

Scalable Bayesian Inference for the Inverse Temperature of a Hidden Potts Model

Matthew Moores^{*,†}, Geoff K. Nicholls[‡], Anthony N. Pettitt[§], and Kerrie Mengersen[§]

Abstract. The inverse temperature parameter of the Potts model governs the strength of spatial cohesion and therefore has a major influence over the resulting model fit. A difficulty arises from the dependence of an intractable normalising constant on the value of this parameter and thus there is no closed-form solution for sampling from the posterior distribution directly. There is a variety of computational approaches for sampling from the posterior without evaluating the normalising constant, including the exchange algorithm and approximate Bayesian computation (ABC). A serious drawback of these algorithms is that they do not scale well for models with a large state space, such as images with a million or more pixels. We introduce a parametric surrogate model, which approximates the score function using an integral curve. Our surrogate model incorporates known properties of the likelihood, such as heteroskedasticity and critical temperature. We demonstrate this method using synthetic data as well as remotely-sensed imagery from the Landsat-8 satellite. We achieve up to a hundredfold improvement in the elapsed runtime, compared to the exchange algorithm or ABC. An open-source implementation of our algorithm is available in the R package **bayesImageS**.

Keywords: approximate Bayesian computation, exchange algorithm, hidden Markov random field, image analysis, indirect inference, intractable likelihood.

MSC 2010 subject classifications: Primary 62M40, 62F15; secondary 62-04.

1 Introduction

Markov random field (MRF) models have seen widespread use in image analysis since their introduction by Besag (1974), as surveyed by Winkler (2003) and Li (2009). A MRF is a generalisation of the Markovian dependence structure to more than one dimension: satellite imagery has two spatial dimensions, while medical images such as computed tomography (CT) are three-dimensional. The hidden Potts (1952) model employs a latent MRF on discrete states to describe spatial dependence between adjacent neighbours. The degree of dependence in the model is governed by a parameter β , known as the inverse temperature due to its origin in statistical physics (Huang, 2010). It is

*Department of Statistics, University of Warwick, Coventry CV4 7AL, UK.

†*Current Address:* National Institute for Applied Statistics Research Australia (NIASRA), School of Mathematics & Applied Statistics, University of Wollongong, NSW 2522, Australia, mmoores@uow.edu.au, url: <https://uow.edu.au/~mmoores>

‡Department of Statistics, University of Oxford, 24-29 St Giles Oxford OX1 3LG, UK, nicholls@stats.ox.ac.uk

§Australian Research Council Centre for Excellence in Mathematical and Statistical Frontiers (ACEMS) and the School of Mathematical Sciences, Queensland University of Technology, Brisbane, Queensland 4001, Australia, a.pettitt@qut.edu.au, k.mengersen@qut.edu.au

Number of pixels	Satellite (900m ² /px)	CT slices (512×512)	HD Video 1080p
4 ⁶	3.7km ²
7 ⁶	105.9km ²	0.4	...
10 ⁶	900.0km ²	3.8	0.5
15 ⁶	10251.6km ²	43.5	5.5

Table 1: Scale of common types of images, including the area covered by a satellite image, the number of axial slices of a CT scan, and the number of frames of HD video.

difficult to set this parameter by trial and error, particularly for noisy images. Rather than using a fixed value, it would be preferable to estimate β as part of the model fitting. However, the intractable normalising constant of the Potts model depends on the value of the inverse temperature. This means that there is no closed-form solution for sampling from its posterior distribution, since the Metropolis-Hastings (M-H) ratio cannot be computed directly.

Møller et al. (2006) introduced an auxiliary-variable method that gives an exact MCMC algorithm for the special case of a 2-component Potts model, also known as an Ising model. The exchange algorithm of Murray et al. (2006) is a variant of this exact method that avoids the need for a fixed estimate of β . A drawback of these algorithms is that they require unbiased sampling from the stationary distribution of the Potts model. This is possible for $k = 2$ or 3 labels using coupling from the past (CFTP; Propp and Wilson, 1996), perfect slice sampling (Mira et al., 2001), or bounding chains (Huber, 2003, 2016). The recursive algorithm of Reeves and Pettitt (2004) can also be used to obtain an exact sample (Friel and Rue, 2007), but only if the lattice is very small. McGrory et al. (2009) reported that the time required for CFTP increased sharply for larger values of β . The available methods for perfect simulation can be computationally prohibitive for practical applications. For this reason, Cucala et al. (2009) substitute 500 iterations of Gibbs sampling on the auxiliary variable to produce an approximate sample from its stationary distribution. This approximate exchange algorithm (AEA) can also be applied for Potts models with $k > 3$, where the state space \mathcal{Z} cannot be partially ordered.

Like the exchange algorithm, approximate Bayesian computation (ABC) uses an auxiliary variable to decide whether to accept or reject the proposed value of β . ABC for the Potts model was introduced by Grelaud et al. (2009). They observed that these models form a natural exponential family and hence possesses a sufficient statistic, $S(\mathbf{z})$. Everitt (2012) combined AEA with particle MCMC and also implemented ABC with sequential Monte Carlo (SMC-ABC) for the Ising model. Everitt found that the computational costs of both AEA and ABC were dominated by simulation of the auxiliary variable. While exact inference is theoretically possible, its applicability is limited to datasets with fewer than a thousand pixels. Comparisons such as McGrory et al. (2009); Moores and Mengersen (2014) and Moores et al. (2015) have shown that auxiliary-variable methods such as AEA and ABC are infeasible for the scale of data that is regularly encountered in image analysis.

Images containing multiple megapixels are now commonplace, from digital photography and high-definition (HD) video to medical imaging and remote sensing. Table 1 illustrates how the number of pixels translates to real-world scale for various types of images. There is often the need to use multiple images to cover the spatial (and temporal) extent of an imaging study. Multispectral images from the Landsat-7 (NASA, 2011) and Landsat-8 (USGS, 2016) satellites have a spatial resolution of 30 metres per pixel (area of 900m^2). A Landsat image covers an area of approximately 170km north-south by 183km east-west. Tomographic reconstructions such as CT scans are usually represented as 3D image stacks, with 512×512 pixels per slice. The pixel resolution and slice width varies depending on the clinical protocol. A single frame of HD video is typically 1920×1080 pixels with 24 frames per second (fps), although higher frame rates and resolutions (such as ultra-high definition, UHD) are available. The volumes of data involved in video processing necessitate specialised methods, which are beyond the scope of this paper. For examples of Bayesian methods for video analysis, see Simoncelli (1999); Minvielle et al. (2010), and the references therein. The remainder of the discussion will focus on static 2D images.

Scalable inference for intractable likelihoods is an active area of research. Some algorithms have been able to improve their runtime while still targeting the exact posterior distribution. For example, both Russian roulette (Lyne et al., 2015) and lazy ABC (Prangle, 2016) used random truncation of the likelihood computation. Liang et al. (2016) used an auxiliary chain to define an importance sampling (IS) distribution for the exchange algorithm. Sherlock et al. (2017) used k -nearest neighbours (k NN) combined with delayed acceptance (Christen and Fox, 2005) to accelerate pseudo-marginal MCMC. Compared to the exchange algorithm or ABC, these “exact-approximate” methods have provided an improvement in computational efficiency of an order of magnitude or less. This is not always sufficient to meet the practical requirements of real-world applications.

For Bayesian inference to be feasible on large image datasets, it is necessary to sacrifice some degree of accuracy in order to achieve further improvements in speed. A variety of approximate methods have recently been proposed, which provide a tradeoff between accuracy and computational cost. One particular class of methods, known as Bayesian indirect likelihood (BIL; Drovandi et al., 2011, 2015), employs a surrogate model to approximate the likelihood function. BIL accelerates computation by interpolating between previous values of the auxiliary variables. It can also take advantage of massively parallel hardware to initialise the surrogate model using a precomputation step. Moores et al. (2015) used linear interpolation to accelerate ABC for the hidden Potts model. Boland et al. (2017) derived a theoretical upper bound on the bias introduced by this and similar piecewise approximations. Conrad et al. (2016) used local polynomial or Gaussian process (GP) models to accelerate MCMC. Wilkinson (2014) and Meeds and Welling (2014) used GP emulation for ABC, while Drovandi et al. (2018) used a GP in a pseudo-marginal algorithm. Järvenpää et al. (2016) used a heteroskedastic GP model and demonstrated how the output of the precomputation step could be used for Bayesian model choice. The IS (Liang et al., 2016) and k NN (Sherlock et al., 2017) approaches could also be viewed as forms of BIL. These nonparametric surrogate models approximate the likelihood under fairly generic assumptions regarding its functional form.

In this paper, we introduce the parametric functional approximate Bayesian (PFAB) algorithm for the Potts model. This algorithm incorporates a surrogate model that approximates the distribution of the sufficient statistic, therefore it is a type of BIL. The advantage over previous BIL algorithms is that we use a parametric integral curve that reflects the known properties of the intractable likelihood function, providing efficiency gains in comparison to nonparametric alternatives. We estimate the parameters of the function using a parallel precomputation step. This leads to a great improvement in runtime, since there is no longer any need to simulate auxiliary variables during model fitting. Using a simulation study, we demonstrate empirically that this approximation error has very little impact on the resulting estimate of the posterior distribution. We also apply our algorithm to images from the Landsat-8 satellite, where the computational cost of AEA or ABC is prohibitive. The superior scalability of PFAB makes it possible to perform fully-Bayesian inference for demanding applications such as this.

The remainder of this article is organised as follows. The synthetic data and satellite imagery are described in Section 2. We define the hidden Potts model in Section 3 and examine the properties of its intractable normalising constant. Our PFAB algorithm for estimating the inverse temperature is described in Section 4. We present our experimental results in Section 5 and conclude the article with a discussion.

2 Example datasets

2.1 Synthetic data

The inverse temperature cannot be directly observed, so the only way to measure the accuracy of a method is to use synthetic data with known parameter values. We employ the simulation-based calibration (SBC) method of Cook et al. (2006); Talts et al. (2018) to evaluate the quality of our PFAB approximation to the posterior for β . The central concept of this approach is to draw a set of values for the model parameters from the prior distribution, then to simulate data from the generative model for each set of parameters, and finally to fit the model and compare the estimates of the posterior distribution to the known values. Simulation-based methods for validating Bayesian posterior coverage have also been proposed by Monahan and Boos (1992); Geweke (2004); and in the context of ABC by Prangle et al. (2014). Cook et al. used posterior quantiles for this comparison, which assumes independent samples (Gelman, 2017). Instead, we compute the rank statistics suggested by Talts et al., which are more applicable to autocorrelated samples obtained from MCMC methods.

