# Bayesian Pairwise Comparison of High-Dimensional Images

Subharup Guha and Peihua Qiu

Department of Biostatistics, University of Florida

December 21, 2024

#### Abstract

A fundamental task in the automated analysis of images is the development of effective image pair comparison techniques. For two high-dimensional images, a statistical method must automatically label them as "similar" or "different" depending on whether random error and spatial dependencies could account for the pixel-wise differences. We develop a Bayesian strategy by constructing a novel extension of Dirichlet processes called the *spatial random partition model* (sRPM). The process groups spatially proximal image pixels with similar intensities into clusters, thereby achieving dimension reduction in the large number of pixels. Next, we apply the sRPM-based analytical procedure to compare two images. The image comparison problem is formulated as a hypothesis test involving a univariate metric adaptive to spatial correlations and robust to random variability in the pixel intensities. To handle the computational burden, we foster a two-stage technique for MCMC analysis and hypothesis testing of image pairs. A simulation study analyzes artificial datasets and finds compelling evidence for the high accuracy of sRPM in image comparison. We demonstrate the effectiveness of the technique by statistically analyzing satellite image data.

*Keywords:* Aggregated attraction function; Bayesian hierarchical model; Differential pixel proportion; Markov Chain Monte Carlo; Nonparametric Bayes; Spatial random partition model.

## 1 Introduction

Due to the rapid advancements in image acquisition techniques, images have emerged as one of the most versatile and information-rich data resources across diverse industries and scientific disciplines. For example, in the manufacturing industry, images play a pivotal role in quality control processes due to their ease of acquisition and cost-effectiveness. Similarly, in the natural sciences, satellite images are indispensable for analyzing Earth's surface and environment across disciplines such as forest science, climate science, agriculture, forecasting, ecology, and fire science.

In many of these contexts, a series of images is often collected over time from a longitudinal process—such as images acquired from a rolling process in metalworking—where a critical goal is to monitor for temporal changes and detect potential "change points." Identifying such change points is essential as they indicate moments where the process diverges beyond acceptable limits of variation, potentially signaling anomalies or shifts in underlying dynamics. This task is closely related to statistical process control (SPC) (Qiu, 2014), but existing SPC methodologies are insufficient for high-dimensional image data due to unrealistic model assumptions and limited practical applicability.

A foundational task underpinning this monitoring process is image pair comparison. By developing effective techniques for comparing high-dimensional images, we can systematically classify pairs of images as "similar" or "different" based on whether observed pixel-wise differences can be attributed to random error. Successful image comparison is often the first step toward more complex objectives, such as inferring the nature of detected changes, isolating regions of interest, or predicting future behaviors of a system under monitoring. For example, in quality control, identifying a shift might guide interventions to rectify the process, while in satellite image analysis, detected changes might signal environmental shifts requiring further study. Thus, image comparison lays the groundwork for meaningful downstream analyses and decision-making in these applications.

In response to the growing need for sophisticated methodologies that combine robust modeling with efficient computation, this paper aims to address the challenges of highdimensional image comparison and provide a practical framework for advancing SPC techniques tailored to modern image data. To be useful, image comparison techniques must have (i) high sensitivity and specificity, (ii) the ability to handle high dimensionality and spatial dependencies through model-based dimension reduction, (iii) flexibility, in the sense that they do not make restrictive parametric assumptions limiting the applicability of the models in real settings, and (iv) computational efficiency, for example, being able to harness the power of efficient computation to analyze high-dimensional images. This paper addresses these critical challenges, providing a practical framework for advancing SPC techniques tailored to modern image data.

To motivate the methodological development, consider three images presented in Figure 1 and consisting of  $100 \times 100$  pixels each. We denote these images by  $\Im_1$ ,  $\Im_2$  and  $\Im_3$ , respectively. Although the images may look identical at the first glance, in reality, all three images are different to varying degrees. A visual inspection reveals that differences between images  $\Im_1$  and  $\Im_2$  are relatively subtle compared to differences between  $\Im_1$  and  $\Im_3$ . Intuitively, it appears that random variation could possibly account for the differences between  $\Im_1$  and  $\Im_2$ . However, although  $\Im_1$  and  $\Im_3$  are very similar in most places, there

are systematic differences, e.g., the dark area on the lower right side of  $\mathfrak{I}_1$  has a slightly different shape.

Relevant background and present state of knowledge There is only limited existing research on monitoring image pairs or sequences, mainly in the chemical and industrial engineering literature, where images have been widely used in recent years (Megahed et al., 2011; Prats-Montalban and Ferrer, 2014; Yan et al., 2015). Existing methods for image comparison often proceed in two main steps. For instance, the methods first extract some features from each image using techniques such as principal component analysis (PCA), and then monitor the extracted features by a conventional control chart (Duchesne et al., 2012; Lin et al., 2008). Some other methods focus on certain prespecified regions in individual images called regions of interest (ROIs), and then monitor the images by a control chart constructed using a summary statistic of the ROIs, e.g., the average image intensity (Jiang et al., 2011; Megahed et al., 2012). The first type of method completely ignores the spatial structure of the images while the second type considers the spatial structure only within the prespecified ROIs. Most of these methods fail to take into account image edges and complicated correlation structures of the spatial image intensities. Consequently, they do not provide a reliable tool for image monitoring applications; see Qiu (2020) for a related discussion. Notable exceptions include Wang and Ye (2010), who compare nonparametric curves under a heteroscedastic model with spatially correlated errors; Koosha et al. (2017), who use nonparametric wavelet basis functions to extract key features and detect change points and fault locations simultaneously; and Roy and Mukherjee (2024), who propose a feature-based image comparison method in which edges and jump points are considered



(a) Image  $\Im_1$ 

(b) Image  $\Im_2$ 



(c) Image  $\Im_3$ 

Figure 1: Are images  $\mathfrak{I}_1$  and  $\mathfrak{I}_2$  significantly different? Are images  $\mathfrak{I}_1$  and  $\mathfrak{I}_3$  significantly different?

as the primary features. As a side note, in the literatures of image processing, fMRI, and machine learning, there exist methods or algorithms for analyzing a set of images obtained in a given time interval (Feng and Qiu, 2018; Guo and Zhang, 2015; Guo et al., 2018; Julea et al., 2011; Lindquist, 2008). However, these methods are retrospective and cannot be used effectively for prospective monitoring of spatiotemporally correlated image sequences.

There are relatively few Bayesian techniques for image comparison. Nonparametric Bayes techniques that are potentially applicable to this problem typically involve extensions of Dirichlet processes (Ferguson, 1973; Ghosal et al., 1999) to spatial settings; see Reich and Fuentes (2015) for a recent review. These strategies include Gelfand et al. (2005), Duan et al. (2007), Dunson and Park (2008), Griffin and Steel (2006), Rodriguez et al. (2010), and Reich and Bondell (2011). Another class of Bayesian methods (e.g., Na et al., 2013; Chatzis and Tsechpenakis, 2010) applies infinite-dimensional spatial extensions of hidden Markov models to the problem of image segmentation. However, these models have certain limitations. One notable challenge is their sensitivity to the predefined definition of a 'neighborhood,' which is assumed to be uniform across the image and known in advance. This assumption can limit the model's flexibility, as the Markov assumption treats pixels in different neighborhoods as independent. As a result, these models may find it difficult to capture more complex or unknown spatial correlation structures beyond first-order dependencies, including medium- or long-range interactions that can vary across the image space.

