

BAYESIAN COMPLEMENTARY CLUSTERING, MCMC AND ANGLO-SAXON PLACENAMES

BY GIACOMO ZANELLA*

University of Warwick

Common cluster models for multi-type point processes model the aggregation of points of the same type. In complete contrast, in the study of Anglo-Saxon settlements it is hypothesized that administrative clusters involving complementary names tend to appear. We investigate the evidence for such an hypothesis by developing a Bayesian Random Partition Model based on clusters formed by points of different types (complementary clustering).

As a result we obtain an intractable posterior distribution on the space of matchings contained in a k -partite hypergraph. We apply the Metropolis-Hastings (MH) algorithm to sample from this posterior. We consider the problem of choosing an efficient MH proposal distribution and we obtain consistent mixing improvements compared to the choices found in the literature. Simulated Tempering techniques can be used to overcome multimodality and a multiple proposal scheme is developed to allow for parallel programming. Finally, we discuss results arising from the careful use of convergence diagnostic techniques.

This allows us to study a dataset including locations and placenames of 1319 Anglo-Saxon settlements dated between 750 and 850 AD. Without strong prior knowledge, the model allows for explicit estimation of the number of clusters, the average intra-cluster dispersion and the level of interaction among placenames. The results support the hypothesis of organization of settlements into administrative clusters based on complementary names.

1. Introduction.

1.1. *The historical problem.* The starting point of this work is a dataset supplied by Professor John Blair of Queen's College, Oxford. The dataset consists of the locations and placenames of 1316 Anglo-Saxon settlements dated between 750 and 850 AD (dataset fully available in [Supplement F](#), for illustrations see [Supplement E](#) or [Figure 12](#)). In the dataset there are 20 different kinds of placenames in total. Placenames form an important source

*Supported by EPSRC through a PhD position under the CRiSM grant EP/D002060/1

Keywords and phrases: Random partition models, Complementary clustering, Data association problems, Metropolis-Hastings algorithm, Efficient proposal distribution, K-cross function, kernel smoothing, bandwidth, Anglo-Saxon placenames locations

of information regarding the Anglo-Saxon civilization (see Gelling and Cole (2000) and Jones and Semple, 2012). Historians think that such placenames indicate specific functions or features of each settlement (for example *Burton* is thought to label a fortified settlement having a military role). Blair’s hypothesis is that the settlements were organized into administrative clusters tending to involve a variety of complementary placenames in each cluster, as illustrated in Figure 1.

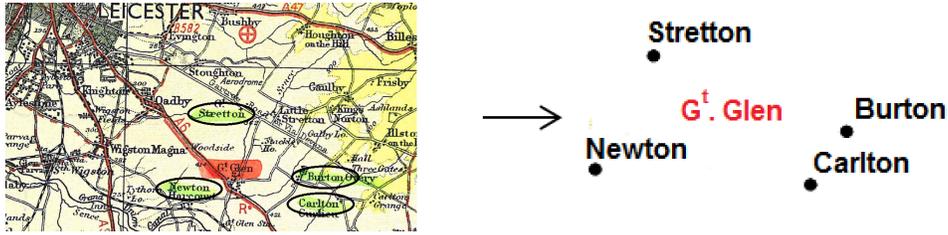


FIG 1. A cluster of four Anglo-Saxon settlements (highlighted in green and circled) in the region of Great Glen (highlighted in red).

The objective of our statistical approach to the study of settlements names and geographical locations is to address the following questions: is there statistical support for Blair’s hypothesis? What is the typical distance between settlements in the same cluster? Which portion of the settlements are clustered together? Which placenames tend to cluster together? Can we provide a list of those clusters which are more strongly supported by the analysis?

Our intention is to provide a useful contribution to historical research on this topic based on a quantitative approach, bearing in mind the scarcity of textual evidences regarding the Anglo-Saxon period.

1.2. *Modeling approach.* By considering the placenames as marks attached to points, we model our data as the realization of a k -type point process (also called k -variate point process), where k is the number of different placenames available (see Baddeley, 2010). We can view our problem as a clustering problem based on aggregations of points of different types. In fact we seek a *complementary clustering*: each cluster may contain at most one settlement for each placename. This requirement is imposed as a result of consultation with the historians involved in the project and is motivated by the assumption that each placename represents a different administrative function (role) within the cluster.

Our intention is to perform explicit inferences on the partition of settlements into clusters. As with hierarchical models, it would be desirable to analyze the dataset all at once, so as not to lose statistical power, and also

to provide inferences at the single cluster level to facilitate visualization and historical interpretation of the results of the analysis.

We employ Random Partition Models (RPMs), often used in the Bayesian Nonparametric literature (see for example [Lau and Green, 2007](#)), as they permit natural inferences on the cluster partition and they have enough flexibility to allow specification of a useful model for complementary clustering.

Standard approaches for point process cluster modeling, like the Log-Gaussian Cox Processes ([Lawson and Denison, 2010](#), ch.3) or the Neyman-Scott model, see for example [Loizeaux and McKeague \(2001\)](#), are not appropriate here, as such models usually provide inferences on the cluster centers or on the point process intensity, while we seek explicit inferences on the cluster partition. Moreover standard cluster methods for marked point process consider the marks as an additional dimension and search for aggregations of points with similar marks. In complete contrast, we seek for aggregations of points of different types.

[Diggle, Eglen and Troy \(2006\)](#) seek evidence for repulsion among points of different types in a bivariate spatial distribution of amacrine cells. They use a pairwise interaction model, which has theoretical limitations which prevent its use for clustering. While this approach could be extended to our case by using area-interaction point processes, which can model clustering ([Baddeley and Van Lieshout, 1995](#)), it would not provide us with explicit estimates of the cluster partition and it would not easily allow complementary clustering specification (at most one point of each type in each cluster).

Multi-target tracking involves the Data Association problem, that is to group together measurements recorded at different time intervals to create objects tracks (see [Oh, Russell and Sastry, 2009](#)). This problem is similar to the problem of performing complementary clustering of a k -type point process. In Data Association problems, however, the interest is to find the best association, while we are interested in assessing the strength of clustering and the level of interaction between different placenames, and in quantifying the uncertainty of our estimates. In fact the modeling aspects we have to be careful about are different from the ones of Data Association problems, while the computational challenges are similar (see Sections 3 and 4).

1.3. Organization of the paper. In Section 2 we define a RPM for complementary clustering and discuss two prior distributions for the cluster partition. Such models lead to intractable posterior distributions. In Section 3 we express such a posterior in terms of matchings contained in hypergraphs. We thus link the problems of sampling from the posterior and finding the posterior mode to the more classical problems of Data Association and Op-

timal Assignment. In Section 4 we design a Metropolis-Hastings algorithm to obtain approximate samples from the posterior. We consider the problem of choosing an efficient proposal distribution, we explore the use of Simulated Tempering to overcome multimodality and we develop a multiple proposal scheme to allow for parallel computation. In Section 5 we analyze the Anglo-Saxon placename location data, first using exploratory spatial statistics tools, such as K -functions, and secondly fitting our RPM. The results support the hypothesis of settlements being organized into administrative clusters. Finally in Section 6 we discuss future directions of research. Supplementary material includes extensive calculations, additional tables and plots, the settlements dataset and R codes to perform the data analysis.

2. The Anglo-Saxon settlements model.

2.1. Random Partition Models. We present Random Partition models (RPMs) in the specific context of planar k -type point processes. For a more general and detailed discussion see [Lau and Green \(2007\)](#) and [Müller and Quintana \(2010\)](#). Let ρ be a partition of an ordered set of marked points $\mathbf{x} = ((x_1, m_1), \dots, (x_{n(\mathbf{x})}, m_{n(\mathbf{x})})) \subseteq \mathbb{R}^2 \times \{1, \dots, k\}$. Thus ρ can be represented as an *unordered* collection $\{C_1, \dots, C_{N(\rho)}\}$ of disjoint non-trivial subsets of the indices $\{1, \dots, n(\mathbf{x})\}$ whose union is the whole set $\{1, \dots, n(\mathbf{x})\}$. A RPM is used to draw inferences on the partition ρ given the observed points \mathbf{x} . Given $C_j = \{i_1^{(j)}, \dots, i_{s_j}^{(j)}\}$ we define $\mathbf{x}_{C_j} = \{(x_{i_1^{(j)}}, m_{i_1^{(j)}}), \dots, (x_{i_{s_j}^{(j)}}, m_{i_{s_j}^{(j)}})\}$, for j running from 1 to $N(\rho)$. We call \mathbf{x}_{C_j} cluster and s_j the size of the cluster. Given the partition ρ , we suppose that locations in each cluster \mathbf{x}_{C_j} are generated independently of locations in other cluster, according to a probability density function $h_{(s_j, \sigma)}(\cdot)$ depending on s_j and on a global intra-cluster dispersion parameter σ . Thus the probability density function of \mathbf{x} conditional on ρ and σ is $\prod_{j=1}^{N(\rho)} h_{(s_j, \sigma)}(\mathbf{x}_{C_j})$.

We assign independent prior distributions to ρ and σ . With a slight abuse of notation, we denote them by $\pi(\rho)$ and $\pi(\sigma)$ respectively. We require $\pi(\rho)$ to be exchangeable with respect to the point indices $\{1, \dots, n(\mathbf{x})\}$ to reflect the fact that point labels are purely arbitrary and have no specific meaning. We obtain the following expression for the posterior density function

$$\pi(\rho, \sigma | \mathbf{x}) \propto \pi(\rho) \pi(\sigma) \prod_{j=1}^{N(\rho)} h_{(s_j, \sigma)}(\mathbf{x}_{C_j}).$$

2.2. Likelihood function. Given ρ and σ , each cluster \mathbf{x}_{C_j} is constructed as follows. First an unobserved center point z_j is sampled from $W \subseteq \mathbb{R}^2$ with

probability density function $g(\cdot)$. Then the observed points $x_{i_1^{(j)}}, \dots, x_{i_{s_j}^{(j)}}$ are given by

$$(2.1) \quad x_{i_l^{(j)}} = z_j + y_{i_l^{(j)}}, \quad l = 1, \dots, s_j,$$

where $y_{i_l^{(j)}}$ is defined as $w_{i_l^{(j)}} - s_j^{-1} \sum_{l=1}^{s_j} w_{i_l^{(j)}}$ with $w_{i_1^{(j)}}, \dots, w_{i_{s_j}^{(j)}}$ being independent bivariate $N(0, \frac{\sigma^2}{\pi} \mathbb{I}_2)$ random vectors, where \mathbb{I}_2 is the 2×2 identity matrix. The variance parametrization $\frac{\sigma^2}{\pi}$ is chosen so that σ equals the expected distance between two points in the same cluster, independently of the value of s_j . Finally the marks $m_{i_1^{(j)}}, \dots, m_{i_{s_j}^{(j)}}$ are sampled uniformly from the set $\{\{m_1, \dots, m_{s_j}\} \subseteq \{1, \dots, k\} \mid m_{l_1} \neq m_{l_2} \text{ for } l_1 \neq l_2\}$. The resulting likelihood function is

$$(2.2) \quad h_{(s_j, \sigma)}(\mathbf{x}_{C_j}) = \frac{\prod_{i \neq j} \mathbb{1}(m_i \neq m_j) g(\bar{x}_{C_j})}{\binom{k}{s_j} s_j (2\sigma^2)^{s_j-1}} \exp\left(-\frac{\pi \delta_{C_j}^2}{2\sigma^2}\right),$$

([Supplement A](#) gives calculations) where \bar{x}_{C_j} is the Euclidean barycenter of \mathbf{x}_{C_j} and $\delta_{C_j}^2 = \sum_{i \in C_j} (x_i - \bar{x}_{C_j})^2$. Here we treat $g(\cdot)$ as a known function. For the purposes of data analysis we will replace g with an estimate using Gaussian kernel smoothing (see for example Figure 4 of [Supplement E](#)) with bandwidth chosen according to the cross-validation method ([Diggle, 2003](#), p.115-118). Note that this replacement commits us to the use of a data-driven prior.