We use a fixed number of 5 mixture components to generate 7 datasets containing 20 images each. The images within a dataset all have the same dimensions: the smallest images have $2^6 = 8 \times 8$ pixels; the other image sizes are $3^6 = 27 \times 27$; $4^6 = 64 \times 64$; $5^6 = 125 \times 125$; $6^6 = 216 \times 216$; $7^6 = 343 \times 343$; and $10^6 = 1000 \times 1000$. This enables us to study how the performance of the algorithms changes as the size of the images increases. For $n = 5^6$ pixels, we also generate datasets with $k = 2, 3$ or 4 hidden states. This gives us 10 simulated datasets in total, containing 200 synthetic images. After discarding the initial, transient phase of the MCMC sampler as burn-in, we thin the

output from PFAB to obtain $L = 175$ approximately-independent samples from the posterior (Sect. 5.1 of Talts et al. explains why this is necessary).

2.2 Satellite remote sensing

Abundant image data are available from Earth observation satellites. The size of satellite images and the frequency with which they are generated create a requirement for automated methods of image processing. Remotely-sensed satellite imagery has been used for a variety of purposes, such as estimating land use (Small, 2001), water quality (McClain, 2009) and economic growth (Henderson et al., 2011). In this study, we aim to classify the pixels in a satellite image according to the type and abundance of vegetation that is present. By labelling the pixels, we can quantify the levels of vegetation in an area and identify contiguous clusters of forest and parkland. These estimates could be used by environmental scientists to estimate global levels of plant biomass. Land-use changes such as vegetation clearing could be detected by monitoring changes in pixel classification over time.

Metrics such as the normalised-difference vegetation index (NDVI) use the ratio of reflectance of red visible light (VIS) and near-infrared (NIR) light as a proxy for the abundance of chlorophyll (Tucker, 1979; Roy et al., 2016):

$$\text{NDVI} = \frac{\text{NIR} - \text{VIS}}{\text{NIR} + \text{VIS}}. \quad (1)$$

Land surface reflectance is estimated from Landsat-8 satellite imagery by correcting for atmospheric effects (Flood, 2014; Vermote et al., 2016). When converted to NDVI, this produces numbers on a scale from -1 to $+1$. NDVI is generally used to detect vegetated areas, with an expected value for vegetation being between 0.3 and 0.8. In this study, we aim to classify pixels as vegetation, developed, or water. Using NDVI, we are able to distinguish between urbanised or industrial areas as one sub-category and suburbs as a separate category. We also use separate categories for deep forest and areas of light vegetation. These five categories span the range of NDVI values. However there is no clear separation between them, as illustrated by the histogram in Figure 1b. A simple thresholding approach would be insufficient to categorise these pixels. Instead, we require a probabilistic approach that takes into account the spatial similarity between pixel labels.

We apply our PFAB algorithm to analyse Landsat-8 imagery of Brisbane, Australia. The Landsat-8 satellite was launched in February 2013 and orbits the Earth every 16 days. To minimise artefacts, we only use imagery with less than 10% cloud cover. There are 18 such images available for 2015 and 2016. These images were cropped to a region of interest 30 km square, containing one million pixels. This roughly corresponds to the greater metropolitan area. An example of one of these images is shown in Figure 1a. Images from Landsat-8 have 16 bits per pixel, with 55,000 possible values. This is an improvement over Landsat-5 and 7, which produce 8 bit images (256 values). The Landsat-8 bands are also narrower, which provides higher spectral resolution. The VIS band contains wavelengths between 636 and 673 nm, while the NIR band is from 851 to 879 nm (USGS, 2016).

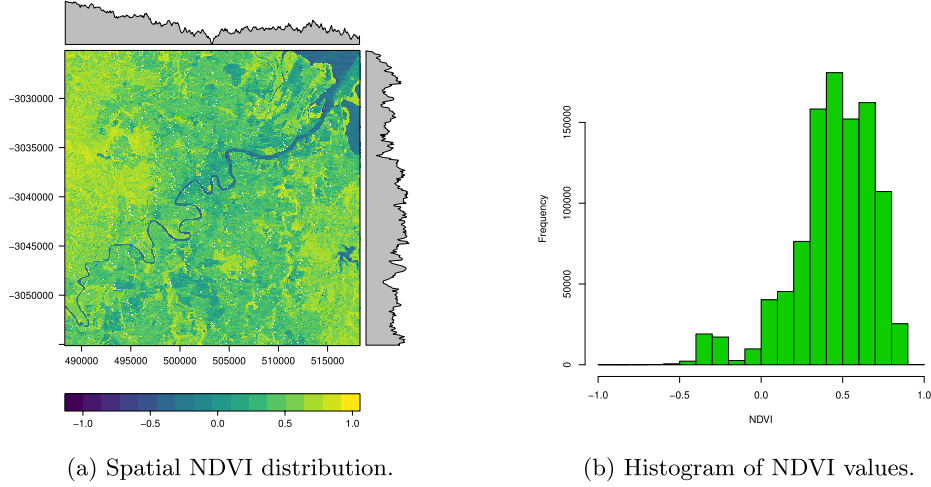


Figure 1: Normalised difference vegetation index (NDVI) computed from a Landsat-8 satellite image of Brisbane, Australia.

3 Hidden Potts model

Image segmentation can be viewed as the task of labelling the observed pixels $y_i \in \mathbb{R}$, $i = 1, \dots, n$ according to a finite set of discrete states $\mathbf{z} \in \{1, \dots, k\}^n$. The hidden Potts model allows for spatial correlation between neighbouring labels in the form of a MRF. The latent labels follow a Gibbs distribution, which is specified in terms of its conditional probabilities:

$$p(z_i | z_{\setminus i}, \beta) = \frac{\exp \left\{ \beta \sum_{\ell \in \partial(i)} \delta(z_i, z_\ell) \right\}}{\sum_{j=1}^k \exp \left\{ \beta \sum_{\ell \in \partial(i)} \delta(j, z_\ell) \right\}}, \quad (2)$$

where β is the inverse temperature, $z_{\setminus i}$ represents all of the labels except z_i , $\partial(i)$ are the neighbouring pixels of i , and $\delta(u, v)$ is the Kronecker delta function. Thus, $\sum_{\ell \in \partial(i)} \delta(z_i, z_\ell)$ is a count of the neighbours that share the same label.

If the pixels in a rectangular 2D lattice with r rows and c columns are indexed row-wise, the nearest (first-order) neighbours $\partial(i)$ are $\ell \in \{i-1, i-c, i+c, i+1\}$, except on the boundary. Pixels situated at the boundary of the image domain have fewer than four neighbours. These neighbourhood relationships are reciprocal, so $h \in \partial(i)$ implies $i \in \partial(h)$. If \mathcal{E} is the set of all unique neighbour pairs or edges $\langle h, \ell \rangle$ in the image lattice, then $\#\mathcal{E} = 2(n - \sqrt{n})$ for a square lattice or $2r \cdot c - r - c$ for a rectangular lattice.

The observation equation links the latent labels to the corresponding pixel values:

$$p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta}) = \prod_{i=1}^n p(y_i | z_i, \theta_{z_i}), \quad (3)$$

where θ_{z_i} are the parameters that govern the distribution of the pixel values with label z_i . The hidden Potts model can thus be viewed as a spatially-correlated generalisation of the finite mixture model (Rydén and Titterton, 1998). Green and Richardson (2002) used a Poisson likelihood for (3), with intensity λ_{z_i} . Instead we follow Geman and Geman (1984); Alston et al. (2007), and many others in assuming that the pixels with label j share a common mean μ_j corrupted by additive Gaussian noise with variance σ_j^2 :

$$y_i | z_i = j, \mu_j, \sigma_j^2 \sim \mathcal{N}(\mu_j, \sigma_j^2). \quad (4)$$

The joint distribution of all of the pixel labels can be expressed in the form of an exponential family, as noted by Grelaud et al. (2009):

$$p(\mathbf{z} | \beta) = \exp\{\beta S(\mathbf{z}) - \log \mathcal{C}(\beta)\}. \quad (5)$$

The augmented likelihood $p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}, \beta)$ can therefore be factorised into $p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta})p(\mathbf{z} | \beta)$, where the second factor does not depend on the observed data, but only on the sufficient statistic:

$$S(\mathbf{z}) = \sum_{\langle i, \ell \rangle \in \mathcal{E}} \delta(z_i, z_\ell). \quad (6)$$

This statistic represents the total number of like neighbour pairs in the image. The joint posterior is then:

$$p(\boldsymbol{\theta}, \beta, \mathbf{z} | \mathbf{y}) \propto p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta})p(\boldsymbol{\theta})p(\mathbf{z} | \beta)p(\beta), \quad (7)$$

where $p(\boldsymbol{\theta})$ is the joint prior for the parameters of the observation equation (3) and $p(\beta)$ is the prior for the inverse temperature. We use informative, conjugate priors for $\mu \sim \mathcal{N}(m_0, s_0^2 I)$ and $\sigma_j^2 \sim \text{InvG}(n_0/2, n_0 v_0/2)$, based on the properties of the NDVI values:

$$m_0 = (-1.0, -0.5, 0.0, +0.5, +1.0)^T, \quad s_0^2 = 0.1^2, \quad (8)$$

$$n_0 = 5, \quad v_0 = 0.024. \quad (9)$$

The $k = 5$ mixture components correspond to different types of land use. In increasing order of their means, these can be classified as: water; urban/industrial; suburban; light vegetation; and dense vegetation. Although these priors are informative for the parameters, there is no direct mapping between the NDVI values and their categorisation. There is too much overlap between the distributions, as illustrated by Figure 1b. This is why we require the probabilistic labelling approach of the hidden Potts model, rather than a simpler method such as thresholding. For the smoothing parameter β , we follow previous authors including Møller et al. (2006); McGrory et al. (2012); Everitt (2012); and Lyne et al. (2015) in using a uniform prior over a bounded, finite domain. This prior is useful for comparing different algorithms, since it places equal weight on a range of possible values for β , including both ordered ($\beta > \beta_{crit}$) and disordered ($\beta < \beta_{crit}$) states, as well as the critical point.

