

Coordinate Based Random Effect Size meta-analysis of neuroimaging studies

Tench CR^{1*}

Christopher.Tench@Nottingham.ac.uk

Radu Tanasescu^{1,2}

Radu.Tanasescu@nottingham.ac.uk

Constantinescu CS¹

Cris.Constantinescu@Nottingham.ac.uk

Auer DP³

dorothee.auer@nottingham.ac.uk

Cottam WJ³

William.cottam@nottingham.ac.uk

¹Division of Clinical Neurosciences, Clinical Neurology, University of Nottingham, Queen's Medical Centre, Nottingham, UK

²Department of Neurology, University of Medicine and Pharmacy Carol Davila Bucharest, Colentina Hospital, Bucharest, Romania

³Division of Clinical Neuroscience, Radiological Sciences, University of Nottingham, Queen's Medical Centre, Nottingham, UK

*Corresponding author

Keywords

Meta-analysis, neuroimaging, voxel based morphometry, functional MRI

1 Abstract

Low power in neuroimaging studies can make them difficult to interpret, and Coordinate based meta-analysis (CBMA) may go some way to mitigating this issue. CBMA has been used in many analyses to detect where published functional MRI or voxel-based morphometry studies testing similar hypotheses report significant summary results (coordinates) consistently. Only the reported coordinates and possibly t statistics are analysed, and statistical significance of clusters is determined by coordinate density.

Here a method of performing coordinate based random effect size meta-analysis and meta-regression is introduced. The algorithm (ClusterZ) analyses both coordinates and reported t statistic or Z score, standardised by the number of subjects. Statistical significance is determined not by coordinate density, but by a random effects meta-analysis of reported effects performed cluster-wise using standard statistical methods and taking account of censoring inherent in the published summary results. Type 1 error control is achieved using the false cluster discovery rate (FCDR), which is based on the false discovery rate. This controls both the family wise error rate under the null hypothesis that coordinates are randomly drawn from a standard stereotaxic space, and the proportion of significant clusters that are expected under the null. Such control is necessary to avoid propagating and even amplifying the very issues motivating the meta-analysis in the first place. ClusterZ is demonstrated on both numerically simulated data and on real data from reports of grey matter loss in multiple sclerosis (MS) and syndromes suggestive of MS, and of painful stimulus in healthy controls. The software implementation is available to download and use freely.

2 Introduction

Neuroimaging studies often involve few subjects and have low statistical power to detect true effects, and with lack of power comes increased risk that significant results are false positives [1]. Add to this the common use of uncorrected p-value thresholds [2], and neuroimaging studies can become difficult to interpret. This situation may be compounded if the data violate the methodological assumptions of the analysis [3]. Meta-analysis can be used to synthesize the evidence across similar neuroimaging studies going some way to mitigating these problems [4], and there are various methods of statistically combining the results [5]. Image based meta-analysis (IBMA) is the most powerful approach, but is currently limited by availability of suitable statistical images. Coordinate based meta-analysis (CBMA), on the other hand, uses just the available summary reports (coordinates and possibly Z scores or t statistics) from functional MRI/PET or voxel-based morphometry studies measuring common effects, and has been utilised in many published studies; the aim is similar to that of IBMA within the limits of the available data [6]. The results of CBMA consists of clusters of coordinates where studies have reported significant effect in similar anatomical locations, representing concordance and indicating relevancy of brain structures, while coordinates not recruited into clusters are considered study specific. Concordance of the reported coordinates is determined statistically relative to a null hypothesis that the coordinates in different studies are uncorrelated, which in practice is simulated by replacing the reported coordinates by random coordinates. Popular CBMA algorithms include the activation likelihood estimate (ALE) [7-11] and the multi-level kernel density (MKDA) algorithm [12]. Signed differential mapping (SDM) [13] is similar to the ALE but incorporating the sign of effect at the reported coordinates to distinguish grey matter loss from grey matter increase, or fMRI activation from deactivation. Effect size SDM (ES-SDM) [14] takes this further and uses the reported t statistic associated with each coordinate, and can also incorporate statistical parametric maps.

