for the mixing measure P , such as the Pitman–Yor process (Pitman and Yor (1997)), also known as the two-parameter Poisson–Dirichlet process, or the normalized generalized Gamma process or more generally, a prior within the class of normalized completely random measures, Poisson–Kingman models (Pitman (2003)), or stick-breaking priors (Ishwaran and James (2001)). See Lijoi and Prünster (2011) for an overview.

In general, the marginal likelihood of the data given the partition or the prior of the partition used to compute the posterior in (1) may not be available in closed form. Moreover, there are

$$S_{N,k} = \frac{1}{k!} \sum_{j=0}^k (-1)^j \binom{k}{j} (k-j)^N,$$

a Stirling number of the second kind, ways to partition the N data points into k groups and

$$B_N = \sum_{k=1}^N S_{N,k},$$

a Bell number, possible partitions of the N data points. Even for small N , this number is very large, which makes computation of the posterior intractable for the simplest choice

of prior and likelihood. Thus, MCMC techniques are typically employed, such as the marginal samplers described by Neal (2000) with extensions in Favaro and Teh (2013) for normalized completely random measures and in Lomellí et al. (2016) for σ -stable Poisson–Kingman models; the conditional samplers described in Ishwaran and James (2001), Papaspiliopoulos and Roberts (2008), or Kalli et al. (2011), with extensions in Favaro and Teh (2013) for normalized completely random measures and in Favaro and Walker (2012) for σ -stable Poisson–Kingman models; or the recently introduced class of hybrid samplers for σ -stable Poisson–Kingman models in Lomellí et al. (2015). These algorithms produce approximate samples $(\mathbf{c}^m)_{m=1}^M$ from the posterior (1). Clearly, describing all the posterior samples is infeasible, and our aim is to develop appropriate summary tools to characterize the posterior.

Extensions of Bayesian nonparametric mixture models are numerous and allow one to model increasingly complex data. These include extensions for partially exchangeable data (Teh et al. (2006)), inclusion of covariates (MacEachern (2000)), time dependent data (Griffin and Steel (2006)), and spatially dependent data (Duan et al. (2007)) to name a few. See Müller and Quintana (2004) and Dunson (2010) for an overview. These extensions also induce latent clusterings of the observations, and the summary tools developed here are applicable for these settings as well.

2.2 Point estimation for clustering

Firstly, we seek a point estimate of the clustering that is representative of the posterior, which may be of direct interest to the researcher or, more generally, important for understanding the behavior of the posterior. From decision theory, a point estimate is obtained by specifying a loss function $L(\mathbf{c}, \hat{\mathbf{c}})$, which measures the loss of estimating the true clustering \mathbf{c} with $\hat{\mathbf{c}}$. Since the true clustering is unknown, the loss is averaged across all possible true clusterings, where the loss associated to each potential true clustering is weighted by its posterior probability. The point estimate \mathbf{c}^* corresponds to the estimate which minimizes the posterior expected loss,

$$\mathbf{c}^* = \operatorname{argmin}_{\hat{\mathbf{c}}} \mathbb{E}[L(\mathbf{c}, \hat{\mathbf{c}}) | y_{1:N}] = \operatorname{argmin}_{\hat{\mathbf{c}}} \sum_{\mathbf{c}} L(\mathbf{c}, \hat{\mathbf{c}}) p(\mathbf{c} | y_{1:N}).$$

A simple choice for the loss function is the 0-1 loss, $L_{0-1}(\mathbf{c}, \hat{\mathbf{c}}) = \mathbf{1}(\mathbf{c} \neq \hat{\mathbf{c}})$, which assumes a loss of 0 if the estimate is equal to the truth and a loss of 1 otherwise. Under the 0-1 loss, the optimal point estimate is the posterior mode. However, this loss function is unsatisfactory because it doesn't take into account similarity between two clusterings; a partition which differs from the truth in the allocation of only one observation is penalized the same as a partition which differs from the truth in the allocation of many observations. Moreover, it is well-known that the mode can be unrepresentative of the center of a distribution. Thus, more general loss functions are needed.

However, constructing a more general loss is not straightforward because, as pointed out by Binder (1978), the loss function should satisfy basic principles such as invariance to permutations of the data point indices and invariance to permutations of the cluster labels for both the true and estimated clusterings. Binder notes that this first condition

implies that the loss is a function of the counts $n_{ij} = |C_i \cap \hat{C}_j|$, which is the cardinality of the intersection between C_i , the set of data point indices in cluster i under \mathbf{c} , and \hat{C}_j , the set of data point indices in cluster j under $\hat{\mathbf{c}}$, for $i = 1, \dots, k_N$ and $j = 1, \dots, \hat{k}_N$; the notation k_N and \hat{k}_N represents the number of clusters in \mathbf{c} and $\hat{\mathbf{c}}$, respectively. He explores loss functions satisfying these principles, starting with simple functions of the counts n_{ij} . The so-called Binder's loss is a quadratic function of the counts, which for all possible pairs of observations, penalizes the two errors of allocating two observations to different clusters when they should be in the same cluster or allocating them to the same cluster when they should be in different clusters:

$$B(\mathbf{c}, \hat{\mathbf{c}}) = \sum_{n < n'} l_1 \mathbf{1}(c_n = c_{n'}) \mathbf{1}(\hat{c}_n \neq \hat{c}_{n'}) + l_2 \mathbf{1}(c_n \neq c_{n'}) \mathbf{1}(\hat{c}_n = \hat{c}_{n'}).$$

If the two types of errors are penalized equally, $l_1 = l_2 = 1$, then

$$B(\mathbf{c}, \hat{\mathbf{c}}) = \frac{1}{2} \left(\sum_{i=1}^{k_N} n_{i+}^2 + \sum_{j=1}^{\hat{k}_N} n_{+j}^2 - 2 \sum_{i=1}^{k_N} \sum_{j=1}^{\hat{k}_N} n_{ij}^2 \right),$$

where $n_{i+} = \sum_j n_{ij}$ and $n_{+j} = \sum_i n_{ij}$. Under Binder's loss with $l_1 = l_2$, the optimal partition \mathbf{c}^* is the partition \mathbf{c} which minimizes

$$\sum_{n < n'} |\mathbf{1}(c_n = c_{n'}) - p_{nn'}|,$$

or equivalently, the partition \mathbf{c} which minimizes

$$\sum_{n < n'} (\mathbf{1}(c_n = c_{n'}) - p_{nn'})^2, \quad (2)$$

where $p_{nn'} = P(c_n = c_{n'} | y_{1:N})$ is the posterior probability that two observations n and n' are clustered together. This loss function was first studied in Bayesian nonparametrics by Lau and Green (2007). We note that in earlier work Dahl (2006) considered minimization of (2) but without the connection to Binder's loss and the decision theoretic approach.

Binder's loss counts the total number of disagreements, D , in the $\binom{N}{2}$ possible pairs of observations. The Rand index (Rand (1971)), a cluster comparison criterion, is defined as the number of agreements, A , in all possible pairs divided by the total number of possible pairs. Since $D + A = \binom{N}{2}$, Binder's loss and the Rand index, denoted $R(\mathbf{c}, \hat{\mathbf{c}})$, are related:

$$B(\mathbf{c}, \hat{\mathbf{c}}) = (1 - R(\mathbf{c}, \hat{\mathbf{c}})) \binom{N}{2},$$

and the point estimate obtained from minimizing the posterior expected Binder's loss is equivalent to the point estimate obtained from maximizing the posterior expected Rand's index. Motivated by this connection, Fritsch and Ickstadt (2009) consider maximizing the adjusted Rand index, introduced by Hubert and Arabie (1985) to correct the Rand index for chance. An alternative loss function is explored by Quintana and Iglesias (2003) specifically for the problem of outlier detection.

3 A comparison of the variation of information and Binder's loss

Meilă (2007) introduces the *variation of information* (VI) for cluster comparison, which is constructed from information theory and compares the information in two clusterings with the information shared between the two clusterings. More formally, the VI is defined as

$$\begin{aligned} \text{VI}(\mathbf{c}, \hat{\mathbf{c}}) &= H(\mathbf{c}) + H(\hat{\mathbf{c}}) - 2I(\mathbf{c}, \hat{\mathbf{c}}) \\ &= -\sum_{i=1}^{k_N} \frac{n_{i+}}{N} \log\left(\frac{n_{i+}}{N}\right) - \sum_{j=1}^{\hat{k}_N} \frac{n_{+j}}{N} \log\left(\frac{n_{+j}}{N}\right) - 2 \sum_{i=1}^{k_N} \sum_{j=1}^{\hat{k}_N} \frac{n_{ij}}{N} \log\left(\frac{n_{ij}N}{n_{i+}n_{+j}}\right), \end{aligned}$$

where \log denotes \log base 2. The first two terms represent the entropy of the two clusterings, which measures the uncertainty in bits of the cluster allocation of an unknown randomly chosen data point given a particular clustering of the data points. The last term is the mutual information between the two clusterings and measures the reduction in the uncertainty of the cluster allocation of a data point in \mathbf{c} when we are told its cluster allocation in $\hat{\mathbf{c}}$. The VI ranges from 0 to $\log(N)$. A review of extensions of the VI to normalize or correct for chance are discussed in Vinh et al. (2010). However, some desirable properties of the VI are lost under these extensions.

