

The Coordinate-Based Meta-Analysis of Neuroimaging Data

Pantelis Samartidis, Silvia Montagna, Timothy D. Johnson and Thomas E. Nichols

Abstract. Neuroimaging meta-analysis is an area of growing interest in statistics. The special characteristics of neuroimaging data render classical meta-analysis methods inapplicable and therefore new methods have been developed. We review existing methodologies, explaining the benefits and drawbacks of each. A demonstration on a real dataset of emotion studies is included. We discuss some still-open problems in the field to highlight the need for future research.

Key words and phrases: Meta-analysis, neuroimaging, functional magnetic resonance imaging.

1. INTRODUCTION

Functional magnetic resonance imaging (fMRI) has experienced a rapid growth over the past two decades and has led to significant advances in our understanding of the human brain, including the differences in brain function between maternal and romantic love (Bartels and Zeki, 2004), the effect of alcohol while performing simulated driving (Calhoun and Pearlson, 2012), or the effect of doing nothing at all (Cole, Smith and Beckmann, 2010). The availability of MRI scanners, inexpensive computational resources and accessible analysis software has made fMRI an ubiquitous tool in psychology, neurology and psychiatry, in addition to new areas like neuromarketing and neuroeconomics.

Nevertheless, there are a variety of factors that limit the interpretability of fMRI results. The principal lim-

itation is the small sample sizes typically used, leading to individual studies that suffer from low power and hence low reproducibility (Button et al., 2013). Some other concerns include high prevalence of false positives (Wager, Lindquist and Kaplan, 2007), poor reproducibility (Raemaekers et al., 2007) and considerable heterogeneity in the analysis pipeline (Carp, 2012). As a result, it is unsurprising that the validity of fMRI is being challenged in both the scientific (Vul et al., 2009) and popular (Shermer, 2008) literature (see Farah, 2014 for an even-handed review).

Meta-analysis provides a way to address all of these limitations. Meta-analysis is the process of combining the results of independently conducted studies to increase power and obtain more reproducible findings than the original studies (Hedges and Olkin, 1985). Meta-analysis of functional neuroimaging data is an active field of research, whose growth is facilitated by the constantly increasing body of literature in fMRI along with the limitations of single experiments. The goal of this paper is to review recent advances in fMRI meta-analysis, evaluate existing methods, and highlight the open problems that need to be investigated. We consider a meta-analytic dataset analysed in a number of ways and make the data freely available to aid other researchers trying a hand at this intriguing area.

The remainder of this manuscript is organised as follows. In Section 2, we provide some background on neuroimaging and explain the special characteristics of neuroimaging studies that make meta-analysis in this setting challenging. In Section 3 we give a detailed

Pantelis Samartidis is Investigator Statistician, MRC Biostatistics Unit, University of Cambridge, Robinson Way, Cambridge, CB2 0SR, United Kingdom (e-mail: pantelis.samartidis@mrc-bsu.cam.ac.uk). Silvia Montagna is Lecturer, School of Mathematics, Statistics and Actuarial Science, University of Kent, Canterbury, CT2 7FS, United Kingdom (e-mail: s.montagna@kent.ac.uk). Timothy D. Johnson is Professor, Biostatistics Department, University of Michigan, Ann Arbor, Michigan 48109, USA (e-mail: tdjtdj@umich.edu). Thomas E. Nichols is Professor, Oxford Big Data Institute, Li Ka Shing Centre for Health Information and Discovery, Nuffield Department of Population Health, University of Oxford, OX3 7LF, United Kingdom (e-mail: thomas.nichols@bdi.ox.ac.uk).

description of the currently most popular methods for fMRI meta-analysis. In Section 4, we proceed to an evaluation of existing methods, which is motivated by application to a real dataset of emotion studies. Finally, we discuss some possible directions for research in Section 5.

2. NEUROIMAGING BACKGROUND

What follows is a very brief review of fMRI and the practical steps involved in a fMRI study. For a more detailed introduction, see Lindquist (2008) for a review of fMRI for statisticians, or Kim and Ogawa (2012) for a detailed, technical review of the meaning of the fMRI signal; Huettel, Song and McCarthy (2009) provide an accessible textbook treatment, while Poldrack, Mumford and Nichols (2011) give a practical, data-analysis-oriented perspective.

The objective of a single fMRI study is to identify the neural correlates of a physical, mental or perceptual process. When neurons in a region of the brain increase their firing rate, there is an increased demand for oxygen which is met by a localised increase in blood flow. The magnetic resonance signature, or susceptibility, of oxygenated and de-oxygenated blood differs, and thus a MRI scanner can capture changes in local oxygenation. This mechanism is known as the *blood oxygenation level-dependent* (BOLD) effect.

During an fMRI acquisition, participants lie flat in the scanner and are asked to perform a series of tasks, such as viewing images or reading texts, while the MRI scanner measures the BOLD signal. For each participant, the data takes the form of a time series of images, 3D snapshots of signal measurements all over the brain. The typical acquisition lasts 6–12 minutes, with data collected every 2 seconds, producing data on a grid with 2 mm × 2 mm spacing in-plane and 2 mm–4 mm slices, producing anywhere from 40,000 to over

100,000 voxels (volume elements) in the brain. Note that this is quite coarse spatial resolution, and separate, fine-resolution images (e.g., 1 mm × 1 mm × 1 mm) are also taken to depict individual's anatomy.

Before the raw data can be analysed, a series of preprocessing steps needs to be undertaken. These include motion correction, which accounts for movements during the acquisition, and spatial smoothing which increases the signal-to-noise ratio. To make data comparable across subjects, a crucial step is “spatial normalisation”, the process of warping all subjects to a standard brain template, or brain atlas. There are different atlases available, but essentially all authors use either the Talairach atlas (Talairach and Tournoux, 1988) or the MNI atlas (see Figure 1).

After spatial normalisation, all subjects' data exist in a common space. Specifically, we can assume that a given voxel corresponds to (roughly) the same region in all subjects' brains. Statistical analysis then proceeds in a mass-univariate approach, fitting a model at each voxel independently of every other voxel. For every subject, time series regression models are fitted in each voxel, where the regression coefficients represent the effect of the different tasks. Task effects or comparisons can be made with the proper “contrasts”, meaning the estimated linear combination of parameter estimates that relates to the effect of interest. These subject-specific contrasts are subsequently analysed in a “second level” population model. The result is a 3D image of T statistics, one for each voxel in the brain, measuring the evidence against the null hypothesis of no effect. The T images are assessed either voxel-by-voxel, or by assessing the size of connected components, or *clusters*, after thresholding the T image at an arbitrary threshold. See Friston et al. (2002), Mumford and Nichols (2006) and Mumford and Nichols (2009) for a detailed review of different approaches for the statistical analysis of fMRI data.




FIG. 1. An average brain in MNI space. Note the directional labels at the edge of each panel: P for Posterior, A for Anterior, S for Superior, I for Inferior, L for Left and R for Right. The origin approximately corresponds to an anatomical structure known as the anterior commissure.

An essential issue in the statistical analysis of fMRI data is *multiple testing*. A T statistic image can have 100,000 or more voxels in the brain, requiring 100,000 simultaneous tests for every contrast of interest. Under a global null hypothesis of no effect in any voxel, we therefore expect around 5000 false positives using the classic significance level of $\alpha = 0.05$. In the early history of fMRI (roughly 1992–2002), arbitrary rule-of-thumb thresholding procedures were common, like a combination of an uncorrected voxel-wise $\alpha = 0.001$ and cluster size threshold $k \geq 10$ (only clusters of size 10 voxels or more). Thresholding methods that controlled the *familywise error* (FWE), the chance of one or more false positives, later became widespread using either random field theory or permutation testing (see Nichols and Hayasaka, 2003 for a review of FWE methods in neuroimaging). More recently, the *false discovery rate* (FDR) was introduced for the thresholding of T images (Genovese, Lazar and Nichols, 2002). The FDR is the expected proportion of false positives among positive findings (Benjamini and Hochberg, 1995).

In any discipline of science, discussion of findings typically includes the point estimates for an effect, the associated standard errors, and the p -values. In neuroimaging, each of these quantities is a 3D image, and sharing such large data files was considered impractical 20 years ago when fMRI was first developed. Yet even today there is general resistance towards sharing the full images. Instead, authors routinely report the x , y , z atlas coordinates of activation peaks. Going forward we will call these coordinates the *foci* (singular *focus*). In other words, the results of an fMRI study are summarised in a list of foci. Based on author preference and software defaults, foci can either be *singletons* that is one focus per significant region or *multiple* that is two or more per significant region.

2.1 Limitations of Individual Studies and Meta-Analysis

There are three aspects of fMRI experiments that challenge the utility of individual studies. First, fMRI studies suffer from low power. The typical sample size of an fMRI study is small, and the majority of experiments involves far less than 20 participants (Carp, 2012). While power depends on the (unknown) true effect size, at least one empirical study supported the notion that fMRI n 's are too small. By sub-sampling from a large sample ($n = 150$), Thirion et al. (2007) found that analyses with 20 or fewer subjects were poor approximations of the full 150-subject result. Further,

Type I error rates are likely to be high, especially for older papers that did not use inference procedures corrected for multiple testing. Using a survey of publications' thresholding methods, Wager, Lindquist and Kaplan (2007) estimated that 17% of all reported foci are false positives. Finally, neuroimaging studies suffer from low test-retest reliability. For example, when scanning a group of subjects twice, once and then 7 days later, Raemaekers et al. (2007) found intra-class correlations for BOLD fMRI activations ranged from 0 to 0.88.

Apart from these inherent limitations, the way fMRI studies are carried out also exhibits great heterogeneity. Each step of a neuroimaging study can be implemented in various ways and there is no standard way to present a stimulus, preprocess the data or construct the linear model for the BOLD response. As a result, there is only a partial agreement in how experiments are conducted. For example, in an analysis of 241 fMRI studies Carp (2012) observed 223 different analytical strategies. Results heavily depend on the type of analysis employed (Button et al., 2013), thus it is not uncommon to observe discrepancies in the outcomes of studies that investigate the same scientific question. Consequently, it is exceptionally hard to yield a conclusion. All reasons combined support the use of meta-analysis to account for these problems and draw more reliable inferences.

A well-performed meta-analysis can tackle the aforementioned issues by modelling the observed heterogeneity between studies, combining the available information to increase power and ultimately separating the consistent findings from those that happened by chance. There exist two broad approaches for meta-analysis of neuroimaging studies: *image-based meta-analysis* (IBMA), if the full T statistic images are available, and *coordinate-based meta-analysis* (CBMA) if only foci are reported. IBMA proceeds by means of some common meta-analytic tools applied to each voxel of the images along with either FWE or FDR corrections for multiple testing. See Hartung, Knapp and Sinha (2008) for an overview of conventional meta-analysis, and Lazar et al. (2002) for a review of IBMA methods.