The conditional posterior distributions $p(\boldsymbol{\theta} | \mathbf{z}, \mathbf{y})$ and $p(z_i | z_{\setminus i}, \beta, y_i, \boldsymbol{\theta}_{z_i})$ can be simulated using Gibbs sampling, but $p(\beta | \mathbf{y}, \mathbf{z}, \boldsymbol{\theta})$ involves an intractable normalising constant:

$$p(\beta | \mathbf{y}, \mathbf{z}, \boldsymbol{\theta}) = \frac{p(\mathbf{z} | \beta)\pi(\beta)}{\int_{\beta} p(\mathbf{z} | \beta)\pi(d\beta)} \quad (10)$$

$$\propto \frac{\exp\{\beta S(\mathbf{z})\}}{\mathcal{C}(\beta)} \pi(\beta). \quad (11)$$

The normalising constant $\mathcal{C}(\beta)$ is also known as a partition function in statistical physics. It has computational complexity of $\mathcal{O}(nk^n)$, since it involves a sum over all possible configurations of the labels $\mathbf{z} \in \mathcal{Z}$:

$$\mathcal{C}(\beta) = \sum_{\mathbf{z} \in \mathcal{Z}} \exp\{\beta S(\mathbf{z})\}. \quad (12)$$

It is infeasible to calculate this value exactly for nontrivial images, thus computational approximations are required.

The conditional expectation of $S(\mathbf{z})$ given β can be expressed in terms of the normalising constant:

$$\mathbb{E}_{\mathbf{z}|\beta}[S(\mathbf{z})] = \frac{d}{d\beta} \log\{\mathcal{C}(\beta)\}. \quad (13)$$

As β approaches infinity, all of the pixels in the image are almost surely assigned the same label, thus the expectation of $S(\mathbf{z})$ approaches the total number of edges $\#\mathcal{E}$ asymptotically, while the variance approaches zero. When $\beta = 0$, (5) simplifies to $(\sum_{\mathbf{z} \in \mathcal{Z}} \exp\{0\})^{-1} = k^{-n}$, hence the labels z_i are independent and uniformly-distributed. We can also express the expectation and variance of $S(\mathbf{z})$ in closed form, in terms of the sum over configuration space \mathcal{Z} , according to the following lemma:

Lemma 1. *The sum over configuration space of the sufficient statistic of the k -state Potts model with first-order neighbours is $\sum_{\mathbf{z} \in \mathcal{Z}} S(\mathbf{z}) = k^{n-1} \#\mathcal{E}$. Similarly, the sum of the square of the sufficient statistic is $\sum_{\mathbf{z} \in \mathcal{Z}} S(\mathbf{z})^2 = \#\mathcal{E}^2 k^{n-2} + \#\mathcal{E} k^{n-1} (1 - k^{-1})$.*

Proof. We can prove the first equation by switching the order of summation:

$$\sum_{\mathbf{z} \in \mathcal{Z}} S(\mathbf{z}) = \sum_{\mathbf{z} \in \mathcal{Z}} \sum_{\langle i, \ell \rangle \in \mathcal{E}} \delta(z_i, z_\ell) = \sum_{\langle i, \ell \rangle \in \mathcal{E}} \sum_{\mathbf{z} \in \mathcal{Z}} \delta(z_i, z_\ell) = \sum_{\langle i, \ell \rangle \in \mathcal{E}} \frac{k^n}{k} = k^{n-1} \#\mathcal{E}. \quad (14)$$

The cardinality of the configuration space $\#\mathcal{Z}$ is k^n . When we reverse the order of summation, we consider only a single pair of indices $\langle i, \ell \rangle$ at a time. The Kronecker delta $\delta(z_i, z_\ell) = 1$ for $\frac{1}{k}$ out of the possible values of z_i and z_ℓ .

By (6), the second equation can be expressed as a double sum:

$$\sum_{\mathbf{z} \in \mathcal{Z}} S(\mathbf{z})^2 = \sum_{\mathbf{z} \in \mathcal{Z}} \sum_{\langle i, j \rangle \in \mathcal{E}} \sum_{\langle k, \ell \rangle \in \mathcal{E}} \delta(z_i, z_j) \delta(z_k, z_\ell) \quad (15)$$

$$= \sum_{\langle i, j, k, \ell \rangle \in \mathcal{E}^2} \sum_{\mathbf{z} \in \mathcal{Z}} \delta(z_i, z_j) \delta(z_k, z_\ell), \quad (16)$$

where $\mathcal{E}^2 = \bigcup_{\langle i, j \rangle \in \mathcal{E}} \bigcup_{\langle k, \ell \rangle \in \mathcal{E}} \{\langle i, j, k, \ell \rangle\}$. Note that the combined indices of the set \mathcal{E}^2 are not unique, since any given pixel can appear twice, so terms involving $\langle i, j, k, \ell \rangle$,

$\langle i, j; i, \ell \rangle$ and $\langle i, j; i, j \rangle$ all appear in the sum. We can partition \mathcal{E}^2 into the sets of all pairs of edges sharing $m = 0, 1$, or 2 vertices:

$$\mathcal{F}_m = \{ \langle i, j; k, \ell \rangle \in \mathcal{E}^2 : |\{i, j\} \cap \{k, \ell\}| = m \}.$$

The pairs of edges in \mathcal{F}_0 share no vertices, so events on these edges are independent. It follows that $P(\delta(z_i, z_j) = 1 \wedge \delta(z_k, z_\ell) = 1) = 1/k^2$ and hence:

$$\sum_{\langle i, j; k, \ell \rangle \in \mathcal{F}_0} \sum_{\mathbf{z} \in \mathcal{Z}} \delta(z_i, z_j) \delta(z_k, z_\ell) = \#\mathcal{F}_0 k^n \frac{1}{k^2}.$$

The pairs of edges in \mathcal{F}_1 share a single pixel location in common. In this case:

$$\sum_{\langle i, j; k, \ell \rangle \in \mathcal{F}_1} \sum_{\mathbf{z} \in \mathcal{Z}} \delta(z_i, z_j) \delta(z_k, z_\ell) = \#\mathcal{F}_1 k^n \frac{1}{k^2},$$

since the shared vertex can take any of the k levels and then the two adjacent vertices share its value in $1/k^2$ of the terms. The pairs of edges in \mathcal{F}_2 share both pixels in common, therefore:

$$\sum_{\langle i, j; k, \ell \rangle \in \mathcal{F}_2} \sum_{\mathbf{z} \in \mathcal{Z}} \delta(z_i, z_j) \delta(z_k, z_\ell) = \sum_{i, j \in \mathcal{E}} \sum_{\mathbf{z} \in \mathcal{Z}} \delta(z_i, z_j)^2 = \#\mathcal{E} k^n \frac{1}{k},$$

as per (14). $\#\mathcal{F}_2 = \#\mathcal{E}$, so it follows that $\#\{\mathcal{F}_0 \cup \mathcal{F}_1\} = \#\mathcal{E}^2 - \#\mathcal{E}$, since $\mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2$ represent a complete partition of \mathcal{E}^2 . We can then rewrite (16) as:

$$\begin{aligned} \sum_{\mathbf{z} \in \mathcal{Z}} S(\mathbf{z})^2 &= \#\mathcal{F}_0 k^n \frac{1}{k^2} + \#\mathcal{F}_1 k^n \frac{1}{k^2} + \#\mathcal{E} k^n \frac{1}{k} \\ &= (\#\mathcal{E}^2 - \#\mathcal{E}) k^{n-2} + \#\mathcal{E} k^{n-1} = \#\mathcal{E}^2 k^{n-2} + \#\mathcal{E} k^{n-1} (1 - k^{-1}). \quad \square \end{aligned}$$

Corollary. *The expectation and variance of the k -state Potts model on a rectangular 2D lattice with inverse temperature $\beta = 0$ are $\mathbb{E}_{\mathbf{z}|\beta=0}[S(\mathbf{z})] = \#\mathcal{E}/k$ and $\mathbb{V}_{\mathbf{z}|\beta=0}[S(\mathbf{z})] = \mathcal{I}(0) = \#\mathcal{E}k^{-1}(1 - k^{-1})$, respectively.*

Proof. The proof follows from Lemma 1 by noting that $p(\mathbf{z}|\beta = 0) = k^{-n}$ and hence:

$$\begin{aligned} \mathbb{E}_{\mathbf{z}|\beta=0}[S(\mathbf{z})] &= \sum_{\mathbf{z} \in \mathcal{Z}} S(\mathbf{z}) p(\mathbf{z} | \beta = 0) \\ &= k^{n-1} \#\mathcal{E} k^{-n} = \frac{\#\mathcal{E}}{k}. \\ \mathbb{V}_{\mathbf{z}|\beta=0}[S(\mathbf{z})] &= \mathbb{E}_{\mathbf{z}|\beta=0}[S(\mathbf{z})^2] - \mathbb{E}_{\mathbf{z}|\beta=0}[S(\mathbf{z})]^2 \\ &= \sum_{\mathbf{z} \in \mathcal{Z}} S(\mathbf{z})^2 p(\mathbf{z} | \beta = 0) - \left(\frac{\#\mathcal{E}}{k} \right)^2 \\ &= \frac{\#\mathcal{E}^2 k^{n-2} + \#\mathcal{E} k^{n-1} (1 - k^{-1})}{k^n} - \frac{\#\mathcal{E}^2}{k^2} = \#\mathcal{E} k^{-1} (1 - k^{-1}). \quad \square \end{aligned}$$

The score function of the Potts model is the difference between the observed sufficient statistic and its expectation:

$$\frac{d}{d\beta} \log\{p(\mathbf{S}(\mathbf{z}) \mid \beta)\} = \mathbf{S}(\mathbf{z}) - \mathbb{E}_{\mathbf{z}|\beta}[\mathbf{S}(\mathbf{z})]. \quad (17)$$

The variance of the score function, also known as the Fisher information, is given by the expectation of the derivative of (13):

$$\mathcal{I}(\beta) = \mathbb{E}_{\mathbf{z}|\beta} \left[\left(\frac{d}{d\beta} \log\{p(\mathbf{S}(\mathbf{z}) \mid \beta)\} \right)^2 \right] \quad (18)$$

$$= \frac{d^2}{d\beta^2} \log\{\mathcal{C}(\beta)\}. \quad (19)$$

The Potts model undergoes a phase transition at the critical value of β , switching from a disordered to an ordered state. Potts (1952) showed that the critical value for a regular 2D lattice can be calculated exactly according to:


$$\beta_{crit} = \log \left\{ 1 + \sqrt{k} \right\}. \quad (20)$$

This is the location of the phase transition for a lattice with r rows and infinite columns, so for example the critical value for the images in Figure 2 is different to (20). However, the error introduced by a finite boundary diminishes as n increases, due to the finite-dimensional scaling property of the Potts model. Figure 3 shows that (20) is very accurate in predicting the behaviour of $\mathbf{S}(\mathbf{z})$ for a 2D image with a maximum value of $\mathbf{S}(\mathbf{z})$ for first-order neighbours of 1,998,000. This corresponds to the satellite image described in Section 2.2, with $k = 5$ mixture components and $\beta_{crit} \approx 1.174$. $\mathbf{S}(\mathbf{z})$ is approximated by simulation using the algorithm of Swendsen and Wang (1987).


