

Statistica Sinica Preprint No: SS-2021-0387

Title	Statistical Inference With Anchored Bayesian Mixture of Regressions Models: An Illustrative Study of Allometric Data
Manuscript ID	SS-2021-0387
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202021.0387
Complete List of Authors	Deborah Kunkel and Mario Peruggia
Corresponding Authors	Deborah Kunkel
E-mails	dekunke@clemson.edu

Notice: Accepted version subject to English editing.

Statistical inference with anchored Bayesian mixture of regressions models: An illustrative study of allometric data

Deborah Kunkel¹ and Mario Peruggia²

1. School of Mathematical & Statistical Sciences, Clemson University, Clemson, SC, USA

2. Department of Statistics, The Ohio State University, Columbus, OH, USA

Abstract:

We present an illustrative study in which we use a mixture of regressions model to improve on an ill-fitting simple linear regression model relating log brain mass to log body mass for 100 placental mammalian species. The slope of the model is of particular scientific interest because it corresponds to a constant that governs a hypothesized allometric power law relating brain mass to body mass. We model these data using an anchored Bayesian mixture of regressions model, which modifies the standard Bayesian Gaussian mixture by pre-assigning small subsets of observations to given mixture components with probability one. These observations (called anchor points) break the relabeling invariance (or label-switching) typical of exchangeable models. In the article, we develop a strategy for selecting anchor points using tools from case influence diagnostics. We compare the performance of three anchoring methods on the allometric data and in simulated settings.

Key words and phrases: Case-deletion weights, Clustering, EM algorithm

1. Introduction

In the natural sciences, allometry studies the relationships between physical and physiological measurements taken on various animal species (Peters, 1983; Gayon, 2000). Of particular interest is to determine how other measurements may be affected by body mass. Examples include the relationships between body mass and brain mass, body mass and metabolic rate, body mass and gestation duration. It is often postulated that pairs (x, y) of such measurements may be related via a power law of the form $y = cx^b$, for some unknown constants c and b , typically assumed to be positive. The estimation of the exponent b is often of primary scientific interest. On a logarithmic scale, the power law turns into the linear relationship $\log y = (\log c) + b \log x$ and the investigative focus shifts toward the estimation of the slope of the regression line.

Given a set of (x, y) pairs of traits measured on a variety of species, it is by now generally accepted that fitting a single linear regression model to the entire data set provides too crude a summary, especially when many species from different taxa and genetically diverse groups are included in the data set (Jerison, 1955; Bennett and Harvey, 1985a,b). More refined approaches rely on the incorporation of evolutionary information (possibly inferred from a taxonomy) to perform an analysis based on models for derived quantities that can be treated as independent, rather than for the original measured traits that exhibit species-related dependencies. For example, this is the case for a popular type of analysis

based on phylogenetically independent contrasts (Felsenstein, 1985; Garland Jr et al., 1992). MacEachern and Peruggia (2002) show that traditional Bayesian variance components models applied directly to allometric data for which taxonomic information is available can produce a good fit and yield easily interpretable inferences.

In this article we reanalyze the data that MacEachern and Peruggia (2002) used to illustrate their methods. The data comprise the body and brain mass measurements on 100 species of placental mammals originally reported by Sacher and Staffeldt (1974) as well as a taxonomy that assigns each species to an order and sub-order based on its morphological and physiological traits. In total, the data contain species that represent 13 orders and 19 sub-orders.

The data are shown in Figure 1. The left panel displays a scatterplot of the centered log body mass and log brain mass and the least-squares fit from a naive simple linear regression model. The residuals for this model are shown in the right panel of Figure 1. This plot raises some concerns about model fit. First, there is a slight increase in residual variability as log body mass grows. Second, other features of the residuals can be traced back to the species orders (distinguished by plotting color): all Primates (red points) have positive residuals, while most Rodentia (blue points) have negative residuals. This structure points to the fact that the least squares line does not properly account for within-order similarities in the allometric relationship.

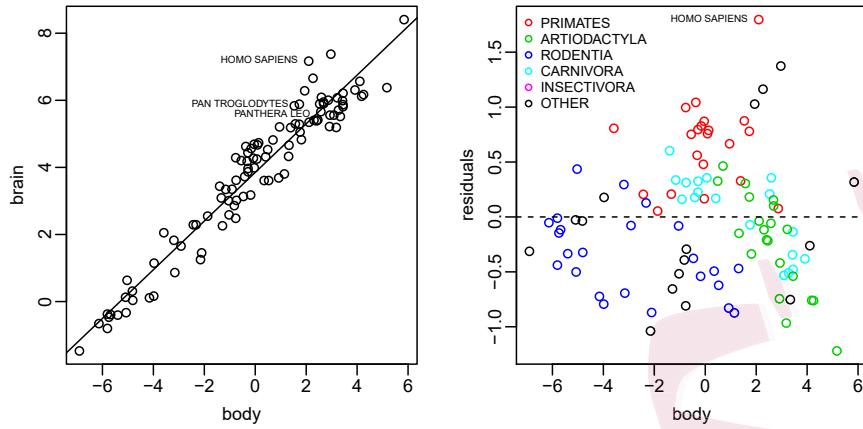


Figure 1: Mammals data (left) with the estimated least-squares regression line and residuals (right) from the least-squares regression fit.

MacEachern and Peruggia (2002) present a detailed evaluation that uncovers the lack-of-fit of this model, introducing Bayesian model diagnostic techniques based on case-deletion importance sampling weights (Geweke, 1989; Bradlow and Zaslavsky, 1997; Peruggia, 1997; Epifani et al., 2008; Thomas et al., 2018). They also present a Bayesian variance components model that includes additive random effects for orders and sub-orders. These random effects induce positive correlations between the residuals of species belonging to the same taxonomic groups and subgroups. The random effect adjustment effectively creates a separate regression line with its own intercept for the species within each subgroup. MacEachern and Peruggia (2002) show that this way of accounting for the taxonomic information significantly ameliorates the quality of the fit.

Suppose now that no information about the taxonomy were available. We

might wish to remedy a suspected lack-of-fit of the simple linear regression model, but a random effects model is not an option when the true group memberships are not known. In such situations, a reasonable modeling alternative is to assume that there exist finitely many subgroups of observations for which separate regression lines are appropriate, and yet it is unknown which observations to associate with which regression. This leads naturally to the formulation of a mixture of regressions model in which, conditional on *unobserved* group membership indicators, observations falling in separate groups follow separate regression models. Using these data as the foundation for an illustrative study, we consider a Bayesian mixture of regressions to account for unobserved heterogeneity in a data set. The fundamental difference between the mixture setting and the random effects setting considered by MacEachern and Peruggia (2002) is that, in the former, group memberships are latent and are the subject of inferential investigation. An additional difference is that we will allow for group-specific slopes in addition to group-specific intercepts.