There are technical limitations with these CBMA algorithms that impact interpretability and specificity. Firstly statistical tests are performed voxel-wise making the relevant cluster-wise type 1 error rates difficult, if not impossible, to assess. Secondly the significance is, at least in part, determined by the density of coordinates from different experiments meaning that coordinates forming a small cluster are more significant than if they formed a larger cluster. Yet it is not clear, for example, that studies reporting thalamic coordinates producing a cluster over the thalamic volume should be less significant than the same studies reporting coordinates producing a smaller cluster in the smaller putamen structure. Finally, the uncorrected p-value threshold employed by both SDM and ES-SDM does not control the type 1 error rate in a principled way [2], and without estimated error rates there is no way to assess the significance of the results given the $\sim 2 \times 10^5$ voxel-wise statistical tests; this may propagate the very problems the MA was employed to mitigate (see [15] for example).

LocalALE is a CBMA algorithm [16, 17] that addresses some of these limitations. It employs an interpretable cluster-level type 1 error rate control scheme, the false cluster discovery rate (FCDR), made possible by performing statistical tests at the coordinate, rather than the voxel, level. The results are such that at-most some specified proportion of the clusters declared significant are expected under the null hypothesis. LocalALE also adjusts its parameters to avoid false negatives when there are few studies and avoid false positives when there are many studies. Furthermore, LocalALE assigns coordinates to clusters in a binary fashion (belonging to a specific cluster, or no cluster), and as a consequence can analyse positive and negative effects (activation and deactivation, for example) simultaneously, allowing post-hoc checks for sign consistency. Nevertheless, LocalALE is unable to utilise the sign or magnitude of the reported effect to perform statistical inference, and the cluster significance is determined by coordinate density biasing the results to smaller clusters.

Here a new coordinate based random effect size (CBRES) meta-analysis (MA), and meta-regression, method (ClusterZ) is detailed. The algorithm deviates from other CBMA methods by performing inference on a standardised effect size, which is related to the Z score or *t* statistic reported by most studies. Consequently the density of coordinates within cluster does not influence statistical significance, so large and small clusters are considered on an equal footing. A random effects meta-analysis approach is taken and model parameters are estimated by maximum likelihood estimation (MLE) and significance assessed by comparing models using a likelihood ratio test (LRT). This is a common approach to meta-analysis, and one that has been applied to neuroimaging studies previously [18], but using a different null hypothesis, that the effect size is zero, to other CBMA methods. Models can be devised to test for evidence of a non-zero effect size, effect size difference between groups, or significant linear regression. ClusterZ also requires consistent spatial effect across studies for significance, and uses this to control the type 1 error rate such that quantifiably more clusters are declared significant than are expected if the studies report uncorrelated spatial effects. Furthermore, it adjusts parameters to avoid false positive and false negative results depending on the number of studies. ClusterZ is similar to traditional MA in that estimates of effect and variance are computed. It provides an alternative to ES-SDM for coordinate based meta-analysis but with the advantage that the type 1 error rate is controlled, quantified, and interpretable. ClusterZ is implemented into NeuRoi, which can be downloaded and used freely: <http://www.nottingham.ac.uk/research/groups/clinicalneurology/neuroi.aspx>.

3 Methods

There are several steps to the ClusterZ algorithm, detailed below. In summary, clusters are formed by reported coordinates that are more densely packed than average. Then, a random effects analysis is performed to give a p-value in each cluster. The same analysis is then performed on many pseudo experiments, in which each coordinate has been replaced by a random one to simulate studies reporting spatially uncorrelated effects. Declaration of significance in ClusterZ has two requirements: 1) that within cluster there is a consistent effect size reported such that the p-value is small, and 2) for a given p-value threshold the number of observed clusters with smaller p-values is quantifiably greater than average for the pseudo experiments. The second requirement indicates how ClusterZ controls the type 1 errors through the false cluster discovery rate.

3.1 Cluster forming

The clustering algorithm is identical to that used by LocalALE, and is detailed in [16] but recapped here. It is based on a popular algorithm: density based spatial clustering of applications with noise (DBSCAN) [19]. The aim is to produce clusters of densely packed coordinates while not recruiting coordinates outside these clusters, which DBSCAN considers noise; in the present application these coordinates are considered study specific effects rather than noise. The initial step is a measure of overlap of coordinates in different studies. A coordinate that overlaps (they are separated by a distance $< \Delta$) coordinates in n other studies has an overlap score of n . For a coordinate to be considered part of a cluster, its overlap score must be at least 3 according to the DBSCAN algorithm, since an overlap score of 2 or less means the coordinate is link in a chain, rather than a cluster, of coordinates. The peak of any cluster is the coordinate, or collection of coordinates, with the highest overlap score. The clustering algorithm proceeds by finding the peak coordinate that is not already assigned to a cluster and assigns it a cluster number. Coordinates overlapping members of this cluster, and have equal or lower overlap score, are recruited to the cluster. This continues until there are no more valid overlapping coordinates to be added to the cluster. The process then continues starting with the coordinate with the highest overlap score that is not already part of a cluster. The result is a set of clusters of coordinates that have a reducing (but not strictly) overlap score moving away from the peak; this can help to prevent close neighbouring clusters merging into one bigger cluster [16].