Figure 2a illustrates how the location of the phase transition changes for different k . Baxter (1973) established that the score function of the Potts model is continuous and smoothly-varying for $k \leq 4$. This is known as a second-order phase transition, since there is a sharp change in the variance as shown in Figure 2c. This heteroskedasticity has important implications for the methods discussed in Section 4. For $k > 4$ there is a first-order phase transition at the critical point, shown by the discontinuity in Figures 2a, 2b, and 3a for $k = 5$. The difference between $\lim_{\beta \nearrow \beta_{crit}} \mathbb{E}_{\mathbf{z}|\beta}[\mathbf{S}(\mathbf{z})]$ and $\lim_{\beta \searrow \beta_{crit}} \mathbb{E}_{\mathbf{z}|\beta}[\mathbf{S}(\mathbf{z})]$ is known as the latent heat. This phase transition only emerges as n becomes large, so for example there is no discontinuity for $n = 2^6$ even when $k = 5$.

4 Bayesian indirect likelihood


In AEA and ABC, the conditional distribution $f(\mathbf{S}(\mathbf{w}) \mid \beta)$ of the auxiliary variables or pseudo-data \mathbf{w} is independent of the observed data \mathbf{y} and the labels \mathbf{z} . BIL involves constructing a suitable surrogate model $f_A(\mathbf{S}(\mathbf{w}) \mid \phi(\beta))$ to approximate this distribution, using a function $\phi(\beta)$ that we estimate by simulating pseudo-data at a fixed set




(a) Expectation for $n = 5^6$ and increasing values of $k \in \{2, 3, 4, 5\}$.



(b) Expectation for $k = 5$ and increasing values of $n \in \{3^6, 4^6, 5^6\}$.



(c) Standard deviation for $n = 5^6$ and increasing values of $k \in \{2, 3, 4, 5\}$.



(d) Standard deviation for $k = 5$ and increasing values of $n \in \{3^6, 4^6, 5^6\}$.

Figure 2: Approximate mean and standard deviation of $S(\mathbf{z})$ using Swendsen-Wang for increasing values of the inverse temperature β , the number of pixels n , and the number of unique labels k . The black dots are empirical estimates computed from random samples, while the fitted curves illustrate the parametric functional approximation introduced in Section 4.

of values $\{\beta_x\}_{x=1}^X$. This function can be reused across multiple datasets, amortising its computational cost. By replacing $S(\mathbf{w})$ with our surrogate model, we avoid the need to simulate auxiliary variables during model fitting. In Moores et al. (2015), it was shown that this approximation could lead to two orders of magnitude improvement in the elapsed runtime for fitting the hidden Potts model, while also improving the convergence properties of the original ABC-SMC algorithm. This is known as nonparametric




Figure 3: Approximate mean and standard deviation of $S(\mathbf{z})$ using Swendsen-Wang for a 2D image with $k = 5$ and $\#\mathcal{E} = 1,998,000$. The vertical, dashed line is the critical value of β . The horizontal, dashed line is $\#\mathcal{E}$. The black dots are empirical estimates computed from random samples, while the fitted curves illustrate the parametric functional approximation of Section 4.

Bayesian indirect likelihood with summary statistics, or nsBIL (Drovandi et al., 2015). Boland et al. (2017) introduced several variants of nsBIL for ABC-MCMC and proved some theoretical bounds on the approximation error.

In this section, we describe our PFAB algorithm in detail. We introduce a parametric function that takes advantage of the properties of the Potts model that were described in Section 3. We utilise this function to perform Bayesian indirect inference, as shown in Algorithm 1. We have implemented this algorithm, as well as AEA and nsBIL, using `RcppArmadillo` (Eddelbuettel and Sanderson, 2014). Our open-source R package for Windows, Linux or macOS is available from the CRAN repository (Moores and Mengersen, 2018).

Figures 2 and 3 illustrate that $\mathbb{E}_{\mathbf{z}|\beta}[S(\mathbf{z})]$ has the form of a sigmoid function, where the shape depends on the number of labels k and edges $\#\mathcal{E}$. The variance of the distribution is equal to the gradient of this curve, according to (13) and (19). Clearly, piecewise-constant approximations such as k NN (Sherlock et al., 2017) or random forests (Pudlo et al., 2016) would be unsuitable as surrogate models. Even the piecewise-linear model of Moores et al. (2015) can create issues for MCMC algorithms, due to sharp changes in the approximate variance. At the other end of the scale, GP approximations (Meeds and Welling, 2014) are too smooth to represent the discontinuity at the first-order phase transition boundary. Similarly, the generalised logistic function (Richards, 1959) is not flexible enough to be used for this purpose.

Figure 4b illustrates three different parametric approximations to the variance of the 2D Potts model with $k = 3$ and $n = 5^6$. The skew-Gaussian approximation is given

Algorithm 1 Parametric Functional Approximate Bayesian (PFAB) Algorithm

-
- 1: Generate $\mathbf{w}_x | \beta_x$ for sample points β_x , where $x = 1, \dots, X$
 - 2: Fit the parametric functions $\hat{\phi}_{\sigma^2}(\beta)$ & $\hat{\phi}_{\mu}(\beta)$
 - 3: **for all** iterations $t = 1, \dots, T$ **do**
 - 4: Update the labels $z_i \sim p(y_i | z_i, \mu_{z_i}, \sigma_{z_i}^2) p(z_i | z_{\setminus i}, \beta) \forall i \in \{1, \dots, n\}$
 - 5: Calculate sufficient statistics $\mathbf{S}(\mathbf{z})$ and $\bar{y}_j, s_j^2 \forall z_i = j, \forall j \in \{1, \dots, k\}$
 - 6: Update the noise parameters $\mu_j, \sigma_j^2 \sim p(\bar{y}_j, s_j^2 | \mu_j, \sigma_j^2) \pi(\mu_j | \sigma_j^2) \pi(\sigma_j^2)$
 - 7: Draw proposed parameter value $\beta^t \sim q(\beta^t | \beta_{t-1})$
 - 8: Approximate the Radon–Nikodým derivative:

$$\rho = \frac{q(\beta_{t-1} | \beta^t) \pi(\beta^t) f_A(\mathbf{S}(\mathbf{z}) | \hat{\phi}_{\mu}(\beta^t), \hat{\phi}_{\sigma^2}(\beta^t))}{q(\beta^t | \beta_{t-1}) \pi(\beta_{t-1}) f_A(\mathbf{S}(\mathbf{z}) | \hat{\phi}_{\mu}(\beta_{t-1}), \hat{\phi}_{\sigma^2}(\beta_{t-1}))} \quad (21)$$

- 9: Draw $u \sim \text{Uniform}[0, 1]$
 - 10: **if** $u < \min(1, \rho)$ **then**
 - 11: $\beta_t \leftarrow \beta^t$ **else** $\beta_t \leftarrow \beta_{t-1}$
 - 12: **end if**
 - 13: **end for**
-

by:

$$\hat{\phi}_{\sigma^2}(\beta) = \begin{cases} \mathbb{V}_0 + (\mathbb{V}_{max} - \mathbb{V}_0) e^{-\vartheta_1(\beta_{crit} - \beta)^2} & : 0 \leq \beta < \beta_{crit}, \\ \mathbb{V}_{max} e^{-\vartheta_2(\beta_{crit} - \beta)^2} & : \beta \geq \beta_{crit}. \end{cases} \quad (22)$$

The tails of the Gaussian decay much faster than the true model. The skew-Laplace approximation is slightly better:

$$\hat{\phi}_{\sigma^2}(\beta) = \begin{cases} \mathbb{V}_0 + (\mathbb{V}_{max} - \mathbb{V}_0) e^{-\vartheta_1(\beta_{crit} - \beta)} & : 0 \leq \beta < \beta_{crit}, \\ \mathbb{V}_{max} e^{-\vartheta_2(\beta - \beta_{crit})} & : \beta \geq \beta_{crit}. \end{cases} \quad (23)$$



(a) $\hat{\phi}_{\mu}(\beta)$ (b) $\hat{\phi}_{\sigma^2}(\beta)$

Figure 4: Parametric, functional approximations to $\mathbf{S}(\mathbf{w}) | \beta$ for $n = 5^6, k = 3$.

$\#\mathcal{E}$	ϑ_1	ϑ_2	$\mathbb{V}_1/\#\mathcal{E}$	$\mathbb{V}_2/\#\mathcal{E}$	$\mathbb{E}_{\beta_{crit}}/\#\mathcal{E}$
112	[5.08; 5.23]	[3.50; 3.59]	[2.75; 2.87]	...	[0.5800; 0.5843]
1404	[4.59; 4.70]	[4.22; 4.30]	[2.83; 2.99]	[4.55; 4.61]	[0.6281; 0.6321]
8064	[4.57; 4.63]	[6.16; 6.27]	[2.92; 2.99]	[5.72; 5.73]	[0.7104; 0.7115]
31000	[4.69; 4.72]	[6.46; 6.51]	[3.14; 3.18]	[5.63; 5.74]	[0.7304; 0.7324]
92880	[4.71; 4.73]	[6.29; 6.33]	[3.19; 3.21]	[4.83; 4.90]	[0.7573; 0.7584]
234612	[4.69; 4.70]	[6.55; 6.59]	[3.18; 3.20]	[5.02; 5.08]	[0.7626; 0.7633]
1998000	[4.72; 4.75]	[6.40; 6.43]	[3.25; 3.31]	[4.56; 4.60]	[0.7764; 0.7767]

Table 2: Posterior estimates for the parameters of the surrogate model for $k = 5$ and $n = 2^6, 3^6, 4^6, 5^6, 6^6, 7^6$, and 10^6 .