A unique challenge arises when modeling with Bayesian mixture models: a typical Bayesian formulation would specify, *a priori*, a fully exchangeable mixture, yielding a fully exchangeable posterior that is indifferent to any arbitrary relabeling of the mixture components. This phenomenon, often referred to as label switching, precludes informative assignment of labels to the mixture components and prevents meaningful interpretation of the component-specific parameter esti-

mates. Kunkel and Peruggia (2020) introduce a modeling device called anchoring that breaks the labeling invariance of the exchangeable version of a Bayesian Gaussian mixture model. The idea is to identify small subsets of representative observations, called the anchor points, and allocate them to separate components before performing the analysis. Kunkel and Peruggia (2020) characterize the equivalence between this process and the specification of a weakly data-dependent prior for the component specific parameters. A shrewd selection of anchor points is essential for the successful implementation of the strategy. They propose a strategy for model specification based on a modified EM algorithm.

Motivated by the mixture of regressions model and its application to our allometric data, we extend the work of Kunkel and Peruggia (2020) in two main directions. First, we generalize the EM anchoring strategy to the case of mixture of regressions models. Second, we introduce a new strategy for the selection of anchor points that builds on the work on case-deletion analysis presented in MacEachern and Peruggia (2002) and Thomas et al. (2018). We compare this new strategy for the selection of anchor points (which we call CDW-reg) to the EM anchoring strategy based on the mixture of regressions model (which we call EM-reg). The strategies differ in the way they extract information from the data to select the points representative of the various mixture components.

In Section 2 we present the mixture of regressions model and its corresponding anchor model. In Section 3 we present the EM method for selecting anchor points

under the mixture of regressions model. In Section 4 we describe the CDW-reg method for choosing anchor points via clustering of case-deletion weights. The results of our simulations and of the analysis of the mammals data are presented in Sections 5 and 6. A few final considerations are discussed in Section 7.

2. Background

2.1 Mixture of regressions model

The k -component mixture of linear regressions model specifies that the observations $(y_i, \mathbf{x}_i), i = 1, \dots, n$, are drawn from k homogeneous subgroups within the population. The response y_i is a scalar and $\mathbf{x}_i \in R^p$ is a row vector of predictors whose first element is equal to one. The data are denoted by $\mathbf{y}' = (y_1, \dots, y_n)$ and by the $n \times p$ matrix \mathbf{X} with i -th row equal to \mathbf{x}_i . The mixture likelihood is

$$f(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta}, \sigma^2, \boldsymbol{\eta}) = \prod_{i=1}^n \sum_{j=1}^k \eta_j \phi(y_i; \mathbf{x}_i \boldsymbol{\beta}_j, \sigma^2), \quad (2.1)$$

where $\boldsymbol{\eta}' = (\eta_1, \dots, \eta_k)$ is a vector of mixture probabilities that satisfy $\sum_{j=1}^k \eta_j = 1$ and $\phi(\cdot; a, b)$ denotes the density function of a normal distribution with mean a and variance b , evaluated at its argument. The component-specific regression parameters in this model are $\boldsymbol{\beta} = (\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_k)$, where $\boldsymbol{\beta}_j \in R^p$ is the vector of regression coefficients associated with the j th component. This model can be written equivalently by introducing latent allocations $\mathbf{s} = (s_1, \dots, s_n)$, where $s_i = j$ indicates that observation i was generated from component j . Specifying that $P(S_i = j) = \eta_j$, for $j = 1, \dots, k$, produces the mixture of regressions model

2.1 Mixture of regressions model

in (2.1). Conditional on $s_i = j$, the mean of y_i is $\mathbf{x}_i\boldsymbol{\beta}_j$ and, with the prior specification given below, the model is a random effects regression.

Assuming conditional independence throughout, we specify the following exchangeable prior:

$$\begin{aligned}\boldsymbol{\beta}_j \mid \boldsymbol{\mu}_\beta, \mathbf{V} &\sim N_p(\boldsymbol{\mu}_\beta, \mathbf{V}), \quad j = 1, \dots, k, \\ \sigma^{-2} \mid a, b &\sim \text{Gamma}(a, b), \\ \boldsymbol{\eta} &\sim \text{Dirichlet}(\alpha \mathbf{1}_k),\end{aligned}\tag{2.2}$$

where \mathbf{V} is a $p \times p$ diagonal matrix whose q -th diagonal element is v_{q-1} and the Gamma distribution is parameterized to have mean a/b . The exchangeable specification is a natural way of expressing prior ignorance about the relation between y and \mathbf{x} within each group: we make no assumptions that induce different prior distributions on the $\boldsymbol{\beta}_j$. A consequence of the exchangeable specification is that the posterior density of $\boldsymbol{\beta}$ is invariant to relabeling: for any permutation of the integers $1, \dots, k$, denoted by $\rho_q(1 : k)$, $p(\boldsymbol{\beta} | \mathbf{y})$ is equal to $p(\boldsymbol{\beta}_{\rho_q(1:k)} | \mathbf{y})$ for all $\boldsymbol{\beta}$. This produces $k!$ symmetric regions in the posterior distribution of $\boldsymbol{\beta}$, each corresponding to one of the $k!$ possible labelings of the mixture components. Inferentially, this is undesirable because marginal distributions of the component specific parameters are identical, making it impossible to use estimates of the labeled parameters $\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_k$ to infer differences among features of the k groups. Computationally, the label switching phenomenon hampers fitting the model by Markov chain Monte Carlo simulation (Jasra et al., 2005).

2.2 Anchor models

2.2 Anchor models

In previous work, (Kunkel and Peruggia, 2020) we have proposed a new class of models called *anchor models* which pre-classify some observations in order to induce a non-exchangeable, data-dependent prior which can alleviate the inferential nuisances caused by the model's posterior exchangeability. The model is defined using k index sets A_1, \dots, A_k , where A_j contains the indices of a small number of observations which will be pre-labeled (or “anchored”) to component j . The number of points in A_j , denoted by m_j , is chosen ahead of time and each observation is anchored to at most one component; i.e., $A_j \cap A_{j'} = \emptyset$ for $j \neq j'$.

Given the anchor points, $A = \bigcup_{j=1}^k A_j$, the anchored version of model (2.1) arises by modifying the distribution on the latent allocations, such that

$$P(S_i = j) = \begin{cases} 1, & i \in A_j, \\ 0, & i \in A_{j'}, j' \neq j, \end{cases} \quad (2.3)$$

and $P(S_i = j) = \eta_j$ for $i \notin A$ as in the model (2.1).