Moreover, these Bayesian methods are either broadly applicable spatial models or are primarily designed for the analysis of single images. As a result, they are not well-suited for image comparison, which is the central focus of this paper. Specifically, they lack the ability to account for spatial correlation structures between non-random pixel locations, identify and transfer relevant information across images using unsupervised, model-based dimensionality reduction techniques, or are inefficient at detecting significant deviations from quality standards while considering the spatial nature of images.

Motivated by these challenges, this article develops, in stages, a Bayesian strategy for comparison of high-dimensional image pairs. Section 2 invents a new nonparametric approach for analyzing a single image with multiple intensity values at each of p pixels or voxels (e.g. a colored image). We construct a novel extension of Dirichlet processes, called the spatial random partition model (sRPM), that groups proximal image pixels with similar intensities into the same cluster. The infinite mixture model adapts the general Bayesian frameworks of random partition models and density estimation in the presence of arbitrary covariates (e.g. Müller et al., 2011) to spatial settings and image analysis in particular. The proposed sRPM model favors a small number of relatively large clusters, thereby achieving dimension reduction in the large number of pixels in a manner that accounts for spatial correlations. A key benefit of this dimension reduction is that the extracted image information is useful for answering inferential questions in image sequence monitoring applications.

Section 3 extends the sRPM-based analytical procedure to compare two images that share the same pixel locations but possess possibly different intensity values at multiple pixels. Image comparison is achieved by an intuitively appealing univariate metric for measuring image differences. This metric is not only robust to random response variability in the large number of pixels but also adaptive to spatial correlations. The image comparison problem is then formulated as a single-parameter hypothesis test. This comparison procedure utilizes only cluster-related information previously extracted from the first image, as developed in Section 2. Depending on the degree to which the two images differ, the marginal posterior of this metric is used to call the two images "similar" or "different" based on the posterior probabilities of competing hypotheses. To handle the computational burden in analyzing high-dimensional images, Section 3.1 fosters a two-stage technique for MCMC analysis and hypothesis testing of image pairs. Section 4 analyzes artificial datasets using the proposed technique and finds compelling evidence for the high accuracy of sRPM in image comparison. Section 5 applies the sRPM technique to detect differences between pairs of satellite images. Finally, several remarks conclude the article in Section 6.

## 2 Bayesian analysis of a single image

For  $d \ge 2$ , imagine that we have a single *d*-dimensional image, denoted by  $\mathfrak{I}_1$  and consisting of *p* pixels. In Figure 1, d = 2 because  $\mathfrak{I}_1$  is two-dimensional. However, we are equally interested in analyzing three-dimensional images, such as brain images used in Alzheimer's disease research. We wish to achieve dimension reduction in the large number of pixels in a manner that accounts for spatial correlation. The extracted information should be subsequently useful for comparisons between  $\mathfrak{I}_1$  and other images in image sequence monitoring.

## 2.1 Spatial random partition model (sRPM)

The proposed sRPM defines a coherent *process*, rather than simply a model. Specifically, the distribution of the responses for any subset of pixels in image  $\Im_1$  can be obtained from the distribution for a larger set of pixels by marginalization. The common domain of the images is assumed to be a design set denoted by  $\Delta_p$  containing p regularly or irregularly arranged d-dimensional pixels. This design set is assumed to belong to a compact convex set in  $\mathcal{R}^d$  denoted by  $\mathcal{D}$ . For pixel  $i = 1, \ldots, p$ , where p is large, and for image indexed by  $t = 1, 2, \ldots$ , the  $i^{th}$  pixel is located at  $\mathbf{s}_i = (x_i, y_i)' \in \Delta_p$  for two-dimensional images (i.e., d = 2) and is associated with a pixel-specific intensity vector  $\mathbf{z}_{it}$  of length r. Color images can be represented by triplets,  $\mathbf{z}_{it} = (z_{itR}, z_{itG}, z_{itB}) \in [0, 1]^3$ , corresponding to RGB values. Monochrome images correspond to scalar intensities, so that  $z_{it} \in [0, 1]$ .

Non-random pixel locations. We regard the pixel locations as non-random design points and the associated intensities as random, image-specific responses. To achieve dimension reduction in the large number of pixels, p, one possibility is to utilize the sparsityinducing property of Dirichlet processes to group a set of proximal pixels with similar intensities in image  $\Im_1$  into  $q_1$  latent clusters, where  $q_1$  is unknown but typically much smaller than p. This would drastically reduce the necessary memory storage space from O(p) to include only the parameters related to the cluster characteristics. The use of Dirichlet processes to achieve dimension reduction has precedence in the literature; see Kim et al. (2006) and Dunson and Park (2008).

However, standard Dirichlet processes are inadequate for image analysis because they

treat all p pixels as exchangeable, and they would allocate the pixels to latent clusters regardless of spatial proximity. To accommodate the intrinsic spatial nature of images, we propose extending the Dirichlet process to a *spatial random partition model* (sRPM) so that only proximal pixels with similar intensities are a posteriori allocated to the same cluster. Applying density estimation regressed on covariates (e.g. Müller et al., 2011) and random partition models, the sRPM process provides a novel framework for image analysis that flexibly adapts to the spatial correlations among the pixels.

Allocation variables  $c_1 = (c_{11}, \ldots, c_{p1})$ . Let the event  $[c_{i1} = k]$  represent pixel *i* of image  $\Im_1$  belonging to cluster *k* in a spatial random partition model, which we will formulate after introducing key concepts and definitions. In the eventual inferential procedure, we will find that an MCMC point estimate of allocation vector  $c_1$  can be computed if desired.

Aggregated attraction function. Each cluster is associated with a latent spatial knot and a cluster-specific bandwidth, with individual pixels in image  $\mathfrak{I}_1$  stochastically allocated to a cluster depending on the pixel-cluster attraction. More specifically, for cluster k, the spatial knot is a d-dimensional vector  $\boldsymbol{\mu}_k$ , and the bandwidth is a  $d \times d$  positive definite matrix,  $\boldsymbol{\Sigma}_k$ . Since the spatial knots and bandwidths are unknown, we assume normalinverse Wishart (NIW) priors for these cluster-specific quantities:  $(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \stackrel{iid}{\sim} \mathcal{W}_d(\boldsymbol{\varphi}_0), k =$  $1, \ldots, q_1$ , for hyperparameter  $\boldsymbol{\varphi}_0 = (\boldsymbol{\mu}_0, \kappa_0, \boldsymbol{\Sigma}_0, \nu_0)$  consisting of positive hyperparameters  $\nu_0$  and  $\kappa_0$ , d-dimensional vector  $\boldsymbol{\mu}_0$ , and  $d \times d$  positive-definite matrix  $\boldsymbol{\Sigma}_0$ .