The approximation can be further improved by taking the square root of the distance between β and its critical value:

$$\hat{\phi}_{\sigma^2}(\beta) = \begin{cases} \mathbb{V}_0 + (\mathbb{V}_{max} - \mathbb{V}_0)e^{-\vartheta_1\sqrt{\beta_{crit}-\beta}} & : 0 \leq \beta < \beta_{crit}, \\ \mathbb{V}_{max}e^{-\vartheta_2\sqrt{\beta-\beta_{crit}}} & : \beta \geq \beta_{crit}. \end{cases} \quad (24)$$

Since the variance of the score function is equal to its derivative (19), it follows that the functional form for the expectation can be obtained as the integral curve of $\hat{\phi}_{\sigma^2}(\beta)$. Given the parametric function (24), the corresponding approximation for the expectation can be obtained as the algebraic solution of this integral (see Moores et al., 2018, Section A):

$$\hat{\phi}_{\mu}(\beta) = \begin{cases} \mathbb{E}_0 + \beta\mathbb{V}_0 + \int_0^\beta (\mathbb{V}_{max} - \mathbb{V}_0)e^{-\vartheta_1\sqrt{\beta_{crit}-t}} dt & : 0 \leq \beta < \beta_{crit}, \\ \mathbb{E}_{\beta_{crit}} + \int_{\beta_{crit}}^\beta \mathbb{V}_{max}e^{-\vartheta_2\sqrt{t-\beta_{crit}}} dt & : \beta \geq \beta_{crit}. \end{cases} \quad (25)$$

Figure 4a illustrates $\hat{\phi}_{\mu}(\beta)$ for a simulated dataset. The initial conditions \mathbb{E}_0 and \mathbb{V}_0 can be calculated exactly for a rectangular lattice, according to Lemma 1. Pickard (1987) gives a formula for calculating \mathbb{V}_{max} , the variance at the critical point, in the asymptotic limit:

$$\lim_{\#\mathcal{E} \rightarrow \infty} \frac{\mathbf{S}(\mathbf{z}) - \mathbb{E}_{\mathbf{z}|\beta_{crit}}[\mathbf{S}(\mathbf{z})]}{\sqrt{\#\mathcal{E} \log\{\#\mathcal{E}\}}} \xrightarrow{d} \mathcal{N}\left(0, \frac{2}{\pi}\right). \quad (26)$$

We have found this formula useful as an upper bound, but a single parameter \mathbb{V}_{max} does not produce a satisfactory fit to the simulated pseudo-data when $k > 4$. Instead, we split it into two separate parameters, $\mathbb{V}_1 = \lim_{\beta \nearrow \beta_{crit}} \mathbb{V}_{\mathbf{z}|\beta}[\mathbf{S}(\mathbf{z})]$ and $\mathbb{V}_2 = \lim_{\beta \searrow \beta_{crit}} \mathbb{V}_{\mathbf{z}|\beta}[\mathbf{S}(\mathbf{z})]$. Due to the discontinuity at the first-order phase transition, we also treat $\mathbb{E}_{\beta_{crit}}$ as a free parameter. This leaves five parameters to estimate, $\Theta = \{\vartheta_1, \vartheta_2, \mathbb{V}_1, \mathbb{V}_2, \mathbb{E}_{\beta_{crit}}\}$, for any given values of $\#\mathcal{E}$ and $k > 4$, as shown in Table 2.

At Step 1 of Algorithm 1 we select a set of sample points $\{\beta_x\}_{x=1}^X$. We know from the Fisher information (18) that we need more points close to β_{crit} , since that is where the variance is at its maximum. We selected $X = 36$ points for the experiments in Section 5, since we are able to generate pseudo-data independently, in parallel on a 36-core Intel Xeon server. This computation could easily be distributed across multiple servers if necessary, or run on a general-purpose graphical processing unit (GPU). However, we

k	ϑ_1	ϑ_2	$\mathbb{V}_{max}/\#\mathcal{E}$
2	[5.35; 5.38]	[5.24; 5.26]	[2.17; 2.18]
3	[5.38; 5.39]	[5.68; 5.70]	[3.60; 3.62]
4	[5.61; 5.63]	[6.10; 6.12]	[4.83; 4.86]

Table 3: Posterior estimates for the parameters of the surrogate model for $n = 5^6$ and $k = 2, 3$, and 4.

found that 36 points were sufficient to fit our parametric surrogate model. For each point, we ran 500 iterations of the Swendsen-Wang algorithm, discarding the first 125 as burn-in. When all of these parallel simulations are combined, we obtain a 36×375 matrix of $S(\mathbf{w})$ values, conditional on β_x , $\#\mathcal{E}$, and k .

At Step 2, we use Stan (Carpenter et al., 2017) to sample from the posterior distribution $\pi(\Theta | S(\mathbf{w}), \beta_x)$, treating our matrix of simulations as observed data. Following the central limit theorem established by Pickard (1987), we approximate the distribution $f(S(\mathbf{w}) | \beta_x)$ using a truncated Gaussian:

$$f_A(S(\mathbf{w}) | \hat{\phi}_\mu(\beta_x), \hat{\phi}_{\sigma^2}(\beta_x)) \sim \mathcal{N}(\hat{\phi}_\mu(\beta_x), \hat{\phi}_{\sigma^2}(\beta_x)) \mathbb{1}_{(0 \leq S(\mathbf{w}) \leq \#\mathcal{E})}. \quad (27)$$

Since $S(\mathbf{w})$ is the count of like neighbours, this continuous model is not suitable for very small images, such as when $n = 2^6$ ($\#\mathcal{E} = 112$). There are also problems for very large β (e.g. $\beta > 1.5$), as the function approaches its horizontal asymptote. However, the fit rapidly improves as n increases.

The 95% highest posterior density (HPD) intervals for Θ when $k = 5$ and using 7 different values of $\#\mathcal{E}$ for each of our simulated datasets are shown in Table 2. Ignoring the outlier in the first row, which is too small for asymptotic properties to apply, there are clear trends in the parameter values as the size of the lattice increases. It might be possible to interpolate between these parameter values for other n , or even to extrapolate to larger images. Table 3 shows HPD intervals for the surrogate model when $k \leq 4$. There are only 3 free parameters to estimate here, since there is no first-order phase transition.

In most cases, the number of iterations T required for the primary chain will depend on the mixing time of the checkerboard Gibbs sampler that is used in Step 4. This depends on both the size of the state vector \mathbf{z} and the value of β . The Gibbs sampler is not guaranteed to be ergodic for $\beta > \beta_{crit}$, so this is an important consideration when analysing real image data. Similarly, the heuristic number of 500 iterations advocated by Cucala et al. (2009); Everitt (2012) for simulating \mathbf{w} in AEA will be insufficient for large images when $\beta > \beta_{crit}$. Liang (2010) showed that convergence of the auxiliary chain could be improved by initialising \mathbf{w} at the current value of \mathbf{z} . Using this warm start, the number of required iterations depends on the autocorrelation time of the Markov chain, which can be an order of magnitude less than the duration of the transient phase or burn-in period. This can be combined with tempered transitions (Neal, 2005) to further improve the acceptance rate. Our approach is to instead use the Swendsen-Wang algorithm to simulate \mathbf{w} because this has better mixing properties for large values of β (Cooper and Frieze, 1999).

We use conjugate priors for the parameters of the observation equation (4), so σ_j^2 and μ_j can be updated by Gibbs sampling from their conditional posterior distributions at Step 6. This requires calculation of the sufficient statistics for each label: $n_j = \sum_{i=1}^n \delta(z_i, j)$; $\bar{y}_j = n_j^{-1} \sum_{\{i:z_i=j\}} y_i$; and $s_j^2 = n_j^{-1} \sum_{\{i:z_i=j\}} (y_i - \bar{y}_j)^2$. The inverse temperature is sampled using an approximate Metropolis-within-Gibbs step. We use adaptive, Gaussian random walk proposals for $q(\beta'|\beta_{t-1}) \sim \mathcal{N}(\beta_{t-1}, \sigma_{MH}^2)$, where the step size σ_{MH}^2 is tuned adaptively using a Robbins–Monro recursion (Andrieu and Thoms, 2008; Garthwaite et al., 2015). The surrogate model is used in place of the intractable likelihood to approximate the Metropolis–Hastings ratio (21).

5 Experimental results

5.1 Simulation study

We generated synthetic data using the R package `PottsUtils` (Feng, 2008; Feng and Tierney, 2011) as explained in Section 2.1. We used informative priors for the $j \in \{1, \dots, 5\}$ mixture components (8) and a uniform prior on the interval $[0, 1.2\beta_{crit}]$ for the inverse temperature. β_{crit} was calculated using Equation (20) to be approximately equal to 1.174. We chose an upper bound of $1.2\beta_{crit} \approx 1.41$ because we found that simulated datasets with β above that value tended to have pixels with fewer than 5 unique labels, so that the effective value of k was different than what was originally specified. This was a particular problem with smaller images, where all of the pixels could be assigned the same label. It is impossible to estimate μ_j or σ_j^2 accurately if there are no pixels with the label j . The maximum value of β also affects the runtime, since the mixing of the Gibbs sampler slows dramatically for $\beta > \beta_{crit}$.

The precomputation step, running Swendsen-Wang for 500 iterations with 36 different values of β , took less than 5 seconds for $n = 8^2$, 27^2 , 64^2 , or 125^2 . This increased slightly to 7.6 seconds for $n = 216^2$, 19.3 seconds for $n = 343^2$, and 169.8 seconds for $n = 1000^2$. We used 4 parallel chains with 10,000 iterations of the no-U-turn sampler (NUTS) per chain to fit the parametric surrogate model in Stan. Half of the iterations were discarded as burn-in. This took 5.5 minutes on a 2.8GHz Intel Core *i7* laptop. The computational cost of fitting the surrogate model only depends on the number of iterations of Swendsen-Wang and NUTS, as well as the number of training points X , not on the sizes of the images involved. The parameters of the surrogate models are shown in tables 2 and 3.