The clustering process depends on the clustering distance Δ , which is analogous to the FWHM parameter used in other CBMA algorithms [7, 13], and the algorithm to compute this has been detailed previously [17]. The choice of Δ is determined by three aims of the clustering algorithm: 1) to allow the true clusters to form, 2) to prevent study specific coordinates forming clusters, and 3) to prevent study specific coordinates being recruited into the true clusters. The first aim requires Δ to be large enough so that the densely packed coordinates within-cluster overlap. The second and third aims conversely require Δ to be small such that the low density coordinates falling between the clusters do

not overlap on average. The density of coordinates within, and between, clusters is unknown, but the density of random coordinates can be estimated, and in the true clusters the coordinates are more densely packed than this and between clusters the coordinates are packed with lower density on average. The algorithm proceeds by redistributing the coordinates randomly (see below) within an anatomical mask, which depending on the problem might be a grey-matter, white-matter, or whole-brain mask. For these coordinates a small value of Δ results in few coordinates having non-zero overlap scores, but this increases for larger Δ . It is helpful to consider the proportion of coordinates with non-zero overlap scores (divided by 2 to avoid coordinate A overlapping coordinate B being considered a second time as B overlapping A) as a function of the clustering distance: $\phi(\Delta)$, the overlap fraction. The clustering distance used is that Δ which, on average, just causes each random coordinate to overlap with another in one other study such that $\phi(\Delta)=0.5$. With this value of Δ the coordinates within the clusters become density reachable [19], satisfying aim 1. The study specific coordinates, on the other hand, are not density reachable satisfying aims 2 & 3. For larger Δ the study specific coordinates will (wrongly) begin to form clusters or make them density reachable from the clusters violating aims 2 & 3, while for smaller Δ the true clusters may not form at all violating aim 1. This method adapts Δ to the number of studies, such that it is small for many studies because the density of coordinates is higher. Consequently the false negatives are reduced when there are fewer studies (the low density of coordinates when few studies are included requires Δ to be large so that clusters still form), while false positives are reduced when there are many studies (because the small Δ prevents the study specific coordinates being recruited into clusters) [17]; this is in contrast to the fixed FWHM parameter used in other algorithms that can paradoxically increase the false positives with an increasing number of studies [17].

A subtlety in this algorithm is that independence of the coordinates within study cannot be guaranteed [9, 12]. Reported coordinates that are very close to each other may be related, and the aim is to preserve this relationship. Independently randomising the coordinates within a mask using a uniform distribution would violate this aim, and therefore a more sophisticated approach is taken and has been described in detail previously [16]. Within-study clusters of coordinates (each coordinate separated by a distance $<\Delta$ from at least one other coordinate from the same cluster) are formed. Each within-study cluster has a centroid, and the mean and standard deviation of the distances from the centroid, of the coordinates belonging to the cluster, computed. Each centroid is then randomly placed, with uniform probability, within the mask, and the cluster coordinates randomly placed about the centroid with the computed mean and standard deviation distance. A randomisation is rejected, and subsequently repeated until successful, if the within-study clusters of coordinates are randomised such that they overlap; see [16] for details.

3.2 Models of effect size

In fMRI a biologically meaningful measure of effect is not available, although %BOLD (blood oxygenation level dependent) signal change has been suggested [20]. In VBM studies a more meaningful measure in terms of volume reduction is possible, but not routinely reported. Typically neuroimaging studies only report effect sizes reflecting statistical significance: the t score or Z score. These are dimensionless and relative effect sizes similar to Cohen's d if scaled [21]; t and Z are both scaled by subject numbers. The meaning, from a biological perspective, of such effects is not directly apparent, and often cannot be inferred from the report [20]. Prospective studies of regions known to be consistently reported as significant could be designed to answer this question, but first the regions must be identified and a statistical effect size estimated so that a sample size might be computed [21]. ClusterZ can be used to provide this detail.