REMARK 1. *This model does not constrain $x_{i_l^{(j)}} = z_j + y_{i_l^{(j)}}$ to lie in the observed window W . To make the model more realistic one could condition the distribution of $y_{i_l^{(j)}}$ in (2.1) on $z_j + y_{i_l^{(j)}} \in W$ (in a sort of edge-correction manner). Nevertheless in our application the values of σ are small (order of 10 kilometers) compared to the size of W and so most correction terms would be negligible. Moreover computing a correction term for each center point z_j would result in a consistent additional computational burden for each step of the MCMC in Section 4. Therefore we avoid such correction terms.*

2.3. *Prior distribution on σ .* History and context suggest some considerations regarding the expected intra-cluster dispersion (σ between 3 and 10km). For example, a basic consideration is that settlements of the same cluster needed to be at no more than a few hours walking distance, in order for the inhabitants of the settlements to interact administratively and politically. Nevertheless we prefer not to impose strong prior information on σ ,

as this allows us the opportunity to see whether our study of geographical location is in accordance with available contextual information. Thus we use a prior for $\sigma^2 \in \mathbb{R}_+$ which is an Inverse Gamma conjugate prior distribution

$$\sigma^2 \sim \text{InvGamma}(\alpha_\sigma, \beta_\sigma).$$

Choices of $\alpha_\sigma = 0.001, 0.01, 0.1, 1$, and $\beta_\sigma = 0.1, 1, 10, 100$ did not significantly change the resulting posterior. In the data analysis of Section 5 we set $\alpha_\sigma = 0.1$ and $\beta_\sigma = 10$.

2.4. *Prior distribution on ρ .* We need to model a partition made up of many small clusters. In fact each cluster can contain at most k points (one for each color), and the historians expect most of the original clusters to have had fewer than 6 settlements. Common RPMs usually result in clusters with many data points and therefore do not seem to be appropriate to our case (see for example Remark 2). Consequently we consider the following two priors $\pi(\rho)$.

2.4.1. *Poisson Model for $\pi(\rho)$.* The number of clusters $N(\rho)$ follows a Poisson distribution with mean λ and each cluster size s_j is sampled from $\{1, \dots, k\}$ according to a probability distribution $\mathbf{p}^{(c)} = (p_1^{(c)}, \dots, p_k^{(c)})$. Note that in such a model the (unobserved) point process of centers $\{z_1, \dots, z_{N(\rho)}\}$ is a Poisson point process with intensity measure $\lambda g(\cdot)$ and the number of observed points need not equal n . Conditioning on having n observed points, the induced prior distribution on ρ is $\pi(\rho|\lambda, \mathbf{p}^{(c)}) \propto \prod_{j=1}^{N(\rho)} \lambda p_{s_j}^{(c)}$. We assign the following conjugate priors to λ and $\mathbf{p}^{(c)}$

$$\lambda \sim \text{Gamma}(k_\lambda, \theta_\lambda), \quad \mathbf{p}^{(c)} = (p_1^{(c)}, \dots, p_k^{(c)}) \sim \text{Dir}(\alpha_1^{(c)}, \dots, \alpha_k^{(c)}).$$

Choices of $k_\lambda = 100, 300, 600$, $\theta_\lambda = 0.5, 1, 3$ and $(\alpha_1^{(c)}, \dots, \alpha_k^{(c)})$ equal to $(1/k, \dots, 1/k)$, $(1, \dots, 1)$ and $(1, 1/(k-1), \dots, 1/(k-1))$ (where k is the number of colors) did not change the posterior significantly. In the data analysis of Section 5 we set $k_\lambda = 300$, $\theta_\lambda = 1$ and $(\alpha_1^{(c)}, \dots, \alpha_k^{(c)}) = (1/k, \dots, 1/k)$.

REMARK 2. *In the RPM literature it is common to assign a Dirichlet Process (DP) prior to ρ : $\pi(\rho|\theta) \propto \prod_{j=1}^{N(\rho)} \theta(s_j - 1)!$, where the concentration parameter θ can be fixed or random. A DP prior (conditioning on having no cluster with more than k points) would be equivalent to use of the Poisson model with fixed $\mathbf{p}^{(c)}$ given by $p_l^{(c)} = \frac{(l-1)!}{\sum_{i=1}^k (i-1)!}$, for $l = 1, \dots, k$. Such a choice would enforce most clusters to have almost k points and thus is not appropriate to this context where we expect most clusters to be smaller.*

2.4.2. *Dirichlet-Multinomial Model for $\pi(\rho)$.* For l running from 1 to k , we define $N_l(\rho)$ as the number of clusters of ρ having size l and $Y_l(\rho) = l \cdot N_l(\rho)$ so that $Y_l(\rho)$ is the total number of points in all the clusters of size l . Note that $\sum_{l=1}^k Y_l(\rho) = n(\mathbf{x})$. In this model the random vector $\mathbf{Y}(\rho) = (Y_1(\rho), \dots, Y_k(\rho))$ follows a Dirichlet-Multinomial distribution conditioned on Y_l being a multiple of l (for l running from 1 to k)

$$(2.3) \quad \Pr(Y_1 = y_1, \dots, Y_k = y_k) \propto \begin{cases} \frac{n!}{y_1! \cdots y_k!} p_1^{y_1} \cdots p_k^{y_k} & \text{if } \sum_{l=1}^k y_l = n \text{ and} \\ & y_l \text{ is a multiple of } l, \\ 0 & \text{otherwise.} \end{cases}$$

We assume that the parameter vector $\mathbf{p} = (p_1, \dots, p_k)$ is unknown with prior distribution $\text{Dir}(\alpha_1, \dots, \alpha_k)$. The resulting prior distribution of ρ given \mathbf{p} , recalling that we want such distribution to be exchangeable, is

$$(2.4) \quad \pi(\rho \mid \mathbf{p}) \propto \frac{1}{\eta(\rho)} \frac{n(\mathbf{x})!}{Y_1(\rho)! \cdots Y_k(\rho)!} p_1^{Y_1(\rho)} \cdots p_k^{Y_k(\rho)},$$

where $\eta(\rho) = \#\{\tilde{\rho} \mid \mathbf{Y}(\rho) = \mathbf{Y}(\tilde{\rho})\} = n! \left(\prod_{l=1}^k (l!)^{\frac{Y_l}{l}} (Y_l/l)! \right)^{-1}$.

REMARK 3. *This model can be seen as a Dirichlet-Multinomial mixture of k classes having Y_1, Y_2, \dots, Y_k points corresponding to singletons, couples, up to k -tuples. We are therefore converting the problem of finding an unknown number (between $\frac{n}{k}$ and n) of small clusters into the problem of finding k big clusters, with k fixed and relatively small (20 in our case).*

REMARK 4. *Note that p_l represents the probability of a point being in a cluster of size l . Since we conditioned Y_l on being a multiple of l , though, this is just an approximation. Nevertheless for $n(\mathbf{x})$ big (e.g. $n(\mathbf{x}) \geq 10$) the approximation error is negligible.*

2.5. *Model parameters and Posterior Distribution.* The two models presented above result in the following unknown elements

$$\begin{aligned} (\text{Poisson}) \quad & (\rho, \sigma, \mathbf{p}^{(c)}, \lambda) \in \mathcal{P}_n \times \mathbb{R}_+ \times [0, 1]^k \times \mathbb{R}_+, \\ (\text{Dirichlet-Multinomial}) \quad & (\rho, \sigma, \mathbf{p}) \in \mathcal{P}_n \times \mathbb{R}_+ \times [0, 1]^k, \end{aligned}$$

where \mathcal{P}_n is the set of all partitions of $\{1, \dots, n\}$. Figure 2 provides graphical representations of the two underlying conditional independence structures. Given the prior and likelihood distributions described in Sections 2.2, 2.3

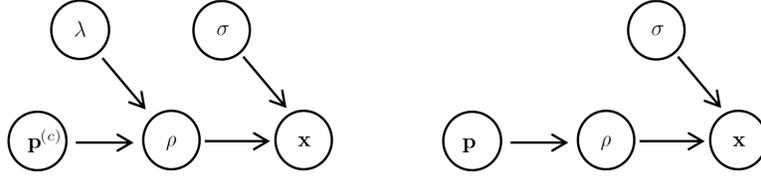


FIG 2. *Conditional independence of the random elements involved in the Poisson Model (left) and Dirichlet-Multinomial model (right).*

and 2.4 we obtain the following conditional posterior distributions for the Poisson Model

$$(2.5) \quad \pi(\rho \mid \mathbf{x}, \sigma, \mathbf{p}^{(c)}, \lambda) \propto \prod_{j=1}^{N(\rho)} \left(\frac{g(\bar{x}_{C_j}) \lambda p_{s_j}^{(c)}}{c_{s_j} \sigma^{2(s_j-1)}} \exp\left(-\frac{\pi \delta_{C_j}^2}{2\sigma^2}\right) \prod_{i,l \in C_j, i \neq l} \mathbb{1}(m_i \neq m_l) \right),$$

$$(2.6) \quad \sigma^2 \mid \mathbf{x}, \rho, \mathbf{p}^{(c)}, \lambda \sim \text{InvGamma} \left(\alpha_\sigma + n - N(\rho), \beta_\sigma + \sum_{j=1}^{N(\rho)} \frac{\pi \delta_{C_j}^2}{2} \right),$$

$$(2.7) \quad \mathbf{p}^{(c)} \mid \mathbf{x}, \rho, \sigma, \lambda \sim \text{Dir} \left(\alpha_1^{(c)} + N_1(\rho), \dots, \alpha_k^{(c)} + N_k(\rho) \right),$$

$$(2.8) \quad \lambda \mid \mathbf{x}, \rho, \sigma, \mathbf{p}^{(c)} \sim \text{Gamma} (k_\lambda + N(\rho), \theta_\lambda / (\theta_\lambda + 1)),$$

where $c_s = \binom{k}{s_j} s_j 2^{s_j-1}$. Similarly for the Dirichlet-Multinomial model

$$(2.9) \quad \pi(\rho \mid \mathbf{x}, \sigma, \mathbf{p}^{(c)}, \lambda) \propto \prod_{l=1}^k \frac{N_l!}{(lN_l)!} \cdot \prod_{j=1}^{N(\rho)} \left(\frac{g(\bar{x}_{C_j}) (p_{s_j})^{s_j}}{c_{s_j} \sigma^{2(s_j-1)}} \exp\left(-\frac{\pi \delta_{C_j}^2}{2\sigma^2}\right) \prod_{i,l \in C_j, i \neq l} \mathbb{1}(m_i \neq m_l) \right),$$

$$(2.10) \quad \sigma^2 \mid \mathbf{x}, \rho, \mathbf{p} \sim \text{InvGamma} \left(\alpha_\sigma + n - N(\rho), \beta_\sigma + \sum_{j=1}^{N(\rho)} \frac{\pi \delta_{C_j}^2}{2} \right),$$

$$(2.11) \quad \mathbf{p} \mid \mathbf{x}, \rho, \sigma \sim \text{Dir} (\alpha_1 + Y_1(\rho), \dots, \alpha_k + Y_k(\rho)).$$

In both models the posterior distribution of ρ is intractable, meaning that we cannot obtain exact inferences from it and even performing approximate

inferences is challenging. In fact the posterior sample space \mathcal{P}_n is too large (of order between $n!$ and n^n) to perform brute force optimization or integration, and the complementary clustering condition makes it not easy to move in the state space. Although we cannot solve the problem in its general form (see Section 3 for more precise statements), Monte Carlo methods can still give satisfactory results in specific applications. In Section 4 we develop MCMC techniques to perform approximate inferences and use careful diagnostic techniques to monitor its convergence.