Given the richness of brain imaging data, with gigabytes collected on each subject, it is lamentable that all this data gets distilled to a list of xyz coordinates (foci). In a comparative study, Salimi-Khorshidi et al. (2009) demonstrated the benefits of using IBMA over CBMA, and observed a dramatic loss of information when only foci were used relative to utilising full image data. Their results highlight the need for sharing

full image data from single experiments (namely, effect magnitude estimates and standard errors, as well as key analysis details), and have encouraged the development of online repositories such as *OpenfMRI*¹ (Poldrack et al., 2013) and *NeuroVault*² (Gorgolewski et al., 2015) for sharing entire statistical maps. But sharing of full data is still rare, the use of such tools is not yet a standard practice and it is still not possible to access image data results for the majority of published research (Poline et al., 2012). On the contrary, coordinate data are easily available through large databases such as *BrainMap*³ (Laird, Lancaster and Fox, 2005) and *NeuroSynth*⁴ (Yarkoni et al., 2011). These repositories include on-line tools for quick data extraction and on-line implementation of CBMA meta-analyses. As a result, CBMA still constitutes the main approach for the meta-analysis of fMRI data, and we will focus solely on a review of CBMA methods hereafter.

3. CBMA METHODS

The limitations of single experiments (see Section 2.1 for a discussion), along with the historical lack of data sharing, quickly presented researchers in the field of fMRI with a challenge. The standard meta-analytic tools used in other fields (see, for example, Hartung, Knapp and Sinha, 2008, for a fairly recent review) could not be applied to the coordinate data and hence there was a need for new methodologies. Early works mainly utilised exploratory data analysis and visualisation techniques to blend the results from different studies (Fox, Parsons and Lancaster, 1998) and it was not until the early 2000s that the first methods for CBMA were proposed (Fox et al., 1997, Turkeltaub et al., 2002, Nielsen and Hansen, 2002, Wager et al., 2003). Since then, many new methods and modifications appeared in the neuroimaging (Laird et al., 2005, Wager, Lindquist and Kaplan, 2007, Radua and Mataix-Cols, 2009, Turkeltaub et al., 2012, Caspers et al., 2014, to name a few) as well as the statistics (Kang et al., 2011, Yue, Lindquist and Loh, 2012, Kang et al., 2014, Montagna et al., 2017) literature.

All of these methods share the same goal: to **identify areas of the human brain that show consistent activation across studies. The different approaches broadly fall into two main categories: kernel-based and model-based methods.** In what follows, we present the most

TABLE 1

A subset of data from a meta-analysis study of emotions

Author	Year	Emotion	X	Y	Z	Participants
Damasio	2000	fear	-10	-62	-17	23
			-1	-66	-1	23
			34	3	32	23
Damasio	2000	anger	-2	-29	-12	23
Philips	2004	disgust	4	-20	15	8
			7	-17	9	8
			4	-63	26	8
Baker	1997	sad	36	20	-8	11
			-44	32	-8	11
Baker	1997	happy	-26	28	0	11
			-34	52	8	11
Williams	2005	anger	7	31	28	13
			7	28	-7	13
...

widely used methods in both categories. We start by setting the notation used throughout the manuscript.

A typical CBMA dataset consists of a list of foci from I independent studies. Each study i , $i = 1, \dots, I$ comes with a set of 3-dimensional coordinates $\mathbf{x}_{ik} \in \mathcal{B}$, where $\mathcal{B} \subset \mathbb{R}^3$ is the standard atlas space and k indexes the multiple foci for a particular study. Table 1 is part of a real dataset from a meta-analysis of emotion studies that will be analysed for the purposes of this review. In this example, \mathbf{x}_{52} would correspond to the second foci ($[-34, 52, 8]^\top$) in the fifth study (Baker 1997, happy). We denote $\mathbf{x}_i = \bigcup_k \mathbf{x}_{ik}$, the complete set of foci reported in study i . Note that some of the studies (e.g., Damasio 2000, fear and Damasio 2000, anger) are obtained from the same experiment; we treat these studies as independent following the standard conventions in the field. Finally, we will denote as $v = [v_x, v_y, v_z] \in \mathcal{B}$ the center location of a particular voxel in the brain atlas, $v = 1, \dots, V$.

3.1 Kernel-Based Methods

The most widely used kernel based methods are the *multilevel kernel density analysis* (Wager, Lindquist and Kaplan, 2007, MKDA), the *activation likelihood estimation* (Eickhoff et al., 2012, ALE) and the *signed differential mapping* (Radua et al., 2012, SDM). All these methods share the same rationale. Briefly, **one starts by creating focus maps: that is, full brain images obtained through smoothing of reported activations with a spatial kernel. Obviously, there are as many focus maps as the total number of foci. Secondly, the focus maps corresponding to a particular study are**

¹<https://openfmri.org> (RRID:SCR_005031)

²<http://neurovault.org> (RRID:SCR_003806)

³<http://www.brainmap.org> (RRID:SCR_003069)

⁴<http://www.neurosynth.org> (RRID:SCR_006798)

combined to create the study-specific maps. These per-study images are subsequently combined into a single image that represents the evidence for consistent activation (clustering). Significance of these images is assessed with a Monte Carlo test under the null hypothesis of complete spatial randomness. We now discuss MKDA, ALE and SDM in detail.

3.1.1 Multilevel kernel density analysis. First, introduced by Wager et al. (2003), MKDA was modified to its current version by Wager, Lindquist and Kaplan (2007). To obtain the focus maps, M_{ik} , one places a sphere of unit intensity and radius r centred at each focus:

$$(3.1) \quad M_{ik}(v) = \mathbf{1}_{\{d(v, \mathbf{x}_{ik}) \leq r\}},$$

where $d(\cdot, \cdot)$ stands for the Euclidian distance. The study specific images, M_i , are then obtained by applying the maximum operator to the focus maps of the study. The procedure can be expressed by the following formula:

$$(3.2) \quad M_i(v) = \begin{cases} 1, & \exists k \text{ s.t. } d(v, \mathbf{x}_{ik}) \leq r, \\ 0, & \text{otherwise.} \end{cases}$$

We call M_i the *comparison indicator maps*. A value of 1 means that there is activation within distance r of a given location. Wager, Jonides and Reading (2004) suggest giving r a value of 10 or 15 mm. The MKDA statistic image m is given as a weighted combination of M_i :

$$(3.3) \quad m(v) = \frac{1}{\sum_i w_i} \sum_{i=1}^I w_i M_i(v).$$

The weights are usually chosen to be proportional to the number of participants in each study thus allowing for studies with larger sample size to contribute more to the value of the statistic. If the weights are all set to 1 then $m(v)$ denotes the proportion of studies that reported activation within distance r to v . Large values of $m(v)$ suggest systematic clustering of foci around its location.

The distribution of the MKDA statistic does not have a closed form and thus Monte Carlo testing is used to assess significance. Multiple synthetic datasets are created by uniformly drawing peak locations from \mathcal{B} , keeping the original number of foci per study fixed. The m statistic map is calculated for these datasets and the maximum value is saved at each replicate. This produces a sample of the maximal statistic under the null hypothesis of random foci allocation.

The sample is then used to obtain FWE-corrected p -values (Kober et al., 2008) as suggested by Nichols and Holmes (2002). Recently, Costafreda, David and Brammer (2009) derived a parametric significance test based on the properties of the spatial Poisson process. For applications of MKDA on real data see Etkin and Wager (2007) and Kober et al. (2008).

3.1.2 Activation likelihood estimation. The idea behind ALE is to model the probability that a voxel v is the true location of a reported focus \mathbf{x}_{ik} as a discretised (over voxels) and truncated (outside the brain) Gaussian distribution centered around the reported location (Turkeltaub et al., 2002). For a certain study i , let L_{ik} be the map based on a single focus \mathbf{x}_{ik} ,

$$(3.4) \quad L_{ik}(v) = c\phi_3(v \mid \mathbf{x}_{ik}, \sigma_i^2 \mathbf{I}),$$

where $\phi_3(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ is the density of a three dimensional Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ evaluated at $\mathbf{x} \in \mathbb{R}^3$, \mathbf{I} is the identity matrix, and c is a normalising constant ensuring that the sum of $\phi_3(\cdot)$ over voxels is equal to one. Thus, $L_{ik}(v)$ represents the probability that voxel v is the true location of focus \mathbf{x}_{ik} . The Gaussian kernel used for ALE is analogous to the uniform kernel used for MKDA, but assigns higher values to voxels closer to the foci. To determine σ_i , Eickhoff et al. (2009) created a mapping between the number of participants in each study, n_i , and the standard deviation σ_i , the intuition being that larger sample sizes lead to improved spatial precisions. However, the mapping is based on an empirical study consisting of 21 subjects and may be unsuitable for experimental paradigms other than the one used by the authors.

The next step for the ALE algorithm is to combine the focus maps L_{ik} into a single study map L_i . Early versions of the algorithm (Turkeltaub et al., 2002, Eickhoff et al., 2009) constructed study maps as the probability that at least one focus is located at a given voxel v .⁵ The drawback of this approach is that it leads to ALE scores that are heavily influenced by studies reporting several foci per region. To address this issue, Turkeltaub et al. (2012) proposed taking the maximum over focus maps as in MKDA, that is,

$$(3.5) \quad L_i(v) = \max_k L_{ik}(v).$$

We call L_i the *modelled activation map*. $L_i(v)$ quantifies the probability that the focus which is closest to v is truly located at v .

⁵This is achieved by taking $L_i(v) = 1 - \prod_k (1 - L_{ik}(v))$.

The ALE statistic ℓ is then computed as

$$(3.6) \quad \ell(v) = 1 - \prod_{i=1}^I (1 - L_i(v)).$$

Expression 3.6 was originally adopted by Turkeltaub et al. (2002) and represents the probability that at least one of the closest activations is truly located in voxel v .

The Monte Carlo significance test of ALE is equivalent yet slightly different to the one of MKDA. In particular, Eickhoff et al. (2009) maintain that the spatial arrangement of activations (focus-to-focus distances) within a study must be preserved when generating synthetic studies, so that convergence across studies (rather than convergence over individual foci) is assessed. Hence for location v , multiple realisations of the ALE statistic $\ell^*(v)$ are created by sampling each activation map from a random location. The null ALE are then obtained as

$$(3.7) \quad \ell^*(v) = 1 - \prod_i (1 - L_i(v^*)),$$

where v^* is drawn uniformly from all possible brain locations. By using (3.7), Eickhoff et al. (2012) showed that it is possible to enumerate exhaustively all the possible outcomes of ℓ , directly obtaining the exact marginal null distribution of the ALE statistic without Monte Carlo. The null can be used to compute (uncorrected) p -values and from these FDR-corrected p -values (Laird et al., 2005). FWE-corrected voxel-wise and cluster size, however, require the Monte Carlo simulation of null distribution of the maximum of ℓ (Eickhoff et al., 2012). Eickhoff et al. (2016) advocated FWE-corrected voxel-wise inference; they found, using simulated data, that this is less susceptible to “spurious” findings compared to voxel-wise FDR correction.