Within clusters the Z scores or t statistics are combined to test if there is consistent statistical significance; publication bias or outlying effects, for example, might be reported by only one study and would not be consistent. Methods combining p-values (or equivalent), reviewed in [5], are appropriate for IBMA, where the test is for significant effect within a voxel given multiple statistical images. In CBMA, however, each reported coordinate has already been declared as significant by the reporting study, and the relevant question is about significant consistency of the reported effects across studies. Combining p-values (known to be significant already) is then not appropriate, so a meta-analytic method of combining results from different studies [5] is employed instead, necessitating an effect size and estimate of variance.

Under certain assumptions the t statistic or Z score can be transformed and used as a dimensionless effect for MA or meta-regression. Assuming that studies report analyses of one or two groups of subjects then the t statistic, which has a student's t distribution, is a measure of effect, or difference in effect between groups, reported in multiples of the standard error. The first difficulty with the t statistic as an effect size is that the standard error depends on the number of subjects, so larger studies appear to report larger effects. Therefore, the effect used in CBRES analysis is the t statistic, or Z score, scaled to obtain an estimate measured in multiples of the sample standard deviation, which does not depend on the number of subjects. A basic analysis is assumed, whereby simple t tests are performed. Then the effect size is computed from the t statistic by dividing by a function of the number of subjects n^*

$$\mathcal{E}'_i = \frac{t}{\sqrt{n^*}}, \quad (1)$$

and the within-study variance is

$$\mathcal{V}'^2_i = \frac{df}{df-2} \frac{1}{n^*_i}, \quad (2)$$

which includes the variance of the Student's t distribution with df degrees of freedom. For a one sample study, for example just healthy control (HC) subjects, then n^* is the number of subjects and $df=n^*-1$. For a two sample study, such as a patient versus HC study, then

$$n_i^* = \frac{n_1 \times n_2}{n_1 + n_2}, \quad (3)$$

where n_1 and n_2 are the numbers of subjects in each sample and $df=n_1+n_2-2$; this assumes equal variance of effect in both groups. These standardised effect size and distributional assumptions assume a random effects approach has been used for analysis.

Under the given assumptions behind equations (1) to (3), they provide an effect size that is standardised against the number of subjects and have an associated variance estimate as required for MA. However, many studies employ a more sophisticated analysis and include multiple regressors [22]. This is partly corrected for by modifying the degrees of freedom of the t distributions by the number of extra regressors included; if the number of regressors is unknown, some underestimation of the effect variance will result. Furthermore, scaling the t statistic by $1/\sqrt{n^*}$ will not standardise to the same sample standard deviation across studies. This is not correctable, due to the limitations of reported effect, and will result in some between-study variance in effect size due to the range of analyses employed. Nevertheless, standardising the reported statistics should result in an effect with approximately known variance, and which tends to be larger for larger effects. Furthermore, ClusterZ attempts to make allowance for this limitation by including a between-study variance as a random effect.

One more assumption is made, for convenience, in what follows: that the subject numbers are sufficiently large that the reported t statistics are well approximated by Z scores such that the effect is

$$\varepsilon_i = \frac{Z}{\sqrt{n^*}} \quad (4)$$

with variance

$$v_i^2 = \frac{1}{n_i^*}. \quad (5)$$

The Z scores are standardised by comparison to the range of possible Student's t distributions making them easier to consider for the purpose of experimentation. However, ClusterZ does have the option to deal with either t or Z for effect sizes using either equations (1) and (2) or equations (4) and (5).

3.3 Random effect model

A random effects approach is employed to combine the effects given by equations (1-5), making some allowance for the often appreciable differences in experimental design and analysis between studies. To perform inference a random effect distribution must be assumed. Typically a normal distribution is assumed [5] so that for study i

$$\varepsilon_i \sim N(\theta_i, \sigma^2 + \nu_i^2). \quad (6)$$

The parameter θ_i is modified to reflect different models, and are estimated using MLE. The total variance is composed of the within-study variance ν_i^2 (equation (5)) and the between-study variance σ^2 , which is also estimated by MLE. With this model estimates are weighted more heavily to the larger studies due to the smaller within-study variance.

Various models are possible using different parameterisations of equation (6). To model the effect sizes by a grand mean the parameterisation is $\theta_i = \mu$. For meta-regression the grand mean is modulated by a study specific covariate so $\theta_i = \mu + \beta c_i$, where β models the change in the grand mean due to the covariate c ; this could be a continuous variable such as age, or a group indicator to investigate differences in effect size between groups.