3. Intractability of the model: relevant literature. We now focus on the intractability of the posterior distribution of ρ . We consider the latter distribution in its simplest form, $\pi(\rho \mid \mathbf{x}, \sigma, \mathbf{p}^{(c)}, \lambda)$ given in (2.5), that in fact factorizes over clusters (i.e. it is a Product Partition Model). For simplicity we will denote $\pi(\rho \mid \mathbf{x}, \sigma, \mathbf{p}^{(c)}, \lambda)$ by $\hat{\pi}(\rho)$.

In this Section we describe $\hat{\pi}(\rho)$ in terms of a hypergraph construction and then we consider the complexity of the following tasks: finding the normalizing constant of $\hat{\pi}(\rho)$, finding the mode $\rho_{max} = \arg \max_{\rho \in \mathcal{P}_n} \hat{\pi}(\rho)$ and sampling from $\hat{\pi}(\rho)$.

We will distinguish between the two-color case ($k = 2$) and the multi-color case ($k \geq 3$) because they present substantially different complexity issues. Motivated by the corresponding literature in the theory of algorithms we often refer to those as two-dimensional and k-dimensional case, even though in both cases our points lie on a plane.

3.1. Formulation of the model in terms of hypergraphs. Hypergraphs are the generalization of graphs where each hyperedge can contain more than two vertices (Berge and Minieka, 1973). In particular the complete k -partite hypergraph induced by k sets V_1, \dots, V_k is defined as $G = (V, E)$ where $V = V_1 \cup \dots \cup V_k$ and $E = \{e \subseteq V : |e \cap V_l| \leq 1 \forall l, |e| \geq 2\}$ (see Figure 3 (a)). A partition $\rho \in \mathcal{P}_n$ of n points into clusters is admissible for our

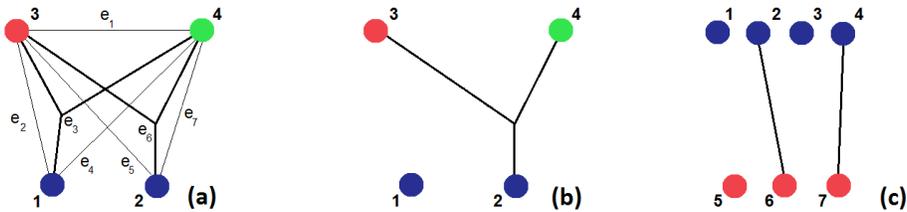


FIG 3. (a): Complete 3-partite hypergraph induced by the sets $V_1 = \{1, 2\}$, $V_2 = \{3\}$ and $V_3 = \{4\}$ corresponding to the colors blue, red and green. (b)-(c): Partial matching corresponding to $\rho = \{\{1\}, \{2, 3, 4\}\}$ and $\rho = \{\{1\}, \{2, 6\}, \{3\}, \{4, 7\}, \{5\}\}$ respectively.

model if and only if no cluster of ρ contains two points of the same type. Therefore a set of points is an admissible cluster if and only if the hyperedge connecting them belongs to the complete k -partite hypergraph induced by the k set of points corresponding to the k types. Every admissible partition ρ can then be interpreted as a partial matching (i.e. hypergraph with at most one hyperedge containing each point) contained in G : each cluster with at least two points corresponds to a hyperedge and each unlinked point is a cluster by itself (see Figure 3 (b)). Moreover we can define for each hyperedge $e = \{x_1, \dots, x_s\} \in E$ a weight

$$(3.1) \quad w(e) = \frac{(c_1)^s \lambda p_s^{(c)} g(\bar{x}) \sigma^{-2(s-1)}}{c_s \left(\lambda p_1^{(c)} \right)^s g(x_1) \cdots g(x_n)} \exp \left(-\frac{\pi \sum_{i=1}^s (x_i - \bar{x})^2}{2\sigma^2} \right),$$

in such a way that $\hat{\pi}(\rho)$ is proportional to the weight of the matching ρ , defined as $\prod_{e \in \rho} w(e)$. In (3.1) \bar{x} denotes the barycenter of x_1, \dots, x_s .

Since the two formulations are equivalent, in the remainder of the paper we will treat ρ indifferently as a partition or as a matching. Note that in the two-color case ρ reduces to a matching in a bipartite graph, see Figure 3 (c).

3.2. Finding the normalizing constant. The normalizing constant of (2.5) is the sum of the weights of all the matchings ρ contained in G , that is the total weight of G . The problem of computing the total weight of a k -partite hypergraph is an $\#P$ -hard counting problem (Valiant, 1979), even in the easiest version where $k = 2$ and the edge weights can only be 0 or 1 (i.e. counting the number of partial matchings in a bipartite graph). The $\#P$ -hard complexity class for counting problems is analogous to the NP -hard complexity class for decision problems. A counting problem y is said to be $\#P$ -hard if and only if every problem in $\#P$ (i.e. every polynomially checkable counting problem) is Cook-reducible to y (see Valiant (1979) for definitions of these terms).

3.3. Finding the posterior mode. Finding $\rho_{max} = \operatorname{argmax}_{\rho} \hat{\pi}(\rho)$ can be reduced to a k -dimensional optimal assignment problem as follows.

PROBLEM 1. (*k -dimensional optimal assignment problem*)

Instance: k sets I_1, \dots, I_k of size n and a cost function $C : I_1 \times \dots \times I_k \rightarrow \mathbb{R}$.

Problem: find an assignment A , i.e. a subset $A \subseteq I_1 \times \dots \times I_k$ containing each point of I_1, \dots, I_k exactly once, that minimizes $\sum_{(i_1, \dots, i_k) \in A} C(i_1, \dots, i_k)$.

First note that $\rho_{max} = \operatorname{argmax}_{\rho} \sum_{e \in \rho} \log(w(e))$, where ρ belongs to the set of matchings of $G = (V, E)$. Suppose now that V is made of n_1, \dots, n_k

vertices of colors $1, \dots, k$ respectively. For each color i add $n - n_i$ auxiliary points that do not contribute to the weight of any edge. Each partial matching ρ of G can then be seen as a complete matching $\tilde{\rho}$ (i.e. a matching connecting all the vertices in V) in the augmented version of G such that ρ and $\tilde{\rho}$ have the same weight. Expressing $\tilde{\rho}$ as an assignment A we obtain Problem 1.

For $k = 2$ this problem is efficiently solvable, in $O(n^3)$ steps, using the Hungarian Algorithm (Kuhn, 1955), which is based on concepts from Optimal Transportation Theory (Villani, 2009). In contrast for $k \geq 3$ this is an NP-hard optimization problem. Even more, unless $P=NP$, there is no deterministic polynomial-time approximation algorithm for a general cost function (i.e. the problem is not in APX). The same holds even if the cost function C is decomposable as $C(x_1, \dots, x_k) = \sum_{i \neq j} d(x_i, x_j)$. Some polynomial time approximation algorithms exist if d satisfies the triangle inequality, but this is not our case (see, for example, Crama and Spieksma (1992) and Bandelt, Crama and Spieksma, 1994). Balas and Saltzman (1991) propose an heuristic algorithm for a general cost function C , but no constant of approximation is provided and only the case $k = 3$ is considered.

Finally De la Vega et al. (2003) propose a polynomial time approximation scheme to partition n points of \mathbb{R}^d in m clusters that minimize the sum of the intra-clusters squared Euclidean distances. This problem is similar to ours but unfortunately the running time of their algorithm is polynomial in n but exponential in m and in our context it seems reasonable to suppose m to be of the same order of n in magnitude.

In conclusion the literature does not appear to provide a generic bounded-complexity method to obtain (or approximate) ρ_{max} . We might therefore resort to heuristics suited to our particular case.

3.4. *Approximate sampling.* In Statistical Physics a monomer-dimer system is a collection of n sites covered by molecules occupying one site, monomers, or two sites, dimers. It can be described with the following model.

MODEL 1. (*monomer-dimer system*)

Instance: A finite undirected graph $G = (V, E)$ with non negative edge weights $w : E \rightarrow [0, \infty)$ such that $w(e) > 0$ for at least one $e \in E$.

State space: the set of partial matchings $\rho \subseteq E$.

Probability distribution: $\hat{\pi}(\rho) \propto \prod_{e \in \rho} w(e)$.

Although monomer-dimer system are usually considered in lattice frameworks, the two-dimensional version of our model can be interpreted as a monomer-dimer system (see Section 3.1). Jerrum and Sinclair (1996) propose

a Metropolis-Hastings (MH) random walk algorithm to obtain approximate samples from monomer-dimer systems distributions in polynomial time. Using a canonical paths argument they prove that for any starting state ρ the mixing time of their Markov Chain satisfy

$$(3.2) \quad \tau_{\rho}(\epsilon) \leq 4(\#E)(\#V)w'^2 (\log(\#E)\#E + \log(\epsilon^{-1})),$$

where $w' = \max\{1, \max_{e \in E} w(e)\}$. [Huber and Law \(2012\)](#) consider the same Markov Chain starting from the mode ρ_{max} (which can be found in $O(\#V)^3$ by the Hungarian algorithm) and improve slightly the bound (3.2) to

$$(3.3) \quad \tau_{\rho_{max}}(\epsilon) \leq 4(\#E)(\#V)w'^2 (\log(2)\#E + \log(\epsilon^{-1})).$$

REMARK 5. *The bounds (3.2) and (3.3) seem to be very conservative in practice. For example in the framework of Section 4.1.2 the bound in (3.3) is of order 10^9 (depending weakly on ϵ). Convergence diagnostic methods, though, suggest that order 10^5 steps are enough to approximate $\hat{\pi}$.*

Can we approximately sample from $\hat{\pi}(\rho)$ in polynomial time for $k \geq 3$ too? This is equivalent to approximately count matchings in hypergraphs in polynomial time (see Chapter 3 of [Jerrum \(2003\)](#) for the relationship between approximate sampling and approximate counting). Unfortunately, as far as we are aware, there are not many results in this field. [Karpinski, Rucinski and Szymanska \(2012\)](#) try to extend the methods of [Jerrum and Sinclair \(1996\)](#) to a hypergraph setting but they managed to do it only for a specific class of sparse hypergraphs, that do not include our case. They also prove a negative result: unless $\text{NP}=\text{RP}$ (RP is the analogous of P but for randomized decision algorithms), there cannot be any FPRAS (Fully Polynomial Random Approximation Scheme, see for example [Jerrum, 2003](#)) to obtain approximate samples from the k -dimensional version of the monomer-dimer system for $k \geq 6$ (see Proposition 3 of [Karpinski, Rucinski and Szymanska, 2012](#)). Strictly speaking, this still does not imply that such a scheme cannot exist for our problem, because our problem is constrained by additional conditions (e.g. our hypergraph is k -partite).

3.5. *Summary of intractability situation.* Finding the normalizing constant of $\hat{\pi}(\rho)$ is NP-hard even for $k = 2$. The posterior mode ρ_{max} can be efficiently computed for $k = 2$, while it is an NP-hard problem for $k \geq 3$, although heuristic algorithms exist. Therefore, assuming $\text{P} \neq \text{NP}$, we should not expect to perform exact posterior inferences in polynomial time.

Polynomial-time algorithms to sample from $\hat{\pi}(\rho)$ exist for $k = 2$ (although they are not practically feasible), while some results suggest that, unless

$P=NP$, they cannot exist for $k \geq 6$ (see Section 3.4). Theoretical results of this kind do not rule out the possibility of obtaining approximate samples in specific situations, but do exclude the possibility of finding a scheme that does so (in polynomial time) for arbitrary instances of a certain class of distributions. Since the problem we consider is by no mean arbitrary it is feasible that special methods may produce good approximate samples.

In Section 4 we propose an MCMC algorithm for the two-color case and one for the k -color case. As a consequence of the results presented in this Section it is clear that additional care is needed when empirically studying MCMC mixing properties.