In this paper, we propose to use the VI as a loss function. Note that since $I(\mathbf{c}, \hat{\mathbf{c}}) = H(\mathbf{c}) + H(\hat{\mathbf{c}}) - H(\mathbf{c}, \hat{\mathbf{c}})$, we can write

$$\begin{aligned} \text{VI}(\mathbf{c}, \hat{\mathbf{c}}) &= H(\mathbf{c}) + H(\hat{\mathbf{c}}) - 2H(\mathbf{c}) - 2H(\hat{\mathbf{c}}) + 2H(\mathbf{c}, \hat{\mathbf{c}}), \\ &= -H(\mathbf{c}) - H(\hat{\mathbf{c}}) + 2H(\mathbf{c}, \hat{\mathbf{c}}), \\ &= \sum_{i=1}^{k_N} \frac{n_{i+}}{N} \log\left(\frac{n_{i+}}{N}\right) + \sum_{j=1}^{\hat{k}_N} \frac{n_{+j}}{N} \log\left(\frac{n_{+j}}{N}\right) - 2 \sum_{i=1}^{k_N} \sum_{j=1}^{\hat{k}_N} \frac{n_{ij}}{N} \log\left(\frac{n_{ij}}{N}\right). \end{aligned}$$

We provide a detailed comparison with an N -invariant version of Binder's loss, defined as

$$\tilde{\text{B}}(\mathbf{c}, \hat{\mathbf{c}}) = \frac{2}{N^2} \text{B}(\mathbf{c}, \hat{\mathbf{c}}) = \sum_{i=1}^{k_N} \left(\frac{n_{i+}}{N}\right)^2 + \sum_{j=1}^{\hat{k}_N} \left(\frac{n_{+j}}{N}\right)^2 - 2 \sum_{i=1}^{k_N} \sum_{j=1}^{\hat{k}_N} \left(\frac{n_{ij}}{N}\right)^2.$$

Both loss functions are considered N -invariant as they only depend on N through the proportions n_{ij}/N . We focus on these two loss functions as they satisfy several desirable properties.

The first important property is that both VI and $\tilde{\text{B}}$ are metrics on the space of partitions.

Property 1. *Both VI and $\tilde{\text{B}}$ are metrics on the space of partitions.*

A proof for VI can be found in Meilă (2007). For $\tilde{\text{B}}$, the proof results from the fact that $\tilde{\text{B}}$ can be derived as the Hamming distance between the binary representation of the clusterings.




Figure 1: Hasse diagram for the lattice of partitions with a sample size $N = 4$. A line is drawn from \mathbf{c} up to $\hat{\mathbf{c}}$ when \mathbf{c} is covered by $\hat{\mathbf{c}}$.

The next properties involve first viewing the space of partitions as a partially ordered set. In particular, consider the space of partitions \mathbf{C} and the binary relation \leq on \mathbf{C} defined by set containment, i.e. for $\mathbf{c}, \hat{\mathbf{c}} \in \mathbf{C}$, $\mathbf{c} \leq \hat{\mathbf{c}}$ if for all $i = 1, \dots, k_N$, $C_i \subseteq \hat{C}_j$ for some $j \in \{1, \dots, \hat{k}_N\}$. The partition space \mathbf{C} equipped with \leq is a partially ordered set.

For any $\mathbf{c}, \hat{\mathbf{c}} \in \mathbf{C}$, \mathbf{c} is *covered* by $\hat{\mathbf{c}}$, denoted $\mathbf{c} \prec \hat{\mathbf{c}}$, if $\mathbf{c} < \hat{\mathbf{c}}$ and there is no $\hat{\hat{\mathbf{c}}} \in \mathbf{C}$ such that $\mathbf{c} < \hat{\hat{\mathbf{c}}} < \hat{\mathbf{c}}$. This covering relation is used to define the *Hasse diagram*, where the elements of \mathbf{C} are represented as nodes of a graph and a line is drawn from \mathbf{c} up to $\hat{\mathbf{c}}$ when $\mathbf{c} \prec \hat{\mathbf{c}}$. An example of the Hasse diagram for $N = 4$ is depicted in Figure 1.

The space of partitions possesses an even richer structure; it forms a lattice. This follows from the fact that every pair of partitions has a *greatest lower bound* and *least upper bound*; for a subset $\mathbf{S} \subseteq \mathbf{C}$, an element $\mathbf{c} \in \mathbf{C}$ is an upper bound for \mathbf{S} if $\mathbf{s} \leq \mathbf{c}$ for all $\mathbf{s} \in \mathbf{S}$, and $\mathbf{c} \in \mathbf{C}$ is the least upper bound for \mathbf{S} , denoted $\mathbf{c} = \text{l.u.b.}(\mathbf{S})$, if \mathbf{c} is an upper bound for \mathbf{S} and $\mathbf{c} \leq \mathbf{c}'$ for all upper bounds \mathbf{c}' of \mathbf{S} . A lower bound and the greatest lower bound for a subset $\mathbf{S} \subseteq \mathbf{C}$ are similarly defined, the latter denoted by $\text{g.l.b.}(\mathbf{S})$. We define the operators \wedge , called the *meet*, and \vee , called the *join*, as $\mathbf{c} \wedge \hat{\mathbf{c}} = \text{g.l.b.}(\mathbf{c}, \hat{\mathbf{c}})$ and $\mathbf{c} \vee \hat{\mathbf{c}} = \text{l.u.b.}(\mathbf{c}, \hat{\mathbf{c}})$. Following the conventions of lattice theory, we will use $\mathbf{1}$ to denote the greatest element of the lattice of partitions, i.e. the partition with every observation in one cluster $\mathbf{c} = (\{1, \dots, N\})$, and $\mathbf{0}$ to denote the least element of the lattice of partitions, i.e. the partition with every observation in its own cluster $\mathbf{c} = (\{1\}, \dots, \{N\})$. See Nation (1991) for more details on lattice theory and the Supplementary Material (Wade and Ghahramani, 2017) for specific details on the lattice of partitions.

A desirable property is that both VI and $\tilde{\mathbf{B}}$ are *aligned* with the lattice of partitions. Specifically, both metrics are *vertically aligned* in the Hasse diagram; if $\hat{\hat{\mathbf{c}}}$ is connected up to $\hat{\mathbf{c}}$ and $\hat{\mathbf{c}}$ is connected up to \mathbf{c} , then the distance between $\hat{\hat{\mathbf{c}}}$ and \mathbf{c} is the vertical sum of the distances between $\hat{\hat{\mathbf{c}}}$ and $\hat{\mathbf{c}}$ and between $\hat{\mathbf{c}}$ and \mathbf{c} (see Property 2). And, both metrics are *horizontally aligned*; the distance between any two partitions is the horizontal sum of the distances between each partition and the meet of the two partitions (see Property 3).




Figure 2: Hasse diagram stretched by VI with a sample size $N = 4$. Note $2 - \frac{3}{4} \log(3) \approx 0.811$. From the VI stretched Hasse diagram, we can determine the distance between any two partitions. Example: if $\mathbf{c} = (\{1, 2\}, \{3, 4\})$ and $\hat{\mathbf{c}} = (\{1\}, \{3\}, \{2, 4\})$, then $\mathbf{c} \wedge \hat{\mathbf{c}} = (\{1\}, \{2\}, \{3\}, \{4\})$ and $d(\mathbf{c}, \hat{\mathbf{c}}) = d(\mathbf{c} \wedge \hat{\mathbf{c}}, \mathbf{1}) - d(\mathbf{c}, \mathbf{1}) + d(\mathbf{c} \wedge \hat{\mathbf{c}}, \mathbf{1}) - d(\hat{\mathbf{c}}, \mathbf{1}) = 2 - 1 + 2 - 1.5 = 1.5$.




Figure 3: Hasse diagram stretched by \tilde{B} with a sample size $N = 4$. From the \tilde{B} stretched Hasse diagram, we can determine the distance between any two partitions. Example: if $\mathbf{c} = (\{1, 2\}, \{3, 4\})$ and $\hat{\mathbf{c}} = (\{1\}, \{3\}, \{2, 4\})$, then $\mathbf{c} \wedge \hat{\mathbf{c}} = (\{1\}, \{2\}, \{3\}, \{4\})$ and $d(\mathbf{c}, \hat{\mathbf{c}}) = d(\mathbf{c} \wedge \hat{\mathbf{c}}, \mathbf{1}) - d(\mathbf{c}, \mathbf{1}) + d(\mathbf{c} \wedge \hat{\mathbf{c}}, \mathbf{1}) - d(\hat{\mathbf{c}}, \mathbf{1}) = 0.75 - 0.5 + 0.75 - 0.625 = 0.375$.

Property 2. For both VI and \tilde{B} , if $\mathbf{c} \geq \hat{\mathbf{c}} \geq \hat{\hat{\mathbf{c}}}$, then

$$d(\mathbf{c}, \hat{\hat{\mathbf{c}}}) = d(\mathbf{c}, \hat{\mathbf{c}}) + d(\hat{\mathbf{c}}, \hat{\hat{\mathbf{c}}}).$$

Property 3. For both VI and \tilde{B} ,

$$d(\mathbf{c}, \hat{\mathbf{c}}) = d(\mathbf{c}, \hat{\mathbf{c}} \wedge \mathbf{c}) + d(\hat{\mathbf{c}}, \hat{\mathbf{c}} \wedge \mathbf{c}).$$

Proofs can be found in the Supplementary Material. These two properties imply that if the Hasse diagram is stretched to reflect the distance between any partition and $\mathbf{1}$, the distance between any two partitions can be easily determined from the *stretched Hasse diagram*. Figures 2 and 3 depict the Hasse diagram for $N = 4$ in Figure 1 stretched according to VI and \tilde{B} respectively.

From the stretched Hasse diagram, we gain several insights into the similarities and differences between the two metrics. An evident difference is the scale of the two diagrams.

Property 4. *A distance on partitions satisfying Properties 2 and 3 has the property that for any two partitions \mathbf{c} and $\hat{\mathbf{c}}$,*

$$d(\mathbf{c}, \hat{\mathbf{c}}) \leq d(\mathbf{1}, \mathbf{0}).$$

Thus,

$$VI(\mathbf{c}, \hat{\mathbf{c}}) \leq \log(N) \quad \text{and} \quad \tilde{B}(\mathbf{c}, \hat{\mathbf{c}}) \leq 1 - \frac{1}{N}.$$

A proof can be found in the Supplementary Material. In both cases, the bound on the distance between two clusterings depends on the sample size N . However, the behavior of this bound is very different; for VI, it approaches infinity as $N \rightarrow \infty$, and for \tilde{B} , it approaches one as $N \rightarrow \infty$. As N grows, the number of total partitions B_N increases drastically. Thus, it is sensible that the bound on the metric grows as the size of the space grows. In particular, $\mathbf{1}$ and $\mathbf{0}$ become more distant as $N \rightarrow \infty$, as there is an increasing number, $B_N - 2$, of partitions between these two extremes; for \tilde{B} , the loss of estimating one of these extremes with the other approaches the fixed number one, while for VI, the loss approaches infinity.