The pixel-cluster attraction is determined by the pixel's position relative to the spatial knot  $\mu_k$ , with the relative distance calibrated by cluster bandwidth  $\Sigma_k$ . In particular, the



Figure 2: Illustration of pixel-cluster attraction for a pixel located at  $s_i$ .

attraction between pixel *i* and cluster *k* is assumed to be equal to the multivariate normal density,  $N_d(\mathbf{s}_i \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ . This ensures that the pixel-cluster attraction is greater for pixels near the spatial knot and for clusters with smaller bandwidths. The concept of pixel-cluster attraction is illustrated in Figure 2 for q = 6 clusters belonging to a rectangular design space in a single image. The areas and orientations of the ellipses are determined by the bandwidths quantified by the positive-definite matrices  $\boldsymbol{\Sigma}_k$ . The sizes of the dots marking the spatial knots are proportional to the attraction between pixel  $\mathbf{s}_i$  and the *k*th cluster.

In image  $\mathfrak{I}_1$ , using the pixel-cluster allocation variables, let  $\mathbf{s}_{k1}^* = \{\mathbf{s}_i : c_{i1} = k, i = 1, \dots, p\}$ denote the pixel locations assigned to the *k*th cluster. The spatial compactness of cluster *k* in image  $\mathfrak{I}_1$  is quantified by an *aggregated attraction function* denoted by  $g(\mathbf{s}_{k1}^*)$ . Larger values of aggregated attraction are associated with tighter spatial clusters, i.e., nearby pixels grouped into smaller clusters. Marginalizing over the unknown cluster-specific knots and bandwidths, and regarding the pixel locations as non-random design points, the aggregated attraction function for cluster k in image  $\Im_1$  is obtained as follows:

$$g(\boldsymbol{s}_{k1}^*) = \int \int \prod_{i:c_{i1}=k} N_d(\boldsymbol{s}_i | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \mathcal{W}_d(\boldsymbol{\varphi}_0) d\boldsymbol{\mu}_k d\boldsymbol{\Sigma}_k = \frac{\Gamma_d(\nu_{k1}/2)}{\pi^{m_{k1}d/2}\Gamma_d(\nu_0/2)} \frac{|\boldsymbol{\Sigma}_0|^{\nu_0/2}}{|\boldsymbol{\Sigma}_{k1}|^{\nu_{k1}/2}} \left(\frac{\kappa_0}{\kappa_{k1}}\right)^{d/2},$$
(1)

where  $\Gamma_d(\cdot)$  denotes the multivariate gamma function. The new quantities appearing in equation (1) have the following definition. Based on the pixels assigned to the *k*th cluster, the updated NIW prior has hyperparameter  $\varphi_{k1} = (\mu_{k1}, \kappa_{k1}, \Sigma_{k1}, \nu_{k1})$ , where

$$\boldsymbol{\mu}_{k1} = \frac{\kappa_0}{\kappa_0 + m_{k1}} \boldsymbol{\mu}_0 + \frac{m_{k1}}{\kappa_0 + m_{k1}} \bar{\boldsymbol{s}}_{k1}, \quad \nu_{k1} = \nu_0 + m_{k1}, \quad \kappa_{k1} = \kappa_0 + m_{k1},$$
$$\boldsymbol{\Sigma}_{k1} = \boldsymbol{\Sigma}_0 + \boldsymbol{V}_{k1} + \frac{\kappa_0 m_{k1}}{\kappa_0 + m_{k1}} (\bar{\boldsymbol{s}}_{k1} - \boldsymbol{\mu}_0) (\bar{\boldsymbol{s}}_{k1} - \boldsymbol{\mu}_0)'. \tag{2}$$

These quantities rely on the following cluster-related summaries for image  $\mathfrak{I}_1$ : (i) number of pixels,  $m_{k1}$ , allocated to cluster k by the allocation vector  $\mathbf{c}_1$ , (ii) pixel location averages,  $\bar{\mathbf{s}}_{k1} = \sum_{i:c_{i1}=k} \mathbf{s}_i/m_{k1}$ , and (iii) matrix error sum of squares for the pixel locations,  $\mathbf{V}_{k1} =$  $\sum_{i:c_{i1}=k} (\mathbf{s}_i - \bar{\mathbf{s}}_{k1})(\mathbf{s}_i - \bar{\mathbf{s}}_{k1})'$ , for  $k = 1, \ldots, q_1$ . Notice that these quantities rely only on the pixel locations but not on the image intensities.

We emphasize that the pixel locations are conceived as non-random design points. By construction, the correlation induced by the spatial knots and bandwidths of the clusters implies that larger values of the attraction function are a priori associated with tighter spatial clusters, i.e., proximal pixels allocated to clusters with smaller areas or volumes, i.e., Lebesgue measures in  $\mathcal{R}^d$ . Spatial random partition model for images. As mentioned, standard Dirichlet processes are inadequate because they ignore the spatiotemporal nature of the images and regard the intensities at the p pixels as a priori exchangeable. Using the cluster-specific aggregated attraction functions in (1), we extend the standard Dirichlet process to a *spatial* random partition model (sRPM) for the pixel-cluster allocations in image  $\Im_1$ :

$$[\boldsymbol{c}_1] \propto \alpha^{q_1} \prod_{k=1}^{q_1} g(\boldsymbol{s}_{k1}^*) \Gamma(m_{k1}), \quad \boldsymbol{c}_1 \in \mathcal{B}_p,$$
(3)

where  $[\cdot]$  denotes density functions,  $\mathcal{B}_p$  is the set of possible partitions of p pixels, and  $\alpha > 0$  is the mass parameter.

In other words, allocation vector  $\mathbf{c}_1 = (c_{11}, \ldots, c_{p1})$  is a priori distributed as the partitions induced on the p image pixels by sRPM. Furthermore, attraction function (1) ensures that a pixel is a priori more likely to join clusters containing a larger number of spatially proximal pixels. Summing over all possible partitions of the p pixels in image  $\mathfrak{I}_1$ , the normalization constant for density (3) is  $\sum_{c_1 \in \mathcal{B}_p} \alpha^{q_1} \prod_{k=1}^{q_1} g(\mathbf{s}_{k1}^*) \Gamma(m_{k1})$ .

Intensities. Finally, we specify the sampling distribution of the intensity or response vector  $\mathbf{z}_{i1} = (z_{i11}, \ldots, z_{ir1})$ , for  $i = 1, \ldots, p$ . Conditional on the pixel-cluster allocations of image  $\mathfrak{I}_1$ , the intensity vectors of the member pixels in cluster k are modeled as a shared cluster-specific mean vector,  $\mathbf{v}_k = (v_{k1}, \ldots, v_{kr})$ , plus Gaussian white noise. More specifically, conditional on the event that the  $i^{th}$  pixel belongs to the kth cluster,  $\mathbf{z}_{i1} \mid$  $\{c_{i1} = k\} \stackrel{iid}{\sim} N_r(\mathbf{v}_k, \sigma^2 \mathbf{I})$ , with  $\sigma^2$  assigned an inverse chi-square prior. Small values of  $\sigma$  guarantee that the clusters consist of pixels with similar intensities. Consequently, the posterior distribution of sRPM tends to group proximal image pixels with similar intensities into the same cluster, and tends to favor a small number of relatively large clusters, thereby achieving dimension reduction.