4. Description of proposed MCMC algorithm. We use the Metropolis-within-Gibbs algorithm to sample from $\pi(\rho, \sigma, \mathbf{p}^{(c)}, \lambda | \mathbf{x})$ and $\pi(\rho, \sigma, \mathbf{p} | \mathbf{x})$ given in (2.5)-(2.11). Direct sampling from the conditional posterior distributions of σ , \mathbf{p} , $\mathbf{p}^{(c)}$ and λ is straightforward. In contrast we can only obtain approximate samples from $\pi(\rho | \mathbf{x}, \sigma, \mathbf{p}^{(c)}, \lambda)$ and $\pi(\rho | \mathbf{x}, \sigma, \mathbf{p})$ (see Section 3). To do this we use the Metropolis-Hastings (MH) algorithm. We consider ways of improving the efficiency and of assessing the convergence of MH algorithms in this framework. For simplicity, in the remainder of Section 4 we will denote both $\pi(\rho | \mathbf{x}, \sigma, \mathbf{p}^{(c)}, \lambda)$ and $\pi(\rho | \mathbf{x}, \sigma, \mathbf{p})$ by $\hat{\pi}(\rho)$.

4.1. 2-color case. We commence by considering the two-color case because there is more known theory than in the general case and because the combinatorial structure of the sample space is simpler. We view ρ as a matching in a bipartite graph with n_1 red points and n_2 blue points (see Section 3.1). We denote the edge connecting the i -th red point and the j -th blue point by the ordered couple $(i, j) \in \{1, \dots, n_1\} \times \{1, \dots, n_2\}$.

The proposal $Q^{2D}(\rho_{old}, \rho_{new})$ for ρ is defined in two steps. First we select an edge (i, j) according to some probability distribution $q_{\rho_{old}}(i, j)$ on $\{1, \dots, n_1\} \times \{1, \dots, n_2\}$. Then, having defined i' as the index such that $(i', j) \in \rho_{old}$, if such an i' exists, and similarly j' as the index such that $(i, j') \in \rho_{old}$, if such a j' exists, we propose a new state $\rho_{new} = \rho_{old} \circ (i, j)$ defined as

$$(4.1) \quad \left\{ \begin{array}{ll} \rho_{old} + (i, j), & \text{if neither } i' \text{ nor } j' \text{ exists,} \\ \rho_{old} - (i, j), & \text{if } (i, j) \in \rho_{old}, \\ \rho_{old} - (i, j') + (i, j), & \text{if } j' \text{ exists and } i' \text{ does not exist,} \\ \rho_{old} - (i', j) + (i, j), & \text{if } i' \text{ exists and } j' \text{ does not exist,} \\ \rho_{old} - (i', j) - (i, j') \\ \quad + (i, j) + (i', j'), & \text{if } i' \text{ and } j' \text{ exist and } (i, j) \notin \rho_{old}, \end{array} \right. \quad \begin{array}{l} \text{(Addition)} \\ \text{(Deletion)} \\ \text{(Switch)} \\ \text{(Switch)} \\ \text{(Double-Switch)} \end{array}$$

where $\rho - (i, j)$ and $\rho + (i, j)$ denote the matchings obtained from ρ by respectively removing or adding the edge (i, j) . Display (4.1) defines the set of allowed moves starting from ρ_{old} and it induces a neighbouring structure on the space of matchings as follows: ρ_{new} is a neighbour of ρ_{old} if $\rho_{new} = \rho_{old} \circ (i, j)$ for some (i, j) . Jerrum and Sinclair (1996) and Oh, Russell and Sastry (2009) consider similar but slightly smaller sets of allowed moves, given by the addition and deletion moves and addition, deletion and switch moves, respectively. It is plausible that increasing the set of allowed moves improves the mixing of the MH Markov chain.

Display (4.1) does not identify uniquely the proposal $Q^{2D}(\rho_{old}, \rho_{new})$ because we still need to choose $q_{\rho_{old}}(\cdot, \cdot)$. Different choices of $q_{\rho_{old}}(\cdot, \cdot)$ will affect the mixing properties of the MH algorithm. Previous works (e.g. Jerrum and Sinclair (1996) and Oh, Russell and Sastry, 2009) chose $q_{\rho_{old}}(i, j)$ to be a uniform measure over the edges $(i, j) \in E$. A naive implementation of such choice leads to poor mixing because most proposed matchings ρ_{new} are improbable and therefore are typically rejected (in our experiments usually less than 1% of the proposed moves were accepted). Some authors overcome this problem using a truncation approximation of the posterior: they force edge weights below a certain threshold δ to be zero, and then choose

$$(P1) \quad q_{\rho_{old}}(i, j) \propto \mathbb{1}_{\{w_{ij} > \delta\}},$$

where w_{ij} is the weight of the edge (i, j) defined in (3.1) and $\mathbb{1}$ denotes the indicator function. See for example the measurement validation step in Oh, Russell and Sastry (2009).

In the following we propose a choice of $q_{\rho_{old}}$ that achieves a better mixing than (P1) and does so without requiring to target an approximation of the posterior.

Firstly note that, especially when $\hat{\pi}(\rho)$ has a factorization in terms of edge weights, it is straightforward to evaluate $\hat{\pi}$ up to a multiplicative constant on the set of neighbours of ρ_{old} defined in (4.1). For example, for the addition move, $\frac{\hat{\pi}(\rho_{old} \circ (i, j))}{\hat{\pi}(\rho_{old})} = w_{ij}$. Thus, one may be tempted to propose proportionally to $\hat{\pi}$ restricted on the set of allowed moves as follows

$$(P2) \quad q_{\rho_{old}}(i, j) \propto \hat{\pi}(\rho_{new}) \quad \text{where } \rho_{new} = \rho_{old} \circ (i, j).$$

Such a choice, however, does not take into account the fact that the normalizing constants of $q_{\rho_{old}}(\cdot, \cdot)$ and $q_{\rho_{new}}(\cdot, \cdot)$ differ for $\rho_{old} \neq \rho_{new}$. As a consequence, for example, detailed balance conditions, $\frac{Q^{2D}(\rho_{old}, \rho_{new})}{Q^{2D}(\rho_{new}, \rho_{old})} = \frac{\hat{\pi}(\rho_{new})}{\hat{\pi}(\rho_{old})}$, are not satisfied, not even approximately. A better choice for $q_{\rho_{old}}(\cdot, \cdot)$ is

$$(P3) \quad q_{\rho_{old}}(i, j) \propto \frac{\hat{\pi}(\rho_{new})}{\hat{\pi}(\rho_{old}) + \hat{\pi}(\rho_{new})}, \quad \text{where } \rho_{new} = \rho_{old} \circ (i, j).$$

Our experiments show that the latter choice leads to a significant improvement in the mixing of the MH Markov chain compared to (P1) and (P2) (see Section 4.1.2). The main reason is that the MH algorithm induced by such proposal has a very high acceptance rate (usually above 99%) without changing the set of allowed moves. It can be shown that, under some regularity assumption on the weights, the proposal given by (P3) satisfies detailed balance condition in the asymptotic regime (i.e. when the number of points tends to infinity), and this helps to explain why the acceptance rate is so high. Similarly one could also derive Peskun ordering arguments in the asymptotic regime. We omit those theoretical results here in favor of demonstrating the mixing improvement given by (P3) using the convergence diagnostic techniques in Section 4.1.2.

There is a trade-off between the complexity of the proposal and the mixing obtained (a complex proposal increases the cost of each step, while a poor proposal increases the number of MCMC steps needed). We seek a compromise with good mixing properties, like (P3), while still requiring little computation, like (P1). In Supplement B we derive the following proposal distribution to try to obtain such goal

$$(P4) \quad q_{\rho_{old}}(i, j) \propto \begin{cases} q^{(add)}(i, j) & \text{if } (i, j) \notin \rho_{old}, \\ q^{(rem)}(i, j) & \text{if } (i, j) \in \rho_{old}, \end{cases}$$

where $q^{(rem)}(i, j) = w_{ij}^{-1/2}$ and

$$q^{(add)}(i, j) = \sqrt{w_{ij}} \left(1 - \sum_{j' \neq j} \frac{w_{ij'} - \sqrt{w_{ij'}}}{1 + \sum_{s \neq i} w_{sj'} + \sum_l w_{il}} \right) \left(1 - \sum_{i' \neq i} \frac{w_{i'j} - \sqrt{w_{i'j}}}{1 + \sum_{s \neq j} w_{i's} + \sum_l w_{lj}} \right).$$

Note that $q^{(rem)}(i, j)$ and $q^{(add)}(i, j)$ do not depend on ρ and can be precomputed at the beginning of the MCMC run. See Section 4.1.2 for discussion of performance.

4.1.1. *Scaling the proposal with a multiple proposal scheme.* When using the MH algorithm on continuous sample spaces one can usually tune the variance of its proposal distribution to improve the efficiency of its algorithm (see for example Roberts, Gelman and Gilks, 1997). Given the very high acceptance rate obtained proposing according to (P3) it is natural to consider the

possibility of scaling our proposal in order to obtain longer-scale moves. The scaling problem for MH algorithms in discrete contexts has been considered, for example, in [Roberts \(1998\)](#). In that case the sample space was $\{0, 1\}^N$, the vertices of the N -dimensional hypercube, and the scaling parameter, say l , was a positive integer representing the number of randomly-chosen bits to be flipped at any given proposal.

Unfortunately, because of the nature of our sample space, it is not so straightforward to scale the proposal distribution $Q^{2D}(\rho_{old}, \rho_{new})$. One possibility is to scale by choosing l edges, $\{(i_h, j_h)\}_{h=1}^l$, and performing l moves defined in (4.1), proposing $\rho_{new} = \rho_{old} \circ (i_1, j_1) \circ \dots \circ (i_l, j_l)$. However the l moves corresponding to $\{(i_h, j_h)\}_{h=1}^l$ cannot be performed independently: consider, for example, the case where $i_1 = i_2$. We would then have to perform l moves sequentially, at a computational cost being roughly l times the one of a single move. Therefore scaling the proposal in such a way does not seem to be effective.

If the l moves could be performed independently, it would be possible to implement a multiple proposal scheme using parallel computation, thus leading to a significant computational gain. This can be obtained by considering an approximation of our model, where points at a distance greater or equal than some r_{max} have probability 0 of being in the same cluster. The latter procedure is equivalent to the truncation procedure cited in Section 4.1 and can be viewed as coming from the use of truncated Gaussian distributions to model points distribution within clusters, see (2.1). Using this truncated model and dividing the observed region into a grid, we defined a multiple proposal scheme where the l moves are proposed and accepted/rejected simultaneously and independently. Therefore, at each MH step, such l moves can be performed in an embarrassingly parallel fashion, meaning that they can be performed without the need for any communication between them. In [Supplement C](#) we give more details on the implementation and we show that in practice the mixing of the resulting MH algorithm improves by a factor roughly equal to l itself (note that the maximum value of l is bounded above, in a way that depends on r_{max} and the size of the window). A parallel-computing implementation of this algorithm would offer significant speed-ups (we anticipate speed ups by a factor around 8 for our dataset, see [Supplement C](#)). Such speed-ups would increase with the size of the dataset and window, making this proposal scheme especially relevant for applications to very large datasets. In [Supplement C](#) this scheme is presented and tested for fixed σ . In case σ is varying, either one requires an upper bound on σ or one needs different square grids for different values of σ .

4.1.2. *Convergence Diagnostics.* We used various convergence diagnostic techniques in order to assess the reliability of our algorithm, to indicate the number of iterations needed, and to compare the efficiency of the four proposals (P1)-(P4) of Section 4.1. We demonstrate such techniques on the posterior $\pi(\rho|\sigma, \mathbf{p}^{(c)}, \lambda, \mathbf{x})$ given by the Poisson model with $k = 2$, $\sigma = 0.3$, $p_1^{(c)} = p_2^{(c)} = 0.5$, $\lambda = 50$ and the center intensity $g(\cdot)$ is uniform over a window $W = [0, 10] \times [0, 10]$. Here \mathbf{x} is a synthetic sample of 44 red and 47 blue points generated according to the Poisson model just defined (see Figure 5 (a)). We set the threshold δ of (P1) to 0.001. R code used to produce the results presented in this Section is available in [Supplement F](#).