An anchored mixture model can be regarded as a hybrid between a random effects model, in which the class membership is known for all observations, and a pure mixture model, in which the class membership is unknown for all observations. The properties of this model depend on which and how many anchor points are selected. Kunkel and Peruggia (2020) show that an anchor model can result in a nearly-unimodal posterior density on the component-specific parameter when the anchor points are judiciously chosen and at least $k - 1$ components have anchor points. They propose a modified expectation-maximization (EM) algorithm

for specifying anchor points for a multivariate Normal mixture model. The next section describes an extension of this method to mixture of regressions models.

3. Anchoring with the EM algorithm

Kunkel and Peruggia (2020) propose the anchored EM algorithm, a modified version of the EM algorithm for maximum a posteriori estimation of anchored mixture models. The standard EM algorithm for mixture models obtains estimates of the model parameters $\boldsymbol{\theta}$ by iterating between an E-step, which estimates the probability distribution $p(\mathbf{s})$ on the latent allocations conditional on the current estimate of $\boldsymbol{\theta}$, and an M-step, which updates $\boldsymbol{\theta}$ to maximize the expected (with respect to $p(\mathbf{s})$) joint posterior density of $\boldsymbol{\theta}$ and \mathbf{s} . This approach can be viewed as iteratively maximizing an objective function $F(q, \boldsymbol{\theta}) = E_q(\log(p(\mathbf{y}, \boldsymbol{\theta}, \mathbf{s})/q(\mathbf{s})))$ with respect to a distribution $q(\mathbf{s})$ on the latent allocations (E step) and $\boldsymbol{\theta}$ (M-step) Neal and Hinton (1998).

The anchored EM algorithm modifies the E-step by constraining the distribution $q(\mathbf{s})$ to correspond to an anchor model; that is, a distribution for which m_j observations are allocated to component j with probability 1, for $j = 1, \dots, k$.

It can be shown that the optimal distribution $q(\cdot)$ is that which is closest to the posterior distribution on \mathbf{s} conditional on $\boldsymbol{\theta}$ and \mathbf{y} (Kunkel and Peruggia, 2020).

The resulting algorithm recovers an anchoring structure that closely approximates the unanchored posterior distribution of $\boldsymbol{\theta}$ near one of its local modes.

3.1 Anchored EM for the mixture of regressions

3.1 Anchored EM for the mixture of regressions

The EM method for the mixture of regressions model (2.1) comprises these steps.

Initialization. Choose a small tolerance > 0 . Set $t = 1$; $\Delta = 100$. Initialize

$$\theta^0 = (\boldsymbol{\beta}^0, \sigma^0, \boldsymbol{\eta}^0).$$

While $\Delta > \text{tolerance}$ **do:**

E-step. Calculate r_{ij}^t for $i = 1, \dots, n$, $j = 1, \dots, k$, where r_{ij} is the posterior probability that $S_i = j$ given $\mathbf{y}, \mathbf{X}, \boldsymbol{\beta}, \sigma, \boldsymbol{\eta}$, and satisfies

$$r_{ij}^t \propto \eta_j^t \phi(y_i; \mathbf{x}_i \boldsymbol{\beta}_j^t, \sigma^{2t}) \quad (3.4)$$

Anchor step. For fixed values m_j , $j = 1, \dots, k$, update the anchor points by

finding $A^t = \bigcup_{j=1}^k A_j^t$ to maximize $\sum_{j=1}^k \sum_{i \in A_j} r_{ij}^t$, subject to $A_j \cap A_{j'} = \emptyset$

and $|A_j| = m_j$ for all $j \neq j'$. Let

$$\tilde{r}_{ij}^t = \begin{cases} r_{ij}^t & \text{if } i \notin A^t \\ 1 & \text{if } i \in A_j^t \\ 0 & \text{if } i \in A_{j'}^t, \quad j' \neq j \end{cases} \quad (3.5)$$

M-step. Update $\theta^t = (\boldsymbol{\beta}^t, \sigma^t, \boldsymbol{\eta}^t)$ to maximize $F^*(q^t, \theta)$, where

$$F^*(q^t, \theta) = \log(p(\boldsymbol{\beta}, \sigma, \boldsymbol{\eta})) + \sum_{j=1}^k \sum_{i=1}^n \tilde{r}_{ij}^t \log(\phi(y_i; \mathbf{x}_i \boldsymbol{\beta}_j, \sigma^2)), \quad (3.6)$$

$p(\boldsymbol{\beta}, \sigma, \boldsymbol{\eta})$ is given in 2.2, and $F^*(q^t, \theta)$ is equal to $F(q^t, \theta)$ plus terms that are constant with respect to θ .

3.2 Anchoring on the mammals data

Update $\Delta = F^*(q^t, \theta^t) - F^*(q^{t-1}, \theta^{t-1})$. Set $t = t + 1$.

End do

Return A_j^t , $j = 1, \dots, k$ and $F^*(q^t, \theta^t)$.

Update steps and recommendations for initialization and remediation of possible non-convergence are given in Supplement 1.

3.2 Anchoring on the mammals data

We now present an analysis of the mammals data using three different anchor models, which differ in the method used to select the anchor points. All models use the same number of mixture components ($k = 3$) and the same number of anchor points in each component ($m = 3$). Selecting the number of mixture components is an open research question, made more difficult by the intractability of the mixture likelihood, poor theoretical properties of standard selection criteria, and the non-identifiability of the model (Roeder, 1994; Nobile, 2004). Several promising strategies that treat k as a random quantity have been developed in recent years (Malsiner-Walli et al., 2016; Miller and Harrison, 2018), and these model-based approaches may be developed in future work into methods that can specify both the anchor points and the number of components. In this study, the choice of three components is motivated by a practical desire to restrain model complexity while retaining sufficient flexibility to capture salient morphological differences between groups of species. Traditional information indices support the

3.3 EM anchoring on the mammals data

choice of a small number of components: the BIC favors $k = 2$ and increases as k ranges from 2 to 6 (BIC = 192, 197, 211, 220, 233, for $k = 2, \dots, 6$), while the AIC is smallest at $k = 3$ with little evidence that $k \geq 4$ aids in model fit (AIC = 175.9, 173.6, 179.3, 180.7, 186.5 for $k = 2, \dots, 6$). Comments on a two-component fit are given in Section 7.