The PFAB algorithm as well as nsBIL accelerated ABC-MCMC (Moores et al., 2015; Boland et al., 2017) were run for 20,000 iterations on each image, discarding the first 10,000 as burn-in. Due to the computational cost of AEA, it was run for only 5,000 iterations on the medium-sized images and 2,000 iterations on the images with 1 million pixels. The differences in elapsed runtime are illustrated by Figure 5d. The precomputation step for both PFAB and nsBIL-MCMC has reduced the elapsed runtime for images with $n = 343^2$ pixels from 5 hours for 5,000 iterations to only 15-23 minutes for 20,000 iterations. The runtime for $n = 1000^2$ pixels improved from 9-12 hours for 2,000 iterations of AEA to 2.5 hours for 20,000 iterations of PFAB or nsBIL. AEA

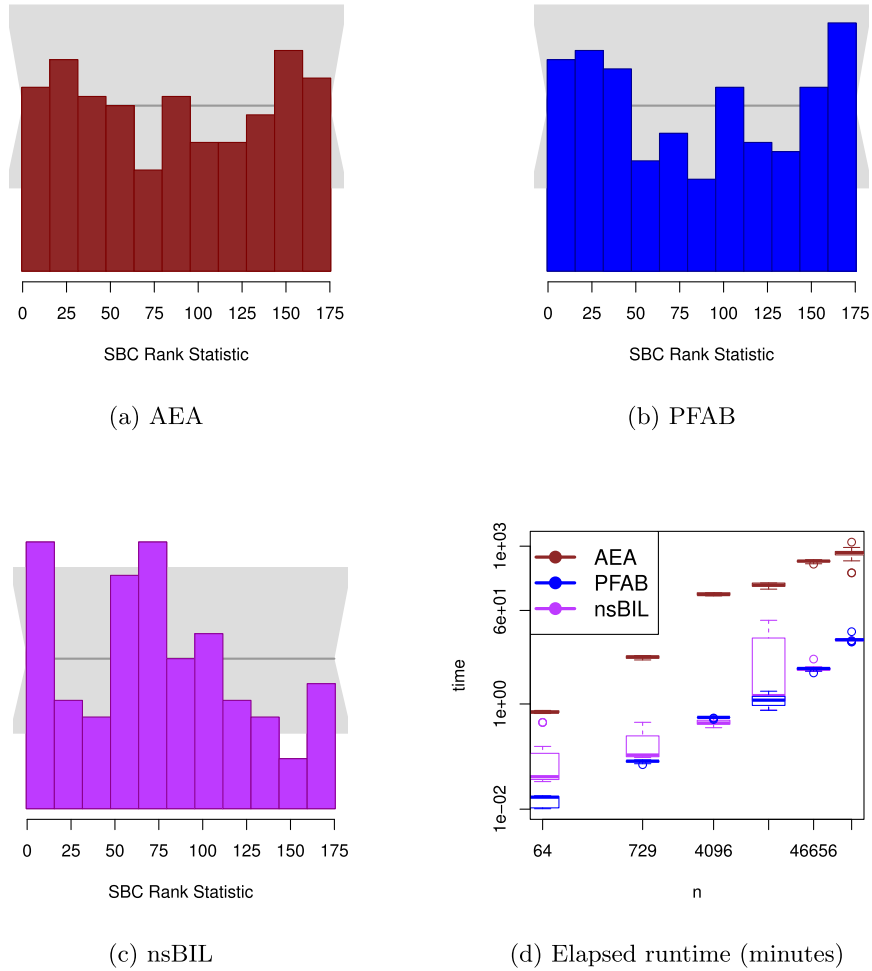


Figure 5: Simulation-based calibration: histograms for the rank statistic of the true parameter value of β under the posterior distribution obtained from (a) the approximate exchange algorithm; (b) our parametric functional approximate Bayesian algorithm; and (c) nonparametric Bayesian indirect likelihood with sufficient statistics. The dark grey, horizontal line shows the expected value and the light grey wedges show the 99% confidence interval. The elapsed runtimes of the 3 algorithms are shown in (d).

takes less than 1 minute to run when $n = 2^6$, which is less time than it takes to fit the surrogate model in Stan. Clearly, our PFAB method is not suited for small images such as these.

Histograms of the rank statistics computed using SBC (Talts et al., 2018) are shown in Figure 5. The variations in these histograms of both AEA and PFAB are within

the expectations of uniformity shown in the grey bars, which suggests that the fits are faithfully representing the true posterior distributions. There will naturally be some bias introduced by the surrogate model in PFAB, as well as by using Swendsen-Wang instead of perfect sampling in AEA. However, this bias is too minor to be detected by the present simulation study. On the other hand, nsBIL shows clear signs of bias towards overestimating β . This is due to a combination of the ABC tolerance ϵ , as well as problems with convergence of the ABC-MCMC algorithm (Lee and Łatuszyński, 2014). We previously found that nsBIL performed much better when incorporated into an ABC-SMC algorithm (Moore et al., 2015), although this increases its computational cost.


5.2 Satellite remote sensing

For the 18 satellite images described in Section 2.2 we used the same priors for μ_j (8), σ_j^2 (9), and β as for the simulation study. Since these images have the same dimensions (1000 × 1000 pixels) as the largest synthetic images, we could reuse the same surrogate models for PFAB and nsBIL. The reusability of surrogate models for many datasets is a major advantage of our approach.

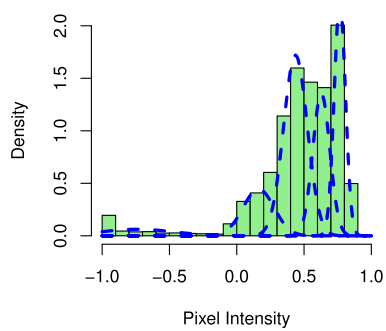
We were unable to evaluate accuracy for these images because the true values of the inverse temperature and the pixel labels are unknown. However, we can still compare the posterior distributions for β obtained from the three algorithms, as well as the time taken to fit the model. These results are summarised in Figure 6. We ran PFAB and nsBIL for 20,000 iterations on each image, while AEA was run for only 2,000 iterations, with 200 iterations of Swendsen-Wang for the auxiliary variables. AEA took 19 hours per image, while PFAB and nsBIL took 1.6 hours. This resulted in an average ESS per hour of 1.88 for AEA, 598 for PFAB, and 692 for nsBIL.

An example segmentation of one of the satellite images is shown in Figure 6a. There is close correspondence between each of the 5 labels and categories of land use around Brisbane. The Brisbane River, Tingalpa Reservoir (in the south-east), and Enoggera Reservoir (in the north-west) are labelled in blue, while urbanised and industrial areas such as the central business district are labelled in black and the suburbs are labelled in grey. Large, forested areas, such as Mount Cotton and Mount Coot-tha, are labelled in green, as well as smaller features such as the City Botanic Gardens and New Farm Park. The histogram of pixel values with corresponding mixture model is shown in Figure 6b.

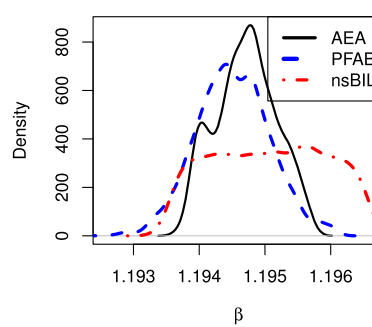
Figure 6c illustrates the posterior density plot for β obtained using the three algorithms. The posteriors from AEA and PFAB were generally in close agreement for the 18 images. As expected, the posterior from nsBIL is over-dispersed and uniformly-distributed over the interval $\pm \epsilon$. In one case, the nsBIL-MCMC algorithm failed to converge to the same value, most likely due to the issues with local modes that were observed in the simulation study.



(a) Image segmentation using 5 labels: water (blue); urban/industrial (black); suburbs (grey); parkland and forest (light/dark green)



(b) Mixture of Gaussians



(c) Posteriors for β

Figure 6: Results for the satellite image of Brisbane, Australia.

6 Concluding remarks

MCMC algorithms that involve simulating auxiliary variables at every iteration are too computationally-expensive for applications in image analysis. Exact inference is infeasible for the scale of data that is commonly encountered in scientific studies, such as satellite remote sensing. Methods such as Russian roulette (Lyne et al., 2015), lazy ABC (Prangle, 2016), or delayed acceptance (Sherlock et al., 2017) can reduce the number of auxiliary iterations that need to be performed, but not enough to compensate for the two orders of magnitude difference in runtime that we observe.

To address this problem, we have introduced a parametric functional approximate Bayesian (PFAB) algorithm. This algorithm incorporates known properties of the Potts model, such as the first-order phase transition for $k > 4$, to obtain a close approximation to the true density function. We fit our surrogate model using a precomputation step, which can be run independently in parallel for selected parameter values. The same surrogate model can be reused across multiple datasets, amortising its computational cost. The resulting posterior distribution for β has been shown to cover the true parameter value in simulation studies. When applied to real data, there is close agreement between the posterior densities obtained from PFAB and the exchange algorithm. Our approximation results in a hundred times speedup, making it feasible to use MCMC for images of realistic size, given reasonable computational power by contemporary standards. PFAB could also be used to accelerate other spatial Markov models, not only for image analysis.

The score function of the Potts model is continuous with respect to n and k , with two notable exceptions. The first is that the phase transition only emerges for sufficiently large n , for example $n \geq 3^6$ when $k = 5$. Secondly, the nature of the phase transition is different for $k \leq 4$ and $k > 4$. This is reflected in the parameters of our surrogate model, creating the possibility of interpolating between surrogate models for a new value of n . PFAB might be used in combination with trans-dimensional algorithms or other methods of inference when k is unknown. Other methods for rescaling the Potts model, such as Cucala and Marin (2013), might also be applied in this context.

For the purpose of comparison, we have used random walk proposals for β with all three algorithms: AEA, PFAB, and nsBIL. However, the parametric surrogate model could be incorporated in an approximate Hamiltonian or Langevin algorithm (Strathmann et al., 2015; Zhang et al., 2017) to improve the efficiency of the MCMC sampler. Our surrogate model also provides an approximation to the Fisher information (18). This could be used in Bayesian optimisation to select design points for iterative refinement of the approximation, in a similar approach to Gutmann and Corander (2016) or Ryan et al. (2016). It could also be used to derive an approximate Jeffreys' prior for β .

Supplementary Material

R package (DOI: [10.1214/18-BA1130SUPPA](https://doi.org/10.1214/18-BA1130SUPPA); .zip). R source package containing code to perform image segmentation using Markov chain Monte Carlo. Includes C++ implementations of all of the methods described in the article: PFAB, ABC, and AEA.

Further Results (DOI: [10.1214/18-BA1130SUPPB](https://doi.org/10.1214/18-BA1130SUPPB); .pdf). Algebraic solution of the integral equation (25), as well as additional figures for the results of the simulation study of Section 5.1 and the satellite imagery of Section 5.2.