For cluster  $k = 1, \ldots, q_1$ , let mean vectors  $\boldsymbol{v}_k \stackrel{iid}{\sim} N_r(\boldsymbol{v}_0, \boldsymbol{\Sigma}_0)$  in the prior. Standard conditionally conjugate priors are imposed on all the remaining hyperparameters to complete the model specification. Notice that a common  $\sigma$  parameter for the r intensity components presupposes that the intensities have comparable scales; this may require standardizing each intensity component before analysis. Despite the apparently strong parametric assumption of homoscedastic Gaussian errors in the sRPM model, Dirichlet processes and many of their extensions, including sRPM, are intrinsically nonparametric and capable of flexibly adapting to complex-structured regression surfaces (Lijoi and Prünster, 2010).

### **2.2** Posterior inferences for image $\Im_1$

The following Proposition establishes the posterior equivalence of two models with very different interpretations. The proof is presented in Supplementary Material. Applying well-established sampling procedures, the sRPM parameters are then updated in a relatively straightforward manner.

**Proposition 1** Consider an alternative model in which (i) sampling distribution  $(\mathbf{s}_i, \mathbf{z}_{i1}) |$  $\{c_{i1} = k\} \stackrel{indep}{\sim} N_d(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \times N_r(\boldsymbol{v}_k, \sigma^2 \mathbf{I}), \text{ and (ii) allocation vector } \mathbf{c}_1 = (c_{11}, \ldots, c_{p1}) \text{ is distributed as the partitions induced by a$ **regular** $Dirichlet process prior with mass parameter <math>\alpha$  and base distribution  $\mathcal{W}_d(\boldsymbol{\varphi}_0) \times N_r(\boldsymbol{v}_0, \boldsymbol{\Sigma}_0)$ . This model regards the pixel locations  $\mathbf{s}_1, \ldots, \mathbf{s}_p$  as random, unlike sRPM, for which  $\Delta_p$  is the design set of pixel locations. Nevertheless, for every p, given the responses and set  $\Delta_p$ , the posterior distribution of the remaining parameters of sRPM and the alternative model are identical.

Intuitively, posterior inferences about the key parameters of interest are equivalent under the alternative model due to the high degree of separability between the pixel locations and image intensities in the two models. Starting with ad hoc initial values, the sRPM model parameters for image  $\Im_1$  are then iteratively generated by Gibbs sampling steps from their full conditionals. The post-burn-in MCMC sample is used for posterior inference. We briefly outline the Gibbs sampling steps below for some of the parameters:

- 1. Allocation vector  $c_1$ . For the allocation vector of the *alternative* posterior distribution described in Proposition 1, this parameter vector is updated using the Gibbs sampler for finite-dimensional Dirchlet process representation of Ishwaran and James (2001).
- 2. Spatial knots and bandwidths. Given the data and the current values of the remaining model parameters including allocation vector  $\boldsymbol{c}_1$ , the spatial knots and bandwidths have independent NIW distributions:  $(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \stackrel{indep}{\sim} \mathcal{W}_d(\boldsymbol{\varphi}_{k1})$ , for  $k = 1, \ldots, q_1$ , with the updated NIW hyperparameter  $\boldsymbol{\varphi}_{k1}$  described after equation (1).

Cluster-related posterior inferences. By using the MCMC sample and applying the least-squares method of Dahl (2006), we can compute  $\hat{c}_1$ , an estimate of the pixel-cluster allocation vector for image  $\Im_1$ . Intuitively, this procedure first uses the MCMC sample to estimate A, the "adjacency matrix" whose elements represent the estimated posterior probabilities that each pixel pair belongs to a common cluster. Next, using an MCMC



Figure 3: Image  $\mathfrak{I}_1$  in Figure 1 and its detected cluster characteristics by the sRPM technique.

sample, estimate  $\hat{c}_1$  is chosen to minimize the Frobenius norm of the difference between  $\hat{A}$ and the adjacency matrix associated with  $\hat{c}_1$ . Since they are straightforward deterministic functions of  $\hat{c}_1$ , the estimated number of clusters,  $\hat{q}_1$ , and the estimated numbers of pixels allocated to the clusters,  $\hat{m}_{11}, \ldots, \hat{m}_{1\hat{q}_1}$ , are immediately available. Furthermore, estimates of the spatial knots,  $\hat{\mu}_1, \ldots, \hat{\mu}_{\hat{q}_1}$ , and bandwidths,  $\hat{\Sigma}_1, \ldots, \hat{\Sigma}_{\hat{q}_1}$ , can be computed. For example, for image  $\mathfrak{I}_1$  of Figure 1, the detected cluster characteristics, along with image  $\mathfrak{I}_1$ for comparison, are graphically presented in Figure 2.2. sRPM detected  $\hat{q}_1 = 57$  spatial clusters, achieving impressive dimension reduction in the 10,000 image pixels.

## 3 Comparison of image pairs

We will extend the sRPM-based analytical procedure of Section 2 to the comparison of two images. Suppose that the images, denoted by  $\mathfrak{I}_1$  and  $\mathfrak{I}_2$ , share the same pixel locations but possess possibly different intensities at multiple pixel locations. Thus, design set  $\Delta_p \subset \mathcal{R}^d$  consists of p regularly or irregularly arranged pixels that are common to both images. Image comparison is achieved by an intuitively appealing measure for image differences that is not only robust to random response variability of the large number of pixels but is also adaptive to spatial correlation. Furthermore, the comparisons utilize only cluster-related information extracted from  $\mathfrak{I}_1$ , as described in Section 2. Depending on the degree to which the two images are different, the marginal posterior of this measure is used to call the images "similar" or "different" using the posterior probabilities of competing hypotheses.