In selecting m , we confine the anchor points to be small subsets of optimally chosen observations. A minimum of one anchor point for $k - 1$ components is required to eliminate the prior exchangeability of the model that leads to label-switching. We exceed this minimum in order to specify data-dependent prior information that can partially characterize the locations and scales of underlying groups. However, we keep m small because our previous work has demonstrated that when the groups are not well-separated, using many anchor points can lead to poor out-of-sample predictive performance (Kunkel and Peruggia, 2020). The sensitivity analysis outlined in Supplement 3 indicates that inferences on the mammals data are similar when we reduce the number of anchor points from $m = 3$ to $m = 2$, but that $m = 1$ appears to provide insufficient prior information to identify three distinct lines.

3.3 EM anchoring on the mammals data

We now apply the anchored EM method outlined in the previous sections to the mammals data. We set the prior hyperparameters in (2.2) to be $a = 5$, $b = 1$,

$\mu_\beta = (3.5, 0.6)'$, and $v_0 = 1$, $v_1 = 0.5$. The left panel of Figure 2 shows the anchor points selected. For ease of exposition, we code mixture grouping by color. The anchored EM essentially performs an approximate maximum a posteriori fit of the mixture model, and it selects anchor points to be those with largest probability of allocation to their respective components. Thus, anchor points assigned to the same component, are, in a sense, those points that would be least likely to be assigned to either of the remaining components. The anchor points for the red and blue groups exhibit some variability in the x-direction and identify approximately parallel lines with different intercepts. The green points are clustered close together in both the x- and y-directions, but suggest a line with a steeper slope than that of the other two groups. The auxiliary information about the orders of the species selected as anchors, shown in the bottom panel of Figure 2, indicates that each component has two anchor points from a particular order. This suggests that the anchor points may be partially identifying the underlying structure that is driven by the taxonomic information.

4. Anchoring with case deletion weights

In a Bayesian context, case-deletion analysis quantifies the influence of an individual observation on the overall analysis by comparing the posterior distribution conditional on the entire data set and the posterior distribution conditional on the reduced data set obtained by omitting the observation under consideration.

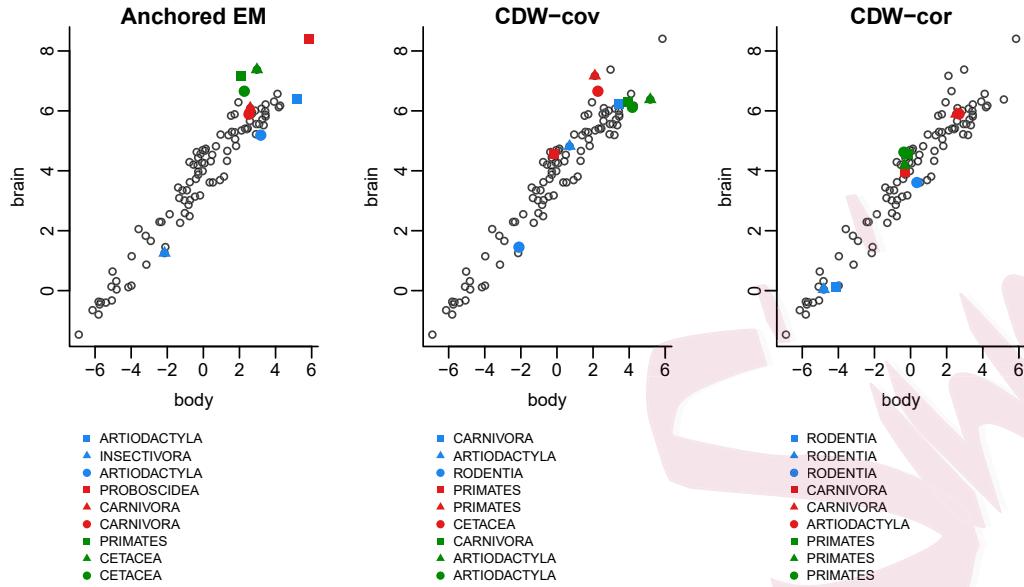


Figure 2: Selected anchor points for the mammals data. The legend indicates the order of each anchor point.

Case-deletion analysis is an effective tool for identifying influential observations in Bayesian models (Bradlow and Zaslavsky, 1997; MacEachern and Peruggia, 2002). It also provides an avenue to assess the similarity of the impact of observations on the inferential conclusions. We now derive a strategy for selecting the anchor points of a Bayesian Gaussian mixture model as “typical” representatives of clusters identified via a preliminary case-deletion analysis based on a model which assumes a single regression line to describe all data points.

Consider a set of observations $\mathbf{y} = (y_1, \dots, y_n)$ following a model with density $f(\mathbf{y}|\boldsymbol{\theta})$ conditional on a set of parameters $\boldsymbol{\theta}$ having prior distribution $\pi(\boldsymbol{\theta})$. The posterior distribution of $\boldsymbol{\theta}$ given \mathbf{y} is proportional to the joint distribution of \mathbf{y}

and $\boldsymbol{\theta}$ which is $p(\mathbf{y}, \boldsymbol{\theta}) = f(\mathbf{y}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})$. Similarly, denoting by $\mathbf{y}_{\setminus i}$ the reduced data set obtained by deleting observation i , the posterior distribution of $\boldsymbol{\theta}$ given $\mathbf{y}_{\setminus i}$ is proportional to $p_{\setminus i}(\mathbf{y}_{\setminus i}, \boldsymbol{\theta}) = f(\mathbf{y}_{\setminus i}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})$.

The ratio between the case-deleted and full posterior reacts to the influence of the deleted case on the inferential conclusions. For this reason it is useful to understand the behavior of the random variables

$$w_i(\boldsymbol{\theta}) = \frac{p_{\setminus i}(\mathbf{y}_{\setminus i}, \boldsymbol{\theta})}{p(\mathbf{y}, \boldsymbol{\theta})}, \quad i = 1, \dots, n, \quad (4.7)$$

when $\boldsymbol{\theta}$ follows the posterior distribution conditional of the entire data set. In practice, we can compute the normalized empirical versions of the ratios in (4.7) using a sample $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_L$ from (approximately) the posterior distribution of $\boldsymbol{\theta}$ given \mathbf{y} . These quantities, known as case-deletion importance sampling weights, are given by

$$\bar{w}_i(\boldsymbol{\theta}_\ell) = \frac{w_i(\boldsymbol{\theta}_\ell)}{\sum_{m=1}^L w_i(\boldsymbol{\theta}_m)}, \quad i = 1, \dots, n; \quad \ell = 1, \dots, L. \quad (4.8)$$

The theoretical variability of the $w_i(\boldsymbol{\theta})$ and the sample variability of the $w_i(\boldsymbol{\theta}_\ell)$ and $\bar{w}_i(\boldsymbol{\theta}_\ell)$, $\ell = 1, \dots, L$, are indicators of the influence of observation i on the posterior distribution, with higher variability indicating larger influence (Bradlow and Zaslavsky, 1997; Peruggia, 1997; Epifani et al., 2008). The covariance matrix (with respect to the full posterior distribution of $\boldsymbol{\theta}$) of the log weights, $\mathbf{C} = [C_{ij}] = [Cov(\log w_i(\boldsymbol{\theta}), \log w_j(\boldsymbol{\theta}))]$, is a particularly useful quantity: C_{ij} can be interpreted as summarizing the degree of similarity between the influence of deletion of case i and case j .