ALE has been used for several analyses including those in Delvecchio et al. (2012) and Konova, Moeller and Goldstein (2013). ALE’s popularity has been supported by the availability of software like *GingerALE*, a broadly used, freely available implementation of the ALE algorithm. However, a recent contribution by Eickhoff et al. (2017) reported the presence of several errors in the ALE code. In particular, errors were traced in the voxel-wise FDR and cluster-wise FWE correction procedures for multiple testing, and have resulted in thresholds that were more liberal compared to those specified by users. For more details, see Eickhoff et al. (2017). As of April 2017, we remark that these issues have been fixed.

3.1.3 Signed differential mapping. SDM (Radua and Mataix-Cols, 2009) is a relatively new method that borrows several characteristics from both MKDA and ALE. The novelty of the method lies in incorporating the T statistic values (when available) from the original studies. To make this point clear, imagine that a study investigates brain activation caused by a given task; in some regions of the brain hyperactivation will be observed while in others there will be underactivation. In both cases, significant values of the T statistic will be recorded; these values will be large and positive in the first case and large and negative values in the second. This case is particularly interesting when difference in activation between tasks is being investigated.

Assume that T_{ik} is the reported T value for focus \mathbf{x}_{ik} . SDM will generate focus maps S_{ik} as

$$(3.8) \quad S_{ik}(v) = \text{sign}(T_{ik}) \exp\left(-\frac{d(v, \mathbf{x}_{ik})^2}{2\sigma^2}\right).$$

Similarly to ALE, SDM employs a Gaussian kernel, and the authors suggest using a standard deviation of approximately 25 mm (Radua and Mataix-Cols, 2009). The study maps are then

$$(3.9) \quad S_i(v) = \begin{cases} -1, & \sum_k S_{ik}(v) \leq -1, \\ \sum_k S_{ik}(v), & -1 \leq \sum_k S_{ik}(v) \leq 1, \\ 1, & \sum_k S_{ik}(v) \geq 1. \end{cases}$$

That is, the study map is obtained as the sum of the corresponding focus maps, but is forced to lie within the interval $[-1, 1]$ in the same way the MKDA study maps M_i are given a maximum value of 1. Finally, the SDM statistic image, s , is calculated as the weighted mean of the study specific maps at each voxel:

$$(3.10) \quad s(v) = \frac{1}{\sum_i w_i} \sum_i w_i S_i(v).$$

Weights are once again proportional to number of participants in the study. Since the method averages both positive and negative findings, voxels that show contradicting results will not appear as significant. Inference is based on the same Monte Carlo scheme of MKDA, and thresholding is done either by setting a highly conservative rejection point ($p < 0.001$) or controlling the FDR (Radua and Mataix-Cols, 2009).

In a newer version of the algorithm, the authors use the T_{ik} values to reconstruct the original T statistic images. That way, it is possible to incorporate both CBMA and IBMA data in the same analysis. For more

details, see Radua et al. (2012). The last contribution made on SDM lies in the use of anisotropic kernels in the analysis (Radua et al., 2014). Anisotropy can be easily incorporated in MKDA and ALE but its superiority to the current practice of using isotropic kernels is only based on empirical findings, thus it should be further investigated. Published work utilising SDM for the analyses includes Richlan, Kronbichler and Wimmer (2011) and Fusar-Poli (2012).

3.2 Model-Based Methods

Recently, there has been growing interest in the development of model-based methodologies to address some of the limitations of kernel-based methods. These methods use ideas from spatial statistics to develop stochastic models for the analysis of foci. Unfortunately the literature on model based methods is still very limited thus our review will be almost exhaustive. In particular, we will outline the *Bayesian hierarchical cluster process* model of Kang et al. (2011), BHICP, the *spatial binary regression* model of Yue, Lindquist and Loh (2012), the *hierarchical Poisson/Gamma random field* model of Kang et al. (2014), and the *spatial Bayesian latent factor regression* model of Montagna et al. (2017). In all cases analyses are performed under the Bayesian paradigm and thus inferences are based on posterior distributions for each model’s parameters.

Some of the methods reviewed here are build upon spatial point processes theory. Spatial point processes are random sets of points in the d -dimensional Euclidian space. A detailed description of the theory behind point processes is beyond the scope of this review, so we refer the reader to Møller and Waagepetersen (2004) and Illian et al. (2008) for details and applications. We now proceed to describe the details of model-based methods.

3.2.1 *A Bayesian hierarchical independent cluster process model (BHICP)*. Kang et al. (2011) proposed a hierarchical model based on an independent cluster process to describe the mechanism generating the foci. The model is structured into 3 levels, of which the lowest level, level 1, contains the observations (foci), while higher levels describe the study and population structure respectively. The distinction between singletons and multiple foci is incorporated into the model. In the outline of the model below, we occasionally suppress the k index so that \mathbf{x}_i is the set of foci reported in study i , that is, $\mathbf{x}_i = \bigcup_k \mathbf{x}_{ik}$.

Figure 2 provides a graphical representation of the model. At level 1, we have the foci (Figure 2 bottom,

coloured circles). We denote with \mathbf{X}_i the underlying process generating the observations \mathbf{x}_i in each study. As discussed in Section 2, we can have both singleton and multiple foci. Thus, process \mathbf{X}_i consists of two mechanisms, one generating the multiple foci (Figure 2, red circles) and one that is giving the singleton foci (Figure 2 bottom, green circles): $\mathbf{X}_i = \mathbf{X}_i^1 \cup \mathbf{X}_i^0$.

Multiple foci \mathbf{X}_i^1 can be viewed as an independent cluster process of points centered around study activation centers \mathbf{y}_i . In particular, for every study center $\psi \in \mathbf{y}_i$ we have a process $\mathbf{X}_{i\psi}^1$ of multiple foci, where total number of offsprings has a $\text{Pois}(\eta_i)$ distribution. Conditional on the total number of offsprings, all the points $\chi \in \mathbf{X}_{i\psi}^1$ are normally distributed around ψ with covariance \mathbf{T}_ψ , that is $\chi \sim \mathcal{N}(\psi, \mathbf{T}_\psi)$, for all $\chi \in \mathbf{X}_{i\psi}^1$. In other words, each $\mathbf{X}_{i\psi}^1$ is a Poisson process defined on the brain with intensity given by

$$(3.11) \quad \lambda_{i\psi}^1(\xi) = \eta_i \phi_3(\xi | \psi, \mathbf{T}_\psi), \quad \xi \in \mathcal{B}.$$

The overall observed pattern of multiple foci is then given as the union of all offsprings, that is, $\mathbf{X}_i^1 = \bigcup_{\psi \in \mathbf{y}_i} \mathbf{X}_{i\psi}^1$, and therefore is also a Poisson process with intensity

$$(3.12) \quad \lambda_i^1(\xi) = \sum_{\psi} \lambda_{i\psi}^1(\xi).$$

Singleton foci \mathbf{X}_i^0 arise directly from the population center process \mathbf{z} . As with multiple foci, \mathbf{X}_i^0 has a total of $\text{Pois}(\theta_i)$ points which are normally distributed around the population center locations $\zeta \in \mathbf{z}$, that is, $\chi \sim \mathcal{N}(\zeta, \mathbf{\Sigma}_\zeta)$ for all $\chi \in \mathbf{X}_{i\zeta}^0$ and covariance matrix $\mathbf{\Sigma}_\zeta$. To add more flexibility, the model allows for some singletons to not cluster around any population center, say $\mathbf{x}_{i\emptyset}$. These foci are assumed to arise from a Poisson process $\mathbf{X}_{i\emptyset}$ of constant intensity ε_{1i} :

$$(3.13) \quad \mathbf{X}_{i\emptyset} | \varepsilon_{1i} \sim \mathcal{PP}(\mathcal{B}, \varepsilon_{1i}).$$

Overall, $\mathbf{X}_i^0 = (\bigcup_{\zeta \in \mathbf{z}} \mathbf{X}_{i\zeta}^0) \cup \mathbf{X}_{i\emptyset}$ are the singleton foci of a study. Hence, the intensity of \mathbf{X}_i^0 is

$$(3.14) \quad \lambda_i^0(\xi) = \varepsilon_{1i} + \theta_i \sum_{\zeta \in \mathbf{z}} \phi_3(\xi | \zeta, \mathbf{\Sigma}_\zeta), \quad \xi \in \mathcal{B}.$$

While kernel-based methods condition on the total number of activations, in the BHICP model the total number of foci is a Poisson random variable with mean $\int_{\mathcal{B}} \lambda_i(\xi) d\xi$, where $\lambda_i = \lambda_i^0 + \lambda_i^1$.

At level 2, we have the unobserved study activation centers \mathbf{y}_i , which are the locations around which the multiple foci of a study cluster. The \mathbf{y}_i are realisations of a point process \mathbf{Y}_i , and may either cluster around




FIG. 2. Realisation of the BHICP model for 3 studies. At level 3 (top) latent population centres (grey, \mathbf{z}) lie. At level 2 (middle) we have centres of multiple foci (black). These come either directly from population centres (squares, \mathbf{y}_i) or from background noise (triangles, $\mathbf{y}_{i\emptyset}$). Level 1 (bottom) contains the data (\mathbf{x}_i). These are multiple (red, \mathbf{x}_i^1) or singleton (green, \mathbf{x}_i^0) foci. Singletons come either directly from population centres (dots, $\mathbf{x}_{i\xi}^0$) or from a background Poisson process (asterisks, $\mathbf{x}_{i\emptyset}^0$).

the population centres \mathbf{z} (Figure 2 middle, squares) or appear in random locations across the brain (Figure 2 middle, triangles). To account for the former, clustered study centers $\mathbf{Y}_{i\zeta}$ are introduced as sets of points normally distributed around population centers $\zeta \in \mathbf{z}$, and with covariance matrix Σ_ζ . The total number of points in each $\mathbf{Y}_{i\zeta}$, that is, the total number of clustered study centers for study i is a Poisson random variable with mean κ_i . As for the latter, noise study centers are modelled as a homogenous Poisson process $\mathbf{Y}_{i\emptyset}$ with intensity ε_{2i} . Overall, $\mathbf{Y}_i = (\bigcup_{\zeta \in \mathbf{z}} \mathbf{Y}_{i\zeta}) \cup \mathbf{Y}_{i\emptyset}$ is Poisson point process with intensity given by

$$(3.15) \quad \rho_i(\xi) = \varepsilon_{2i} + \kappa_i \sum_{\zeta \in \mathbf{z}} \phi_3(\xi | \zeta, \Sigma_\zeta), \quad \xi \in \mathcal{B}.$$

At the highest level (level 3), we have the population activation centres (Figure 2 top, gray crosses). These are unobserved realisations \mathbf{z} of an *a priori* homogenous Poisson process \mathbf{Z} of intensity ε_3 :

$$(3.16) \quad \mathbf{Z} | \varepsilon_3 \sim \mathcal{PP}(\mathcal{B}, \varepsilon_3).$$

Population centres are the locations around which study activation centers and singleton foci scatter. As such, they can be viewed as locations in the brain where an overall population effect exists.