3.3.1 Maximum likelihood estimation with censoring

The maximum likelihood estimates of parameters θ and σ are generally straight forward to compute. However, a subtlety is that not all studies will report a coordinate and effect in every cluster, and those that don't are censored by the statistical threshold used [18]; for example a study applying an uncorrected p-value threshold of 0.0001 reports only Z scores exceeding 3.72 in magnitude. Left, right, and interval censoring need to be considered. The log likelihood, which is maximised for parameter estimation, is a sum of contributions from uncensored and censored terms.

For uncensored data the log likelihood contribution for a single study is given by the probability (density) that effect size $E = \varepsilon_i$ given the parameters θ_i and σ ($P(E = \varepsilon_i | \theta_i, \sigma)$)

$$L_i = -\log(\sqrt{2\pi}\Sigma_i) - \frac{(\varepsilon_i - \theta_i)^2}{2\Sigma_i^2}, \quad (7)$$

where

$$\Sigma_i^2 = \sigma^2 + \nu_i^2. \quad (8)$$

Left censoring occurs when a study reports a significant negative effect (for example deactivation) within a cluster but no effect size is given. If study i reports an effect size threshold such that $E \leq -T_i$, where T_i is the threshold magnitude, then the contribution to the likelihood is the probability $P(E \leq -T_i | \theta_i, \sigma)$. Given the random effects model (equation (6)) this probability can be computed accurately using the error function (*erf*) [23]. The contribution to the log likelihood for a single left censored study is therefore

$$L_{left_i} = \log\left(\frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{-T_i - \theta_i}{\sqrt{2}\Sigma_i}\right)\right]\right). \quad (9)$$

With right censoring study i reports an effect size threshold such that $E \geq T_i$ and the contribution to the likelihood is the probability $P(E \geq T_i | \theta, \sigma)$, and log likelihood

$$L_{right_i} = \log \left(1 - \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{T_i - \theta_i}{\sqrt{2}\Sigma_i} \right) \right] \right). \quad (10)$$

Interval censoring occurs when a study does not report a significant result within a cluster. In this case all that is known is that $|E| \leq T_i$, and the likelihood is the probability $P(|E| \leq T_i | \theta, \sigma)$. In this case the contribution to the log likelihood is

$$L_{int_i} = \log \left(\frac{1}{2} \left[\operatorname{erf} \left(\frac{T_i - \theta_i}{\sqrt{2}\Sigma_i} \right) - \operatorname{erf} \left(\frac{-T_i - \theta_i}{\sqrt{2}\Sigma_i} \right) \right] \right). \quad (11)$$

Summing these contributions over all studies gives the log likelihood

$$\begin{aligned} Ll(\theta, \sigma) = & \sum_{i \in \text{uncensored}} L_i \\ & + \sum_{i \in \text{left censored}} L_{left_i} \\ & + \sum_{i \in \text{right censored}} L_{right_i} \\ & + \sum_{i \in \text{interval censored}} L_{int_i} \end{aligned} \quad (12)$$

which is maximised for parameter estimation.

Effect size thresholds are often given in studies that employ uncorrected p-values. However, they are sometimes omitted. In the absence of a stated threshold, ClusterZ estimates them from the smallest magnitude effect size reported by the study. If no effect size is reported, then conservative low magnitude default threshold of 3.09 (corresponding to $p < 0.001$) is used by default; a high default threshold might overestimate the effect and drive significance.

3.3.2 Inference using the likelihood ratio test

Without the censoring, a one sample t -test or simple linear regression (SLR) model can be used for inference. However, neuroimaging study reports are censored so a likelihood ratio test, a general scheme for performing a statistical test by approximating the difference of two log likelihoods to a chi squared distribution, is used. The two likelihoods in question are the maximum likelihoods computed under the null and alternative hypotheses. To test the hypothesis that coordinates in a cluster have a non-zero grand mean effect size μ compute

$$D_\mu = 2 \left(\max_{\mu, \sigma_1} Ll(\mu, \sigma_1) - \max_{\sigma_0} Ll(0, \sigma_0) \right) \quad (13)$$

and compare D_μ to a chi square distribution with one degree of freedom to compute the p-value. For the regression model, where $\theta_i = \mu - \beta c_i$, a test for non-zero regression coefficient β compares

$$D_\beta = 2 \left(\max_{\mu_1, \beta, \sigma_1} Ll(\mu_1, \beta, \sigma_1) - \max_{\mu_0, \sigma_0} Ll(\mu_0, \sigma_0) \right) \quad (14)$$

to a chi square distribution with one degree of freedom to compute the p-value.