Further, for models of conditional independence, Thomas et al. (2018) detail the existence of a close relationship between \mathbf{C} and measures of influence based on infinitesimal geometric perturbations of the multiplicative contributions of each observation to the overall likelihood. In the perturbed likelihood, the multiplicative factor corresponding to each observation is raised to a power ω_i , $i = 1 \dots, n$. The original likelihood is recovered by setting $\omega_i = 1$, $i = 1 \dots, n$. Thomas et al. (2018) show that, in a neighborhood of $\boldsymbol{\omega} = (1, \dots, 1)'$, \mathbf{C} characterizes the curvature of the n -dimensional surface representing the Kullback-Leibler divergence of the posterior based on the original likelihood from the posterior based on the geometrically perturbed likelihood. Thus, \mathbf{C} contains information about the directions in the n -dimensional real hyperplane along which the Kullback-Leibler divergence surface changes more rapidly in response to geometric likelihood perturbations. This insight can be used to assess the directional influence of cases.

As an exploratory tool for assessing such influence, Thomas et al. (2018) recommend to compute the sample covariance matrix, $\hat{\mathbf{C}}$, and the sample correlation matrix, $\hat{\mathbf{R}}$, based on the sample $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_L$ and to perform a principal component analysis (PCA) based on the eigendecompositions of these matrices. The first several components are often sufficient to explain most of the observed variability in the log weights and well-summarize the main directions of influence. A PCA display consisting of a scatterplot of the first two or three normalized eigenvectors helps to reveal structure in the data: points with high loadings in one or more

4.1 CDW anchoring on the mammals data

components are particularly influential, and points with similar loadings in all components have similar influence.

We leverage these ideas to form a strategy for choosing anchor points based on the case-deletion weights. First, we fit a single simple linear regression model that does not accommodate latent heterogeneity in the data. Graphical summaries based on a PCA eigendecomposition of $\widehat{\mathbf{C}}$ or $\widehat{\mathbf{R}}$ can help to visualize clusters with similar directional influence on the base model. To specify an appropriate mixture of regressions model, we choose anchor points to be representatives of these clusters. In this study we use k means applied to the rows of $\widehat{\mathbf{C}}$ or $\widehat{\mathbf{R}}$ to aid in identifying clusters and representative points.

4.1 CDW anchoring on the mammals data

To apply the proposed technique to the analysis of the mammals data, we began by fitting a Bayesian simple linear regression model with the following hyperparameters: $a = 5$, $b = 1$, $\boldsymbol{\mu}_\beta = (3.5, 0.6)'$, and $v_0 = 1$, $v_1 = 0.5$. A Gibbs sampler was run to obtain 5,000 posterior samples after burn-in and thinning of chains. The log case-deletion weights were computed from the sampled residuals. The left panels of Figure 3 shows the PCA displays based on the first two eigenvalues of the eigendecomposition of $\widehat{\mathbf{C}}$ and $\widehat{\mathbf{R}}$.

Again postulating three mixture components, we ran the k means clustering algorithm, as implemented in R, on the rows of $\widehat{\mathbf{C}}$ or $\widehat{\mathbf{R}}$. Doing so, we identified

4.1 CDW anchoring on the mammals data

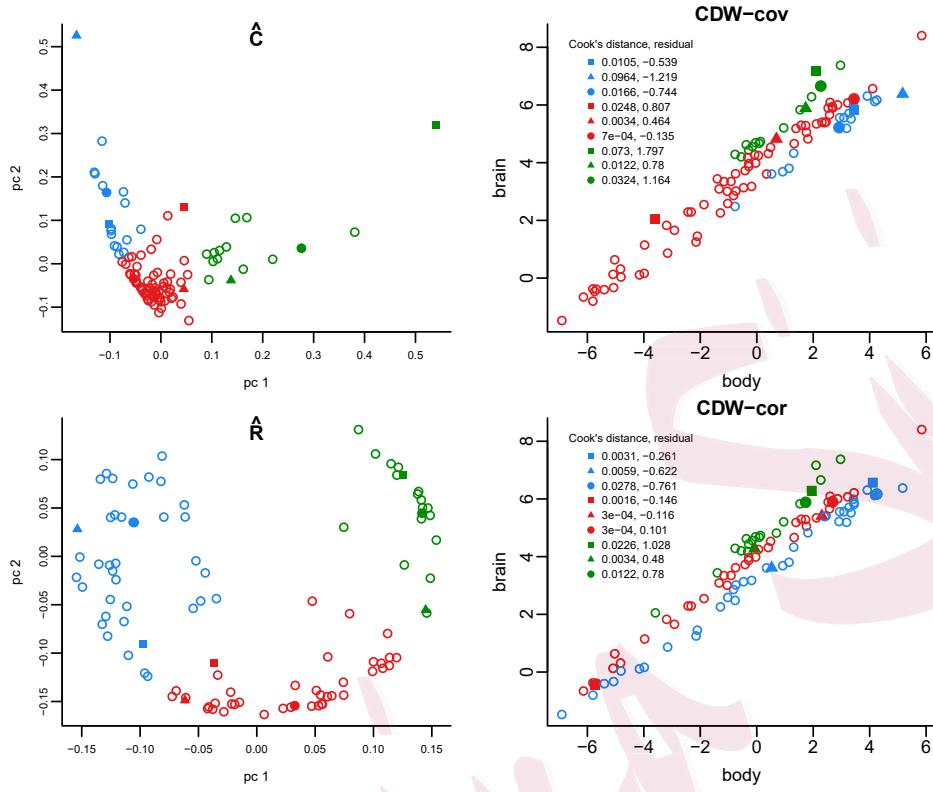


Figure 3: PCA display based on the first two eigenvalues from the eigendecomposition of $\hat{\mathbf{C}}$ (top) and $\hat{\mathbf{R}}$ (bottom) of the log case-deletion weights. The resulting k means clustering with highlighted anchor points is shown in the right panel.

the red, green, and blue clusters in the PCA display. The scatterplots in the right panels of Figure 3 shows how the identified clusters map back into observation space. The $\hat{\mathbf{C}}$ -based clustering appears to be reacting primarily to the size and sign of the residuals from the fitted simple linear regression line, with the green cases representing a small number of values tending to have large, positive residuals and the blue cases representing a small number of values with negative residuals. These green and blue cases stand apart from most points in the PCA space, with

4.1 CDW anchoring on the mammals data

the points with largest residuals appearing well-separated from all other points.