From the stretched Hasse diagram in Figures 2 and 3, we can determine the closest partitions to any \mathbf{c} . For example, the closest partitions to $\mathbf{1}$ are the partitions which split $\mathbf{1}$ into two clusters, one singleton and one containing all other observations; and the closest partitions to $(\{1\}, \{2\}, \{3, 4\})$ are the partition which merges the two smallest clusters $(\{1, 2\}, \{3, 4\})$ and the partition which splits the cluster of size two $(\{1\}, \{2\}, \{3\}, \{4\})$.

Property 5. *For both metrics VI and \tilde{B} , the closest partitions to a partition \mathbf{c} are:*

- *if \mathbf{c} contains at least two clusters of size one and at least one cluster of size two, the partitions which merge any two clusters of size one and the partitions which split any cluster of size two.*
- *if \mathbf{c} contains at least two clusters of size one and no clusters of size two, the partitions which merge any two clusters of size one.*
- *if \mathbf{c} contains at most one cluster of size one, the partitions which split the smallest cluster of size greater than one into a singleton and a cluster with the remaining observations of the original cluster.*

A proof can be found in the Supplementary Material. This property characterizes the set of estimated partitions which are given the smallest loss. Under both loss functions, the smallest loss of zero occurs when the estimated partition is equal to the truth. Otherwise, the smallest loss occurs when the estimated clustering differs from the truth by merging two singleton clusters or splitting a cluster of size two, or, if neither is possible, splitting the smallest cluster of size n into a singleton and a cluster of size $n - 1$. We further note that the loss of estimating the true clustering with a clustering which merges two singletons or splits a cluster of size two, is $\frac{2}{N}$ and $\frac{2}{N^2}$ for VI and \tilde{B} respectively, which converges to 0 as $N \rightarrow \infty$ for both metrics, but at a faster rate for \tilde{B} .

Next, we note that the Hasse diagram stretched by \tilde{B} in Figure 3 appears asymmetric, in the sense that $\mathbf{1}$ is more separated from the others when compared to the Hasse diagram stretched by VI in Figure 2.

Property 6. *Suppose N is divisible by k , and let \mathbf{c}_k denote a partition with k clusters of equal size N/k .*

$$\tilde{B}(\mathbf{1}, \mathbf{c}_k) = 1 - \frac{1}{k} > \frac{1}{k} - \frac{1}{N} = \tilde{B}(\mathbf{0}, \mathbf{c}_k).$$

$$VI(\mathbf{1}, \mathbf{c}_k) = \log(k) \leq \log(N) - \log(k) = VI(\mathbf{0}, \mathbf{c}_k), \quad \text{for } k \leq \sqrt{N},$$

and

$$VI(\mathbf{1}, \mathbf{c}_k) = \log(k) \geq \log(N) - \log(k) = VI(\mathbf{0}, \mathbf{c}_k), \quad \text{for } k \geq \sqrt{N}.$$

Property 6 reflects the asymmetry apparent in Figure 3. In particular, for \tilde{B} , a partition with two clusters of equal size \mathbf{c}_2 will always be closer to the extreme $\mathbf{0}$ of each data point in its own cluster than the extreme $\mathbf{1}$ of everyone in one cluster. However, as the sample size increases, \mathbf{c}_2 becomes equally distant between the two extremes. For all other values of k , the extreme $\mathbf{0}$ will always be closer. This behavior is counter-intuitive for a loss function on clusterings. VI is much more sensible in this regard. If $k = \sqrt{N}$, $\mathbf{0}$ and $\mathbf{1}$ are equally good estimates of \mathbf{c}_k . For $k < \sqrt{N}$, \mathbf{c}_k is better estimated by $\mathbf{1}$ and for $k > \sqrt{N}$, \mathbf{c}_k is better estimated by $\mathbf{0}$; as the sample size increases, these preferences become stronger. In particular, note that loss of estimating \mathbf{c}_2 with $\mathbf{1}$ will always be smaller than estimating it with $\mathbf{0}$ for $N > 4$.

Additionally, we observe from Figure 3 that the partitions with two clusters of sizes one and three are equally distant between the two extremes under \tilde{B} . The following property generalizes this observation.

Property 7. *Suppose N is an even and square integer. Then, the partitions with two clusters of sizes $n = \frac{1}{2}(N - \sqrt{N})$ and $N - n$ are equally distant from $\mathbf{1}$ and $\mathbf{0}$ under \tilde{B} .*

This property is unappealing for a loss function, as it states that the loss of estimating a partition consisting of two clusters of sizes $\frac{1}{2}(N - \sqrt{N})$ and $\frac{1}{2}(N + \sqrt{N})$ with the partition of only one cluster or with the partition of all singletons is the same. Intuitively, however, $\mathbf{1}$ is a better estimate. The behavior of VI is much more reasonable, as partitions with two clusters will always be better estimated by $\mathbf{1}$ than $\mathbf{0}$ for $N > 4$ and partitions with \sqrt{N} clusters of equal size are equally distant from $\mathbf{0}$ and $\mathbf{1}$.

Finally, we note that as both VI and \tilde{B} are metrics on the space of clusterings, we can construct a ball around \mathbf{c} of size ϵ , defined as:

$$B_\epsilon(\mathbf{c}) = \{\hat{\mathbf{c}} \in \mathbf{C} : d(\mathbf{c}, \hat{\mathbf{c}}) \leq \epsilon\}.$$

From Property 5, the smallest non-trivial ball will be the same for the two metrics. When considering the next smallest ball, differences emerge; a detailed example is provide in the Supplementary Material. In the authors' opinions, the VI ball more closely reflects our intuition of the closest set of partitions to \mathbf{c} .

4 Point estimation via the variation of information

As detailed in the previous section, both VI and \tilde{B} share several desirable properties including being aligned with the lattice of partitions and coinciding in the smallest non-trivial ball around any clustering. However, in our comparison, differences also emerged. Particularly, we find that \tilde{B} exhibits some peculiar asymmetries, preferring to split clusters over merging, and we find that the VI ball more closely reflects our intuition of the neighborhood of a partition. In light of this, we propose to use VI as a loss function in Bayesian cluster analysis. Under the VI, the optimal partition \mathbf{c}^* is

$$\begin{aligned} \mathbf{c}^* &= \underset{\hat{\mathbf{c}}}{\operatorname{argmin}} \mathbb{E}[\text{VI}(\mathbf{c}, \hat{\mathbf{c}}) | \mathcal{D}] \\ &= \underset{\hat{\mathbf{c}}}{\operatorname{argmin}} \sum_{n=1}^N \log\left(\sum_{n'=1}^N \mathbf{1}(\hat{c}_{n'} = \hat{c}_n)\right) - 2 \sum_{n=1}^N \mathbb{E}\left[\log\left(\sum_{n'=1}^N \mathbf{1}(c_{n'} = c_n, \hat{c}_{n'} = \hat{c}_n)\right) | \mathcal{D}\right], \end{aligned} \quad (3)$$

with \mathcal{D} denoting the data. For a given $\hat{\mathbf{c}}$, the second term in (3) can be approximated based on the MCMC output, and evaluating this term is of order $O(MN^2)$ (recall M is the number of MCMC samples). This may be computationally demanding if the number of MCMC samples is large and if (3) must be evaluated for a large number of $\hat{\mathbf{c}}$. Alternatively, one can use Jensen's inequality, swapping the log and expectation, to obtain a lower bound on the expected loss which is computationally more efficient to evaluate:

$$\underset{\hat{\mathbf{c}}}{\operatorname{argmin}} \sum_{n=1}^N \log\left(\sum_{n'=1}^N \mathbf{1}(\hat{c}_{n'} = \hat{c}_n)\right) - 2 \sum_{n=1}^N \log\left(\sum_{n'=1}^N P(c_{n'} = c_n | \mathcal{D}) \mathbf{1}(\hat{c}_{n'} = \hat{c}_n)\right). \quad (4)$$

Similar to minimization of the posterior expected Binder's loss, minimization of (4) only depends on the posterior through the posterior similarity matrix, which can be pre-computed based on the MCMC output. In this case, computational complexity for a given $\hat{\mathbf{c}}$ is reduced to $O(N^2)$.

Due to the huge dimensions of the partition space, computing the lower bound in (4) for every possible $\hat{\mathbf{c}}$ is practically impossible. A simple technique to find the optimal partition \mathbf{c}^* restricts the search space to some smaller space of partitions. The **R** package 'mcclust' (Fritsch (2012)), which contains tools for point estimation in Bayesian cluster analysis and cluster comparison, includes a function `minbinder()` that finds the partition minimizing the poster expected Binder's loss among the subset of partitions 1) visited in the MCMC chain or 2) explored in a hierarchical clustering algorithm with a distance of $1 - P(c_n = c_{n'} | \mathcal{D})$ and average or complete linkage. An alternative search algorithm developed in Lau and Green (2007), which is based on binary integer programming, is also implemented.

We propose a greedy search algorithm to locate the optimal partition \mathbf{c}^* based on the Hasse diagram, which can be used for both VI and \tilde{B} . In particular, given some partition $\hat{\mathbf{c}}$, we consider the l closest partitions that cover $\hat{\mathbf{c}}$ and the l closest partitions that $\hat{\mathbf{c}}$ covers. Here, the distance used to determine the closest partitions corresponds to the selected loss of VI or \tilde{B} . Next, the posterior expected loss $\mathbb{E}[L(\mathbf{c}, \hat{\mathbf{c}}) | \mathcal{D}]$ is computed for all proposed partitions $\hat{\mathbf{c}}$, and we move in the direction of minimum posterior expected loss, that is the partition \mathbf{c}' with minimal $\mathbb{E}[L(\mathbf{c}, \mathbf{c}') | \mathcal{D}]$ is selected. The algorithm stops when

no reduction in the posterior expected loss is obtained or when a maximum number of iterations has been reached. At each iteration, the computational complexity is $O(lN^2)$.