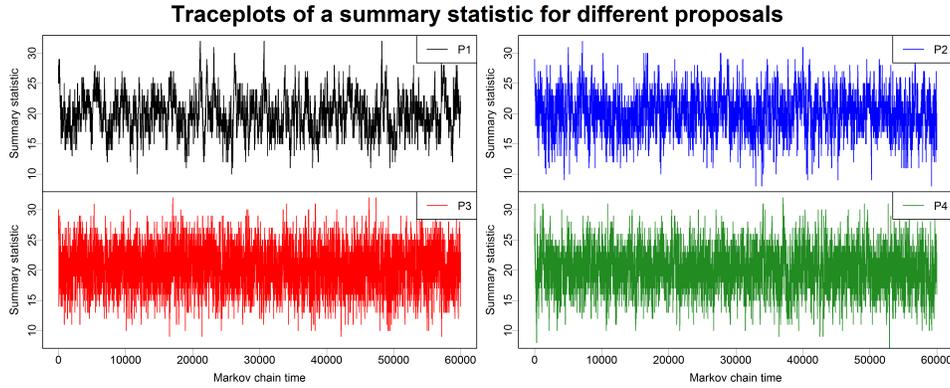


FIG 4. Traceplots of the number of differences from a reference matching.

We *first* performed some qualitative output analysis by looking at summary plots of the MCMC samples of the partition (as the one in Figure 5 (a)). Such plots can be helpful to spot when mixing has not yet occurred (see Section 4.1.3).

Secondly we considered different real valued summary statistics of the chain state (typically the number of different edges from a fixed reference matching). We plotted time series (see Figure 4) and empirical distributions of such real valued functions for different runs of the MCMC starting from different configurations. We estimated the autocorrelation functions (see Figure 5 (b)), the Integrated Autocorrelation Time (IAT) and the Effective Sample Size (ESS) of such real-valued time series using the *R* package *coda* (see [Plummer et al., 2005](#)) in order to compare different versions of the algorithm (see Table 1).

Thirdly we used some standard convergence diagnostic techniques (see [Brooks and Roberts \(1998\)](#) and [Cowles and Carlin \(1996\)](#) for an overview of the techniques available). In particular we used the multivariate version of

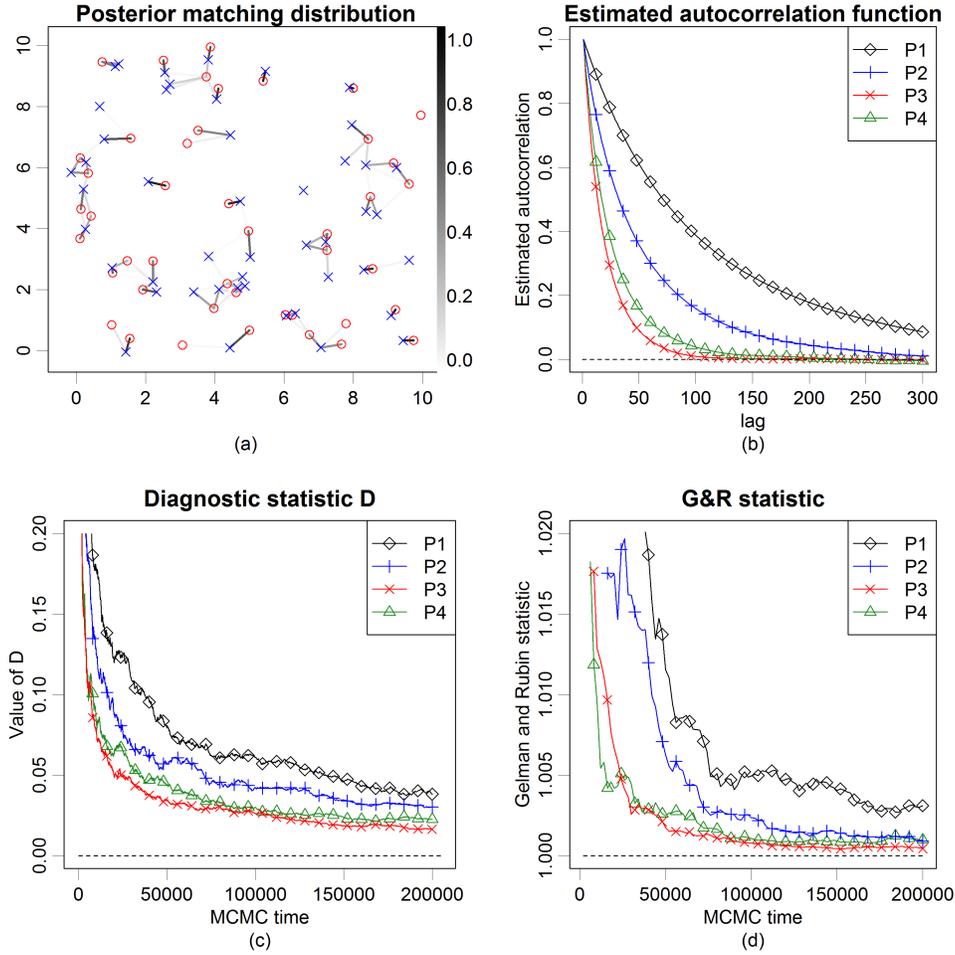


FIG 5. Four convergence diagnostic techniques described in Section 4.1.2.

Gelman and Rubin diagnostic (see [Gelman and Rubin \(1992\)](#) and [Brooks and Gelman, 1998](#)). Figure 5 (d) shows the results obtained by using a 10-dimensional summary statistic of ρ . In this context univariate summary statistics are not sufficiently informative and therefore misleading results can be obtained if these are used as the sole basis for convergence diagnostics.

Finally we compared two independent runs of the algorithm (with different starting states) by looking at estimates of the association probabilities $p_{ij} = \Pr((i, j) \in \rho)$ with $\rho \sim \hat{\pi}$. We consider the following measure of proximity

$$(4.2) \quad D = \sup_{(i,j) \in E} |\hat{p}_{ij}^{(1)} - \hat{p}_{ij}^{(2)}|,$$

where $\hat{p}_{ij}^{(1)}$ and $\hat{p}_{ij}^{(2)}$ denote the proportion of time that (i, j) was present in the two MCMC runs. As starting states we considered the empty matching (each point is a cluster), the posterior mode (obtained with the Hungarian algorithm) and matchings obtained as the output of the MCMC itself. Since equation (4.2) considers each link individually, we expect the resulting convergence diagnostic indicator D to be more severe than the ones obtained from one or few summary statistics. Results are shown in Figure 5 (d).

None of the convergence methods just presented indicate convergence issues except in the complete matching case (when the parameter p_1 or $p_1^{(c)}$ is equal or very close to 0), that is considered in the next subsection.

	mean acc.rate	Estimated IAT	ESS for 10^4 steps [for 1 sec]	steps [sec] to $D < .05$	steps [sec] to $GR < .005$
P1	17%	206	262 [270]	1.4e05 [7.3]	7.6e04 [13.5]
P2	41%	108	544 [40]	7.1e04 [84.6]	6.2e04 [97]
P3	97%	40	1358 [99]	2.0e04 [32.7]	2.4e04 [27.3]
P4	68%	55	1038 [747]	3.4e04 [2.2]	1.6e04 [4.8]

TABLE 1

Performances of the four proposals of Section 4.1 on configuration in Figure 5 (a) averaged over 5 independent runs for each proposal. GR denotes the multivariate Gelman and Rubin statistic (potential scale reduction factor). The running time indicated in brackets is evaluated using R software on a desktop computer with Intel i7 processor.

All convergence diagnostic techniques agree in indicating that proposal (P3) gives the best mixing; however in terms of real computation time the most efficient proposal is (P4). Note that such performances depend on the measure being targeted and, when running time is considered, on the computer implementation of such proposals. For the case considered in this Section, proposal (P4) gives a 3-4 times speed-up over the commonly used choice (P1). Depending on the configuration such speed-up may vary. According to our experiments, for “flatter” distributions (e.g. increasing σ to 1 and $p_1^{(c)}$ to 0.9, while keeping the other parameters unchanged) the speed-up almost disappears, while for “rougher” distributions (e.g. decreasing both σ and $p_1^{(c)}$ to 0.1, while keeping the other parameters unchanged) the speed-up increases and (P4) can be to 10 times faster than (P1).

4.1.3. Multimodality and Simulated Tempering. In the complete matching case the posterior distribution of ρ presents a strongly multimodal behavior. Cycle-like configurations like the one in Figure 6 (a) are local maxima for $\hat{\pi}(\rho)$. In fact in order to reach a higher probability configuration (i.e. shorter links) from such a “cycle” configuration, with the set of allowed moves defined by (4.1), the chain needs to pass through lower probability

configurations (i.e. longer links). If we consider extreme cycle-like configurations (e.g. Figure 6 (b)), then the MCMC will typically get stuck in such local maxima. In order to overcome this potential multimodality problem we

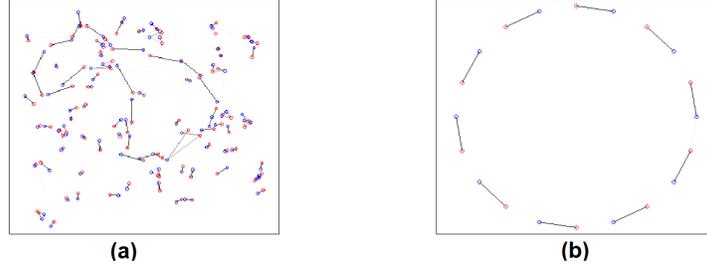


FIG 6. Configurations corresponding to local maxima of $\pi(\rho|\mathbf{x})$ for (a) a synthetic sample and (b) an artificially designed points configuration.

implemented a simulated tempered version of our MCMC algorithm, see for example [Geyer and Thompson \(1995\)](#) or [Marinari and Parisi \(1992\)](#) for references. This technique manages to overcome local maxima for the complete matching case even when extreme cycle-like configurations are present (as in Figure 6 (b)). Nevertheless our specific application do not present a complete matching case and therefore we have a milder multimodality and the MCMC algorithm exhibits sufficient mixing without the use of Simulated Tempering. Therefore Simulated Tempering is not used for the real data analysis, as convergence diagnostic tools do not show suspicious behavior.

We note that [Dellaert et al. \(2003\)](#) deal with multimodality in a similar posterior space (made of perfect matchings in a bipartite graph) arising from the Structure from Motion problem. In order to allow the MH algorithm to overcome local maxima like the one in Figure 6 (b) they allow the MH proposal to include “long” moves that they call “chain flipping”.

4.2. *k-color case.* We now define an MCMC algorithm that targets $\hat{\pi}(\rho)$ when $k \geq 3$. This case is harder than the two-dimensional one because it involves clusters with different dimensions and not just pairwise interaction.

4.2.1. *Description of proposed Gibbs projection MCMC algorithm.* We define the transition kernel P of our MCMC algorithm as a mixture of $\binom{k}{\lfloor k/2 \rfloor}$ MH transition kernels, each of which corresponds to a group A of $\lfloor k/2 \rfloor$ colors

$$(4.3) \quad P(\rho_{old}, \rho_{new}) = \binom{k}{\lfloor k/2 \rfloor}^{-1} \sum_{A \subset \{1, \dots, k\}, |A| = \lfloor k/2 \rfloor} P^{(A)}(\rho_{old}, \rho_{new}),$$

where $\binom{\cdot}{\cdot}$ denotes the binomial coefficient and $\lfloor k/2 \rfloor$ denotes the integer part of $k/2$. Here $P(\cdot, \cdot)$ selects a set of colors A , “projects” the k -color configuration to a 2-colors configuration where the new two colors correspond to A and $A^c = \{1, \dots, k\} \setminus A$ and then acts on the two-colors configuration. More precisely the action of $P^{(A)}$ is the following (see Figure 7):

1. reduce the k -color configuration (\mathbf{x}, ρ_{old}) to a two-color one $(\mathbf{x}^{2D}, \rho_{old}^{2D})$ by replacing the points having colors in A and A^c respectively with their cluster centroids. We denote by d_i the number of points merged together into the i -th point x_i^{2D} ,
2. obtain ρ_{new}^{2D} from $(\mathbf{x}^{2D}, \rho_{old}^{2D})$ with one or more MH moves using the proposal Q^{2D} of Section 4.1 on a target measure $\hat{\pi}^{2D}$ being the two-dimensional version of $\hat{\pi}$ (modified to take account of the multiplicity of the points d_i , see Supplement D),
3. obtain the k -color configuration (\mathbf{x}, ρ_{new}) from $(\mathbf{x}^{2D}, \rho_{new}^{2D})$ by the inverse operation of Step 1.