The red cases, which comprise most of the data set, are closely clustered in the PCA space and appear to be well described by the simple linear regression line.

The $\hat{\mathbf{R}}$ -based PCA plot shows a nearly-circular shape and fewer points with high loadings in a particular direction. In the observation space, we see that the green and blue clusters correspond to points having rather large least squares residuals of the same sign within each cluster.

To specify the CDW anchor models, we identified three representative observations within each cluster to be used as the three anchor points. To do this, for each of the initial clusters, we again ran the k means clustering algorithm to identify three sub-clusters. A representative point from each sub-cluster was chosen as the point nearest to the sub-cluster centroid. The selected points (three for each cluster) are shown as solid dots in the PCA displays and scatterplots of Figure 3. In addition, the Cook's distances and residuals for the anchors points were obtained from the least-squares fit depicted in Figure 1. The same nine points are also displayed in the middle and right panels of Figure 2, along with the orders to which the nine species belong. The model whose anchor points are estimated from $\hat{\mathbf{C}}$ is referred to as CDW-cov, and the model whose anchor points are estimated from $\hat{\mathbf{R}}$ is referred to as CDW-cor.

In the CDW-cov model, we see that, within each cluster, the set of anchor points comprises some points with apparently large influence (the blue triangle

and green square, for example, which have large Cook's distance values and large residuals, relative to the other points) and points that are less influential (such as the red triangle, whose x-value is near the sample average and which falls close to the regression line). This feature follows from the way in which k means decides to allocate member species to the various sub-clusters, with some sub-clusters comprising mostly "usual" observations and other, typically smaller, sub-clusters comprising mostly "unusual" observations. Nonetheless, the anchor points for the blue and green groups tend to be somewhat close together in the x-y space; the larger residuals of points in these clusters produces more distance in the PCA space.

In contrast, the CDW-cor anchor points do not stand out as unusual in the observation space and their Cook's distances tend to be smaller, with only two exceeding 0.02. These selected points are widely separated in the x-direction, with each set of three points hinting at a clearly distinct line (in a least squares sense). These groups seem to be driven primarily by similar directions of influence.

5. Simulation

We performed a simulation study to evaluate the performance of the anchored EM and CDW methods in estimating parameters of three-component mixture of regressions models. The details of the simulation design and its results are presented in Supplement 2. In summary, we found that the anchored EM models

and CDW-cor models tend to produce better accuracy in parameter estimation and cluster estimation than CDW-cov. The weaker performance of CDW-cov can be explained by the tendency for the points whose case-deletion weights have the highest variance to be separated from the others in the $\widehat{\mathbf{C}}$ space. These points, which are the most unusual observations in one or both of the x- and y- directions, are often selected as anchor points by the proposed k means method. However, they may not represent behavior that would be typical of any other points, leading to situations where some components describe only the most unusual cases. This phenomenon is alleviated when using the CDW-cor method because the $\widehat{\mathbf{R}}$ matrix captures a normalized version of the similarity of the influence of cases and is less sensitive to the overall variability of the case deletion weights. Future work may investigate alternatives to the use of the k means clusters and sub-clusters to better select anchor points using $\widehat{\mathbf{C}}$.

An interesting case considered in the simulation is one where data are generated from three parallel lines (Setting B in Supplement 2). This corresponds to a random-intercepts regression when the source of heterogeneity is unobserved. The CDW-cor method resulted in the most accurate estimation in this case, suggesting that $\widehat{\mathbf{R}}$ can be a useful tool in uncovering this type of latent heterogeneity. In contrast, when the lines differ in slope and intercepts, anchored EM performed better than the CDW methods.

6. Analysis of mammals

We now present the inferential results obtained by fitting three anchored mixture of regressions models to the mammals data: the anchored EM (A-EM), CDW-cov, and CDW-cor models. We specify the prior in (2.2) with the same hyperparameters specified in the anchored EM selection method: $a = 5$, $b = 1$, $\mu_\beta = (3.5, 0.6)'$, and $v_0 = 1, v_1 = 0.5$. For each anchor model we ran a Gibbs sampler to obtain $L = 7,500$ posterior samples (after thinning and burn-in) of the model parameters and assessed convergence using trace plots. To sample from an anchor model with the Gibbs sampler, the latent allocation variables s_i are fixed for the anchor points and sampled only for the unanchored points.

6.1 Parameter and cluster estimates

We used the posterior means of the model parameters to estimate component regression lines and pointwise credible intervals for the component-specific mean functions, $\mathbf{X}\beta_j$, for each anchor model. These lines are shown in Figure 4, with Components 1, 2, and 3 drawn in blue, red, and green, respectively. We also estimated group membership for each species by finding a maximum a posteriori estimate of its latent allocation, s_i . Each data point in Figure 4 is color-coded according to its estimated allocation, \hat{s}_i . The anchor points, whose group assignments are assumed to be known, are shown as solid symbols to distinguish them from the remaining observations.

6.1 Parameter and cluster estimates

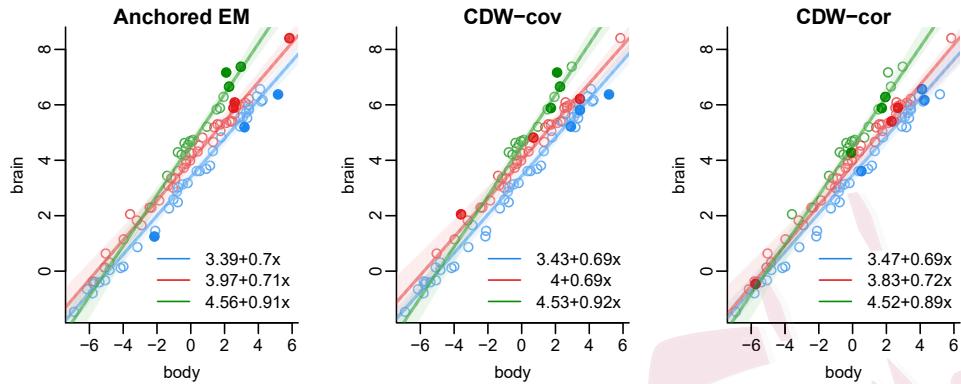


Figure 4: Posterior mean regression lines and 90% pointwise posterior credible intervals for the mean function for each mixture component. The data are color-coded according to their estimated allocation.