We have developed an **R** package ‘mcclust.ext’ (Wade (2015)), expanding upon the ‘mcclust’ package, that is currently available on the author’s website¹ and includes functions `minbinder.ext()` and `minVI()` to find the partition minimizing the posterior expected Binder’s loss and VI, respectively. In addition to implementing the search algorithms of `minbinder()` in ‘mcclust’ described previously, the greedy search algorithm is also included. As is common in greedy search algorithms, results are sensitive to both the starting value of $\hat{\mathbf{c}}$ and the step size l . In practice, we recommend multiple restarts, for example, at different MCMC samples or the best partition found by the other search algorithms. A larger value of l will allow more exploration and reduce the need for multiple restarts, and we have chosen a default value of $l = 2N$ as this showed good exploration in the examples considered with little sensitivity to the initial value of $\hat{\mathbf{c}}$. However, for larger datasets, this may be too expensive and multiple restarts with smaller l may be preferred.

An advantage of the greedy search algorithm over simply restricting to partitions visited in the chain is that partitions not explored in the MCMC algorithm can be considered; in fact, in almost all simulated and real examples, the clustering estimate is not among the sampled partitions and results in a lower expected loss than any sampled partition.

5 Credible balls of partitions

To characterize the uncertainty in the point estimate \mathbf{c}^* , we propose to construct a credible ball of a given credible level $1 - \alpha$, $\alpha \in [0, 1]$, defined as

$$B_{\epsilon^*}(\mathbf{c}^*) = \{\mathbf{c} : d(\mathbf{c}^*, \mathbf{c}) \leq \epsilon^*\},$$

where ϵ^* is the smallest $\epsilon \geq 0$ such that $P(B_{\epsilon}(\mathbf{c}^*)|\mathcal{D}) \geq 1 - \alpha$. The credible ball is the smallest ball around \mathbf{c}^* with posterior probability at least $1 - \alpha$. It reflects the posterior uncertainty in the point estimate \mathbf{c}^* ; with probability $1 - \alpha$, we believe that the clustering is within a distance of ϵ^* from the point estimate \mathbf{c}^* given the data. It can be defined based on any metric on the space of partitions, such as VI and $\tilde{\mathbf{B}}$. If the smallest non-trivial ball under VI or $\tilde{\mathbf{B}}$ has posterior probability of at least $1 - \alpha$, the credible balls under the two metrics will coincide (see Property 5). Typically, however, they will be different.

From the MCMC output, we can obtain an estimate of ϵ^* , and thus the credible ball of level $1 - \alpha$. First, the distance between all MCMC samples $\{\mathbf{c}^m\}$ and \mathbf{c}^* is computed. For any $\epsilon \geq 0$,

$$P(B_{\epsilon}(\mathbf{c}^*)|\mathcal{D}) = \mathbb{E}[\mathbf{1}(d(\mathbf{c}^*, \mathbf{c}) \leq \epsilon)|\mathcal{D}] \approx \frac{1}{M} \sum_{m=1}^M \mathbf{1}(d(\mathbf{c}^*, \mathbf{c}^m) \leq \epsilon),$$

and ϵ^* is the smallest $\epsilon \geq 0$ such that $\frac{1}{M} \sum_{m=1}^M \mathbf{1}(d(\mathbf{c}^*, \mathbf{c}^m) \leq \epsilon) \geq 1 - \alpha$.

¹<https://www2.warwick.ac.uk/fac/sci/statistics/staff/academic-research/wade/>.

To characterize the credible ball, we define the *vertical* and *horizontal bounds* of the credible ball. The vertical upper bounds consist of the partitions in the credible ball with the smallest number of clusters that are most distant from \mathbf{c}^* . The vertical lower bounds consist of the partitions in the credible ball with the largest number of clusters that are most distant from \mathbf{c}^* . The horizontal bounds consist of the partitions in the credible ball that are most distant from \mathbf{c}^* . The bounds are defined more formally below, where the notation $k(\mathbf{c})$ is used for the number of clusters in \mathbf{c} .

Definition 1 (Vertical upper bounds). *The vertical upper bounds of the credible ball $B_{\epsilon^*}(\mathbf{c}^*)$, denoted $v_{\epsilon^*}^u(\mathbf{c}^*)$, are defined as*

$$v_{\epsilon^*}^u(\mathbf{c}^*) = \{\mathbf{c} \in B_{\epsilon^*}(\mathbf{c}^*) : k(\mathbf{c}) \leq k(\mathbf{c}') \forall \mathbf{c}' \in B_{\epsilon^*}(\mathbf{c}^*) \text{ and} \\ d(\mathbf{c}, \mathbf{c}^*) \geq d(\mathbf{c}'', \mathbf{c}^*) \forall \mathbf{c}'' \in B_{\epsilon^*}(\mathbf{c}^*) \text{ with } k(\mathbf{c}) = k(\mathbf{c}'')\}.$$

Definition 2 (Vertical lower bounds). *The vertical lower bounds of the credible ball $B_{\epsilon^*}(\mathbf{c}^*)$, denoted $v_{\epsilon^*}^l(\mathbf{c}^*)$, are defined as*

$$v_{\epsilon^*}^l(\mathbf{c}^*) = \{\mathbf{c} \in B_{\epsilon^*}(\mathbf{c}^*) : k(\mathbf{c}) \geq k(\mathbf{c}') \forall \mathbf{c}' \in B_{\epsilon^*}(\mathbf{c}^*) \text{ and} \\ d(\mathbf{c}, \mathbf{c}^*) \geq d(\mathbf{c}'', \mathbf{c}^*) \forall \mathbf{c}'' \in B_{\epsilon^*}(\mathbf{c}^*) \text{ with } k(\mathbf{c}) = k(\mathbf{c}'')\}.$$

Definition 3 (Horizontal bounds). *The horizontal bounds of the credible ball $B_{\epsilon^*}(\mathbf{c}^*)$, denoted $h_{\epsilon^*}(\mathbf{c}^*)$, are defined as*

$$h_{\epsilon^*}(\mathbf{c}^*) = \{\mathbf{c} \in B_{\epsilon^*}(\mathbf{c}^*) : d(\mathbf{c}, \mathbf{c}^*) \geq d(\mathbf{c}', \mathbf{c}^*) \forall \mathbf{c}' \in B_{\epsilon^*}(\mathbf{c}^*)\}.$$

These bounds describe the extremes of the credible ball and with $1 - \alpha$ posterior probability, how different we believe the partition may be from \mathbf{c}^* . An example is provided in the Supplementary Material. In practice, we define the vertical and horizontal bounds based on the partitions in the credible ball with positive estimated posterior probability.

In existing literature, quantification of uncertainty in the clustering structure is typically described through a heat map of the estimated posterior similarity matrix. However, as opposed to the credible ball of Bayesian confidence level $1 - \alpha$, there is no precise quantification of how much uncertainty is represented by the posterior similarity matrix. Moreover, in the examples of Section 6, we find that in a comparison with the 95% credible balls, the uncertainty is under-represented by the posterior similarity matrix. Additionally, the credible balls have the added desirable interpretation of characterizing the uncertainty around the point estimate \mathbf{c}^* .

6 Examples

We provide both simulated and real examples to compare the point estimates from VI and Binder's loss and describe the credible ball representing uncertainty in the clustering estimate.




Figure 4: The data is simulated from a mixture of four normals with locations $(\pm 2, \pm 2)'$ and colored by cluster membership. In (b) components having varying standard deviations.

6.1 Simulated examples

Two datasets of size $n = 200$ are simulated from:

$$X_i \stackrel{iid}{\sim} \sum_{j=1}^4 \frac{1}{4} N \left(\begin{bmatrix} (-1)^{\lfloor \frac{j-1}{2} \rfloor} 2 \\ (-1)^{j-1} 2 \end{bmatrix}, \begin{bmatrix} \sigma_j^2 & 0 \\ 0 & \sigma_j^2 \end{bmatrix} \right).$$

In the first example, $\sigma_j = 1$ for all components, while in the second example, components have varying standard deviations; $\sigma_j = 1$ for the two components located in the first and third quadrants, $\sigma_j = 0.5$ in the second quadrant, and $\sigma_j = 1.5$ in the fourth quadrant. The datasets for both examples are depicted in Figure 4 and colored by cluster membership.

We consider a Dirichlet process (DP) mixture model:

$$X_i | P \stackrel{iid}{\sim} \int N \left(\begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}, \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix} \right) dP(\mu, \Sigma) \quad \text{and} \quad P \sim \text{DP}(\alpha P_0), \quad (5)$$

where $\mu = (\mu_1, \mu_2)'$ and Σ is a diagonal matrix with diagonal elements (σ_1^2, σ_2^2) . The base measure of the DP is the conjugate product of normal inverse gamma priors with parameters $(\mu_{0,i}, c_i, a_i, b_i)$ for $i = 1, 2$, i.e. P_0 has density

$$p_0(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2) \propto \prod_{i=1}^2 \sqrt{\frac{c_i}{\sigma_i^2}} \exp \left(-\frac{c_i}{2\sigma_i^2} (\mu_i - \mu_{0,i})^2 \right) (\sigma_i^2)^{-a_i-1} \exp \left(-\frac{b_i}{\sigma_i^2} \right).$$

The parameters were fixed to $\mu_{0,i} = 0$, $c_i = 1/2$, $a_i = 2$, $b_i = 1$ for $i = 1, 2$. The mass parameter α is given a $\text{Gam}(1, 1)$ hyperprior.