The BHICP can be viewed as a random effects model as it allows for both within-study and between-study variability. Samples from the posterior distributions are obtained via MCMC. Several interesting quantities can

be inferred upon such as regions of consistent activations (through the posterior distribution of populations centers), the uncertainty in the location of study centers around the population centers (through Σ_ζ) and the variability of the foci within studies (through Ψ_ξ).

3.2.2 A Bayesian spatially adaptive binary regression model. Yue, Lindquist and Loh (2012) use spatial logistic regression for a meta-analysis of emotion studies. For study i and voxel v , let $y_i(v)$ be the binary outcome defined as

$$(3.17) \quad y_i(v) = \begin{cases} 1, & \text{at least one focus at voxel } v, \\ 0, & \text{no foci at voxel } v. \end{cases}$$

Note that the binary study images $\{y_i(v)\}_{v=1}^V$ are identical to the MKDA study maps $M_i(v)$. Logistic regression can be used to model the probability that a voxel is reported as a focus, $p_i(v) = \mathbb{P}(y_i(v) = 1)$. It is assumed that

$$(3.18) \quad p_i(v) = H(z(v)),$$

where $H(\cdot)$ is the link function. The authors use the standard probit and logit link functions.

Spatial correlation is induced through the prior on $\{z(v)\}_{v=1}^V$. In particular, we assume that the process $z(v)$ is an adaptive Gaussian Markov random field (Yue and Speckman, 2010, aGMRF). The aGMRF model defines the conditional distribution of $z(v)$ through a specific dependence with neighbouring voxels. A significant merit of the method is the inclusion of a local smoothness parameter $\gamma(v)$ for the aGMRF. This allows the method to automatically choose the amount of smoothing required depending on the amount of information available.

Authors further introduce a process $\psi(v)$, an indicator of whether the outcome variable $y_i(v)$ is miscoded; the case $\psi(v) = 1$ can either refer to both false positives, voxels that were falsely found as activated, and false negatives, voxels that were not reported as foci even though they were activated. The process $\psi(v)$ is not observed and hence is estimated along with the remaining model parameters.

Posterior probabilities of activation at each voxel are obtained through an auxiliary variable MCMC algorithm. Voxels with high posterior probabilities of being reported as foci are more likely to show an effect. A potential drawback of the method is that it can be currently applied only in two dimensions. In three dimensions, the value of one of the axes is held fixed, for example $z = c$, while the model is fitted for all available observations of the form $\mathbf{x}_{ik} = [x1_{ik}, x2_{ik}, c]$. Authors however, maintain that extending the model to three dimensions is possible.

3.2.3 A hierarchical Poisson/Gamma random field model (HPGRF). A neuroimaging meta-analysis will typically consider several subtypes of tasks. For example, a meta-analysis of emotion may classify the studies according to experiments on “happiness”, “sadness”, “pain”, etc. Yet, the methods described previously are for a single homogeneous group of studies. Kang et al. (2014) propose a model that models each type of foci separately, allowing simultaneously for dependence between the J different types.

Let \mathbf{x}_{ij} be the set of foci reported by study i for task type j . Suppose that \mathbf{x}_{ij} are realisations of a Cox Process \mathbf{X}_j driven by a random intensity measure $\Lambda_j(d\xi)$. Conditional on $\Lambda_j(d\xi)$, \mathbf{X}_j are Poisson processes on the brain \mathcal{B} :

$$(3.19) \quad \mathbf{X}_j \mid \Lambda_j(d\xi) \sim \mathcal{PP}(\mathcal{B}, \Lambda_j(d\xi)).$$

In other words, we have J underlying Cox processes, each one contributing a specific type of foci in some/all of the studies. The intensity measures $\Lambda_j(d\xi)$ arise from a convolution of a finite kernel measure $\mathbf{K}_j(d\xi, \zeta)$ and a Gamma random field $\mathbf{G}_j(d\zeta)$:

$$(3.20) \quad \Lambda_j(d\xi) = \int_{\mathcal{B}} \mathbf{K}_j(d\xi, \zeta) \mathbf{G}_j(d\zeta).$$

The model arising from (3.19)–(3.20) is similar to the Poisson/Gamma random field model of Wolpert and Ickstadt (1998), who first introduced the idea of convolving a Gamma random field with a Poisson process. To introduce dependence between the different tasks, it is assumed that $\mathbf{G}_j(d\zeta)$ are independent realisations of a Gamma random field with common shape measure $\mathbf{G}_0(d\zeta)$ and inverse scale parameter β :

$$(3.21) \quad \mathbf{G}_j(d\zeta) \sim \mathcal{GRF}(\mathbf{G}_0(d\zeta), \beta).$$

Again, $\mathbf{G}_0(d\zeta)$ is a Gamma random field:

$$(3.22) \quad \mathbf{G}_0(d\zeta) \sim \mathcal{GRF}(\alpha(d\zeta), \beta_0).$$

An MCMC scheme is used for posterior computation. The HPGRF model allows for the detection of overall effects based on the posterior intensity $\mathbf{G}_0(d\zeta)$ or task-specific effects based on $\Lambda_j(d\xi)$. Inference on types with fewer observations can be done by borrowing information from the remaining types through correlation under the common base intensity $\mathbf{G}_0(d\zeta)$. A significant benefit of the model is that it requires the specification of very few hyperparameters.

3.2.4 *A spatial Bayesian latent factor regression model.* In a recent contribution, [Montagna et al. \(2017\)](#) generalise the model proposed in [Montagna et al. \(2012\)](#) to the case where observations are spatial point patterns reported from different neuroimaging studies. The authors regard the foci from each study \mathbf{x}_i as a Cox process \mathbf{X}_i driven by a nonnegative random intensity function μ_i :

$$(3.23) \quad \mathbf{X}_i \mid \mu_i \sim \mathcal{PP}(\mathcal{B}, \mu_i).$$

[Montagna et al. \(2017\)](#) consider a functional representation for the (log) intensity function and write $\log \mu_i$ in terms of a collection of basis functions:

$$(3.24) \quad \log \mu_i(\mathbf{v}) = \sum_{m=1}^p \theta_{im} b_m(\mathbf{v}) = \mathbf{b}(\mathbf{v})^\top \boldsymbol{\theta}_i.$$

This specification implies that $\log \mu_i$ belongs to the span of a (fixed) set of basis functions, $\{b_m(\cdot)\}_{m=1}^p$, with $\boldsymbol{\theta}_i$ denoting a vector of study-specific coefficients. For a discussion on the choice of the bases (e.g., B-splines, Gaussian kernels, etc.) and their number p , we refer to [Montagna et al. \(2017\)](#). A low dimensional representation of $\log \mu_i$ is achieved by placing a sparse latent factor model ([Arminger and Muthén, 1998](#)) on the basis coefficients:

$$(3.25) \quad \boldsymbol{\theta}_i = \mathbf{A} \boldsymbol{\eta}_i + \boldsymbol{\zeta}_i \quad \text{with } \boldsymbol{\zeta}_i \sim N_p(0, \boldsymbol{\Sigma}),$$

where $\boldsymbol{\theta}_i = [\theta_{i1}, \dots, \theta_{ip}]^\top$, \mathbf{A} is a $p \times k$ factor loading matrix with $k \ll p$, $\boldsymbol{\eta}_i = (\eta_{i1}, \dots, \eta_{ik})^\top$ is a vector of latent factors for study i , and $\boldsymbol{\zeta}_i = (\zeta_{i1}, \dots, \zeta_{ip})^\top$ is a residual vector that is independent with the other variables in the model and is normally distributed with mean zero and diagonal covariance matrix $\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \dots, \sigma_p^2)$.

Two attractive features of this approach are the ability to accommodate covariate information (*meta-regression*) and perform *reverse inference*. Both goals are achieved by putting the low dimensional vectors of latent factors $\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n$ in any flexible joint model with other variables of interest. For example, information from covariates \mathbf{Z}_i can be incorporated through a simple linear model,

$$(3.26) \quad \boldsymbol{\eta}_i = \boldsymbol{\beta}^\top \mathbf{Z}_i + \boldsymbol{\Delta}_i \quad \text{with } \boldsymbol{\Delta}_i \sim N_k(0, \mathbf{I}),$$

where $\boldsymbol{\beta}$ is a $r \times k$ matrix of unknown coefficients, and r denotes the dimension of \mathbf{Z}_i .

Finally, reverse inference refers to inferring which cognitive process or task generated an observed activation in a certain brain region. In mathematical terms,

it corresponds to estimating $\Pr[\text{Task}|\text{Activation}]$. Suppose for simplicity that the meta-analysis dataset consists of studies that can be categorized as either type A (e.g., happy) or type B (e.g., sad). Let y_i denote the study type, with

$$y_i = \begin{cases} 1, & \text{if study } i \text{ is type A,} \\ 0, & \text{if study } i \text{ is type B.} \end{cases}$$

The interest is in estimating the probability that newly observed point pattern data arose from either a type A or B experiment. Because the study type can be represented as a binary response, the authors build a probit model for study type and predict the posterior probability that a new point pattern data arose from either type. Specifically, they model $p_{y_i} = \Pr(y_i = 1 | \alpha, \boldsymbol{\gamma}, \boldsymbol{\eta}_i) = \Phi(\alpha + \boldsymbol{\gamma}^\top \boldsymbol{\eta}_i)$, where $\Phi(\cdot)$ denotes the standard normal distribution function. Parameter α can be interpreted as the baseline probability that study i is of type A, and $\boldsymbol{\gamma}^\top \boldsymbol{\eta}_i$ accounts for study-specific random deviations. Notice that the latent factors $\boldsymbol{\eta}_i$ are used as a vehicle to link the random intensities (thus, the foci) to the study-type. We remark that the probit model can be easily replaced by an appropriate predictive model for categorical, nominal, or continuous study features.