**Propagation of information from image**  $\mathfrak{I}_1$  to  $\mathfrak{I}_2$ . One of the key features of the Bayesian paradigm is its natural ability to borrow strength from comparable or somewhat similar data. Applied to the image comparison problem, this implies that conditional on image  $\mathfrak{I}_1$ , the posterior of  $\mathfrak{I}_1$  plays the role of the prior for  $\mathfrak{I}_2$ . For example, the pixels in image  $\mathfrak{I}_2$  may join any of the  $q_1$  clusters already discovered in image  $\mathfrak{I}_1$  (if the pixels belong to regions of design space  $\Delta_p$  where the two images are somewhat similar) or may join a new set of clusters unique to image  $\mathfrak{I}_2$  (if the pixels belong to regions where the two images are very different). If  $\mathfrak{I}_2$  is associated with  $q_2$  clusters, then  $q_2 \ge q_1$  because we condition on the  $q_1$  clusters previously discovered in  $\mathfrak{I}_1$ . The first  $q_1$  of these clusters are shared with  $\mathfrak{I}_1$  whereas the remaining  $(q_2 - q_1)$  clusters are unique to  $\mathfrak{I}_2$ . In the extreme situation where the images are very different,  $q_2 > q_1$ , and none of the  $q_1$  clusters of  $\mathfrak{I}_1$  are occupied in  $\mathfrak{I}_2$ . At the other extreme, when the two images are nearly identical,  $q_2 = q_1$ , and only the  $q_1$ clusters of  $\mathfrak{I}_1$  are occupied by the p pixels of  $\mathfrak{I}_2$ . Consequently, the latent spatial knots and bandwidths of the image  $\mathfrak{I}_2$  clusters have the following independent normal-inverse Wishart (NIW) conditional priors:

$$(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \mid \mathfrak{I}_1 \overset{indep}{\sim} \begin{cases} \mathcal{W}_d(\boldsymbol{\varphi}_{k1}), & k = 1, \dots, q_1, \\ \mathcal{W}_d(\boldsymbol{\varphi}_0), & k = (q_1 + 1), \dots, q_2, \end{cases}$$
(4)

where the components of the NIW hyperparameter  $\varphi_{k1}$  is described in equation (2) and estimated using the Gibbs sampler for image  $\Im_1$ . Extending the earlier notation for allocation variables, for i = 1, ..., p, let  $[c_{i2} = k]$  represent the event that pixel i of image  $\Im_2$ belongs to cluster k, and let  $\mathbf{s}_{k2}^* = \{\mathbf{s}_i : c_{i2} = k\}$  be the set of pixel locations belonging to cluster k. It can be shown that the conditional prior (4) gives the following expression for the aggregated attraction function:

$$g(\boldsymbol{s}_{k2}^{*}) = \frac{1}{\pi^{m_{k2}d/2}} \times \begin{cases} \frac{\Gamma_{d}(\nu_{k2}/2)}{\Gamma_{d}(\nu_{k1}/2)} \frac{|\boldsymbol{\Sigma}_{k1}|^{\nu_{k1}/2}}{|\boldsymbol{\Sigma}_{k2}|^{\nu_{k2}/2}} \left(\frac{\kappa_{k1}}{\kappa_{k2}}\right)^{d/2}, & k = 1, \dots, q_{1}, \\ \frac{\Gamma_{d}(\nu_{k2}/2)}{\Gamma_{d}(\nu_{0}/2)} \frac{|\boldsymbol{\Sigma}_{0}|^{\nu_{0}/2}}{|\boldsymbol{\Sigma}_{k2}|^{\nu_{k2}/2}} \left(\frac{\kappa_{0}}{\kappa_{k2}}\right)^{d/2}, & k = (q_{1}+1), \dots, q_{2}. \end{cases}$$
(5)

The new quantities in equation (5) are defined as follows. Based on the assigned pixels, the updated NIW prior has hyperparameters  $\varphi_{k2} = (\mu_{k2}, \kappa_{k2}, \Sigma_{k2}, \nu_{k2})$  with components

$$\boldsymbol{\mu}_{k2} = \frac{m_{k1}}{m_{k1} + m_{k2}} \boldsymbol{\mu}_{k1} + \frac{m_{k2}}{m_{k1} + m_{k2}} \bar{\boldsymbol{s}}_{k2}, \quad \boldsymbol{\nu}_{k2} = \boldsymbol{\nu}_{k1} + m_{k2}, \quad \boldsymbol{\kappa}_{k2} = \boldsymbol{\kappa}_{k1} + m_{k2},$$

$$\boldsymbol{\Sigma}_{k2} = \boldsymbol{\Sigma}_{k1} + \boldsymbol{V}_{k2} + \frac{\kappa_{k1} m_{k2}}{\kappa_{k1} + m_{k2}} (\bar{\boldsymbol{s}}_{k2} - \boldsymbol{\mu}_{k2}) (\bar{\boldsymbol{s}}_{k2} - \boldsymbol{\mu}_{k2})'. \quad (6)$$

These quantities rely on the following cluster-related summaries from image  $\mathfrak{I}_2$ : (i) number of pixels,  $m_{k2}$ , allocated to cluster k by the allocation vector  $\mathbf{c}_2$ , (ii) pixel location averages,  $\bar{s}_{k2} = \sum_{i:c_{i2}=k} s_i/m_{k2}$ , and *(iii)* matrix error sum of squares for the pixel locations,  $V_{k2} = \sum_{i:c_{i2}=k} (s_i - \bar{s}_{k2})(s_i - \bar{s}_{k2})'$ , for  $k = 1, \ldots, q_2$ . Then, the conditional sRPM for the pixelcluster allocations of image  $\Im_2$  becomes

$$[\boldsymbol{c}_2 \mid \boldsymbol{c}_1] \propto \alpha^{q_2} \prod_{k=1}^{q_2} g(\boldsymbol{s}_{k2}^*) \, \Gamma(m_{k1} + m_{k2}), \quad \boldsymbol{c}_2 \in \mathcal{B}_p, \tag{7}$$

where  $m_{k1} = 0$  for  $k > q_1$ , i.e., the new clusters of  $\mathfrak{I}_2$  were empty in  $\mathfrak{I}_1$ . Analogously as in image  $\mathfrak{I}_1$ , the intensity or response vector  $\boldsymbol{z}_{i2} = (z_{i12}, \ldots, z_{ir2})$  is distributed as  $\boldsymbol{z}_{i2} \mid \{c_{i2} = k\} \stackrel{iid}{\sim} N_r(\boldsymbol{v}_k, \sigma^2 \boldsymbol{I})$ . For  $k = 1, \ldots, q_1$ , the prior for mean vector  $\boldsymbol{v}_k$  coincides with its posterior from image  $\mathfrak{I}_1$ . For the new clusters indexed by  $k = (q_1 + 1), \ldots, q_2$ , the mean vectors have the previously specified no-data prior,  $\boldsymbol{v}_k \stackrel{iid}{\sim} N_r(\boldsymbol{v}_0, \boldsymbol{\Sigma}_0)$ .