A marginal Gibbs sampler is used for inference (Neal (2000)) with 10,000 iterations after a burn in period of 1,000 iterations. Trace plots and autocorrelation plots (not

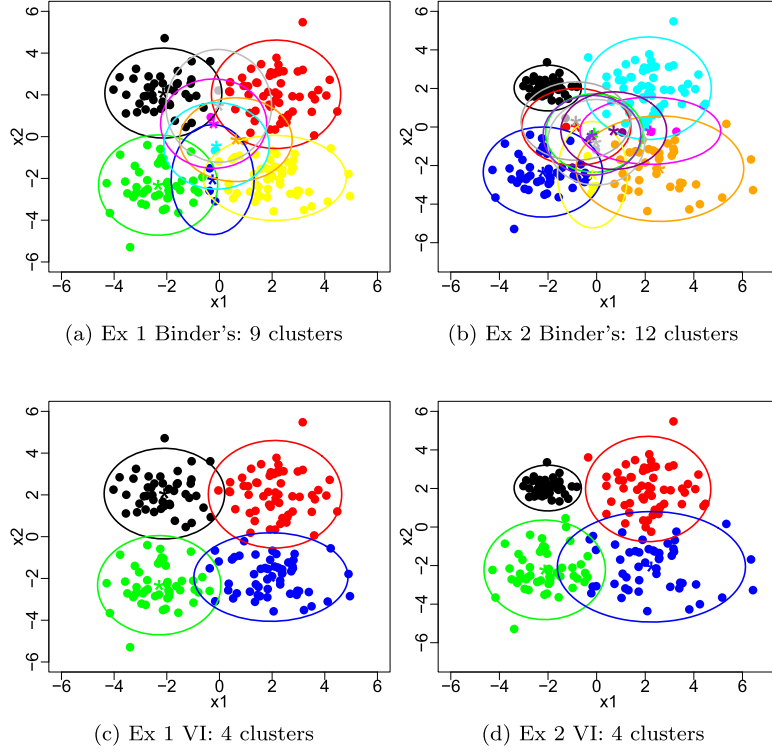


Figure 5: Clustering estimate with color representing cluster membership for Binder's loss (first row) and VI (second row) with columns corresponding to examples.

shown) suggest convergence. Among partitions sampled in the MCMC, only one is visited twice and all others are visited once in the first example, while no partitions are visited more than once in the second example.

Figure 5 depicts the partition estimate found by the greedy search algorithm for Binder's loss and VI and for both examples (with multiple restarts and the default value of $l = 2N$); colors represent cluster membership with the posterior expected cluster-specific mean and variance represented through stars and ellipses, respectively. Tables in the Supplementary Material provide a comparison of the true partition with the estimates through a cross tabulation of cluster labels. In all examples, the four true clusters are visible; however, Binder's loss creates new small clusters for observations on the border between clusters where cluster membership is uncertain, overestimating the number of clusters. This effect is most extreme for the second example, where the fourth cluster (blue in Figure 4b) has increased overlap with the second and third clusters (red and green in Figure 4b), while the first cluster (black in Figure 4b) with decreased variance is well separated from the other clusters and identified in both estimates.

A further comparison of the true partition with the estimates under Binder's loss and VI, for both examples, is provided in Table 1. As expected, the \tilde{B} estimate and

	Loss	k_N^*	N_I	$\mathbb{E}[\tilde{B} \mathcal{D}]$	$\tilde{B}(\mathbf{c}_t, \mathbf{c}^*)$	$\mathbb{E}[\text{VI}_{\text{LB}} \mathcal{D}]$	$\mathbb{E}[\text{VI} \mathcal{D}]$	$\text{VI}(\mathbf{c}_t, \mathbf{c}^*)$
Ex 1:	\tilde{B}	9	13	0.062	0.045	0.545	0.816	0.643
	VI	4	9	0.064	0.044	0.426	0.77	0.569
Ex 2:	\tilde{B}	12	18	0.088	0.056	0.846	1.068	0.764
	VI	4	10	0.093	0.049	0.668	0.99	0.561

Table 1: A comparison of the clustering estimate with \tilde{B} or VI in terms of 1) number of clusters k_N^* ; 2) number of data points incorrectly classified, denoted N_I ; 3) expected \tilde{B} ; 4) \tilde{B} between the optimal and true clusterings; 5) expected lower bound of VI; 6) expected VI; and 7) VI between the optimal and true clusterings for both examples.

Ex 1	Loss	k_N^*	N_I	$\mathbb{E}[\tilde{B} \mathcal{D}]$	$\tilde{B}(\mathbf{c}_t, \mathbf{c}^*)$	$\mathbb{E}[\text{VI}_{\text{LB}} \mathcal{D}]$	$\mathbb{E}[\text{VI} \mathcal{D}]$	$\text{VI}(\mathbf{c}_t, \mathbf{c}^*)$
$N = 200$:	\tilde{B}	9	13	0.062	0.045	0.545	0.816	0.643
	VI	4	9	0.064	0.044	0.426	0.77	0.569
$N = 400$:	\tilde{B}	17	31	0.068	0.052	0.674	1.0	0.769
	VI	4	18	0.073	0.044	0.505	0.933	0.54
$N = 800$:	\tilde{B}	24	62	0.068	0.061	0.615	1.016	0.903
	VI	4	47	0.069	0.056	0.477	0.943	0.742
$N = 1600$:	\tilde{B}	41	93	0.058	0.044	0.551	0.898	0.719
	VI	4	49	0.0596	0.045	0.403	0.814	0.629

Table 2: Example 1 with increasing sample size: a comparison of the clustering estimate with \tilde{B} or VI in terms of 1) number of clusters k_N^* ; 2) number of data points incorrectly classified, denoted N_I ; 3) expected \tilde{B} ; 4) \tilde{B} between the optimal and true clusterings; 5) expected lower bound of VI; 6) expected VI; and 7) VI between the optimal and true clusterings.

VI estimate achieve the lowest posterior expected loss for \tilde{B} and VI, respectively, but interestingly, the VI estimate has the smallest distance from the truth for both \tilde{B} and VI in both examples, with the greatest improvement in the second example. Furthermore, the number of incorrectly classified data points is greater for the \tilde{B} estimate than the VI estimate.

Additional simulated experiments were performed to analyze the effect of increasing the sample size in the first example. The results are succinctly summarized in Table 2. As the sample size increases, more points are located on the border where cluster membership is uncertain. This results in an increasing number of clusters in the \tilde{B} estimate (up to 41 clusters for $N = 1600$), while the VI estimate contains only four clusters for all sample sizes. In both estimates, the number of incorrectly classified data points increases with the sample size, however this number is smaller for the VI estimate in all sample sizes, with the difference between this number for Binder’s and VI growing with the sample size. Furthermore, the VI estimate has improved VI distance with truth and improved or comparable \tilde{B} distance with truth when compared with the \tilde{B} estimate.

Further experiments were carried out to consider highly unbalanced clusters. In this case, the conclusions continue to hold; Binder’s loss overestimates the number of clusters present, placing uncertain observations in new small clusters, and this effect becomes more pronounced with increased overlap between clusters (results not shown).

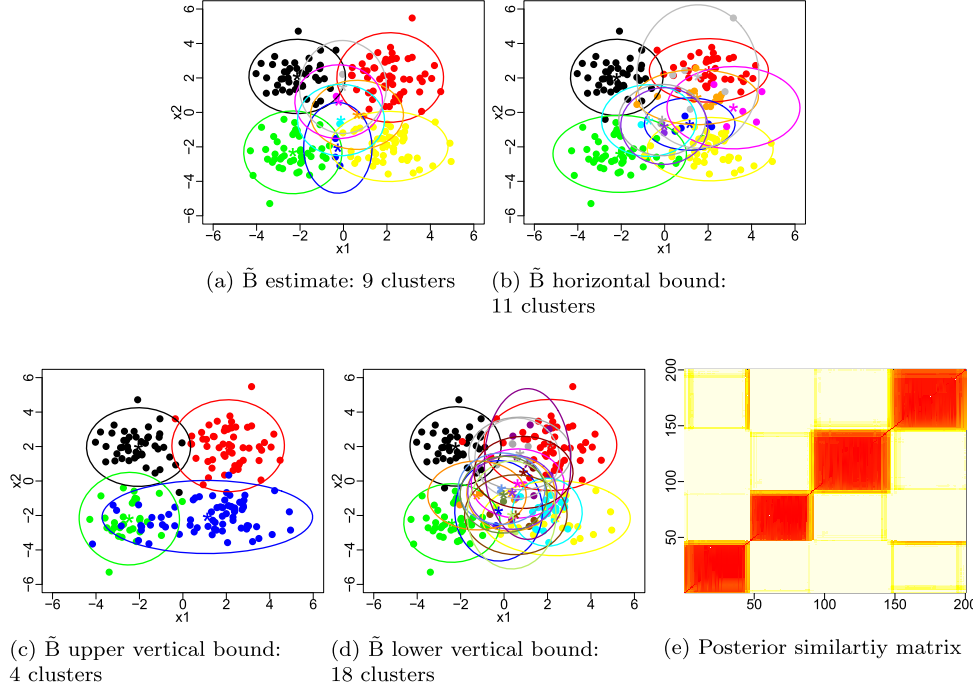


Figure 6: Example 1: 95% credible ball with Binder's loss around \mathbf{c}^* (a) represented by the (b) horizontal bound, (c) upper vertical bound, and (d) lower vertical bound, where color denotes cluster membership, and a heat map of the posterior similarity matrix (e).

For the first example, Figures 6 and 7 represent the 95% credible ball around the optimal partition for $\tilde{\mathbf{B}}$ and VI, respectively, through the upper vertical bound, lower vertical bound, and horizontal bound, with data points colored according to cluster membership. Analogous plots for the second example are found in Figures 8 and 9. The Supplementary Material provides tables comparing the bounds with the true clustering through a cross tabulation of the true cluster labels with the cluster labels for each bound.