References

- Alston, C. L., Mengersen, K. L., Robert, C. P., Thompson, J. M., Littlefield, P. J., Perry, D., and Ball, A. J. (2007). “Bayesian mixture models in a longitudinal setting for analysing sheep CAT scan images.” *Computational Statistics & Data Analysis*, 51: 4282–4296. MR2364445. doi: <https://doi.org/10.1016/j.csda.2006.05.013>. 7
- Andrieu, C. and Thoms, J. (2008). “A tutorial on adaptive MCMC.” *Statistics and Computing*, 18(4): 343–373. MR2461882. doi: <https://doi.org/10.1007/s11222-008-9110-y>. 16
- Baxter, R. J. (1973). “Potts model at the critical temperature.” *Journal of Physics. C. Solid State Physics*, 6(23): L445. 10
- Besag, J. (1974). “Spatial interaction and the statistical analysis of lattice systems.” *Journal of the Royal Statistical Society. Series B. Methodological*, 36(2): 192–236. MR0373208. 1
- Boland, A., Friel, N., and Maire, F. (2017). “Efficient MCMC for Gibbs Random Fields using pre-computation.” *arXiv preprint arXiv:1710.04093 [stat.CO]*. URL <https://arxiv.org/abs/1710.04093>. 3, 12, 16
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., and Riddell, A. (2017). “Stan: A Probabilistic Programming Language.” *Journal of Statistical Software*, 76(1): 1–32. doi: <https://doi.org/10.18637/jss.v076.i01>. 15
- Christen, J. A. and Fox, C. (2005). “Markov chain Monte Carlo Using an Approximation.” *Journal of Computational and Graphical Statistics*, 14(4): 795–810. MR2211367. doi: <https://doi.org/10.1198/106186005X76983>. 3
- Conrad, P. R., Marzouk, Y. M., Pillai, N. S., and Smith, A. (2016). “Accelerating Asymptotically Exact MCMC for Computationally Intensive Models via Local Approximations.” *Journal of the American Statistical Association*, 111(516): 1591–1607. MR3601720. doi: <https://doi.org/10.1080/01621459.2015.1096787>. 3
- Cook, S. R., Gelman, A., and Rubin, D. B. (2006). “Validation of Software for Bayesian Models Using Posterior Quantiles.” *Journal of Computational and Graphical Statistics*, 15(3): 675–692. MR2291268. doi: <https://doi.org/10.1198/106186006X136976>. 4
- Cooper, C. and Frieze, A. M. (1999). “Mixing properties of the Swendsen-Wang process on classes of graphs.” *Random Structures & Algorithms*, 15(3–4): 242–261. MR1716764. doi: [https://doi.org/10.1002/\(SICI\)1098-2418\(199910/12\)15:3/4<242::AID-RSA4>3.3.CO;2-3](https://doi.org/10.1002/(SICI)1098-2418(199910/12)15:3/4<242::AID-RSA4>3.3.CO;2-3). 15

- Cucala, L. and Marin, J.-M. (2013). “Bayesian Inference on a Mixture Model With Spatial Dependence.” *Journal of Computational and Graphical Statistics*, 22(3): 584–597. MR3173732. doi: <https://doi.org/10.1080/10618600.2013.805652>. 20
- Cucala, L., Marin, J.-M., Robert, C. P., and Titterton, D. M. (2009). “A Bayesian Reassessment of Nearest-Neighbor Classification.” *Journal of the American Statistical Association*, 104(485): 263–273. MR2663042. doi: <https://doi.org/10.1198/jasa.2009.0125>. 2, 15
- Drovandi, C. C., Moores, M. T., and Boys, R. J. (2018). “Accelerating pseudo-marginal MCMC using Gaussian processes.” *Computational Statistics & Data Analysis*, 118: 1–17. MR3715260. doi: <https://doi.org/10.1016/j.csda.2017.09.002>. 3
- Drovandi, C. C., Pettitt, A. N., and Faddy, M. J. (2011). “Approximate Bayesian computation using indirect inference.” *Journal of the Royal Statistical Society. Series C, Applied Statistics*, 60(3): 317–337. MR2767849. doi: <https://doi.org/10.1111/j.1467-9876.2010.00747.x>. 3
- Drovandi, C. C., Pettitt, A. N., and Lee, A. (2015). “Bayesian indirect inference using a parametric auxiliary model.” *Statistical Science*, 30(1): 72–95. MR3317755. doi: <https://doi.org/10.1214/14-STS498>. 3, 12
- Eddelbuettel, D. and Sanderson, C. (2014). “RcppArmadillo: Accelerating R with high-performance C++ linear algebra.” *Computational Statistics & Data Analysis*, 71: 1054–63. MR3132026. doi: <https://doi.org/10.1016/j.csda.2013.02.005>. 12
- Everitt, R. G. (2012). “Bayesian Parameter Estimation for Latent Markov Random Fields and Social Networks.” *Journal of Computational and Graphical Statistics*, 21(4): 940–960. MR3005805. doi: <https://doi.org/10.1080/10618600.2012.687493>. 2, 7, 15
- Feng, D. (2008). “Bayesian hidden Markov normal mixture models with application to MRI tissue classification.” Ph.D. thesis, University of Iowa. MR2712000. 16
- Feng, D. and Tierney, L. (2011). *PottsUtils: Utility Functions of the Potts Models*. R package version 0.2-2. URL <http://CRAN.R-project.org/package=PottsUtils> 16
- Flood, N. (2014). “Continuity of Reflectance Data between Landsat-7 ETM+ and Landsat-8 OLI, for Both Top-of-Atmosphere and Surface Reflectance: A Study in the Australian Landscape.” *Remote Sensing*, 6(9): 7952–7970. doi: <https://doi.org/10.3390/rs6097952>. 5
- Friel, N. and Rue, H. (2007). “Recursive computing and simulation-free inference for general factorizable models.” *Biometrika*, 94(3): 661–672. MR2410015. doi: <https://doi.org/10.1093/biomet/asm052>. 2
- Garthwaite, P. H., Fan, Y., and Sisson, S. A. (2015). “Adaptive optimal scaling of Metropolis-Hastings algorithms using the Robbins-Monro process.” *Communications in Statistics. Theory and Methods*, 45(17): 5098–5111. MR3520297. doi: <https://doi.org/10.1080/03610926.2014.936562>. 16

- Gelman, A. (2017). “Correction to Cook, Gelman, and Rubin (2006).” *Journal of Computational and Graphical Statistics*, 26: 940. MR3765358. doi: <https://doi.org/10.1080/10618600.2017.1377082>. 4
- Geman, S. and Geman, D. (1984). “Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6: 721–41. 7
- Geweke, J. (2004). “Getting it Right: Joint Distribution Tests of Posterior Simulators.” *Journal of the American Statistical Association*, 99: 799–804. MR2090912. doi: <https://doi.org/10.1198/016214504000001132>. 4
- Green, P. J. and Richardson, S. (2002). “Hidden Markov models and disease mapping.” *Journal of the American Statistical Association*, 97: 1055–1070. MR1951259. doi: <https://doi.org/10.1198/016214502388618870>. 7
- Grelaud, A., Robert, C. P., Marin, J.-M., Rodolphe, F., and Taly, J.-F. (2009). “ABC likelihood-free methods for model choice in Gibbs random fields.” *Bayesian Analysis*, 4(2): 317–336. MR2507366. doi: <https://doi.org/10.1214/09-BA412>. 2, 7
- Gutmann, M. U. and Corander, J. (2016). “Bayesian Optimization for Likelihood-free Inference of Simulator-based Statistical Models.” *Journal of Machine Learning Research*, 17(1): 4256–4302. MR3555016. 20
- Henderson, V., Storeygard, A., and Weil, D. N. (2011). “A Bright Idea for Measuring Economic Growth.” *The American Economic Review*, 101(3): 194–199. doi: <https://doi.org/10.1257/aer.101.3.194>. 5
- Huang, K. (2010). *Introduction to Statistical Physics*. Boca Raton: Chapman & Hall/CRC Press, 2nd edition. 1
- Huber, M. L. (2003). “A bounding chain for Swendsen-Wang.” *Random Structures & Algorithms*, 22(1): 43–59. MR1943858. doi: <https://doi.org/10.1002/rsa.10071>. 2
- Huber, M. L. (2016). *Perfect Simulation*, volume 148 of *Monographs on Statistics & Applied Probability*. Boca Raton, FL: Chapman & Hall/CRC Press. MR3443710. 2
- Järvenpää, M., Gutmann, M., Vehtari, A., and Marttinen, P. (2018). “Gaussian process modeling in approximate Bayesian computation to estimate horizontal gene transfer in bacteria.” *The Annals of Applied Statistics*, 12(4): 2228–2251. MR3875699. doi: <https://doi.org/10.1214/18-AOAS1150>. 3
- Lee, A. and Latuszyński, K. (2014). “Variance bounding and geometric ergodicity of Markov chain Monte Carlo kernels for approximate Bayesian computation.” *Biometrika*, 101(3): 655–671. MR3254907. doi: <https://doi.org/10.1093/biomet/asu027>. 18
- Li, S. Z. (2009). *Markov Random Field Modeling in Image Analysis*. Dordrecht: Springer, 3rd edition. MR2493908. 1
- Liang, F. (2010). “A double Metropolis Hastings sampler for spatial models with intractable normalizing constants.” *Journal of Statistical Computation*

- and Simulation*, 80(9): 1007–1022. MR2742519. doi: <https://doi.org/10.1080/00949650902882162>. 15
- Liang, F., Jin, I. H., Song, Q., and Liu, J. S. (2016). “An Adaptive Exchange Algorithm for Sampling from Distributions with Intractable Normalizing Constants.” *Journal of the American Statistical Association*, 111(513): 377–393. MR3494666. doi: <https://doi.org/10.1080/01621459.2015.1009072>. 3
- Lyne, A.-M., Girolami, M., Atchadé, Y., Strathmann, H., and Simpson, D. (2015). “On Russian Roulette Estimates for Bayesian Inference with Doubly-Intractable Likelihoods.” *Statistical Science*, 30(4): 443–467. MR3432836. doi: <https://doi.org/10.1214/15-STS523>. 3, 7, 20
- McClain, C. R. (2009). “A Decade of Satellite Ocean Color Observations.” *Annual Review of Marine Science*, 1: 19–42. doi: <https://doi.org/10.1146/annurev.marine.010908.163650>. 5
- McGrory, C. A., Pettitt, A. N., Reeves, R., Griffin, M., and Dwyer, M. (2012). “Variational Bayes and the Reduced Dependence Approximation for the Autologistic Model on an Irregular Grid With Applications.” *Journal of Computational and Graphical Statistics*, 21(3): 781–796. MR2970919. doi: <https://doi.org/10.1080/10618600.2012.632232>. 7
- McGrory, C. A., Titterton, D., Reeves, R., and Pettitt, A. N. (2009). “Variational Bayes for estimating the parameters of a hidden Potts model.” *Statistics and Computing*, 19(3): 329–340. MR2516223. doi: <https://doi.org/10.1007/s11222-008-9095-6>. 2
- Meeds, E. and Welling, M. (2014). “GPS-ABC: Gaussian Process Surrogate Approximate Bayesian Computation.” In *Proc. 30th Conf. UAI*, 593–602. Quebec City, Canada: AUAI Press. 3, 12
- Minvielle, P., Doucet, A., Marrs, A., and Maskell, S. (2010). “A Bayesian approach to joint tracking and identification of geometric shapes in video sequences.” *Image and Vision Computing*, 28(1): 111–123. doi: <https://doi.org/10.1016/j.imavis.2009.05.002>. 3
- Mira, A., Møller, J., and Roberts, G. O. (2001). “Perfect slice samplers.” *Journal of the Royal Statistical Society. Series B. Methodological*, 63(3): 593–606. MR1858405. doi: <https://doi.org/10.1111/1467-9868.00301>. 2
- Møller, J., Pettitt, A. N., Reeves, R., and Berthelsen, K. K. (2006). “An Efficient Markov Chain Monte Carlo Method for Distributions with Intractable Normalising Constants.” *Biometrika*, 93(2): 451–458. MR2278096. doi: <https://doi.org/10.1093/biomet/93.2.451>. 2, 7
- Monahan, J. F. and Boos, D. D. (1992). “Proper Likelihoods for Bayesian Analysis.” *Biometrika*, 79(2): 271–278. MR1185129. doi: <https://doi.org/10.1093/biomet/79.2.271>. 4
- Moores, M. T., Drovandi, C. C., Mengersen, K., and Robert, C. P. (2015). “Pre-processing for approximate Bayesian computation in image analysis.” *Statist-*