All three models have identified a subgroup whose slope is considerably steeper than those of the other groups, arbitrarily labeled Component 3 (shown in green) for all models. This group contains species whose brain masses show a large increase as body masses increase. The estimated regression lines are similar across the methods, with CDW-cor estimating the lowest slope (0.89) and CDW-cov estimating the largest (0.92). The three methods assign many of the same species to Component 3, identifying a string of points in the upper-right end of the scatterplot to be those arising from the steep regression line. Fewer species are allocated to this group than to the other two, and most are those with larger bodies from the Primate or Cetacea orders. In addition to these, the CDW-cor model assigns several smaller species to Component 3, which appear in the middle-left portion of the left panel of Figure 4 in an area near the point where the estimated regres-

6.2 Validation and sensitivity

sion lines for Components 2 and 3 intersect. These same points are assigned to Component 2 in the CDW-cov model, and are split between Components 2 and 3 in the A-EM model.

Component 1, plotted in blue in Figure 4, represents, broadly speaking, species with brains that are small relative to species of the same size. Of the 24 species in the Rodentia order, 18, 20, and 15 are allocated to this component by the A-EM, CDW-cov, and CDW-cor models, respectively. For all models, this group has the smallest intercept, and the estimated slopes are near 0.70 for all three anchor models. Component 2, with regression lines drawn in red, has an estimated slope nearly equal to that of Component 1. This group represents the species which differ from the Component 1 species mostly in their average brain size given their body size. The orders of species allocated to Component 2 are varied, with Artiodactyla and Carinorva being the most prevalent for all anchor models.

6.2 Validation and sensitivity

For this case study, we specified a mixture model in order to accommodate heterogeneity in the data due to taxonomic differences. Because the auxiliary information about the species' orders is available, we can evaluate the similarity in estimated groups from the mixture model to the true orders of the species. In Supplement 4, we demonstrate that all three mixture models have estimated grouping that have some correspondence with the known taxonomies: Group 1 tends to de-

6.2 Validation and sensitivity

scribe Rodentia, while Group 3 seems to describe Primates and Cetacea.

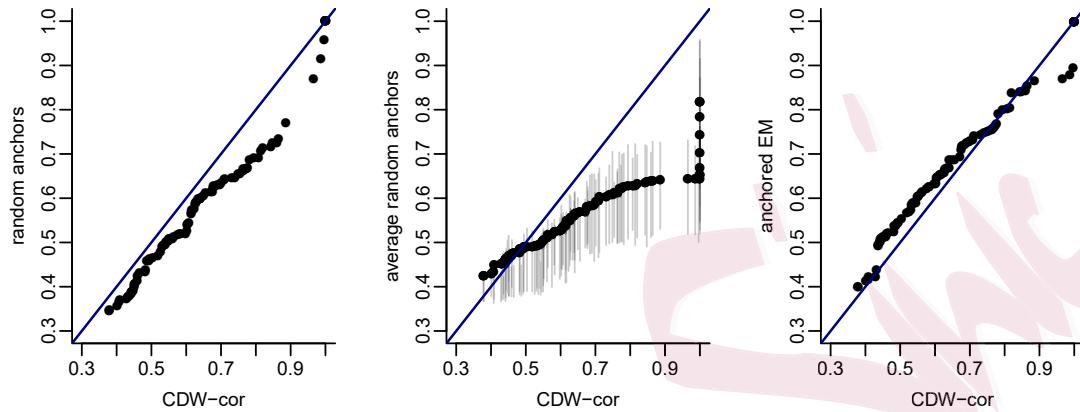


Figure 5: Maximum allocation probabilities sorted by magnitude. The left panel compares the CDW-cor model to one randomly-selected anchor model. The middle panel displays the sorted maximum average allocation probabilities over 100 randomly-selected anchor models. The shaded lines show the 25th and 75th quantiles. The right panel compares the CDW-cor model to the A-EM model.

In the absence of auxiliary information, it may be desirable to compare competing anchor models in terms of their effectiveness in identifying distinct, well-separated mixture components. Kunkel and Peruggia (2020) show that goodness-of-fit in an anchor model is closely related to the degree of separation among the component distributions: a well-fitting anchor model will produce estimated mixture components with distinctive features. In contrast, a poorly-fitting anchor model will exhibit features similar to the “label-switching” seen in full exchangeable models: densities of component-specific parameters may be multimodal and

6.2 Validation and sensitivity

similar in shape, and estimated allocation probabilities will tend towards equal probabilities for each component. Motivated by these considerations, we propose to perform model-checking by summarizing the maximum estimated posterior allocation probabilities calculated for each of the observed data points; that is,

$$p_i^{max} = \max_j \hat{P}(S_i = j|y), i = 1, \dots, n.$$

Figure 5 demonstrates three ways to assess the model using the values of p_i^{max} . The right panel shows the order statistics of the p_i^{max} values from the CDW-cor model on the x-axis and from the A-EM model on the y-axis. A straight line would indicate that both models result in similar overall fit across the two models. The fact that most points fall slightly above the identity line indicates that most probabilities are higher for the A-EM model than for CDW-cor, suggesting a better fit for the A-EM model. If a competing model is not considered, randomization can be used for checking one anchor model. The left panel shows a plot that compares the CDW-cor model to an anchor model whose anchor points are selected completely at random. The CDW-cor probabilities are consistently higher, resulting in a pattern with points falling underneath the identity line. The middle panel is similar, but the y-axis shows the average p_i^{max} values over 100 randomized anchor models with shaded lines indicating the 25th and 75th sample percentiles of the p_i^{max} . This plot gives an approximation of the expected p_i^{max} and their variability across random anchor models.

A sensitivity analysis was performed to investigate the effect of the hyperpa-

rameters a , b , v_0 , and v_1 in (2.2). The analysis, which is detailed in Supplement 3, revealed parameter estimates for models with weaker prior information on σ^{-2} and β that are similar to those in the analysis described in Section 6.1.

7. Discussion

When we specify a finite mixture model, we assume the existence of distinct subgroups in the population. The group membership of any individual observation is unknown and can be estimated a posteriori. A random effects model similarly allows inference on group-specific features, but requires auxiliary, deterministic information indicating which observations are to be grouped together, and the similarities among these observations drive the estimated features of the various groups. In an anchored mixture model, the anchored observations affect the model fit in the same way as labeled observations in a random effects model affect inferences on their groups: they inform the distinct features of their mixture components. These features then influence the probabilities of group membership of the remaining unanchored observations. Thus, the similarities among a component's anchor points are a key driver of what types of groups the mixture model identifies.