A quantitative measure for image comparisons. Using the parameters of the conditional model for image  $\mathfrak{I}_2$ , specifically, the random number of pixels of  $\mathfrak{I}_2$  allocated to the spatial clusters, we define the *differential pixel proportion* of image  $\mathfrak{I}_2$  to be

$$\psi = \frac{1}{p} \sum_{i=1}^{p} \mathcal{I}(c_{i1} \neq c_{i2}), \tag{8}$$

where  $\mathcal{I}(\cdot)$  is the indicator function. This quantity measures the relative area or volume of the regions in  $\mathcal{R}^d$  where images  $\mathfrak{I}_1$  and  $\mathfrak{I}_2$  are different;  $\psi$  values near 0 (1) are indicative of similar (dissimilar) image pairs. Since the pixel-cluster allocations are a denoised version of the image after accounting for spatial correlation, this advantage is inherited by the measure  $\psi$ . Image comparison problem can be formulated as a single-parameter hypothesis test:

$$H_0: \psi \le \psi^*$$
 ("similar images") versus  $H_1: \psi > \psi^*$  ("different images"), (9)

for some prespecified threshold  $\psi^* \in (0, 1)$ . For image comparisons, the key parameter of interest is  $P(\psi > \psi^* | \mathfrak{I}_1, \mathfrak{I}_2)$ , i.e. the posterior probability of hypothesis  $H_1$ . The two images are declared different (similar) if this posterior probability does (not) exceed 0.5.

#### 3.1 Posterior inferences for differential proportion $\psi$

Due to the computational costs of analyzing high-dimensional images, we perform combined MCMC analysis and hypothesis testing of the two images in separate stages:

- <u>Stage 1:</u> First, image  $\mathfrak{I}_1$  is analyzed using the computationally efficient MCMC strategy described in Section 2.2. In this manner, we obtain  $\hat{c}_1$ , an estimate of the pixel-cluster allocation vector of  $\mathfrak{I}_1$ . From the estimated allocations, we compute the estimated number of clusters  $\hat{q}_1$  and the estimated cluster sizes  $\hat{m}_{11}, \ldots, \hat{m}_{1\hat{q}_1}$ . For estimating these important parameters, a detailed description is presented in Section 2.2.
- <u>Stage 2a:</u> Assuming all parameters specific to image  $\Im_1$  to be equal to their posterior estimates, image  $\Im_2$  is analyzed using the Gibbs sampler of Ishwaran and James (2001) applied to the conditional sRPM (7). The Monte Carlo scheme is analogous to that for  $\Im_1$ , and the computational costs of the two MCMC samplers are identical.
- <u>Stage 2b:</u> For the  $l^{th}$  Gibbs sampling update of the  $\mathfrak{I}_2$  parameters, where  $l = 1, \ldots, L$ , let  $m_{12}^{(l)}, \ldots, m_{q_2^{(l)}2}^{(l)}$  denote the sizes of the  $q_2^{(l)}$  clusters. Then,  $\psi^{(l)} = \sum_{k=\hat{q}_1+1}^{q_2^{(l)}} m_{k2}^{(l)}/p_{k2}^{(l)}$

represents the  $l^{(th)}$  MCMC draw for the parameter  $\psi$ . The values  $\psi^{(1)}, \ldots, \psi^{(L)}$  estimate the marginal posterior density of  $\psi$  and the estimated posterior probability  $\hat{P}(\psi > \psi^* \mid \mathfrak{I}_1, \mathfrak{I}_2) = \frac{1}{L} \sum_{l=1}^{L} I(\psi^{(l)} > \psi^*)$ , is immediately available. Images  $\mathfrak{I}_1$  and  $\mathfrak{I}_2$  are called different or similar depending on whether or not  $\hat{P}(\psi > \psi^* \mid \mathfrak{I}_1, \mathfrak{I}_2)$  exceeds 0.5.

## 4 Simulation studies

To investigate the accuracy of sRPM procedure in labeling images as "similar" or "different," we generated 500 sets of truly similar and truly different image pairs (i.e. 1,000 images in total), with each monochrome image consisting of  $100 \times 100$  pixels each. For each image, the p = 10,000 pixel signals were generated from an adaptation of the Potts model (Hurn et al., 2003) with  $q^* = 45$  hidden states on the square lattice. A Potts model induces spatial correlation in the pixel intensities in a very different manner than sRPM. For hidden states  $\nu_1, \ldots, \nu_p \in \{1, \ldots, q^*\}^p$ , the model induces spatial correlation in the hidden states by borrowing information from neighboring pixels. Let  $\mathcal{N}_i$  be the pixel indices surrounding pixel i, so that  $\mathcal{N}_i$  has 3, 5, and 8 elements depending on whether i belongs to the image interior, edge, or corner. Specifically,  $\left[\nu_1, \ldots, \nu_p\right] = \frac{1}{\mathcal{C}(\Lambda)} \exp\left|\sum_{i=1}^p \sum_{u \in \mathcal{N}_i} \lambda_{\nu_u \nu_i}\right|$ , where  $\nu_i \in$  $\{1,\ldots,q^*\}, i=1,\ldots,p$ , with  $\lambda_{\nu_u\nu_i}>0$  being a similarity measure between hidden states  $\nu_u$  and  $\nu_i$  in a diagonally dominant  $q^* \times q^*$  matrix,  $\mathbf{\Lambda} = ((\lambda_{\nu_i \nu_i}))$ . The similarity measures were related to the pairwise distances of q randomly generated points on the square lattice. The normalizing constant is denoted by  $\mathcal{C}(\boldsymbol{\lambda})$ . Due to its intractability, we generated a sample from the Potts density as the last draw of the corresponding Gibbs sampler.

The  $q^*$  hidden signals associated with the hidden states were generated i.i.d. from the standard normal distribution. Finally, zero-mean Gaussian noise was added to the pixel-specific hidden signals to obtain the continuous image intensities, which were then standardized to belong to the interval [0, 1]. Truly similar images shared the same hidden signals but differed in the pixel-specific noise. Truly different images not only differed in pixel-specific noise, but also had slightly different similarity measures in the Potts densities.

Following image generation, the posterior inference strategy outlined in Section 3 was implemented to analyze each image pair using the sRPM technique and call each image pair as similar or different. Averaging over the 500 datasets, the blue line in Figure 4 displays the receiver operating characteristic (ROC) curve of accurate calls, along with 95% confidence bands, as the threshold similarity  $\psi^*$  was increased from 0 to 1. As is well known, the area under the curve (AUC) close to 1 is indicative of a method's reliability. The estimated AUC for sRPM was 0.928 and a 95% confidence interval was (0.914, 0.942).

Minimizing the distance from the ROC curve to the upper left corner of the square, we obtained the optimal threshold value of  $\psi^* = 0.183$ . The percentage of correct and incorrect calls for the optimal  $\psi^*$  is displayed in Table 1. image comparison. In extensive simulation studies, we have obtained reliable inferences for threshold values belonging to the interval, [0.1, 0.2]. For image comparison in real datasets, we recommend selecting the threshold  $\psi^*$  of hypothesis test (9) in this range. For comparison, we analyzed the same datasets using the image comparison approach based on continuity regions (Feng and Qiu, 2018). Averaging over the 500 datasets, the red line of Figure 4 displays the ROC curve with 95% confidence bands. For the Feng and Qiu (2018) method, the estimated AUC



Figure 4: Estimated ROC curves of the sRPM and Feng and Qiu (2018) approaches for the 500 artificial datasets of the simulation study. The dotted lines represent 95% confidence bands for the ROC curves.

was 0.801, with a 95% confidence interval of (0.770, 0.831), demonstrating the significantly greater effectiveness of sRPM in image comparisons. Minimizing the distance of the competing method's ROC curve to the upper left corner of the square in Figure 4, Table 1 displays the optimal percentage of correct and incorrect image pairs calls for the Feng and Qiu (2018) approach. These results reveal that the sRPM method has approximately balanced levels of sensitivity (i.e., ability to detect differences) and specificity (i.e., ability to detect similarities). In contrast, the Feng and Qiu (2018) strategy has significantly lower sensitivity and specificity. These findings reveal that, at least for the types of images investigated here, the sRPM method is more conservative than the Feng and Qiu (2018) strategy in rejecting the null hypothesis of image similarity.