4. EVALUATION OF EXISTING METHODS

One of the aims of this paper is to evaluate CBMA methods. A head-to-head comparison of existing methodologies is unfeasible, because the statistics described earlier have very different interpretations. Instead, we examine some characteristics of CBMA methods that show the drawbacks and merits of each. In what follows, we focus on the comparison between kernel-based and model-based methods. In Section 4.1, we conduct a series of simulations to study the sensitivity properties of the ALE algorithm that we think characterise other kernel-based methods as well. In Section 4.2, we apply the methods for which available software exist on a real dataset and compare the outputs. Finally, in Section 4.3, we proceed to a discussion.

4.1 ALE Simulation Study

Even though kernel-based methods have been extensively used for the analysis of neuroimaging data, their power properties have not been thoroughly investigated on synthetic datasets. We perform a simulation study to assess the power properties of the ALE method. In particular, we want to assess how the power of the algorithm evolves with respect to the number of studies

in the meta-analysis and whether the method is robust to the inclusion of low quality studies. We choose ALE for three main reasons. First, ALE is currently the most broadly used method for CBMA (based on a PubMed search for ALE, MKDA and SDM). Secondly, a recent review of kernel-based methods (Radua and Mataix-Cols, 2012) reported that the three kernel-based methods provide qualitatively similar results, thus we expect that our findings are indicative of MKDA and SDM methods as well. Finally, we strongly believe that the current version of ALE (Eickhoff et al., 2012) provides the best approximation to the Monte Carlo test null distribution upon which inference is based.

We create meta-analytic datasets based on the following setup. Each simulated dataset consists of I studies; of these, $I p$ are valid while the rest $I(1 - p)$ are noise, $0 \leq p \leq 1$. For the valid studies, we assume there exist 8 population centers around which foci cluster. For each center, a valid study reports no foci with probability 0.35 (that is, each center is detected with probability 0.65), a singleton focus with probability 0.5, two multiple foci with probability 0.1 and three multiple foci with probability 0.05. Hence, the expected number of foci per valid study is 6.8, similar to the average number of foci in the application of Section 4.2. The foci are drawn from a three dimensional Gaussian distribution centered at the corresponding population center. As for the noise studies, we simply sample foci uniformly from the brain mask. The expected number of foci for the noise studies is the same as for valid studies. Our setup for a power analysis of the ALE algorithm is similar to the one in Eickhoff et al. (2016). Nevertheless, their study involves a single activated region, and therefore fewer measures of power are considered. Further, their data generating mechanism is different to the one employed in this work.

As sample size, we consider I equal to 20, 40, 60, 80, 100 and 120. For a given I , we successively set $p = 0, 0.05, 0.10, 0.15, \dots, 0.95, 1$. For each distinct combination of I and p , we create $B = 1000$ datasets as described above, and apply the ALE algorithm (Eickhoff et al., 2012) to each dataset. The normal kernel standard deviation is set to $\sigma = 4$ mm, constant across studies; the value is chosen according to the sample size-standard deviation mapping used by *GingerALE*, and is used for studies with ≈ 12 participants. We use an $\alpha = 0.05$ FDR-corrected threshold to assess significance of the ALE statistic images. The following power-related quantities are recorded: (1) the probability that at least one of the 8 population centers

is detected; (2) the probability that all 8 centers are detected; (3) the mean number of centers detected in 1000 runs; (4) the mean voxel-wise true positive rate, where “truly” active voxels are defined by the 95% probability spheres around the population centers.

Our findings are summarised in Figure 3 where quantities 1 – 4 are plotted against the proportion of valid studies. One can observe that all 4 power measures increase monotonically to their maximal values of 1, 1, 8 and 1, respectively, as the number of studies grows. For a given proportion of valid studies, we observe that the bigger the total sample size is, the higher the power. For a fixed sample size, the power increases with the proportion of valid studies. Therefore, ALE is a consistent test. In Figure 4, we plot quantities 1 – 4 versus the total number of valid studies, that is, $I p$ instead of p . We see that the curves for different I tend to coincide. This is a key robustness property of the ALE algorithm, that is, adding pure-noise studies does not degrade power detection. Our results are consistent with the findings of Eickhoff et al. (2016), who observe a similar behaviour of the ALE algorithm in their setup.

In the online supplementary materials (Samartsidis et al., 2017), we repeat the simulation study above with both MKDA and SDM kernels, and using ALE’s exact enumeration procedure to threshold the statistic images (5%, FDR-corrected). This extension aims at investigating the effect of choosing three different kernels under a common inferential procedure. We find that both MKDA and SDM kernels exhibit similar behaviour to ALE. For more details, see the online supplementary materials (Samartsidis et al., 2017).

4.2 Analysis of a Real Dataset

In this section, we perform a meta-analysis of emotion studies. The dataset we consider here was previously analysed by Kober et al. (2008), Kang et al. (2011), Yue, Lindquist and Loh (2012) and Kang et al. (2014), and is now publicly available.⁶ Our goal here is not to extend the analyses in these papers, but rather to facilitate a comparison among the methods described in Section 3 and, more generally, show how CBMA can be used to identify areas of consistent brain activation. Due to lack of software availability, however, we only apply MKDA,⁷ ALE,⁸ SDM,⁹ and the BHICP¹⁰ to the data.

⁶<https://osf.io/3vn9c/>

⁷https://github.com/canlab/Canlab_MKDA_MetaAnalysis

⁸<http://www.brainmap.org/ale/> (RRID:SCR_014921)

⁹<http://www.sdmproject.com/software/> (RRID:SCR_002554)

¹⁰<http://www-personal.umich.edu/~jiankang/software.html>

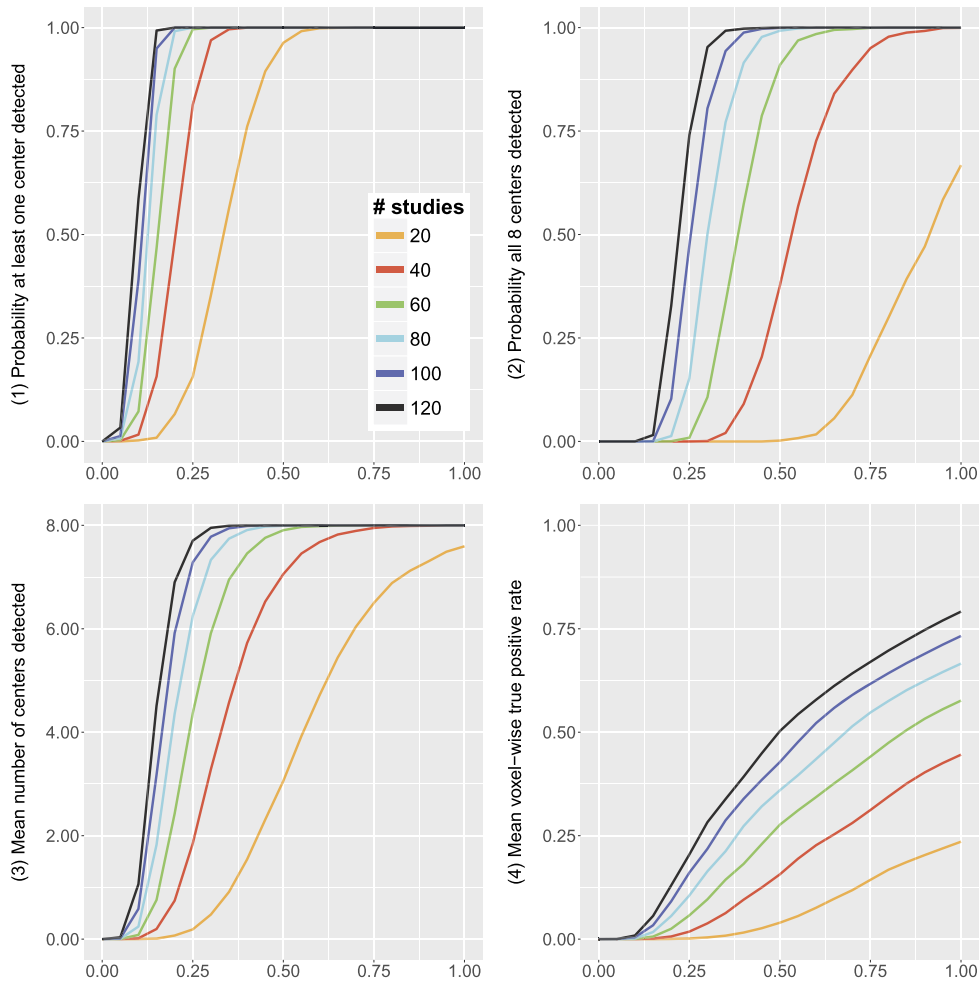


FIG. 3. Results of the simulation study. Power properties of the ALE algorithm are plotted against the proportion of valid studies p . Top left: probability that at least one center is detected. Top right: probability that all 8 centers are detected. Bottom left: mean number of centers detected. Bottom right: mean voxel-wise true positive rate.

The dataset consists of 164 experiments conducted between 1993 and 2005. Most of these experiments investigated multiple contrasts (2.7 contrasts per study on average), and each contrast involved a different emotion. It is also not unlikely that a given experiment used the same group of subjects to investigate multiple emotions, usually in the same scanning session. It is important to note that the CBMA methodologies reviewed in Section 3 do not account for within-study correlation, thus effectively treat contrasts within each study (and across studies) as independent. Accommodating for within-study correlation may be infeasible due to the sparsity of the point patterns, as well as for computational considerations. While we recognise this may be a limitation, extending CBMA methods to account for within-study correlation is beyond the scope of this paper. Thus, we following standard neu-

roimaging convention and treat contrasts as independent hereafter. Hence, our sample consists of $I = 437$ contrasts and a total of 2478 foci, with an average of roughly 6 foci per contrast. Eight emotion types appear in the dataset: affective, anger, disgust, fear, happy, mixed, sad and surprise. In general, the goals of a meta-analysis are twofolds: (1) Identify consistent activation regions (aggregation of foci) across studies of the same emotion; and (2) Identify consistent activation regions across all emotion types. Since methods for which software is available can not accommodate a multi-type meta-analysis, we focus on the second objective here. This in turn enables us to pool all 437 contrasts together and work with a bigger sample size.

Note that 58 experiments used *positron emission tomography* (PET) imaging instead of fMRI. Both PET




FIG. 4. Results of the simulation study. Power properties of the ALE algorithm are plotted against the total number of valid studies l_p . Top left: probability that at least one center is detected. Top right: probability that all 8 centers are detected. Bottom left: mean number of centers detected. Bottom right: mean voxel-wise true positive rate.

and fMRI use blood flow to detect activation signals, but differ fundamentally in the way they exploit changes in blood flow to detect activations (for details on PET acquisition see, e.g., Bailey et al., 2006). Consequently, results may differ depending on which acquisition method was used. For example, it is well known that fMRI scanners provide better spatial resolution compared to PET scanners, but are also more sensitive to motion artefacts. In the emotions literature, there are both studies suggesting that results obtained from the two different modalities are similar (Wager et al., 2008, for example), as well as studies revealing differences between the two (see, e.g., Costafreda et al., 2008). While we chose to include both PET and fMRI studies in our analysis hereafter, we highlight the need for extending existing CBMA methodologies to

account for nuisance effects such as different imaging modalities.