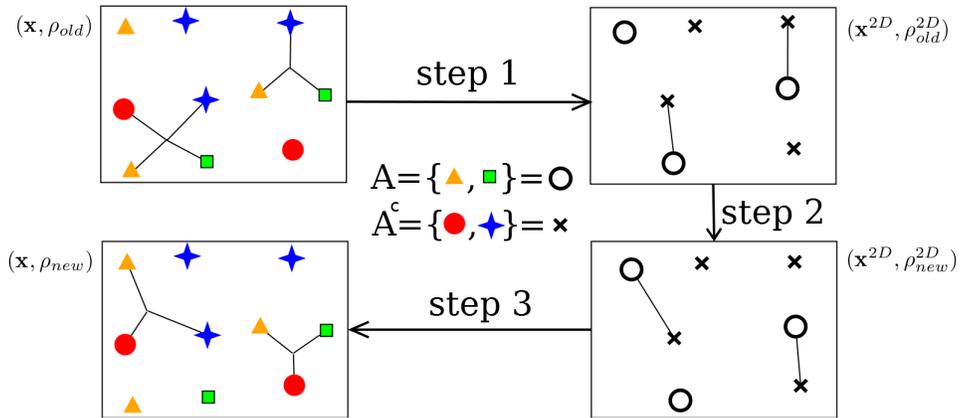


FIG 7. The action of a transition kernel $P^{(A)}$ for a given A .

In order for this algorithm to be correct $\hat{\pi}^{2D}$ must be proportional to $\hat{\pi}$ on the collection of possible moves of $P^{(A)}$, so that $P^{(A)}$ satisfies detailed balance conditions with respect to $\hat{\pi}$. This follows from basic properties of the Gaussian density function and is proven in Supplement D. Note that, when k is even, choosing A or A^c leads to the same transition kernel. This is not an issue and it is indeed equivalent to never choosing A^c and choosing A twice more often.

This scheme allows us to re-use the two-color algorithm. Moreover, once the projection from (\mathbf{x}, ρ_{old}) to $(\mathbf{x}^{2D}, \rho_{old}^{2D})$ has been done, it is faster to per-

form MH moves in the two-color matching space rather than in the original k -color one (note that the $\hat{\pi}^{2D}$ depends on \mathbf{x} only through \mathbf{x}^{2D}).

By merging colors together we allow proposals which move many points at the same time from one cluster to another. Therefore the induced set of allowed moves is broader than, for example, the one of a scheme that moves one point at a time. [Oh, Russell and Sastry \(2009\)](#) consider also, for example, “birth” moves that propose to create a cluster from three or more single points in one step. Such moves are likely to be useful to speed up mixing in applications where there appear clusters with many points.

The k -dimensional sample space is more complicated than the two-dimensional one. Therefore additional care and longer MCMC runs are needed. We implement analogous convergence diagnostic techniques to the ones in [Section 4.1.2](#). As might be expected, the number of MCMC steps needed to reach stationarity and to obtain mixing is much higher than in the two-color case (see [Section 5](#)). Nevertheless our experiments suggest that, as in the two-color case, the MCMC manages to mix properly unless we are in a case close to complete matching (see [Section 4.1.3](#)).

5. Application to the real data. We now analyze the Anglo-Saxon settlements dataset supplied by Prof. John Blair, firstly by using common Spatial Statistic tools and then by applying the model described in [Section 2](#). Such dataset is available in [Supplement F](#).

5.1. Preliminary analysis. The dataset is made of 20 different groups, each of which contains the list of settlements having one of the 20 place-names (see [Table 1](#) of [Supplement E](#)). The historians involved in the project expect the clustering behaviour to involve in particular 13 of those place-names (indicated in [Table 1](#) of [Supplement E](#)). We refer to the settlements relative to those 13 placenames as *reduced dataset*, and to all the settlements recorded as *full dataset*. We will perform the analysis on both datasets.

For each settlements the following variables are given: County, place, Parish / Township, grid ref, date of first evidence (see [Table 2](#)).

The locations are expressed through the Ordnance Survey (OS) National Grid reference system. A set of OS National Grid coordinates, like *SU230870*, identify a $100\text{m} \times 100\text{m}$ square on a grid covering Great Britain. Some locations in the dataset are given with just 2 letters and 4 digits, for example *SU2387*, and they identify a $1\text{km} \times 1\text{km}$ square. In the dataset some coordinates have a letter *c* in front of them, for example *c.SU2387*, to indicate that the location is less accurate.

COUNTY	PLACE	PARISH/ TOWNSHIP	GRID REF	DATE OF FIRST EVIDENCE
BRK	Bourton	Bourton	SU 230870	c. 1200
BUC	Bierton	Bierton with Broughton	SP 836152	DB
BUC	Bourton	Buckingham	SP 710333	DB
CHE	Burton	Burton (T)	SJ 509639	DB
CHE	Burton	Burton (T)	SJ 317743	1152
CHE	Buerton	Buerton (T)	SJ 682433	DB

TABLE 2

Data available regarding the first 6 settlement with the name Burton.

5.1.1. *Data cleaning and data assumptions.* Our analysis is concerned with placenames and geographical locations. We need to express the placenames (variable “place”) as k types and to convert the locations (variable “Grid reference”) to two-dimensional Euclidean coordinates. Such data cleaning process entails historical assumptions on the dataset and thus we have been guided by the judgment of the subject-specific historians involved in this project in doing so.

Placenames: we ignore minor variation in placenames. For example we consider the settlements of Table 2 as having placename *Burton*: their actual recorded placenames vary amongst *Burton*, *Bourton*, *Bierton*, *Buerton*.

Four groups (out of 20) are made up of two subgroups each with similar placenames: *Aston/Easton*, *Charlton/Charlcot*, *Drayton/Draycot* and *Walton/Walcot*. We consider such subgroups to be the same, for example *Charlton* and *Charlcot* are treated as the same placename.

Locations: each settlement is assumed to be located at the center of the square identified by its OS National Grid coordinates.

“Multiple” records: it is sometimes indicated in the original dataset that some couples (or triples) of settlements, with same placename and very close locations, have to be considered as multiple records of the same settlement. We replaced such couples (or triples) of settlements with one settlement located at their midpoint. Moreover there are some other pairs of settlements that have very close locations and the same placename. It is primarily a matter of historical interpretation whether these couples have to be considered as single settlements. We performed the analysis under both hypothesis (keeping them separated and merging them) without seeing significant change in the results. The analysis presented here is made with those settlements merged together (3 km is the threshold distance below which we identify two records of settlements with the same placename).

Table 1 of Supplement E records the numbers of imprecise and multiple records for each placename.

5.2. *K-cross functions and null hypothesis testing.* Second moment functions are a useful tool to investigate interpoint interaction (Chiu et al., 2013). In particular, given a multivariate point pattern, bivariate (or cross-type) K-functions provide good summary functions of the interaction across points of different types. The bivariate K-function $K_{ij}(r)$ is the expected number of points of type j closer than r to a typical point of type i , divided by the intensity λ_j of the type j subpattern of points \mathbf{x}_j (Baddeley, 2010, Sec. 6). For testing and displaying purposes we define a single summary function, a multivariate K-function $K_{cross}(r)$, as the weighted average of $K_{ij}(r)$ for $i \neq j$, where the weights are the product of the intensities $\lambda_i \lambda_j$.

Classical K -functions, however, rely strongly on the assumption that the point pattern is stationary, which is not appropriate for our dataset. Therefore we use the inhomogeneous version of the K -functions, where the contribution coming from each couple of points is reweighted to take into account for spatial inhomogeneity (Baddeley, Møller and Waagepetersen, 2000). Standard estimates of the inhomogeneous bivariate K -functions \hat{K}_{ij} are obtained using the *spatstat R* package (Baddeley and Turner, 2005).

In order to test whether the interaction shown by K -functions is significant or not we need to define a null hypothesis (representing no-interaction among placenames). Section 8 of Baddeley (2010) describes three classical null hypotheses for multivariate point processes: random labeling (given the locations the point types are i.i.d.), Complete Spatial Randomness and Independence (CSRI, the locations are a uniform Poisson point process and the point types are i.i.d.) and independence of components (points of different types are independent). The random labeling and the CSRI hypotheses are unrealistic assumptions for our dataset because our point pattern is clearly not stationary and the distribution of placenames is not spatially homogeneous (some placenames are more concentrated in the South, some in the North and so on). The independence of components hypothesis is realistic but, in order to test it, stationarity of the points pattern is usually assumed. Instead we define the following no-interaction null hypothesis: each point pattern \mathbf{x}_j is an inhomogeneous Poisson point process (with intensity function $\lambda_j(\cdot)$ potentially varying over j).

This leads to an approximate Monte Carlo test as follows. First we estimate the intensities $\lambda_j(\cdot)$ with $\hat{\lambda}_j(\cdot)$ (see Figures 4 and 5 of Supplement E) obtained through standard Gaussian kernel smoothing (bandwidth chosen according to the cross-validation method, see for example Diggle (2003), p.115-118). Secondly we sample 99 independent multivariate inhomogeneous Poisson point patterns according to $\{\hat{\lambda}_j(\cdot)\}_{j=1}^k$. Finally we use those samples both to plot simulation envelopes and to perform a deviation test with sig-

nificance $\alpha = 0.05$ using as a summary function a centered version of the L -function $\hat{L}_{cross}(r) = \sqrt{\frac{\hat{K}_{cross}(r)}{\pi}}$ for $r \in (0, r_{max})$, with $r_{max} = 15$ km. The de-

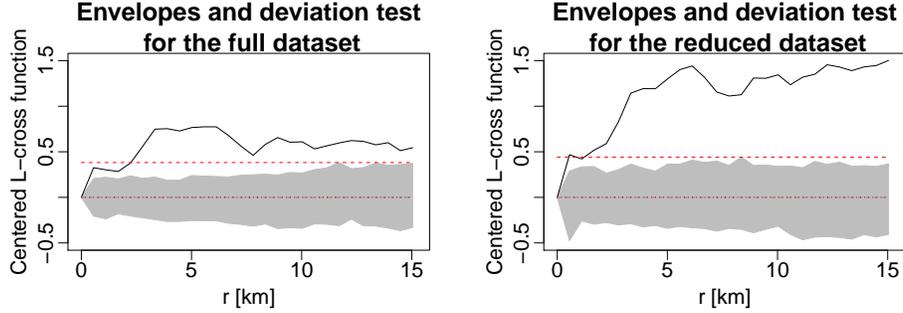


FIG 8. Black solid lines represent $\hat{L}_{cross}(r) - \mathbb{E}[\hat{L}_{cross}(r)]$ for the observed pattern, the 95% envelopes (gray areas) are obtained using 99 simulated patterns and the red dashed lines indicate the upper deviations. Deviation test: if the black solid line rises above the red dashed line then the interaction can be considered significant at significance level $\alpha = 0.05$. The values of $\mathbb{E}[\hat{L}_{cross}(r)]$ are estimated using simulated point patterns generated according to the null hypothesis.

viation test (Grabarnik, Myllymäki and Stoyan, 2011) summarizes the summary function with a single value $D = \max_{r \in (0, r_{max})} \hat{L}_{cross}(r) - \mathbb{E}[\hat{L}_{cross}(r)]$ and compares it to the ones obtained from the 99 simulated samples. The null hypothesis is rejected for both the full and the reduced dataset (see Figure 8 for the results). For the reduced dataset this provides evidence of a stronger clustering effect. The R code used to perform this test and produce Figure 8 is given in Supplement F. Application of the same deviation test on the bivariate L -functions $\hat{L}_{ij}(r)$ provides an indication of which couples of placenames exhibit significant interaction (see Figure 6 of Supplement E).