In this case study, we selected three groups in order to detect hypothesized subgroups tied to unobserved random effect indicators. A model with two mixture components could also have been specified to accommodate some of the

unexplained structure in the residuals from the simple linear regression. Two components, while unable to capture many latent subgroups, can nonetheless distinguish between the most prominent inhomogeneous subsets of data points: for example, the primates such as *homo sapiens* are likely to be separated from the large-bodied, small-brained rodents, regardless of the number of components used. The Supplementary materials gives details regarding the anchor points and estimated regression lines under a two-component model. A two component model may also be a first step in an iterative process, in which mixture components are added and model fit is re-evaluated to determine whether additional components are necessary.

In our illustrative study, the EM method chose anchor points to fall near the straight lines assumed by the mixture of regression lines, and its accompanying mixture model defined groups based on proximity to these lines. For example, if we truly believe the mixture model we have specified, a method such as anchored EM, which subsumes this model at the anchoring stage, is perfectly appropriate. This method is flexible and extensions to other regression models, including non-Gaussian mixtures, can be readily derived within this framework.

The CDW methods are based on the principle that a mixture model is appropriate because the simple model is inadequate; the mixture components are identified as groups exhibiting similar misfit. In our study, the CDW-cov method often selected points whose case deletion weights have high variance. These points

tend to have unusual x-values and/or large residuals.

The CDW-cor method selected anchor points that were representative of groups of observations exerting similar influence on the posterior inferences from a naive model, and our simulation results indicated that this method has promise in using a mixture model to introduce random effects when auxiliary grouping information is unobserved. Further, because CDW requires only a base model and ability to calculate log case-deletion weights, this method can readily be extended to non-Gaussian models, such as logistic or Poisson regression, and hierarchical models. Future work will investigate how the theoretical properties of the $\hat{\mathbf{C}}$ and $\hat{\mathbf{R}}$ matrices relate to the optimal determination of the number of components and selection of anchor points.

In the multiple regression setting, we have demonstrated in simulations that the anchored mixture of regressions models can produce accurate parameter estimates in low dimensions, even if some predictors are collinear. Further investigation is needed to evaluate how well the model scales when the number of predictors grows. In addition, real data sets are likely to exhibit features that affect the relative performance of the methods for selecting anchor points. For example, the current study illustrates that anchored-EM and the CDW methods differ in the variability among selected anchor points: CDW enforces some separation among the points, while anchored-EM may often select anchor points that are close together. We have found that in real data sets where some predictors

may have low variability, discretized values, or inflation at zero, the differences among anchor models estimated from these methods may be magnified.

In sum, to select a method for finding anchor points, it is important to consider which types of distinguishing features the underlying model uses to identify similar observations and to determine how such features become relevant to answer the scientific questions of interest.

8. Supplementary Materials

The online Supplement contains the following: (1) details on the anchored EM algorithm for the anchored mixture of regressions model; (2) simulation study design and results; (3) sensitivity analyses; (4) ; and (5) analysis that explores the relationship between known taxonomy and the model-based clustering.

Acknowledgments

This material is based on work supported by the National Science Foundation under grant no. SES-1424481 and SES-1921523.

References

Bennett, P. and P. Harvey (1985a). Relative brain size and ecology in birds. *Journal of Zoology* 207(2), 151–169.

Bennett, P. M. and P. H. Harvey (1985b). Brain size, development and metabolism in birds and mammals. *Journal of Zoology* 207(4), 491–509.

REFERENCES

Bradlow, E. T. and A. M. Zaslavsky (1997). Case influence analysis in Bayesian inference. *Journal of Computational and Graphical Statistics* 6(3), 314–331.

Epifani, I., S. N. MacEachern, M. Peruggia, et al. (2008). Case-deletion importance sampling estimators: Central limit theorems and related results. *Electronic Journal of Statistics* 2, 774–806.

Felsenstein, J. (1985). Phylogenies and the comparative method. *The American Naturalist* 125(1), 1–15.

Garland Jr, T., P. H. Harvey, and A. R. Ives (1992). Procedures for the analysis of comparative data using phylogenetically independent contrasts. *Systematic Biology* 41(1), 18–32.

Gayon, J. (2000, 08). History of the Concept of Allometry. *American Zoologist* 40(5), 748–758.

Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica* 57(6), 1317–1339.

Jasra, A., C. C. Holmes, and D. A. Stephens (2005). Markov chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling. *Statistical Science* 20(1), 50–67.

Jerison, H. J. (1955). Brain to body ratios and the evolution of intelligence. *Science* 121(3144), 447–449.

Kunkel, D. and M. Peruggia (2020). Anchored Bayesian Gaussian mixture models. *Electronic Journal of Statistics* 14(2), 3869 – 3913.

MacEachern, S. N. and M. Peruggia (2002). Bayesian tools for EDA and model building: A brainy study. In *Case Studies in Bayesian Statistics*, pp. 345–362. Springer.

Malsiner-Walli, G., S. Frühwirth-Schnatter, and B. Grün (2016, Jan). Model-based clustering based on sparse finite Gaussian mixtures. *Statistics and Computing* 26, 303–324.

Miller, J. W. and M. T. Harrison (2018). Mixture models with a prior on the number of components. *Journal of the American Statistical Association* 113(521), 340–356.

Neal, R. M. and G. E. Hinton (1998). *A View of the EM Algorithm that Justifies Incremental, Sparse, and other Variants*, pp. 355–368. Dordrecht: Springer Netherlands.

Nobile, A. (2004). On the posterior distribution of the number of components in a finite mixture. *The Annals of Statistics* 32(5), 2044–2073.

Peruggia, M. (1997). On the variability of case-deletion importance sampling weights in the Bayesian

REFERENCES

linear model. *Journal of the American Statistical Association* 92(437), 199–207.

Peters, R. H. (1983). *The ecological implications of body size*. Cambridge University Press.

Roeder, K. (1994). A graphical technique for determining the number of components in a mixture of normals. *Journal of the American Statistical Association* 89(426), 487–495.

Sacher, G. A. and E. F. Staffeldt (1974). Relation of gestation time to brain weight for placental mammals: implications for the theory of vertebrate growth. *The American Naturalist* 108(963), 593–615.

Thomas, Z. M., S. N. MacEachern, and M. Peruggia (2018). Reconciling curvature and importance sampling based procedures for summarizing case influence in Bayesian models. *Journal of the American Statistical Association* 113(524), 1669–1683.