$\mathbf{sRPM}$			
		Truth	
		Similar	Different
Detected	Similar	86.4% (1.5%)	13.6% (1.5%)
	Different	14.0%~(1.6%)	86.0%~(1.6%)
Feng and Qiu (2018)			
		Truth	
		Similar	Different
Detected	Similar	78.4% (1.8%)	21.6% (1.8%)
	Different	20.8% (2.0%)	70.2% (2.0%)

Table 1: Averaging over the 500 artificial image pairs of the simulation study, the optimal detection accuracies of the sRPM and Feng and Qiu (2018) methods for image comparison. See the text for further explanation. Shown in parentheses are standard errors.

# 5 Analysis of satellite images

#### Chicago area images

Figure 5 displays monochrome satellite images of the Chicago area taken in 1990 (image  $\Im_1$ ) and in 1999 (image  $\Im_2$ ). The data are available at available at https://users.phhp.ufl.edu/ pqiu/research/book/data/index.html. From an informal visual comparison, we find that the two images have a few differences. Following the 9-year interval between the images, there appear to be more dark spots in image  $\Im_2$ . These changes may be due to environmental changes or new construction. Some dark spots in image  $\Im_1$  also appear in image  $\Im_2$ , although their sizes have changed, e.g., in the lower portions of the images. Most of the differences between the images are fairly small in magnitude, difficult to describe using a parametric model, and are mostly local, affecting only small portions of the images.

For calling the two images similar or different, we focused on the differential pixel proportion,  $\psi$ , defined in equation (8), and performed hypothesis test (9) to obtain an MCMC



(a) Image  $\mathfrak{I}_1$ 



(b) Image  $\mathfrak{I}_2$ 

Figure 5: Reference satellite image  $\mathfrak{I}_1$  was taken in 1990 in the Chicago area. Satellite image  $\mathfrak{I}_2$  of the same area taken in 1999.

estimate of posterior probability  $P(\psi > \psi^* | \mathfrak{I}_1, \mathfrak{I}_2)$ . As mentioned earlier, extensive simulation studies have suggested a threshold value  $\psi^*$  belonging to the interval [0.1, 0.2] to obtain reliable inferences. In general, two images were declared to different or similar depending on whether or not the estimate  $\hat{P}(\psi > \psi^* | \mathfrak{I}_1, \mathfrak{I}_2)$  exceeded 0.5.

For the Chicago area images, the posterior inference procedure outlined in Section 3.1 was implemented to obtain the marginal posterior of differential parameter  $\psi$ . An MCMC estimate of the marginal posterior is shown in the left panel of Figure 6. The right panel of Figure 6 plots the estimated posterior probability  $\hat{P}(\psi > \psi^* | \mathfrak{I}_1, \mathfrak{I}_2)$  as  $\psi^*$ varies over the range [0, 0.5]. We find that as  $\psi^*$  increases, the posterior probability stays level at approximately 1 before falling sharply as  $\psi^*$  approaches 0.5. That is, irrespective of the prespecified threshold  $\psi^* \in [0.1, 0.2]$ , the estimated posterior probability exceeds 0.99, which is a Bayes factor greater than 100 if we assign equal prior probabilities to the competing hypotheses. This is decisive evidence that the Chicago area satellite images are different.



Figure 6: For the Chicago area images, the left panel displays an MCMC estimate of the marginal posterior of differential parameter  $\psi$ . The right panel displays an estimate of the posterior probability  $P(\psi > \psi^* | \mathfrak{I}_1, \mathfrak{I}_2)$  for different values of the threshold  $\psi^*$ .

#### Bay Area images

Figure 7 displays two satellite images of the San Francisco Bay Area in 1990 and 1999. Figure 1 of Supplementary Material displays the estimated marginal posterior of differential parameter  $\psi$ . We find negligible posterior mass assigned to  $\psi$  smaller than 0.48. Consequently, the estimated probability  $\hat{P}(\psi > \psi^* | \mathfrak{I}_1, \mathfrak{I}_2)$  was equal to 1 for any prespecified threshold  $\psi^* \in [0.1, 0.2]$ , and is overwhelming evidence that the two satellite images are different.

# 6 Discussion

There is a critical need for statistical techniques that can reliably identify "similar" or "different" image pairs depending on whether random measurement error is able to account for pixel-wise differences between the images. We have developed a Bayesian strategy rely-



Figure 7: In the left panel, satellite image  $\mathfrak{I}_1$  of the San Francisco bay area, taken in 1990. In the right panel, satellite image  $\mathfrak{I}_2$  of the same region, taken in 1999.

ing on a novel extension process called the spatial random partition model (sRPM). This process achieves dimension reduction while accounting for spatial correlation by grouping proximal image pixels with similar intensities into clusters. The extracted image information is then utilized for image comparison via a single-parameter hypothesis test based on a univariate metric that is adaptive to pixel intensity variation and spatial correlation. We foster a computationally efficient two-stage MCMC technique. A simulation study analyzes artificial datasets and finds compelling evidence for the high sensitivity and specificity of sRPM. The success of sRPM is demonstrated by analyzing satellite image data. R code implementing the method is available at https://github.com/sguha-lab/sRPM.

The nonparametric nature of sRPM makes it highly effective in analyzing images with diffuse spatially differentiated regions. However, the method may be challenged by images with sharply demarcated features (e.g., text images) or jagged images with steep slopes. In applications where a sequence of multiple images is available from a longitudinal process, interest often focuses on detecting temporal change points, if any, at which the longitudinal process diverges beyond acceptable limits of variation. We are currently developing natural extensions of the proposed sRPM analytical framework in such a manner that, at any point in the image sequence, only cluster-related information that has been progressively extracted from the preceding images is used to determine whether the current time point is a change point. It is often known that the first few images of the sequence correspond to an "in-control" condition of the longitudinal process. In these situations, we are devising strategies that exploit this information to further improve the inferential accuracy of monitoring image sequences. The MCMC strategy of Ishwaran and James (2001) is adequate for the image sizes of this paper. However, for ultra high-dimensional images, these algorithms are computationally intensive. We are exploring whether the Metropolis-Hastings algorithm proposed by Guha (2010) could potentially be applied to significantly reduce computational times while ensuring that the posterior mixing rates and ESS in the sense of Turek et al. (2017) are satisfactory.