- tics and Computing*, 25(1): 23–33. MR3304900. doi: <https://doi.org/10.1007/s11222-014-9525-6>. 2, 3, 11, 12, 16, 18
- Moores, M. T. and Mengersen, K. (2014). “Bayesian approaches to spatial inference: modelling and computational challenges and solutions.” *AIP Conference Proceedings*, 1636: 112–117. doi: <https://doi.org/10.1063/1.4903719>. 2
- Moores, M. T. and Mengersen, K. (2018). “bayesImageS: Bayesian methods for image segmentation using a Potts model.” R package version 0.5-3. URL <https://CRAN.R-project.org/package=bayesImageS> 12
- Moores, M. T., Nicholls, G. K., Pettitt, A. N., and Mengersen, K. (2018). Supplements to “Scalable Bayesian Inference for the Inverse Temperature of a Hidden Potts Model.” doi: <https://doi.org/10.1214/18-BA1130SUPPA>, <https://doi.org/10.1214/18-BA1130SUPPB>. 14
- Murray, I., Ghahramani, Z., and MacKay, D. J. C. (2006). “MCMC for Doubly-intractable Distributions.” In *Proc. 22nd Conf. UAI*, 359–366. Arlington, VA: AUAI Press. 2
- NASA (2011). “Landsat 7 Science Data Users Handbook.” Technical report, National Aeronautics and Space Administration, Greenbelt, MD. URL <http://landsathandbook.gsfc.nasa.gov/> 3
- Neal, R. M. (2005). “Taking Bigger Metropolis Steps by Dragging Fast Variables.” arXiv preprint *arXiv:math/0502099 [math.ST]*. URL <https://arxiv.org/abs/math/0502099> 15
- Pickard, D. K. (1987). “Inference for Discrete Markov Fields: The Simplest Nontrivial Case.” *Journal of the American Statistical Association*, 82(397): 90–96. MR0883337. doi: <https://doi.org/10.1080/01621459.1987.10478394>. 14, 15
- Potts, R. B. (1952). “Some generalized order-disorder transformations.” *Mathematical Proceedings of the Cambridge Philosophical Society*, 48: 106–9. MR0047571. doi: <https://doi.org/10.1017/S0305004100027419>. 1, 10
- Prangle, D. (2016). “Lazy ABC.” *Statistics and Computing*, 26(1): 171–185. MR3439366. doi: <https://doi.org/10.1007/s11222-014-9544-3>. 3, 20
- Prangle, D., Blum, M. G. B., Popovic, G., and Sisson, S. A. (2014). “Diagnostic tools for approximate Bayesian computation using the coverage property.” *Australian & New Zealand Journal of Statistics*, 56(4): 309–329. MR3300163. doi: <https://doi.org/10.1111/anzs.12087>. 4
- Propp, J. G. and Wilson, D. B. (1996). “Exact sampling with coupled Markov chains and applications to statistical mechanics.” *Random Structures & Algorithms*, 9(1–2): 223–252. MR1611693. doi: [https://doi.org/10.1002/\(SICI\)1098-2418\(199608/09\)9:1/2<223::AID-RSA14>3.3.CO;2-R](https://doi.org/10.1002/(SICI)1098-2418(199608/09)9:1/2<223::AID-RSA14>3.3.CO;2-R). 2
- Pudlo, P., Marin, J.-M., Estoup, A., Cornuet, J.-M., Gautier, M., and Robert, C. P. (2016). “Reliable ABC model choice via random forests.” *Bioinformatics*, 32(6): 859–866. MR3304913. doi: <https://doi.org/10.1007/s11222-014-9530-9>. 12

- Reeves, R. and Pettitt, A. N. (2004). “Efficient Recursions for General Factorisable Models.” *Biometrika*, 91(3): 751–757. MR2090636. doi: <https://doi.org/10.1093/biomet/91.3.751>. 2
- Richards, F. J. (1959). “A Flexible Growth Function for Empirical Use.” *Journal of Experimental Botany*, 10(2): 290–301. 12
- Roy, D. P., Kovalsky, V., Zhang, H. K., Vermote, E. F., Yan, L., Kumar, S. S., and Egorov, A. (2016). “Characterization of Landsat 7 to Landsat 8 reflective wavelength and normalized difference vegetation index continuity.” *Remote Sensing of Environment*, 185: 57–70. doi: <https://doi.org/10.1016/j.rse.2015.12.024>. 5
- Ryan, C. M., Drovandi, C. C., and Pettitt, A. N. (2016). “Optimal Bayesian Experimental Design for Models with Intractable Likelihoods Using Indirect Inference Applied to Biological Process Models.” *Bayesian Analysis*, 11(3): 857–883. MR3543911. doi: <https://doi.org/10.1214/15-BA977>. 20
- Rydén, T. and Titterton, D. M. (1998). “Computational Bayesian Analysis of Hidden Markov Models.” *Journal of Computational and Graphical Statistics*, 7(2): 194–211. MR1649362. doi: <https://doi.org/10.2307/1390813>. 7
- Sherlock, C., Golightly, A., and Henderson, D. A. (2017). “Adaptive, delayed-acceptance MCMC for targets with expensive likelihoods.” *Journal of Computational and Graphical Statistics*, 26(2): 434–444. MR3640199. doi: <https://doi.org/10.1080/10618600.2016.1231064>. 3, 12, 20
- Simioncelli, E. P. (1999). “Bayesian Multi-Scale Differential Optical Flow.” In Jähne, B., Haussecker, H., and Geissler, P. (eds.), *Handbook of computer vision and applications*, volume 2, chapter 14, 397–422. San Diego: Academic Press. 3
- Small, C. (2001). “Estimation of urban vegetation abundance by spectral mixture analysis.” *International Journal of Remote Sensing*, 22(7): 1305–1334. doi: <https://doi.org/10.1080/01431160151144369>. 5
- Strathmann, H., Sejdinovic, D., Livingstone, S., Szabo, Z., and Gretton, A. (2015). “Gradient-free Hamiltonian Monte Carlo with Efficient Kernel Exponential Families.” In *Advances in Neural Information Processing Systems*, volume 28, 955–963. 20
- Swendsen, R. H. and Wang, J.-S. (1987). “Nonuniversal critical dynamics in Monte Carlo simulations.” *Physical Review Letters*, 58: 86–88. doi: <https://doi.org/10.1103/PhysRevLett.58.86>. 10
- Talts, S., Betancourt, M., Simpson, D., Vehtari, A., and Gelman, A. (2018). “Validating Bayesian Inference Algorithms with Simulation-Based Calibration.” arXiv preprint [arXiv:1804.06788](https://arxiv.org/abs/1804.06788) [stat.ME]. URL <https://arxiv.org/abs/1804.06788> 4, 5, 17
- Tucker, C. J. (1979). “Red and photographic infrared linear combinations for monitoring vegetation.” *Remote Sensing of Environment*, 8(2): 127–150. 5
- USGS (2016). “Landsat 8 Data Users Handbook.” Technical Report LSDS-1574, United States Geological Survey, Sioux Falls, SD. Version 2.0. URL <https://landsat.usgs.gov/landsat-8-18-data-users-handbook> 3, 5

- Vermote, E., Justice, C., Claverie, M., and Franch, B. (2016). “Preliminary analysis of the performance of the Landsat 8/OLI land surface reflectance product.” *Remote Sensing of Environment*, 185: 46–56. doi: <https://doi.org/10.1016/j.rse.2016.04.008>. 5
- Wilkinson, R. D. (2014). “Accelerating ABC methods using Gaussian processes.” In *Proc. 17th Int. Conf. AISTATS*, volume 33 of *JMLR W&CP*, 1015–1023. Reykjavik, Iceland: MIT Press. 3
- Winkler, G. (2003). *Image Analysis, Random Fields and Markov chain Monte Carlo Methods: A Mathematical Introduction*. Berlin Heidelberg: Springer-Verlag, 2nd edition. MR1950762. doi: <https://doi.org/10.1007/978-3-642-55760-6>. 1
- Zhang, C., Shahbaba, B., and Zhao, H. (2017). “Precomputing Strategy for Hamiltonian Monte Carlo Method Based on Regularity in Parameter Space.” *Computational Statistics*, 32(1): 253–279. MR3604217. doi: <https://doi.org/10.1007/s00180-016-0683-1>. 20

Acknowledgments

MTM was supported by the UK EPSRC as part of the “*i*-Like” programme grant (ref: EP/K014463/1). KLM was supported by an ARC Laureate Fellowship. KLM and ANP received funding from the Australian Government via the Australian Research Council Centre of Excellence for Mathematical and Statistical Frontiers (project number CE140100049). The authors thank David Firth, Christian Robert, Jean-Michel Marin, Daniel Simpson, Clair Alston-Knox, and Natalie Moores for helpful conversations during the preparation of this manuscript. We also thank the associate editor and two reviewers for their thoughtful comments and suggestions. Landsat imagery courtesy of NASA Goddard Space Flight Center and U.S. Geological Survey. Computational resources and services used in this work were provided by the HPC and Research Support Group, Queensland University of Technology, Brisbane, Australia.