In the first example, we observe that elements of the 95% credible ball with positive estimated posterior probability have at least four clusters for both metrics and at most 18 clusters for $\tilde{\mathbf{B}}$ or 16 clusters for VI, while the most distant elements contain 11 clusters for $\tilde{\mathbf{B}}$ and VI (Table 3). For both metrics, these bounds reallocate uncertain data points on the border with these points either added to one of the four main clusters or to new small to medium-sized clusters. For example, in the $\tilde{\mathbf{B}}$ upper bound, 19 elements of the third cluster (green in Figure 4a) are added to the fourth cluster (blue in Figure 4a) and in the $\tilde{\mathbf{B}}$ lower bound, the fourth cluster (blue in Figure 4a) is split in two medium-sized clusters and several small clusters.

In the second example, the first cluster (black in Figure 4b) is stable in all bounds, while the 95% credible ball reflects posterior uncertainty on whether to divide the re-

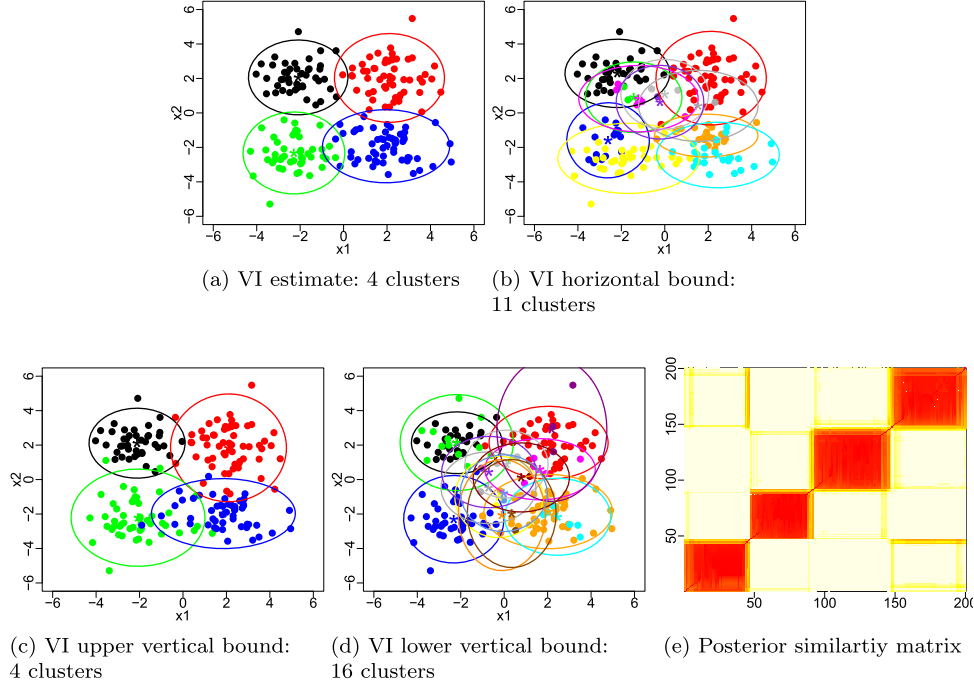


Figure 7: Example 1: 95% credible ball with VI around \mathbf{c}^* (a) represented by the (b) horizontal bound, (c) upper vertical bound, and (d) lower vertical bound, where color denotes cluster membership, and a heat map of the posterior similarity matrix (e).

	Loss	Upper		Lower		Horizontal	
		k_N^u	$d(\mathbf{c}^*, \mathbf{c}_u)$	k_N^l	$d(\mathbf{c}^*, \mathbf{c}_l)$	k_N^h	$d(\mathbf{c}^*, \mathbf{c}_h)$
Ex 1:	$\tilde{\mathbf{B}}$	4	0.097	18	0.097	11	0.097
	VI	4	1.02	16	1.152	11	1.213
Ex 2:	$\tilde{\mathbf{B}}$	4	0.137	19	0.131	10	0.137
	VI	3	1.043	16	1.342	6	1.403

Table 3: A summary of the credible bounds with $\tilde{\mathbf{B}}$ or VI in terms of the number of clusters and distance to the clustering estimate for the upper vertical, lower vertical, and horizontal bounds and for both examples.

maintaining data points into 3 to 18 clusters for $\tilde{\mathbf{B}}$ and 2 to 15 clusters for VI (Table 3). Notice the high uncertainty in the fourth cluster with increased variance (blue in Figure 4b). Additionally, note the greater uncertainty around the optimal estimate in Example 2, as the horizontal distance in Table 3 is greater for Example 2 for both metrics.

Figures 6–9 also present heat maps of the posterior similarity matrix for both examples. In general, the posterior similarity matrix appears to under-represent the uncertainty; indeed, one would conclude from the similarity matrix that there is only

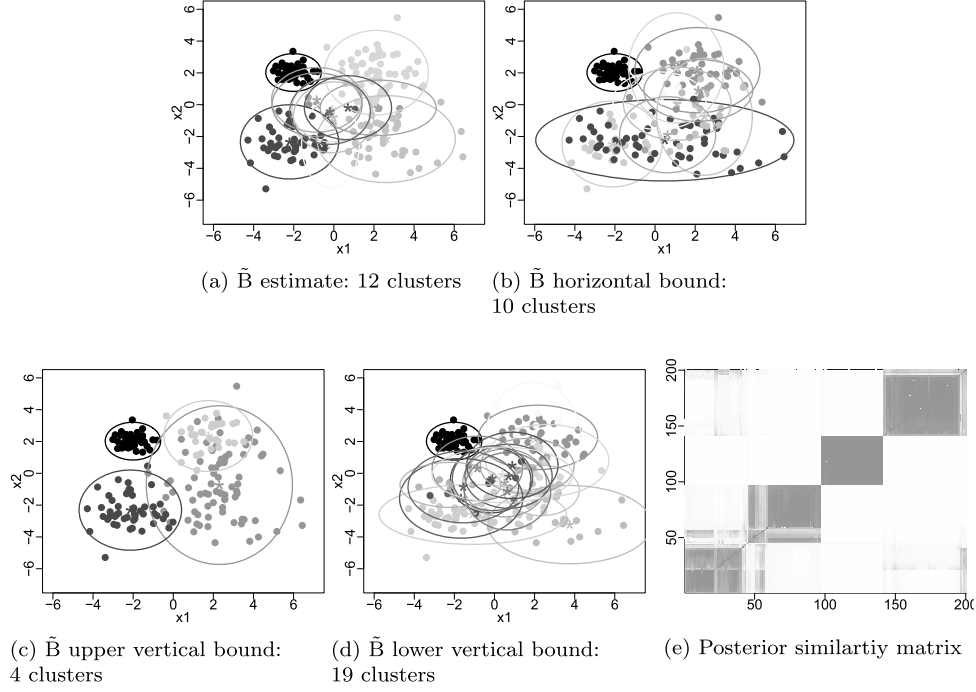


Figure 8: Example 2: 95% credible ball with Binder's loss around \mathbf{c}^* (a) represented by the (b) horizontal bound, (c) upper vertical bounds (only one of two shown for conciseness), and (d) lower vertical bound, where color denotes cluster membership, and a heat map of the posterior similarity matrix (e).

uncertainty in allocation of a few data points in Example 1. Moreover, the 95% credible ball gives a precise quantification of the uncertainty.

6.2 Galaxy example

We consider an analysis of the galaxy data (Roeder (1990)), available in the MASS package of **R**, which contains measurements of velocities in km/sec of 82 galaxies from a survey of the Corona Borealis region. The presence of clusters provides evidence for voids and superclusters in the far universe. The data is modeled with a DP mixture (5). The parameters were selected empirically with $\mu_0 = \bar{x}$, $c = 1/2$, $a = 2$, $b = s^2$, where \bar{x} represents the sample mean and s^2 represents the sample variance. The mass parameter α is given a $\text{Gam}(1, 1)$ hyperprior.

With 10,000 samples after 1,000 burn in, the posterior mass is spread out over 9,636 partitions, emphasizing the need for appropriate summary tools. Figure 10 plots the point estimate of the partition found by the greedy search algorithm for Binder's loss and VI (with multiple restarts and the default value of $l = 2N$). The data values are

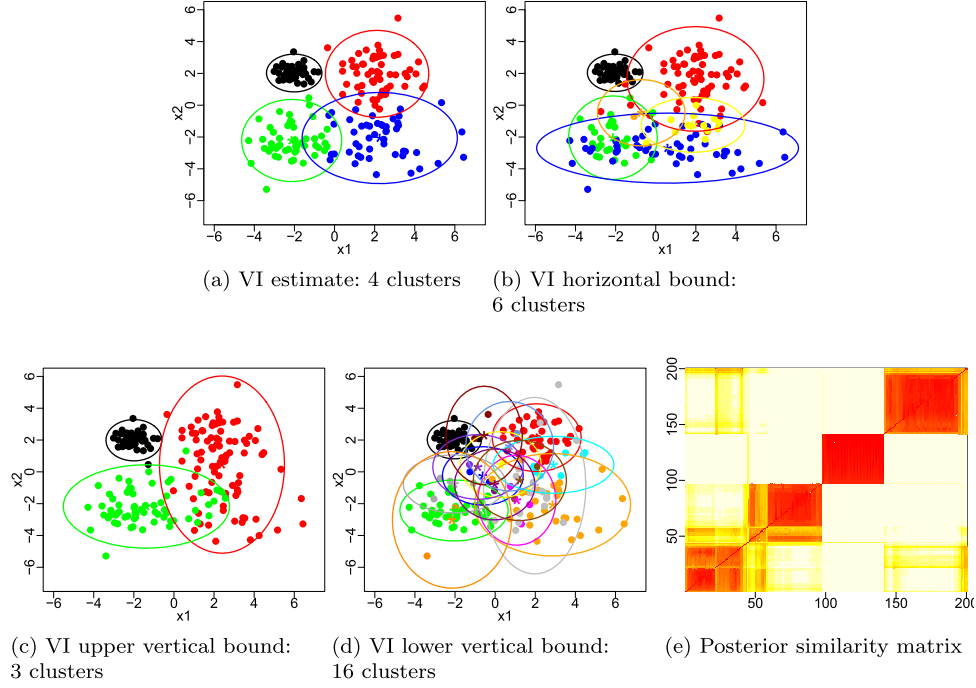


Figure 9: Example 2: 95% credible ball with VI around \mathbf{c}^* (a) represented by the (b) horizontal bound, (c) upper vertical bound, and (d) lower vertical bound, where color denotes cluster membership and a heat map of the posterior similarity matrix (e).

plotted against the estimated density values from the DP mixture model and colored according to cluster membership, with correspondingly colored stars and bars along the x-axis representing the posterior mean and variance within cluster. Again, we observe that Binder's loss places observations with uncertain allocation into singleton clusters, with a total of 7 clusters, 4 of which are singletons, while the VI solution contains 3 clusters. Table 4 compares the point estimates in terms of the posterior expected \tilde{B} , lower bound of VI, and VI; as anticipated, the \tilde{B} solution has the smallest posterior expected \tilde{B} and the VI solution has the smallest posterior expected VI.