3.4 Type 1 error control

The null hypothesis of the LRT's is that there is zero mean effect, which has been used previously for the meta-analysis of neuroimaging studies [18]. The null hypothesis used by ClusterZ and CBMA algorithms, however, is that studies report no common spatial effect. The clustering algorithm detailed above does nothing to statistically preclude clusters forming even for studies reporting different spatial effects, and results of the LRT are no indication of significant spatial clustering. Therefore, the peak summary effects reported by studies, which by definition already surpass a study dependent threshold for significance, may well produce significant meta-analytic results in any incidentally formed clusters; the same is also true of meta-regression if reported Z scores all correlate with the regressor. Consistency in reported effect size, while necessary for significance in CBRES, is not sufficient and a further step is required to prevent studies testing different hypotheses producing such incidental significant results frequently i.e. to control the family wise error rate (FWER). Furthermore, when studies test related hypotheses controlling the type 1 error rate such that the significant clusters quantifiably outnumber those expected under the null makes interpretation simple. In ClusterZ, this is performed using the FCDR.

The concept is that results declared significant by ClusterZ should be more significant than incidental results that might arise if the studies were measuring different effects. A null hypothesis based on an unbiased sample of unrelated studies with similar characteristics (subject numbers, effect sizes etc) is needed, but such samples are not readily available; indeed it is not obvious how unbiased might be defined for a sample of neuroimaging studies. In CBMA approximations to unrelated studies are computed by spatial randomisation of the reported coordinates; ClusterZ uses the algorithm detailed in section 3.1. To be specific, ClusterZ produces 4000 pseudo experiments by replacing the reported coordinates with random coordinates; while preserving the reported effect sizes, subject numbers, censoring thresholds, and covariate. For each of these, clusters are formed and inference on the effect sizes, using the likelihood ratio test, performed. If a total of N_0 clusters are formed by these pseudo experiments, then there are an associated set of p-values: p_{0i} ($1 \leq i \leq N_0$). Similarly, using the reported coordinates a set of N clusters are formed having an associated set of p-values: p_j (sorted such that $p_1 \leq p_2 \leq p_3 \leq \dots \leq p_N$). To control the type 1 error rate, the false cluster discovery rate is limited to a level α such that

$$k = \max_j \frac{1}{4000} \frac{\sum_{1 \leq i \leq N_0} I(p_{0i} \leq p_j)}{j} \leq \alpha. \quad (15)$$

In equation (15) $I(E)$ is an indicator function that equals one if E is true, and equals zero otherwise. The interpretation of this is that k is the maximum number of significant clusters such that the expected number of clusters per pseudo experiment is at most $\alpha \times k$.

The FCDR is very similar to the more familiar false discovery rate (FDR) method [24] on which it is based, and is exactly the same as that employed by LocalALE [16]. It imposes control such that at most a specified proportion of clusters declared significant would be expected from the pseudo experiments used as surrogate experiments involving unrelated studies. Furthermore, it controls the FWER under the null hypothesis [16] just as FDR [24].

3.4.1 Family wise error rate control in ClusterZ

Beyond FCDR, it is also straight forward to control the FWE rate in ClusterZ. By taking the minimum p-value for each of the 4000 pseudo experiments, sorting them into ascending order, then picking the $\alpha \times 4000^{\text{th}}$ p-value in the sorted list as the threshold for significance, the family wise error rate will be controlled at a level α . This is an option in the ClusterZ software, but it is conservative and has the disadvantage that there is no indicator of the proportion of clusters declared significant that are to be expected under the null, which is why FCDR is the default and recommended option.