The simulation parameters are set as following. For MKDA, we use a kernel size of $r = 10$ mm, which is also the software default. A total of 10,000 Monte Carlo datasets are generated under the null hypothesis and used to threshold the MKDA statistic image $m(v)$ at $\alpha = 0.05$, FWE-corrected. ALE automatically assigns a kernel size for each study based on the total number of participants and uses the method of Eickhoff et al. (2012) to calculate the distribution of the statistic under the null hypothesis. The significance of the statistic image $\ell(v)$ is assessed with an FDR corrected $\alpha = 0.05$ threshold. For SDM we use an isotropic kernel of 20 mm since it is the software default and do 500 Monte Carlo randomisations. For the BHICP, we use the same hyperparameter values as

in Kang et al. (2011), and run the MCMC for 120,000 iterations saving once every 100 iterations. This results in a total sample size of 1200 posterior draws, of which we discard the first 200 as a burnin. The run length and burn-in are chosen following Kang et al. (2011). We assess convergence visually by inspecting the traceplots of some of the model parameters. Further, we conduct a second run of the model initialised at different starting values. A comparison of the results from the two runs

reveals no significant differences. For more details, see the online supplementary materials (Samartsidis et al., 2017). We now summarise the results.

Figure 5 shows statistic images obtained from the four methods presented above, conditional on several values of the z dimension. Note that for ALE, MKDA and SDM we show $\ell(v)$, $m(v)$ and $s(v)$, respectively, whereas for the BHICP we show the voxel-wise posterior mean of the activation study intensity function,




FIG. 5. Qualitative comparison of the CBMA methods in regions of high clustering of foci. Column 1 is the MKDA statistic $m(v)$, column 2 is the ALE statistic $\ell(v)$, column 3 is the SDM statistic $s(v)$, and column 4 is the study activation center intensity $\sum_{i=1}^I [\lambda_i^0(v) + \rho_i(v)]$ for the BHICP. Rows 1–3 correspond to axial slices $z = -22$, $z = -16$ and $z = -2$, respectively. Note that differences in brain shapes across methods are simply due to different masks used by the various algorithms.

namely $\sum_{i=1}^I [\lambda_i^0(v) + \rho_i(v)]$.¹¹ We see that all of the methods provide qualitatively similar results. More specifically, the regions of the brain that are mostly engaged in emotion processing are the right and left amygdala (Figure 5, top and middle row). This finding is consistent with previous analyses of the same dataset (Kober et al., 2008, Kang et al., 2011, Yue, Lindquist and Loh, 2012, Kang et al., 2014) as well as results of previous studies (Phelps and LeDoux, 2005, Costafreda et al., 2008). Other regions with moderately high values are the right and left cerebral cortex (Figure 5, bottom row). For the BHICP, this pattern is only noticeable in the right cerebral cortex.

4.3 Discussion

In this section, we build on our real data analysis to discuss the differences between kernel-based and model-based approaches in greater detail.

Interpretation. The meta-analysis of emotions suggests that the results obtained with model-based methods are qualitatively similar to those obtained with kernel-based methods. However, it is difficult to compare these results quantitatively as they have very different interpretations. At each voxel v , the MKDA statistic $m(v)$ represents the weighted proportion of studies that reported an activation within 10 mm of that given voxel. For the emotion dataset, the maximum value of $m(v)$, say $m(v^*)$, is 0.177 and v^* is located in the left amygdala. The ALE statistic $\ell(v)$ estimates the probability that at least one of the foci closest to v in each study is truly located at v . We observe a maximum of 0.158 for $\ell(v)$ in our analysis. SDM uses ALE's normal kernel but combines study maps as MKDA, hence interpretation is difficult. However, $s(v)$ can be still viewed as a qualitative measure of foci clustering around v . The maximum value for SDM is 0.181, and v^* , the voxel where the maximum occurs, is also located near the left amygdala.

In model-based methods, one obtains posterior distributions for all of the model parameters rather than a single statistic. Given the posterior samples, various estimates can be constructed. For the point process models of Kang et al. (2011, 2014) and Montagna et al. (2017), one mainly utilises the posterior draws of the intensity functions, say $\lambda^t(v)$, where t indexes the MCMC samples. Conditional on $\lambda^t(v)$, the expected

number of foci at any voxel v has a Poisson distribution with mean $8\lambda^t(v)$, where 8 is the volume of each voxel on a standard $2 \times 2 \times 2$ mm mask. With the BHICP, one can infer the location of the activation centers or the location of the population centres. Instead, the PGRF and the latent factor model of Montagna et al. (2017) provide the expected number of foci for a given experimental paradigm. Further, the PGRF will produce one intensity per study-type j whereas the latent factor model will produce one intensity map per study. Population results can be obtained by averaging the study-specific posterior intensity maps. Finally, the model of Yue, Lindquist and Loh (2012) produces posterior draws $z^t(v)$ and $\psi^t(v)$, which can be used to estimate the probability that a voxel is activated, $H(z^t(v))$, and the probability that a voxel v is falsely reported, $\mathbb{P}(\psi^t(v) = 1)$.

Implementation & computational considerations. Software for kernel-based methods can be applied to any dataset and will produce a pair of brain images, one with the value of the statistic at each and every voxel and one containing the corresponding p -values. By directly comparing these two images, one can easily identify significant voxels. Instead, it is not straightforward to implement an MCMC scheme for the Bayesian model-based methods discussed in Section 3.2. Prior specifications that are suitable for one dataset may be completely inappropriate for another. Further, it is not possible to know in advance how many iterations are required for the MCMC algorithm to converge, and convergence needs to be assessed as well.

In terms of computing time, kernel-based methods outperform model-based methods (with the exception of SDM). For the emotion dataset in Section 4.2, ALE required approximately 15 minutes of run time whereas MKDA required around 3 hours for 10,000 Monte Carlo replications. On the contrary, the BHICP model took roughly 16 hours for 120,000 MCMC iterations, Kang et al. (2014) needed 20 hours to complete the analysis, and it is yet not possible to run the spatial binary regression model on the full brain.

From a purely computational perspective, therefore, it may not seem appealing to a neuroimaging practitioner to adopt model-based methods, and this might explain why kernel-based methods are generally preferred. Kernel-based methods, however, suffer from some other serious limitations.

Explicit spatial modelling. Kernel-based methods are based on a *mass univariate approach* (MUA) that lacks an explicit spatial approach to the modelling of the foci. Specifically, MUA consists of fitting a model

¹¹Kang et al. (2011) note that singleton foci can also be viewed as activation centers since they come directly from the population centers.

at each voxel independently of every other voxel. Even though the activation of nearby voxels is correlated, estimation with the MUA ignores the spatial correlation, but inference later accounts for it when random field theory or permutation procedures define a threshold for significant activation. By failing to capture the spatial nature of the data, kernel-based methods do not provide an accurate representation of the true data generating mechanism and can not jointly characterise randomness of the number and locations of activations within each study. With point process models instead, one can obtain the expected number of foci in any region of interest simply by integrating the intensity function over that region. For example, we can use the posterior draws of the BHICP model to infer that the expected number of activation centers in the amygdala is 2.1 (95% CI [1.8, 2.6]) for the emotion dataset, and that the median number of population centers in the entire brain is 7 (95% CI [5, 10]).

Several other quantities of interest can be obtained from model-based methods. For the BHICP model, it is possible to derive $(1 - \alpha)\%$ credible ellipses for both population and study activation centers, thus returning an estimate of within-study and between study-variability as in a random effects meta-analysis model. By introducing the latent process $\psi(v)$, [Yue, Lindquist and Loh \(2012\)](#) estimate the probability of a voxel being miscoded. In the HPGRF model, the authors provide correlation estimates between the different emotions. All these quantities can not be obtained by any of the kernel-based methods.

By lacking an explicit spatial model, kernel-based methods can not deliver, for example, spatial confidence intervals and an arbitrary kernel size parameter (r for MKDA and σ for ALE and SDM) must be set (though [Eickhoff et al., 2009](#) offer heuristics). Typically, its value is specified based on previous studies rather than being estimated from the data, and it remains constant across the brain regardless of the amount of smoothing required in each region. However, a bad choice for the kernel size can potentially affect the results. For example, in the third column of [Figure 5](#) we observe bigger clusters resulting from a choice for a larger kernel size than that used for MKDA and ALE. [Yue, Lindquist and Loh \(2012\)](#) automatically choose the required amount of smoothing by introducing an extra smoothness parameters in their GMRF.

Quantification of uncertainty. There is no measure of uncertainty associated with the effect estimate in kernel-based methods, thus conclusions could be misleading. For example, in [Section 4.1](#) we found that

power properties of ALE do not degrade with the inclusion of poor quality studies (see [Figure 4](#)). Since inferences remain roughly unchanged, it is not possible to distinguish between cases with strong signal (few poor quality studies) and weak signal (many poor quality studies). Note that this is a fixed effects model property, where a small proportion of the data drives the inference. Model-based methods tackle this problem by providing standard errors obtained directly from the posterior distribution of any parameter of interest. When the signal is strong there will be small variability in the posterior estimates, whereas when signal is weak uncertainty will be higher.

Mechanisms for reverse inference. The Bayesian framework upon which model-based methods are built facilitates the construction of predictive distributions over new studies. This helps producing the so-called *reverse inferences* ([Poldrack, 2011](#)), a topic of growing interest in the fMRI community. Traditionally, fMRI studies produce *forward inferences*: for a given task or paradigm, inference is made on the location of the brain response to the task. Reverse inference consists of using the pattern of brain activation to infer which task is most likely to have produced the data. If a neuroscientist has developed a new behavioural experiment (say, on emotion), he/she may indeed want to know whether their task engages the brain's emotion processing system. In such case, the researcher would want an estimate of the probability that the data arose from a population of emotion studies. [Kang et al. \(2014\)](#) show that classification based on the HPGRF model outperforms a naive classifier based on MKDA (MKDA + NBC). Reverse inference is naturally embedded into the spatial Bayesian latent factor model of [Montagna et al. \(2017\)](#), and classification performance outperforms MKDA + NBC on a meta-analysis of emotion and executive control studies. These results suggest that spatial models can capture information in the data that can not be learnt with an MUA.

To conclude, kernel-based methods can be more efficient than model-based methods when there is need for quick exploratory analyses, and for quick comparisons between several different datasets. Nevertheless, we believe that model-based methods have significant merits compared to kernel-based methods because of the extensive inference that can be performed when using these methodologies. Of course, there are still several open problems for model-based methods and CBMA in general. We discuss these open problems in the following section.