Loss	k_N^*	$\mathbb{E}[\tilde{B} \mathcal{D}]$	$\mathbb{E}[\text{VI}_{\text{LB}} \mathcal{D}]$	$\mathbb{E}[\text{VI} \mathcal{D}]$
B	7	0.218	0.746	1.014
VI	3	0.237	0.573	0.939

Table 4: Galaxy example: a comparison of the optimal partition with Binder's loss and VI in terms of posterior expected \tilde{B} , lower bound to VI, and VI.

The 95% VI credible ball contains all partitions with a VI distance less than 1.832. Figure 11 depicts the 95% credible ball through the upper vertical, lower vertical, and horizontal bounds, which are further described and summarized in Table 5 and




Figure 10: Galaxy example: optimal clustering estimate with color representing cluster membership for Binder's loss and VI, with correspondingly colored stars and bars along the x-axis representing the posterior mean and variance within cluster.

	Upper		Lower		Horizontal	
	k_N^u	$d(\mathbf{c}^*, \mathbf{c}_u)$	k_N^l	$d(\mathbf{c}^*, \mathbf{c}_l)$	k_N^h	$d(\mathbf{c}^*, \mathbf{c}_h)$
Galaxy	2	1.364	15	1.669	8	1.832

Table 5: Galaxy example: a summary of the credible bounds with VI in terms of the number of clusters and distance to the clustering estimate for the upper vertical, lower vertical, and horizontal bounds.

in cross tabulation tables in the Supplementary Material. We observe a large amount of variability around the optimal partition. With 95% posterior probability, we believe that, on one extreme, the data could be modeled using only 2 components, one with a large variance to account for outliers (black cluster in Figure (11a)). On the other extreme, the data could be further split into one medium sized cluster and many, 14 to be precise, smaller clusters. The horizontal bound, the most extreme partition in the credible ball, splits the largest cluster (red in Figure 10b) into two medium sized clusters and four small clusters and reallocates some of its data points to the first cluster (black in Figure 10b). Figure 11d emphasizes that the posterior similarity matrix under-represents the uncertainty around the point estimate in comparison to the credible ball.

7 Discusssion

Bayesian cluster analysis provides an advantage over classical cluster analysis, in that the Bayesian procedure returns a posterior distribution over the entire partition space, reflecting uncertainty in the clustering structure given the data, as opposed to returning a single solution or conditioning on the parameter estimates and number of clusters. This allows one to assess statistical properties of the clustering given the data. However, due to the huge dimension of the partition space, an important problem in Bayesian cluster analysis is how to appropriately summarize the posterior. To address this problem, we

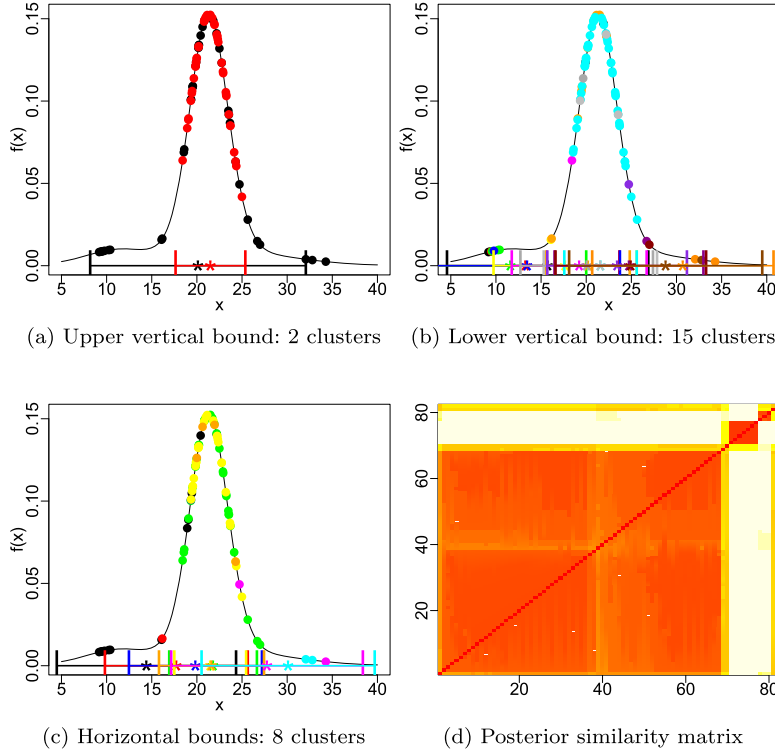


Figure 11: Galaxy example: 95% credible ball with VI represented by the (a) upper vertical bound, (b) lower vertical bound, and (c) horizontal bound, where color denotes cluster membership, with correspondingly colored stars and bars along the x-axis representing the posterior mean and variance within cluster, and (d) a heat map of the posterior similarity matrix.

have developed tools to obtain a point estimate of clustering based on the posterior and describe uncertainty around this estimate via the 95% credible ball.

Obtaining a point estimate through a formal decision theory framework requires the specification of a loss function. Previous literature focused on Binder’s loss. In this work, we propose to use an information theoretic measure, the variation of information, and provide a detailed comparison of the two metrics. We find that Binder’s loss exhibits peculiar asymmetries, preferring to split over merge clusters, and the variation of information is more symmetric in this regard. This behavior of Binder’s loss causes the optimal partition to overestimate the number of clusters, allocating uncertain data points to small additional clusters. In addition, we have developed a novel greedy search algorithm to locate the optimal partition, allowing one to explore beyond the space of partitions visited in the MCMC chain.

To represent uncertainty around the point estimate, we construct 95% credible balls around the point estimate and depict the credible ball through the upper vertical, lower

vertical, and horizontal bounds. In addition to a heat map of the posterior similarity matrix, which is often reported in literature, the 95% credible ball enriches our understanding of the uncertainty present. Indeed, it provides a precise quantification of the uncertainty present around the point estimate, and in examples, we find that an analysis based on the posterior similarity matrix leads one to be over certain in the clustering structure. The developed posterior summary tools for Bayesian cluster analysis are available² through an **R** package ‘mcclust.ext’ (Wade (2015)), expanding upon the existing **R** package ‘mcclust’ (Fritsch (2012)).

In future work, we aim to extend these ideas to Bayesian feature allocation analysis, an extension of clustering which allows observations to belong to multiple clusters (Griffiths and Ghahramani (2011)). A further direction of research will be to explore posterior consistency for the number of clusters based on the VI estimate for Bayesian nonparametric mixture models; this is of particular interest in light of the negative results of Miller and Harrison (2013) and Miller and Harrison (2014) and the positive results in our simulation studies (Table 2). Finally, scalability issues of Bayesian nonparametric mixture models are an important concern for very large datasets. To scale with large sample sizes, a number of papers have avoided exploration of the posterior on partitions through MCMC and focused on finding a point estimate of the partition, often through MAP inference (Heller and Ghahramani (2005), Dahl (2009), Raykov et al. (2014)) or the DP-means algorithm and its extensions (Kulis and Jordan (2012), Jiang et al. (2012), Broderick et al. (2013)). One direction of future research is to develop an algorithm to find the point estimate which minimizes the posterior expected VI that avoids MCMC. Of course, while gaining in scalability, we lose the uncertainty in the clustering structure.

Supplementary Material

Supplementary material for Bayesian cluster analysis: Point estimation and credible balls (DOI: [10.1214/17-BA1073SUPP](https://doi.org/10.1214/17-BA1073SUPP); .pdf).

References

- Binder, D. (1978). “Bayesian Cluster Analysis.” *Biometrika*, 65: 31–38. [MR0501592](#). doi: <http://dx.doi.org/10.1093/biomet/65.1.31>. 3, 6
- Broderick, T., Kulis, B., and Jordan, M. (2013). “MAD-Bayes: MAP-based asymptotic derivations from Bayes.” In *Proceedings of the 30th International Conference on Machine Learning*, 226–234. 25
- Dahl, D. (2006). “Model-based clustering for expression data via a Dirichlet process mixture model.” In Do, K., Müller, P., and Vannucci, M. (eds.), *Bayesian Inference for Gene Expression and Proteomic*, 201–218. Cambridge University Press. [MR2706330](#).

7

²Through the author’s website <https://www2.warwick.ac.uk/fac/sci/statistics/staff/academic-research/wade/>.