K -functions can indicate interaction and suggest which pairs of placenames interact more than others. Nevertheless they do not provide explicit estimates. We regard K -functions as a useful exploratory tool to guide us in the data analysis and to double-check the results we obtain from our model.

We note that Dr. Stuart Brookes from UCL has already used second moment functions to do some preliminary analysis on the Anglo-Saxon settlements dataset presented here (personal communication by Prof. John Blair).

5.3. Analysis of the Anglo-Saxon settlements dataset using our model.

We present now the main results obtained by analyzing the Anglo-Saxon settlements dataset with the Random Partition Model described in Section 2. The analysis gives modest support to the historian hypothesis that these settlements are clustered, and it permits inference about ranges of values for

relevant parameters.

Here the no-clustering null hypothesis corresponds to $p_1^{(c)} = 1$ or $p_1 = 1$ for the Poisson and Dirichlet-Multinomial models respectively (see Section 2). As shown in Figure 9, such a hypothesis clearly lies outside the region

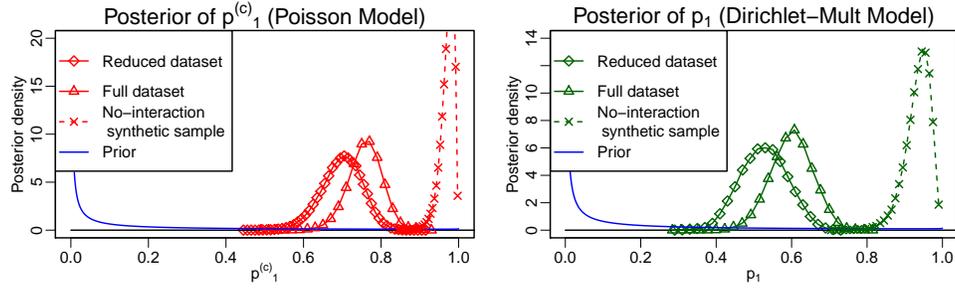


FIG 9. Estimated posterior distribution of $p_1^{(c)}$ and p_1 (see Section 2) for the reduced and full dataset (13 and 20 placenames respectively). The no clustering hypothesis ($p_1^{(c)} = 1$ and $p_1 = 1$ respectively) lies outside the support of the posterior for the real dataset.

where the posterior distribution is concentrated. As a sanity check we also simulated synthetic samples generated according to the no-clustering null hypothesis (as defined in Section 5.2) and plotted the posterior distribution of $p_1^{(c)}$ and p_1 . As one would expect, in this case, $p_1^{(c)} = 1$ or $p_1 = 1$ are included in the posterior support (see Figure 9).

Figure 10(a) shows the estimated posterior distribution of σ for the reduced dataset. Such posteriors are all clearly peaked around 5 km. For the

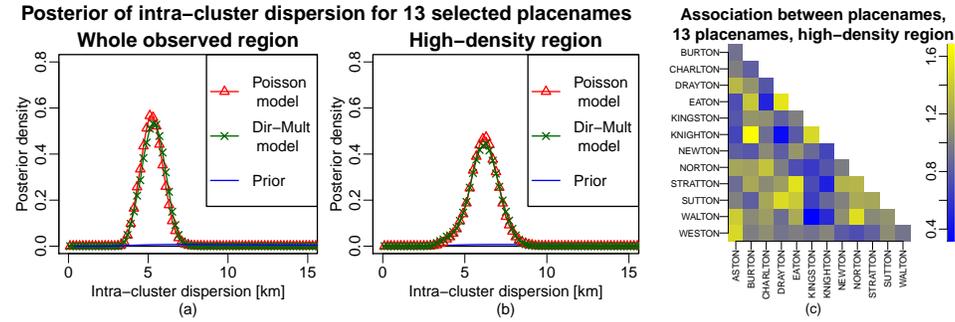


FIG 10. (a) $\pi(\sigma|\mathbf{x})$ for the reduced dataset. (b) $\pi(\sigma|\mathbf{x})$ considering only a high-density region (see Section 6). (c) Measure of association between placenames (see Section 5.3).

Poisson model the 95% Highest Posterior Density interval is (3.5, 6.1) km and the posterior mean is 4.7 km. Therefore, according to the fit given by our model, the clustering behavior consists of clusters with settlements hav-

ing distance being approximately 5 km on average. It is satisfying to note that this value is in accordance with the value suggested by the historians involved in the project and coherent with the historical interpretation (see Section 2.3).

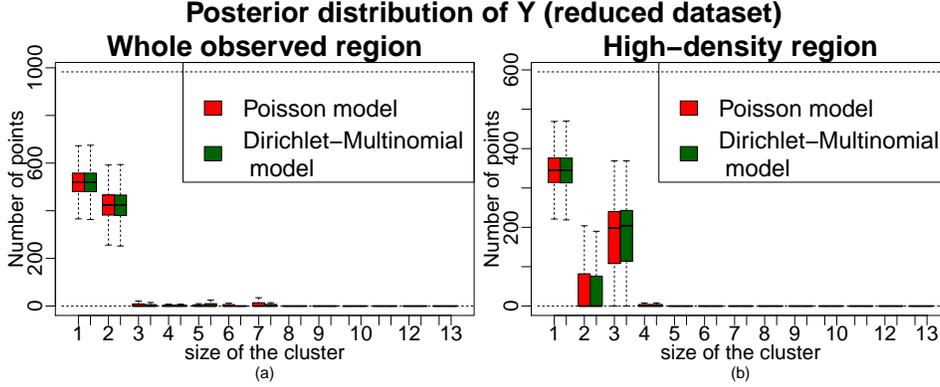


FIG 11. (a) Posterior distribution of $\mathbf{Y} = (Y_1, \dots, Y_k)$ for the reduced dataset. (b) Same but considering only the settlements in a high density region (see Section 6).

Figure 11(a) shows a box plot representation of the posterior distribution of (Y_1, \dots, Y_k) , where Y_l is the number of settlements in clusters of size l (i.e. with l settlements). Note that approximately half of the settlements are not clustered (i.e they belong to clusters of size 1). Moreover most of the clustered settlements belongs to clusters of size 2. Historians expected to see more clusters involving three or four settlements than what was reported by our model. Inspection shows that model-fitting, and the requirement to fit clusters in the low-density region (which mostly contain couples with a high posterior probability), forces all the clusters in the high-density region to be couples too. In fact when the high-density region is analyzed separately (approximately 600 settlements) more triples appear and the posterior of σ shifts slightly towards bigger values, see Figures 10(b) and 11(b). This suggests that there might be an heterogeneity in the clustering behaviour between high and low-density regions which is not captured in the model when applied to the whole region. This indicates a possible direction for future work (see Section 6).

Figure 12 shows a graphical representation of the posterior distribution of the partition ρ (for the reduced dataset analyzed with the Poisson model). This representation is of considerable use since it provides a visual understanding of how the model is fitting the data and enables comparison with contextual information.

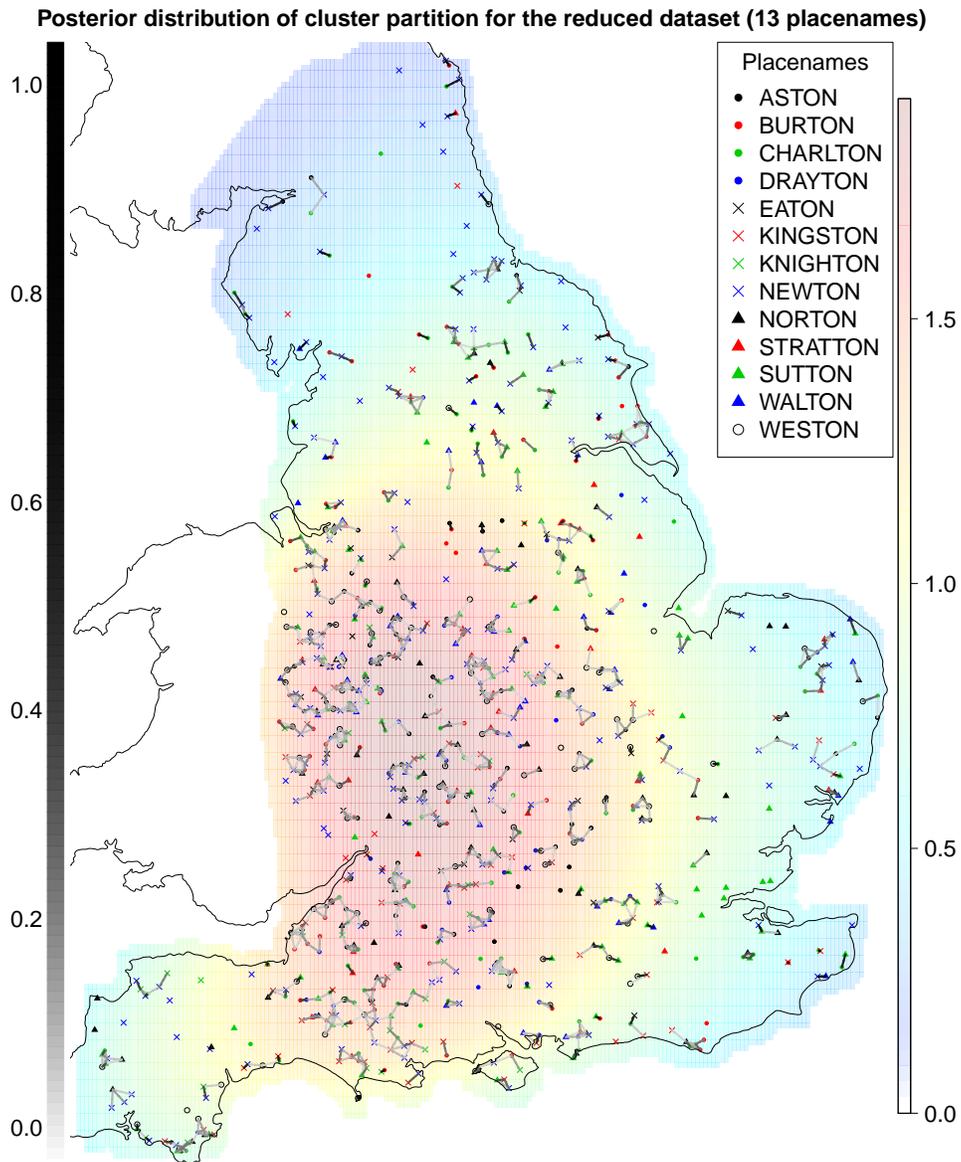


FIG 12. Graphical representation of $\pi(\rho|\mathbf{x})$, where \mathbf{x} is the reduced dataset (13 placenames) in the whole observed region. The intensity of gray corresponds to the estimated posterior probability of the cluster. The truncated kernel density estimation of g is plotted in the background, with values express in relative terms with respect to the uniform measure.

From the previous results it can be seen that the Poisson and the Dirichlet-Multinomial model lead to almost equivalent posterior distributions (e.g.

Figures 10 and 11). The models differ only in the way the prior distribution of the partition ρ is built, so this appears to indicate insensitivity to the specification of this component of the prior. Moreover the posterior did not seem to be much sensitive to the values of the hyperparameters of σ , λ , $\mathbf{p}^{(c)}$ and \mathbf{p} .