## Acknowledgements

This work was supported by NSF grants DMS-1854003 awarded to SG and DMS-1914639 awarded to PQ, and by NIH awards R01 CA269398 and U01 CA209414 to SG. We thank the Associate Editor and two anonymous referees for many insightful remarks that improved the paper's focus and content.

## **Conflict of Interest**

The authors report there are no competing interests to declare.

## References

- Chatzis, Sotirios P and Tsechpenakis, Gabriel (2010), "The infinite hidden Markov random field model," *IEEE Transactions on Neural Networks*, 21, 6, 1004–1014.
- Dahl, D. B. (2006), Model-Based Clustering for Expression Data via a Dirichlet Process Mixture Model. Cambridge University Press.
- Duan, J. A., Guindani, M., and Gelfand, A. E. (2007), "Generalized spatial Dirichlet process models," *Biometrika*, 94, 809–825.
- Duchesne, C., Liu, J., and MacGregor, J. (2012), "Multivariate image analysis in the process industries: A review," *Chemometrics and Intelligent Laboratory Systems*, 117, 116–128.
- Dunson, D. B. and Park, J.-H. (2008), "Kernel stick-breaking processes," *Biometrika*, 95, 307–323.
- Feng, Long and Qiu, Peihua (2018), "Difference detection between two images for image monitoring," *Technometrics*, 60, 3, 345–359.
- Ferguson, Thomas S (1973), "A Bayesian analysis of some nonparametric problems," *The* Annals of Statistics, 209–230.
- Gelfand, A., Kottas, A., and MacEachern, S. (2005), "Bayesian nonparametric spatial modeling with Dirichlet processes mixing," *Journal of the American Statistical Association*, 100, 1021–1035.
- Ghosal, S., Ghosh, J. K., and Ramamoorthi, R. V. (1999), "Posterior consistency of Dirichlet mixtures in density estimation," *The Annals of Statistics*, 27, 143–158.
- Griffin, J. E. and Steel, M. F. J. (2006), "Order-based dependent Dirichlet processes," Journal of the American Statistical Association, 101, 179–194.
- Guha, S. (2010), "Posterior simulation in countable mixture models for large datasets," Journal of the American Statistical Association, 105, 775–786.
- Guo, X. and Zhang, C.M. (2015), "Difference detection between two images for image monitoring," *Statistica Sinica*, 25, 475–498.
- Guo, Y., Jia, X., and Paull, D. (2018), "Effective sequential classifier training for SVMbased multitemporal remote sensing image classification," *IEEE Transactions on Image Processing*, 27, 475–498.
- Hurn, M.A., Husby, O.K., and Rue, H. (2003), A Tutorial on Image Analysis. Springer, New York, NY.
- Ishwaran, Hemant and James, Lancelot F (2001), "Gibbs sampling methods for stickbreaking priors," Journal of the American Statistical Association, 96, 453, 161–173.

- Jiang, W., Han, S.W., Tsui, K.L., and Woodall, W.H. (2011), "Spatiotemporal surveillance methods in the presence of spatial correlation," *Statistics in Medicine*, 30, 569–583.
- Julea, A., Meger, N., Bolon, P., Rigotti, C., Doin, M.P., Lasserre, C., Trouve, E., and Lazarescu, V.N. (2011), "Unsupervised spatiotemporal mining of satellite image time series using grouped frequent sequential patterns," *IEEE Transactions on Geoscience* and Remote Sensing, 49, 1417–1430.
- Kim, S., Tadesse, M. G., and Vannucci, M. (2006), "Variable selection in clustering via Dirichlet process mixture models," *Biometrika*, 93, 877–893.
- Koosha, Mehdi, Noorossana, Rassoul, and Megahed, Fadel (2017), "Statistical process monitoring via image data using wavelets," *Quality and Reliability Engineering Interna*tional, 33, 8, 2059–2073.
- Lijoi, A. and Prünster, I. (2010), "Models beyond the Dirichlet process," In Bayesian Nonparametrics, eds. N. L. Hjort, C. Holmes, P. Müller, and S. G. Walker, Cambridge, U.K.: Cambridge Series in Statistical and Probabilistic Mathematics, pp. 80–136.
- Lin, H.D., Chung, C.Y., and Lin, W.T. (2008), "Principal component analysis based on wavelet characteristics applied to automated surface defect inspection," WSEAS Transactions on Computer Research, 3, 193–202.
- Lindquist, M.A. (2008), "The statistical analysis of fMRI data," *Statistical Science*, 23, 439–464.
- Megahed, F.M., Wells, L.J., Camelio, J.A., and Woodall, W.H. (2012), "A spatiotemporal method for the monitoring of image data," *Quality Reliability and Engineering Interna*tional, 28, 967–980.
- Megahed, F.M., Woodall, W.H., and Camelio, J. A. (2011), "A review and perspective on control charting with image data," *Journal of Quality Technology*, 43, 83–98.
- Müller, P., Quintana, F., and Rosner, G. L. (2011), "A product partition model with regression on covariates," *Journal of Computational and Graphical Statistics*, 20, 260–278.
- Na, In Seop, Cho, Wan Hyun, Kim, Soo Hyung, Kang, Soon Ja, Nguyen, Trung Quy, and Choi, Jun Yong (2013), Unsupervised color images segmentation using spatial hidden MRF GDPM model. in Proceedings of the 7th International Conference on Ubiquitous Information Management and Communication, 1–8.
- Prats-Montalban, J.M. and Ferrer, A. (2014), "Statistical process control based on multivariate image analysis: A new proposal for monitoring and defect detection," *Computers* and Chemical Engineering, 71, 501–511.
- Qiu, P. (2014), Introduction to Statistical Process Control, Chapman & Hall/CRC.

- (2020), "Big data? Statistical process control can help," *The American Statistician*, DOI: 10.1080/00031305.2019.1700163.
- Reich, B.J. and Fuentes, M. (2015), *Spatial Bayesian Nonparametric Methods*. Springer International Publishing.
- Reich, Brian J and Bondell, Howard D (2011), "A spatial Dirichlet process mixture model for clustering population genetics data," *Biometrics*, 67, 2, 381–390.
- Rodriguez, A., Dunson, D. B., and Gelfand, A. E. (2010), "Latent stick-breaking processes," Journal of the American Statistical Association, 105, 647–659.
- Roy, Anik and Mukherjee, Partha Sarathi (2024), "Image comparison based on local pixel clustering," *Technometrics*, 1–12.
- Turek, Daniel, de Valpine, Perry, Paciorek, Christopher J, and Anderson-Bergman, Clifford (2017), "Automated parameter blocking for efficient Markov chain Monte Carlo sampling," *Bayesian Analysis*, 12, 2, 465–490.
- Wang, Xiao-Feng and Ye, Deping (2010), "On nonparametric comparison of images and regression surfaces," *Journal of statistical planning and inference*, 140, 10, 2875–2884.
- Yan, H., Paynabar, K., and Shi, J. (2015), "Image-based process monitoring using low-rank tensor decomposition," *IEEE Transactions on Automation Science and Engineering*, 12, 216–227.