- Dahl, D. (2009). “Modal clustering in a class of product partition models.” *Bayesian Analysis*, 4: 243–264. MR2507363. doi: <http://dx.doi.org/10.1214/09-BA409>. 3, 25
- Duan, J., Guindani, M., and Gelfand, A. (2007). “Generalized spatial Dirichlet processes.” *Biometrika*, 94: 809–825. MR2416794. doi: <http://dx.doi.org/10.1093/biomet/asm071>. 6
- Dunson, D. (2010). “Nonparametric Bayes applications to biostatistics.” In Hjort, N., Holmes, C., Müller, P., and Walker, S. (eds.), *Bayesian nonparametrics*. Cambridge University Press. MR2730665. 6
- Favaro, S. and Teh, Y. (2013). “MCMC for normalized random measure mixture models.” *Statistical Science*, 28: 335–359. MR3135536. doi: <http://dx.doi.org/10.1214/13-STS422>. 6
- Favaro, S. and Walker, S. (2012). “Slice sampling σ -stable Poisson–Kingman mixture models.” *Journal of Computational and Graphical Statistics*, 22: 830–847. MR3173745. doi: <http://dx.doi.org/10.1080/10618600.2012.681211>. 6
- Ferguson, T. (1973). “A Bayesian analysis of some nonparametric problems.” *Annals of Statistics*, 1: 209–230. MR0350949. 5
- Fraley, C. and Raftery, A. (2002). “Model-based clustering, discriminant analysis, and density estimation.” *Journal of the American Statistical Association*, 97: 611–631. MR1951635. doi: <http://dx.doi.org/10.1198/016214502760047131>. 1
- Fritsch, A. (2012). *mcclust: Process an MCMC Sample of Clusterings*. URL <http://cran.r-project.org/web/packages/mcclust/mcclust.pdf> 13, 25
- Fritsch, A. and Ickstadt, K. (2009). “Improved criteria for clustering based on the posterior similarity matrix.” *Bayesian Analysis*, 4: 367–392. MR2507368. doi: <http://dx.doi.org/10.1214/09-BA414>. 3, 7
- Griffin, J. and Steel, M. (2006). “Order-based dependent Dirichlet processes.” *Journal of the American Statistical Association*, 10: 179–194. MR2268037. doi: <http://dx.doi.org/10.1198/016214505000000727>. 6
- Griffiths, T. and Ghahramani, Z. (2011). “The Indian buffet process: An introduction and review.” *Journal of Machine Learning Research*, 12: 1185–1224. MR2804598. 25
- Hartigan, J. and Wong, M. (1979). “Algorithm AS 136: A k-means clustering algorithm.” *Journal of the Royal Statistical Society, Series C*, 28: 100–108. MR0405726. 1
- Heard, N., Holmes, C., and Stephens, D. (2006). “A quantitative study of gene regulation involved in the immune response of anopheline mosquitos: An application of Bayesian hierarchical clustering of curves.” *Journal of the American Statistical Association*, 101: 18–29. MR2252430. doi: <http://dx.doi.org/10.1198/016214505000000187>. 3
- Heller, K. and Ghahramani, Z. (2005). “Bayesian hierarchical clustering.” In *Proceedings of the 22nd International Conference on Machine Learning*, 297–304. 3, 25

- Hubert, L. and Arabie, P. (1985). “Comparing partitions.” *Journal of Classification*, 2: 193–218. 7
- Ishwaran, H. and James, L. (2001). “Gibbs sampling methods for stick-breaking priors.” *Journal of the American Statistical Association*, 96: 161–173. MR1952729. doi: <http://dx.doi.org/10.1198/016214501750332758>. 5, 6
- Jiang, K., Kulis, B., and Jordan, M. (2012). “Small-variance asymptotics for exponential family Dirichlet process mixture models.” In *Advances in Neural Information Processing Systems*, 3158–3166. 25
- Kalli, M., Griffin, J., and Walker, S. (2011). “Slice sampling mixture models.” *Statistics and Computing*, 21: 93–105. MR2746606. doi: <http://dx.doi.org/10.1007/s11222-009-9150-y>. 6
- Kulis, B. and Jordan, M. (2012). “Revisiting K-means: New algorithms via Bayesian nonparametrics.” In *Proceedings of the 29th International Conference on Machine Learning*, 513–520. 25
- Lau, J. and Green, P. (2007). “Bayesian model-based clustering procedures.” *Journal of Computational and Graphical Statistics*, 16: 526–558. MR2351079. doi: <http://dx.doi.org/10.1198/106186007X238855>. 3, 7, 13
- Lijoi, A. and Prünster, I. (2011). “Models beyond the Dirichlet process.” In Hjort, N., Holmes, C., Müller, P., and Walker, S. (eds.), *Bayesian Nonparametrics*, 80–136. Cambridge, UK: Cambridge University Press. MR2730661. 5
- Lo, A. (1984). “On a class of Bayesian nonparametric estimates: I. Density estimates.” *Annals of Statistics*, 12: 351–357. MR0733519. doi: <http://dx.doi.org/10.1214/aos/1176346412>. 2
- Lomellí, M., Favaro, S., and Teh, Y. (2015). “A hybrid sampler for Poisson–Kingman mixture models.” In Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 28*. 6
- Lomellí, M., Favaro, S., and Teh, Y. (2016). “A marginal sampler for σ -stable Poisson–Kingman mixture models.” *Journal of Computational and Graphical Statistics*. To appear. 6
- MacEachern, S. (2000). “Dependent Dirichlet processes.” *Technical Report, Department of Statistics, Ohio State University*. 6
- Medvedovic, M. and Sivaganesan, S. (2002). “Bayesian infinite mixture model based clustering of gene expression profiles.” *Bioinformatics*, 18: 1194–1206. 3
- Medvedovic, M., Yeung, K., and Bumgarner, R. (2004). “Bayesian mixture model based clustering of replicated microarray data.” *Bioinformatics*, 20: 1222–1232. 3
- Meilă, M. (2007). “Comparing clusterings – an information based distance.” *Journal of Multivariate Analysis*, 98: 873–895. MR2325412. doi: <http://dx.doi.org/10.1016/j.jmva.2006.11.013>. 3, 8

- Miller, J. and Harrison, M. (2013). “A simple example of Dirichlet process mixture inconsistency for the number of components.” In Burges, C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. (eds.), *Advances in Neural Information Processing Systems 26*. Curran Associates, Inc. 25
- Miller, J. and Harrison, M. (2014). “Inconsistency of Pitman–Yor process mixtures for the number of components.” *Journal of Machine Learning Research*, 15: 3333–3370. MR3277163. 25
- Molitor, J., Papathomas, M., Jerrett, M., and Richardson, S. (2010). “Bayesian profile regression with an application to the national survey of children’s health.” *Biostatistics*, 11: 484–498. 3
- Müller, P. and Quintana, F. (2004). “Nonparametric Bayesian data analysis.” *Statistical Science*, 19: 95–110. MR2082149. doi: <http://dx.doi.org/10.1214/088342304000000017>. 6
- Nation, J. (1991). *Notes on Lattice Theory*. <http://www.math.hawaii.edu/~jb/books.html>. 9
- Neal, R. (2000). “Markov chain sampling methods for Dirichlet process mixture models.” *Journal of Computational and Graphical Statistics*, 9: 249–265. MR1823804. doi: <http://dx.doi.org/10.2307/1390653>. 6, 16
- Papaspiliopoulos, O. and Roberts, G. (2008). “Retrospective Markov chain Monte Carlo methods for Dirichlet process hierarchical models.” *Biometrika*, 95(1): 169–186. MR2409721. doi: <http://dx.doi.org/10.1093/biomet/asm086>. 6
- Pitman, J. (2003). “Poisson Kingman partitions.” In *Statistics and Science: a Festschrift for Terry Speed*, 1–34. Beachwood: IMS Lecture Notes. MR2004330. doi: <http://dx.doi.org/10.1214/lnms/1215091133>. 5
- Pitman, J. and Yor, M. (1997). “The two-parameter Poisson–Dirichlet distribution derived from a stable subordinator.” *Annals of Probability*, 25: 855–900. MR1434129. doi: <http://dx.doi.org/10.1214/aop/1024404422>. 5
- Quintana, F. (2006). “A predictive view of Bayesian clustering.” *Journal of Statistical Planning and Inference*, 136: 2407–2429. MR2279815. doi: <http://dx.doi.org/10.1016/j.jspi.2004.09.015>. 2
- Quintana, F. and Iglesias, P. (2003). “Bayesian clustering and product partition models.” *Journal of the Royal Statistical Society: Series B*, 65: 557–574. MR1983764. doi: <http://dx.doi.org/10.1111/1467-9868.00402>. 3, 7
- Rand, W. (1971). “Objective criteria for the evaluation of clustering methods.” *Journal of the American Statistical Association*, 66: 846–850. 7
- Rasmussen, C., De la Cruz, B., Ghahramani, Z., and Wild, D. (2009). “Modeling and visualizing uncertainty in gene expression clusters using Dirichlet process mixtures.” *Computational Biology and Bioinformatics, IEEE/ACM Transactions on*, 6: 615–628. 3

- Raykov, Y., Boukouvalas, A., and Little, M. (2014). “Simple approximate MAP Inference for Dirichlet processes.” Available at <https://arxiv.org/abs/1411.0939>. MR3572859. doi: <http://dx.doi.org/10.1214/16-EJS1196>. 3, 25
- Roeder, K. (1990). “Density estimation with confidence sets exemplified by superclusters and voids in galaxies.” *Journal of the American Statistical Association*, 85: 617–624. 21
- Teh, Y., Jordan, M., Beal, M., and Blei, D. (2006). “Hierarchical Dirichlet process.” *Journal of the American Statistical Association*, 101: 1566–1581. MR2279480. doi: <http://dx.doi.org/10.1198/016214506000000302>. 6
- Vinh, N., Epps, J., and Bailey, J. (2010). “Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance.” *Journal of Machine Learning Research*, 11: 2837–2854. MR2738784. 8
- Wade, S. (2015). *mcclust.ext: Point estimation and credible balls for Bayesian cluster analysis*. URL https://www.researchgate.net/publication/279848500_mcclustext-manual. 14, 25
- Wade, S. and Ghahramani, Z. (2017). “Supplementary material for Bayesian cluster analysis: Point estimation and credible balls.” *Bayesian Analysis*. doi: <http://dx.doi.org/10.1214/17-BA1073SUPP>. 9

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/I036575/1].