5. OPEN PROBLEMS

Some aspects of CBMA are still being overlooked by both kernel-based and kernel-based methods. Arguably the most important is publication bias, which occurs when publicly accessible studies are not a representative subset of the total population of studies. A special case of publication bias is the so-called file drawer, that is, studies with significant findings are more likely to get published. If present, publication bias can affect the outcome of a meta-analysis and lead to false conclusions. While there is evidence for the existence of publication biases in fMRI (David et al., 2013), there has been no attempt to quantify neither the extent of the phenomenon nor the effect these biases may have on meta-analysis estimates.

Another area of emerging interest is that of meta-regression, that is, the use of study-specific characteristics as explanatory variables in a meta-analysis model (Greenland, 1994). Meta-regression is an important facet of meta-analysis, especially when there is appreciable heterogeneity between studies. As noted by Carp (2012), fMRI experiments exhibit strong heterogeneity, hence it is essential to explore the effect that study characteristics have on the outcome of the analyses. Of both kernel and model-based approaches revised in Section 3, only the spatial Bayesian latent factor regression model of Montagna et al. (2017) can accommodate covariate information collected on the different studies. It is therefore a breakthrough contribution in neuroimaging meta-analysis research. By incorporating covariates in a meta-analysis on emotions, the authors observe, for example, that failing to adjust for multiple hypothesis testing results into a higher expected number of foci over the brain. The other CBMA methods described in Section 3 would assume that any deviance from the true population effect is due to sampling error. While current implementations of other model-based methods can not accommodate covariates, it is certainly possible to extend them to address this point. Thus, meta-regression remains an important topic that should be explored further in future research.

Another interesting avenue for future research consists in combining CBMA with image data from new fMRI studies. For example, there are no models that combine CBMA with IBMA when full 3D statistic images are available. Further, it is currently not possible to use meta-analysis data to improve estimation in small, underpowered group fMRI studies. A plausible direction is to build a point process model for single

studies where the prior distribution for the study centers corresponds to the posterior intensity as obtained from a point process CBMA and IBMA model.

Often, additional information is available about the foci, such as the corresponding p -values or T scores. A convenient way to model these quantities is via mark processes, where these values become marks of the existing point patterns to improve estimation of the intensity function. Such an approach can enrich the inferences obtained from a meta-analysis by characterising the magnitude of activation at each voxel as opposed to modelling the location (and number) of the activations only, which is the question that current (model-based) methods address.

Finally, there is little work on functional connectivity for model-based methods. Functional connectivity refers to the dependency between one or more regions of the brain. In CBMA functional connectivity is implied by co-activation, that is, when two regions consistently report activations. In a recent work, Xue et al. (2014) use a multivariate Poisson model to induce correlation among the foci in several regions on interest. However, it would be interesting to extend the spatial models in order to capture these correlations as well.

ACKNOWLEDGEMENTS

The authors thank Jian Kang for helpful discussions and gratefully acknowledge the efforts of the labs of Tor Wager and Lisa Feldman Barrett for creating the emotion meta-analysis database and making it freely available. The authors also thank the two referees, the Associate Editor and the Editor for their constructive comments that helped improve the manuscript substantially. This work was largely completed while PS, SM and TEN were at the University of Warwick, Department of Statistics. PS, TEN & TDJ were supported by NIH grant 5-R01-NS-075066; TEN was supported by a Wellcome Trust fellowship 100309/Z/12/Z and NIH grant R01 2R01EB015611-04. The work presented in this paper represents the views of the authors and not necessarily those of the NIH or the Wellcome Trust Foundation.

SUPPLEMENTARY MATERIAL

Supplement to “The coordinate-based meta-analysis of neuroimaging data” (DOI: [10.1214/17-STS624SUPP](https://doi.org/10.1214/17-STS624SUPP); .pdf). The supplementary material includes the results of the simulation study of Section 4.1 using the MKDA and SDM kernels. MCMC convergence diagnostics for the analysis of the emotions data in Section 4.2 with the BHICP are also presented.

REFERENCES

- ARMINGER, G. and MUTHÉN, B. O. (1998). A Bayesian approach to nonlinear latent variable models using the Gibbs sampler and the Metropolis-Hastings algorithm. *Psychometrika* **63** 271–300.
- BAILEY, D. L., TOWNSEND, D. W., VALK, P. E. and MAISEY, M. N., eds. (2006). *Positron Emission Tomography: Basic Sciences*. Springer, Berlin.
- BARTELS, A. and ZEKEI, S. (2004). The neural correlates of maternal and romantic love. *NeuroImage* **21** 1155–1166.
- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B* **57** 289–300. [MR1325392](#)
- BUTTON, K. S., IOANNIDIS, J. P. A., MOKRYSZ, C., NOSEK, B. A., FLINT, J., ROBINSON, E. S. J. and MUNAFÒ, M. R. (2013). Power failure: Why small sample size undermines the reliability of neuroscience. *Nat. Rev., Neurosci.* **14** 365–76.
- CALHOUN, V. D. and PEARLSON, G. D. (2012). A selective review of simulated driving studies: Combining naturalistic and hybrid paradigms, analysis approaches, and future directions. *NeuroImage* **59** 25–35.
- CARP, J. (2012). The secret lives of experiments: Methods reporting in the fMRI literature. *NeuroImage* **63** 289–300.
- CASPERS, J., ZILLES, K., BEIERLE, C., ROTTSCHY, C. and EICKHOFF, S. B. (2014). A novel meta-analytic approach: Mining frequent co-activation patterns in neuroimaging databases. *Neuroimage* **90** 390–402.
- COLE, D. M., SMITH, S. M. and BECKMANN, C. F. (2010). Advances and pitfalls in the analysis and interpretation of resting-state FMRI data. *Frontiers in Systems Neuroscience* **4** 8.
- COSTAFREDA, S. G., DAVID, A. S. and BRAMMER, M. J. (2009). A parametric approach to voxel-based meta-analysis. *NeuroImage* **46** 115–122.
- COSTAFREDA, S., BRAMMER, M., DAVID, A. and FU, C. (2008). Predictors of amygdala activation during the processing of emotional stimuli. *Brains Res. Rev.* **58** 57–70.
- DAVID, S. P., WARE, J. J., CHU, I. M., LOFTUS, P. D., FUSAR-POLI, P., RADUA, J., MUNAFÒ, M. R. and IOANNIDIS, J. P. A. (2013). Potential reporting bias in fMRI studies of the brain. *PLoS ONE* **8**.
- DELVECCHIO, G., FOSSATI, P., BOYER, P., BRAMBILLA, P., FALKAI, P., GRUBER, O., HIETALA, J., LAWRIE, S. M., MARTINOT, J.-L., MCINTOSH, A. M., MEISENZAHN, E. and FRANGOU, S. (2012). Common and distinct neural correlates of emotional processing in bipolar disorder and major depressive disorder: A voxel-based meta-analysis of functional magnetic resonance imaging studies. *Eur. Neuropsychopharmacol.* **22** 100–113.
- EICKHOFF, S. B., LAIRD, A. R., GREFKES, C., WANG, L. E., ZILLES, K. and FOX, P. T. (2009). Coordinate-based activation likelihood estimation meta-analysis of neuroimaging data: A random-effects approach based on empirical estimates of spatial uncertainty. *Hum. Brain Mapp.* **30** 2907–2926.
- EICKHOFF, S. B., BZDOK, D., LAIRD, A. R., KURTH, F. and FOX, P. T. (2012). Activation likelihood estimation meta-analysis revisited. *NeuroImage* **59** 2349–2361.
- EICKHOFF, S. B., NICHOLS, T. E., LAIRD, A. R., HOFFSTAEDTER, F., AMUNTS, K., FOX, P. T., BZDOK, D. and EICKHOFF, C. R. (2016). Behavior, sensitivity, and power of activation likelihood estimation characterized by massive empirical simulation. *Neuroimage* **137** 70–85.
- EICKHOFF, S. B., LAIRD, A. R., FOX, P. M., LANCASTER, J. L. and FOX, P. T. (2017). Implementation errors in the GingerALE software: Description and recommendations. *Hum. Brain Mapp.* **38** 7–11.
- ETKIN, A. and WAGER, T. (2007). Functional neuroimaging of anxiety: A meta-analysis of emotional processing in PTSD, social anxiety disorder, and specific phobia. *Am. J. Psychiatry* **164** 1476–1488.
- FARAH, M. J. (2014). Brain images, babies, and bathwater: Critiquing critiques of functional neuroimaging. *Hastings Cent. Rep.* **44** S19–30.
- FOX, P. T., PARSONS, L. M. and LANCASTER, J. L. (1998). Beyond the single study: Function/location metanalysis in cognitive neuroimaging. *Curr. Opin. Neurobiol.* **8** 178–187.
- FOX, P. T., LANCASTER, J. L., PARSONS, L. M., XIONG, J. H. and ZAMARRIPA, F. (1997). Functional volumes modeling: Theory and preliminary assessment. *Hum. Brain Mapp.* **5** 306–311.
- FRISTON, K. J., PENNY, W., PHILLIPS, C., KIEBEL, S., HINTON, G. and ASHBURNER, J. (2002). Classical and Bayesian inference in neuroimaging: Theory. *NeuroImage* **16** 465–483.
- FUSAR-POLI, P. (2012). Voxel-wise meta-analysis of fMRI studies in patients at clinical high risk for psychosis. *J. Psychiatry Neurosci.* **37** 106–112.
- GENOVESE, C., LAZAR, N. and NICHOLS, T. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *NeuroImage* **15** 870–878.
- GORGOLEWSKI, K. J., VAROQUAUX, G., RIVERA, G., SCHWARZ, Y., GHOSH, S. S., MAUMET, C., SOCHAT, V. V., NICHOLS, T. E., POLDRACK, R. A., POLINE, J.-B., YARKONI, T. and MARGULIES, D. S. (2015). NeuroVault.org: A web-based repository for collecting and sharing unthresholded statistical maps of the human brain. *Front. Neuroinform* **9** 8.
- GREENLAND, S. (1994). Invited commentary: A critical look at some popular meta-analytic methods. *Am. J. Epidemiol.* **140** 290–296.
- HARTUNG, J., KNAPP, G. and SINHA, B. K. (2008). *Statistical Meta-Analysis with Applications*. Wiley, Hoboken, NJ. [MR2435836](#)
- HEDGES, L. V. and OLKIN, I. (1985). *Statistical Methods for Meta-Analysis*. Academic Press, Inc., Orlando, FL. [MR0798597](#)
- HUETTEL, S. A., SONG, A. W. and MCCARTHY, G. (2009). *Functional Magnetic Resonance Imaging Second Edition*. Sinauer Associates, Inc, Sunderland, MA.
- ILLIAN, J., PENTTINEN, A., STOYAN, H. and STOYAN, D. (2008). *Statistical Analysis and Modelling of Spatial Point Patterns*. Wiley, Chichester. [MR2384630](#)
- KANG, J., JOHNSON, T. D., NICHOLS, T. E. and WAGER, T. D. (2011). Meta analysis of functional neuroimaging data via Bayesian spatial point processes. *J. Amer. Statist. Assoc.* **106** 124–134. [MR2816707](#)
- KANG, J., NICHOLS, T. E., WAGER, T. D. and JOHNSON, T. D. (2014). A Bayesian hierarchical spatial point process model for multi-type neuroimaging meta-analysis. *Ann. Appl. Stat.* **8** 1800–1824. [MR3271354](#)
- KIM, S.-G. and OGAWA, S. (2012). Biophysical and physiological origins of blood oxygenation level-dependent fMRI signals. *J. Cereb. Blood Flow Metab.* **32** 1188–1206.