Figure 11(c) represents a measure of association between placenames. Given two placenames, say a and b , the measure is defined as

$$(5.1) \quad \frac{\Pr[A|B]}{\Pr[A]} = \frac{\Pr[A \cap B]}{\Pr[A] \cdot \Pr[B]} = \frac{\Pr[B|A]}{\Pr[B]},$$

where A and B are the events of observing placename a and b respectively in a cluster chosen uniformly at random from the clusters of ρ , with ρ distributed according to $\pi(\rho|\mathbf{x})$. In Figure 11(c) we plot the value of (5.1) in relative terms with respect to a null hypothesis. In the null hypothesis we first choose a cluster from ρ as before and then, denoting the number of settlements in the cluster by s , we sample s placenames independently of each other with placename probabilities proportional to their numerosity in the dataset, conditioning on having pairwise different placenames. The expected values of interest under the null distribution have been estimated using standard Monte Carlo methods. High values in Figure 11(c) suggest positive interaction between placenames, while low values suggest negative interaction. Most of the positive associations suggested by Figure 11(c), such as *Knighton-Burton*, *Eaton-Drayton* or *Knighton-Kingston*, are coherent with the current historians hypothesis. We note that, for a fixed ρ , the measure in (5.1) reduces to the *coefficient of association* used by ecologists to measure association between species (Dice, 1945). Many different measures of association have been proposed in the ecological literature (see e.g. Janson and Vegelius, 1981). We chose (5.1) because it is symmetric, clearly interpretable and our experiments suggest that (5.1) is not much influenced by the numerosity of placenames a or b , unlike most measures proposed in Janson and Vegelius (1981).

In order to obtain the results presented in this section, the MCMC algorithm of Section 4.2 was run for 10^6 steps, where at each step 100 moves of the two-color configuration $(\mathbf{x}^{2D}, \rho^{2D})$ were proposed. We assessed convergence using the methods described in Section 4.1.2 (e.g. the value of D in (4.2) was approximately 0.02). The time needed for such runs using a basic R implementation (available in Supplement F) on a desktop computer with an Intel *i-7* processor is approximately 40 hours when using the Poisson model and 160 hours when using the Dirichlet-Multinomial model (in the latter case the posterior of ρ does not factorize and so more computation is needed at each MH step).

6. Discussion. We have designed a Random Partition Model (RPM) that is able to capture the clustering behaviour expected by the historians who produced this dataset. With no strong prior information, the model produces estimates that are meaningful for the historical context and in accordance with contextual information (e.g. see the posterior distribution of σ and the association between placenames in Figure 10). We also defined two flexible prior distributions for clusters partition that are designed for a “small clusters” framework (where each cluster has at most k points with k small). In doing so we developed a RPM to perform complementary clustering which is applicable to other contexts where one needs to find aggregations of elements of different types. For example Professor Susan Holmes from Stanford University suggests that, in biological contexts, species living in the same geographical area assemble by dissimilarity as they fill different ecological niches, resulting in clusters of complementary species.

We carefully considered the computational aspects of this problem. After considering related problems in the complexity theory literature (see Section 3) we employed the Metropolis-Hastings (MH) algorithm. We proposed a choice of MH proposal distributions that, compared to the usual choices found in the literature, achieves a significantly better mixing by approximating detailed balance conditions (see Section 4.1). We developed a multiple proposal scheme to allow for parallel computation that could be relevant for applications to bigger datasets (see Section 4.1.1). Regarding convergence diagnostic we note that, when monitoring the convergence of the MCMC in the partition space, univariate summary statistics appear to be not sufficiently informative to be used as a basis for convergence diagnostics. Diagnostics based on multivariate summary statistics or on the matrix of the estimated association probabilities seem to give more robust results (see Section 4.1.2).

Although the proposed model manages to capture the pattern we were looking for, there is much room for improvement. For example, a direction for future work could be to extend the model in order to capture the heterogeneity in the clustering behaviour between high and low-density regions (see Section 5). One could try to do this by allowing the parameters $\mathbf{p}^{(c)}$ (or \mathbf{p}) and σ to vary over different regions, maybe as a function of the points density, while taking care not to over-parametrize the problem (the amount of data is limited). One could also try to model the dispersion of the settlements in the same cluster with a non-Gaussian distribution having heavier tails.

Another extension that could result in a better fit is to relax the exchangeability assumption on the distribution of placenames. In fact in our model the probability of choosing a certain placename does not depend on the lo-

cation, while the data suggest that different placenames are more likely to be chosen in different regions.

The context suggests that we are observing a thinned version of the original settlements distribution. Nevertheless it is not obvious how to incorporate missing data in this model without making further assumptions that do not seem realistic and are not supported by the historical informations available (e.g. that in each cluster there is a settlement for each type).

An interesting direction for future work is to try to incorporate other sources of data in the model. For example topographical information seem to be related to the settlements clustering (e.g. historians think that settlements named Burton are related to good viewpoints) it would be interesting to find an efficient way to incorporate them in the model.

Acknowledgments. Prof. Wilfrid Kendall for PhD supervision. Prof. John Blair for collaboration and arranging supply of data. EPSRC for funding through the CRiSM grant EP/D002060/1.

SUPPLEMENTARY MATERIAL

Supplement A: Derivation of likelihood expression
(http://www2.warwick.ac.uk/gzanella/compclust_suppa.pdf).

Supplement B: Derivation of the proposal distribution in (P4)
(http://www2.warwick.ac.uk/gzanella/compclust_suppb.pdf).

Supplement C: Multiple Proposal Scheme
(http://www2.warwick.ac.uk/gzanella/compclust_suppc.pdf).

Supplement D: Correctness of the k-dimensional algorithm
(http://www2.warwick.ac.uk/gzanella/compclust_suppd.pdf).

Supplement E: Additional tables and plots on the data analysis
(http://www2.warwick.ac.uk/gzanella/compclust_suppe.pdf).

Supplement F: Anglo-Saxon settlements dataset and R codes
(http://www2.warwick.ac.uk/gzanella/compclust_suppf.zip).

References.

- BADDELEY, A. (2010). Multivariate and marked point processes. *Handbook of spatial statistics* 371–402.
- BADDELEY, A., MOLLER, J. and WAAGEPETERSEN, R. (2000). Non- and semi-parametric estimation of interaction in inhomogeneous point patterns. *Statistica Neerlandica* **54** 329–350.
- BADDELEY, A. and TURNER, R. (2005). Spatstat: an R package for analyzing spatial point patterns. *Journal of Statistical Software* **12** 1–42. ISSN 1548-7660.

- BADDELEY, A. and VAN LIESHOUT, M. (1995). Area-interaction point processes. *Annals of the Institute of Statistical Mathematics* **47** 601–619.
- BALAS, E. and SALTZMAN, M. (1991). An algorithm for the three-index assignment problem. *Operations Research* **39** 150–161.
- BANDELT, H., CRAMA, Y. and SPIEKSMAN, F. (1994). Approximation algorithms for multi-dimensional assignment problems with decomposable costs. *Discrete Applied Mathematics* **49** 25–50.
- BERGE, C. and MINIEKA, E. (1973). *Graphs and hypergraphs*. Amsterdam: North-Holland publishing company.
- BROOKS, S. and GELMAN, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics* **7** 434–455.
- BROOKS, S. and ROBERTS, G. (1998). Assessing convergence of Markov chain Monte Carlo algorithms. *Statistics and Computing* **8** 319–335.
- CHIU, S., STOYAN, D., KENDALL, W. and MECKE, J. (2013). *Stochastic geometry and its applications*. John Wiley & Sons.
- COWLES, M. and CARLIN, B. (1996). Markov chain Monte Carlo convergence diagnostics: a comparative review. *Journal of the American Statistical Association* **91** 883–904.
- CRAMA, Y. and SPIEKSMAN, F. (1992). Approximation algorithms for three-dimensional assignment problems with triangle inequalities. *European Journal of Operational Research* **60** 273–279.
- DE LA VEGA, W., KARPINSKI, M., KENYON, C. and RABANI, Y. (2003). Approximation Schemes for Clustering Problems in Finite Metrics and High Dimensional Spaces. *Proceedings on the thirty-fifth annual ACM symposium on Theory of Computing* 50–58.
- DELLAERT, F., SEITZ, S. M., THORPE, C. E. and THRUN, S. (2003). EM, MCMC, and chain flipping for structure from motion with unknown correspondence. *Machine Learning* **50** 45–71.
- DICE, L. R. (1945). Measures of the amount of ecologic association between species. *Ecology* **26** 297–302.
- DIGGLE, P. J. (2003). *Statistical analysis of spatial point patterns*. Edward Arnold.
- DIGGLE, P., EGLEN, S. and TROY, J. (2006). Modelling the bivariate spatial distribution of amacrine cells. In *Case Studies in Spatial Point Process Modeling* 215–233. Springer.
- GELLING, M. and COLE, A. (2000). *The landscape of place-names*. Shaun Tyas.
- GELMAN, A. and RUBIN, D. (1992). Inference from Iterative Simulation using Multiple Sequences. *Statistical Science* **4** 457–511.
- GEYER, C. and THOMPSON, E. (1995). Annealing Markov chain Monte Carlo with applications to ancestral inference. *Journal of the American Statistical Association* **90** 909–920.
- GRABARNIK, P., MYLLYMÄKI, M. and STOYAN, D. (2011). Correct testing of mark independence for marked point patterns. *Ecological Modelling* **222** 3888–3894.
- HUBER, M. and LAW, J. (2012). Simulation reduction of the Ising model to general matchings. *Electronic Journal of Probability* **17** 1–15.
- JANSON, S. and VEGELIUS, J. (1981). Measures of ecological association. *Oecologia* **49** 371–376.
- JERRUM, M. (2003). *Counting, Sampling and Integrating: Algorithms and Complexity. Lectures in Mathematics ETH Zürich*. Birkhäuser Verlag, Basel.
- JERRUM, M. and SINCLAIR, A. (1996). The Markov chain Monte Carlo method: an approach to approximate counting and integration. *Approximation algorithms for NP-hard problems* 482–520.
- JONES, R. and SEMPLE, S. (2012). *Sense of Place in Anglo-Saxon England*. Shaun Tyas.
- KARPINSKI, M., RUCINSKI, A. and SZYMANSKA, E. (2012). Approximate Counting of

- Matchings in Sparse Uniform Hypergraphs. *arXiv preprint arXiv:1204.5335* 1–13.
- KUHN, H. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* **2** 83–97.
- LAU, J. and GREEN, P. (2007). Bayesian model-based clustering procedures. *Journal of Computational and Graphical Statistics* **16** 526–558.
- LAWSON, A. and DENISON, D. (2010). *Spatial cluster modelling*. CRC press.
- LOIZEAUX, M. and MCKEAGUE, I. (2001). Perfect sampling for posterior landmark distributions with an application to the detection of disease clusters Bayesian cluster models. *Selected Proceedings of the Symposium on Inference for Stochastic Processes* **37** 321–332.
- MARINARI, E. and PARISI, G. (1992). Simulated tempering: a new Monte Carlo scheme. *EPL (Europhysics Letters)* **c** 1–12.
- MÜLLER, P. and QUINTANA, F. (2010). Random partition models with regression on covariates. *Journal of Statistical Planning and Inference* **140** 2801–2808.
- OH, S., RUSSELL, S. and SASTRY, S. (2009). Markov chain Monte Carlo data association for multi-target tracking. *Automatic Control, IEEE Transactions on* **54** 481–497.
- PLUMMER, M., BEST, N., COWLES, K. and VINES, K. (2005). Output analysis and diagnostics for MCMC. *R package version 0.10-3*, URL <http://cran.rproject.org>.
- ROBERTS, G. (1998). Optimal Metropolis algorithms for product measures on the vertices of a hypercube. *Stochastics and Stochastic Reports* **62** 275–283.
- ROBERTS, G. O., GELMAN, A. and GILKS, W. R. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability* **7** 110–120.
- VALIANT, L. (1979). The complexity of enumeration and reliability problems. *SIAM Journal on Computing* **8** 410–421.
- VILLANI, C. (2009). *Optimal transport: old and new*. Springer New York.

DEPARTMENT OF STATISTICS
 UNIVERSITY OF WARWICK
 COVENTRY, CV4 7AL
 UNITED KINGDOM
 E-MAIL: g.zanella@warwick.ac.uk