- KOBER, H., BARRETT, L. F., JOSEPH, J., BLISS-MOREAU, E., LINDQUIST, K. and WAGER, T. D. (2008). Functional grouping and cortical and subcortical interactions in emotion: A meta-analysis of neuroimaging studies. *NeuroImage* **42** 998–1031.
- KONOVA, A. B., MOELLER, S. J. and GOLDSTEIN, R. Z. (2013). Common and distinct neural targets of treatment: Changing brain function in substance addiction. *Neurosci. Biobehav. Rev.* **37** 2806–2817.
- LAIRD, A. R., LANCASTER, J. J. and FOX, P. T. (2005). BrainMap: The social evolution of a human brain mapping database. *Neuroinformatics* **3** 65–77.
- LAIRD, A. R., FOX, P. M., PRICE, C. J., GLAHN, D. C., UECKER, A. M., LANCASTER, J. L., TURKELTAUB, P. E., KOCHUNOV, P. and FOX, P. T. (2005). ALE meta-analysis: Controlling the false discovery rate and performing statistical contrasts. *Hum. Brain Mapp.* **25** 155–164.
- LAZAR, N. A., LUNA, B., SWEENEY, J. A. and EDDY, W. F. (2002). Combining brains: A survey of methods for statistical pooling of information. *NeuroImage* **16** 538–550.
- LINDQUIST, M. A. (2008). The statistical analysis of fMRI data. *Statist. Sci.* **23** 439–464. [MR2530545](#)
- MØLLER, J. and WAAGEPETERSEN, R. P. (2004). *Statistical Inference and Simulation for Spatial Point Processes*. **100**. Chapman & Hall/CRC, Boca Raton, FL. [MR2004226](#)
- MONTAGNA, S., TOKDAR, S. T., NEELON, B. and DUNSON, D. B. (2012). Bayesian latent factor regression for functional and longitudinal data. *Biometrics* **68** 1064–1073. [MR3040013](#)
- MONTAGNA, S., WAGER, T. D., BARRETT, L. F., JOHNSON, T. D. and NICHOLS, T. E. (2017). Spatial Bayesian latent factor regression modeling of coordinate-based meta-analysis data. *Biometrics*. To appear.
- MUMFORD, J. A. and NICHOLS, T. (2006). Modeling and inference of multisubject fMRI data. *IEEE Eng. Med. Biol. Mag.* **25** 42–51.
- MUMFORD, J. A. and NICHOLS, T. (2009). Simple group fMRI modeling and inference. *NeuroImage* **47** 1469–1475.
- NICHOLS, T. and HAYASAKA, S. (2003). Controlling the family-wise error rate in functional neuroimaging: A comparative review. *Stat. Methods Med. Res.* **12** 419–446. [MR2005445](#)
- NICHOLS, T. E. and HOLMES, A. P. (2002). Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Hum. Brain Mapp.* **15** 1–25.
- NIELSEN, F. Å. and HANSEN, L. K. (2002). Modeling of activation data in the BrainMap™ database: Detection of outliers. *Hum. Brain Mapp.* **15** 146–156.
- PHELPS, E. A. and LEDOUX, J. E. (2005). Contributions of the amygdala to emotion processing: From animal models to human behavior. *Neuron* **48** 175–187.
- POLDRACK, R. (2011). Inferring mental states from neuroimaging data: From reverse inference to large-scale decoding. *Neuron* **72** 692–697.
- POLDRACK, R. A., MUMFORD, J. A. and NICHOLS, T. E. (2011). *Handbook of Functional MRI Data Analysis*. Cambridge Univ. Press, Cambridge. [MR2839490](#)
- POLDRACK, R. A., BARCH, D. M., MITCHELL, J. P., WAGER, T. D., WAGNER, A. D., DEVLIN, J. T., CUMBA, C., KOYEJO, O. and MILHAM, M. P. (2013). Toward open sharing of task-based fMRI data: The OpenfMRI project. *Front. Neuroinform* **7** 12.
- POLINE, J.-B., BREEZE, J. L., GHOSH, S., GORGOLEWSKI, K., HALCHENKO, Y. O., HANKE, M., HASELGROVE, C., HELMER, K. G., KEATOR, D. B., MARCUS, D. S., POLDRACK, R. A., SCHWARTZ, Y., ASHBURNER, J. and KENNEDY, D. N. (2012). Data sharing in neuroimaging research. *Front. Neuroinform* **6** 9.
- RADUA, J. and MATAIX-COLS, D. (2009). Voxel-wise meta-analysis of grey matter changes in obsessive-compulsive disorder. *Br. J. Psychiatry* **195** 393–402.
- RADUA, J. and MATAIX-COLS, D. (2012). Meta-analytic methods for neuroimaging data explained. *Biol. Mood Anxiety Disord.* **2** 6.
- RADUA, J., MATAIX-COLS, D., PHILLIPS, M. L., EL-HAGE, W., KRONHAUS, D. M., CARDONER, N. and SURGULADZE, S. (2012). A new meta-analytic method for neuroimaging studies that combines reported peak coordinates and statistical parametric maps. *Eur. Psychiatry* **27** 605–611.
- RADUA, J., RUBIA, K., CANALES-RODRÍGUEZ, E. J., POMAROL-CLOTET, E., FUSAR-POLI, P. and MATAIX-COLS, D. (2014). Anisotropic kernels for coordinate-based meta-analyses of neuroimaging studies. *Front. Psychiatry* **5** 13.
- RAEMAEKERS, M., VINK, M., ZANDBELT, B., VAN WEZEL, R. J. A., KAHN, R. S. and RAMSEY, N. F. (2007). Test-retest reliability of fMRI activation during prosaccades and antisaccades. *NeuroImage* **36** 532–542.
- RICHLAN, F., KRONBICHLER, M. and WIMMER, H. (2011). Meta-analyzing brain dysfunctions in dyslexic children and adults. *NeuroImage* **56** 1735–1742.
- SALIMI-KHORSHIDI, G., SMITH, S. M., KELTNER, J. R., WAGER, T. D. and NICHOLS, T. E. (2009). Meta-analysis of neuroimaging data: A comparison of image-based and coordinate-based pooling of studies. *NeuroImage* **45** 810–823.
- SAMARTSIDIS, P., MONTAGNA, S., NICHOLS, T. E. and JOHNSON, T. D. (2017). Supplement to “The coordinate-based meta-analysis of neuroimaging data.” DOI:10.1214/17-STS624SUPP.
- SHERMER, M. (2008). Why you should be skeptical of brain scans. *Sci. Am. Mind* **19** 66–71.
- TALAIRACH, J. and TOURNOUX, P. (1988). *Co-Planar Stereotaxic Atlas of the Human Brain*. Thieme, Stuttgart.
- THIRION, B., PINEL, P., MERIAUX, S., ROCHE, A., DEHAENE, S. and POLINE, J.-B. (2007). Analysis of a large fMRI cohort: Statistical and methodological issues for group analyses. *NeuroImage* **35** 105–120.
- TURKELTAUB, P. E., EDEN, G. F., JONES, K. M. and ZEFFIRO, T. A. (2002). Meta-analysis of the functional neuroanatomy of single-word reading: Method and validation. *NeuroImage* **16** 765–780.
- TURKELTAUB, P. E., EICKHOFF, S. B., LAIRD, A. R., FOX, M., WIENER, M. and FOX, P. (2012). Minimizing within-experiment and within-group effects in activation likelihood estimation meta-analyses. *Hum. Brain Mapp.* **33** 1–13.
- VUL, E., HARRIS, C., WINKIELMAN, P. and PASHLER, H. (2009). Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspect. Psychol. Sci.* **4** 274–290.
- WAGER, T. D., JONIDES, J. and READING, S. (2004). Neuroimaging studies of shifting attention: A meta-analysis. *NeuroImage* **22** 1679–1693.

- WAGER, T. D., LINDQUIST, M. and KAPLAN, L. (2007). Meta-analysis of functional neuroimaging data: Current and future directions. *Soc. Cogn. Affect Neurosci.* **2** 150–158.
- WAGER, T. D., PHAN, K. L., LIBERZON, I. and TAYLOR, S. F. (2003). Valence, gender, and lateralization of functional brain anatomy in emotion: A meta-analysis of findings from neuroimaging. *NeuroImage* **19** 513–531.
- WAGER, T. D., BARRETT, L. F., BLISS-MOREAU, E., LINDQUIST, K., DUNCAN, S., KOBER, H., JOSEPH, J., DAVIDSON, M. and MIZE, J. (2008). The neuroimaging of emotion. In *Handbook of Emotions*, 3rd ed. (M. Lewis, J. M. Haviland-Jones and L. F. Barrett, eds.) 249–271. Guilford Press, New York, NY.
- WOLPERT, R. L. and ICKSTADT, K. (1998). Poisson/gamma random field models for spatial statistics. *Biometrika* **85** 251–267. [MR1649114](#)
- XUE, W., KANG, J., BOWMAN, F. D., WAGER, T. D. and GUO, J. (2014). Identifying functional co-activation patterns in neuroimaging studies via Poisson graphical models. *Biometrics* **70** 812–822. [MR3295742](#)
- YARKONI, T., POLDRACK, R. A., NICHOLS, T. E., VAN ESEN, D. C. and WAGER, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nat. Methods* **8** 665–670.
- YUE, Y. R., LINDQUIST, M. A. and LOH, J. M. (2012). Meta-analysis of functional neuroimaging data using Bayesian non-parametric binary regression. *Ann. Appl. Stat.* **6** 697–718. [MR2976488](#)
- YUE, Y. and SPECKMAN, P. L. (2010). Nonstationary spatial Gaussian Markov random fields. *J. Comput. Graph. Statist.* **19** 96–116. [MR2654402](#)