4 Experiments

4.1 Estimation

The first necessary step in performing CBRES meta-analysis is to estimate the effects, only then can inference using the likelihood ratio test be valid. To demonstrate the utility of equations (7-12) numerical simulation was performed. Samples representative of effects in a single cluster were generated using the model

$$\varepsilon_i = \mu + \beta c_i + \eta_i + \rho_i, \quad (16)$$

where the within-study error is

$$\eta_i \sim N(0, \sigma_i^2)$$

and the between-study error is

$$\rho_i \sim N(0, \sigma^2).$$

Samples of parameters μ and β were selected at random (independently) from uniform distributions with range -1 to +1, and σ was selected from a standard uniform random distribution but constrained to be $0.1 \leq \sigma \leq 1$. To simulate N_s studies, N_s samples of η_i and ρ_i were generated using the Box-Muller

algorithm [25]. For the purpose of this experiment the covariate c_i was set equally spaced between $-a$ and $+a$

$$c_i = -a + 2 \frac{(i-1)}{N_s-1} a, \quad (17)$$

which is centred on zero and has a standard deviation of

$$\sigma_c = \sqrt{a^2 \left[\frac{1}{3} + \frac{2}{3} \frac{1}{(N_s-1)} \right]}. \quad (18)$$

To simulate censoring any sample ε_i with magnitude $< T$ was removed.

A pragmatic (given knowledge of previous coordinate based meta-analyses) set of model values was employed: $n^*=20$, $N_s=20$, and censoring threshold $T=3.79/\sqrt{n^*}$ (which represents a commonly used uncorrected p-value threshold of 0.0001). The covariate range ($-a$ to a) was set such that $\sigma_c=1$, which according to the SLR model makes the standard errors of parameters μ and β equal. The model was generated 100 times and the estimated parameters plotted against the true parameters. It was expected that the number of studies improves the estimates, therefore a further similar experiment with $N_s=100$ studies, representing a large MA, was simulated and results plotted.

4.1.1 Estimation Results

Plots of grand mean, between-study standard deviation, and regression coefficient against their maximum likelihood estimates are shown in figure (1). MLE has been successful, over a pertinent range of parameter values, in the presence of censoring typical of whole-brain neuroimaging studies. The standard errors on the estimates are, as expected, smaller for 100 studies than for just 20.

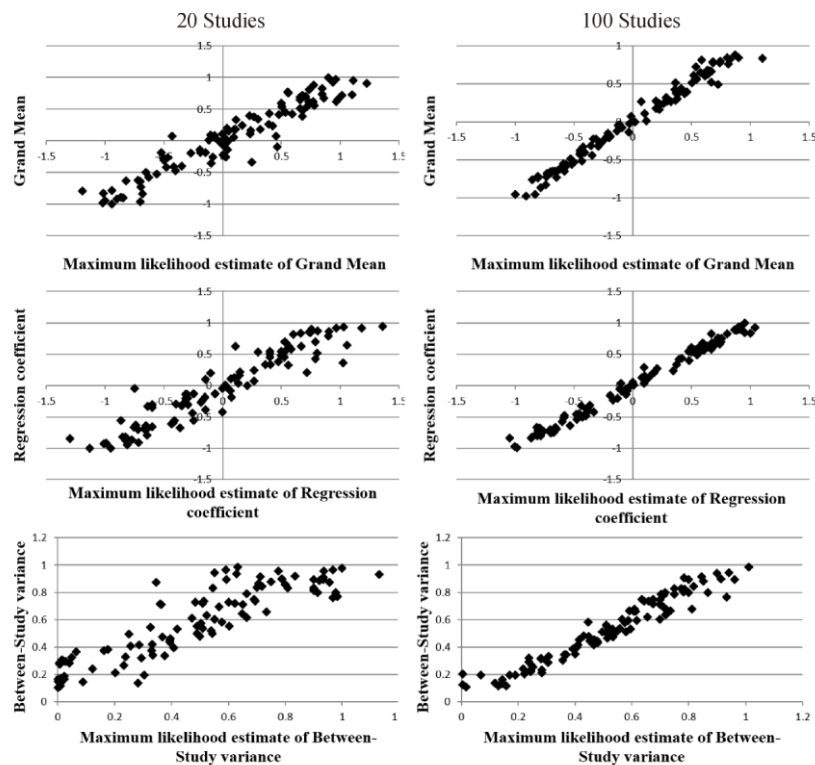


Figure 1. Utility of MLE to estimate the model parameters in the presence of censoring at $Z > 3.79$ ($p \leq 0.0001$). Estimates, as expected, are more accurate for many studies.

4.2 Validation of clustering distance

To consider how sensitive ClusterZ is to Δ , numerical simulation was used. Experiments with 20 subjects were simulated with activation Z scores normally distributed with a mean of 3.79 and standard deviation 1.0; truncated to a maximum of 6.0 to avoid unrealistically large values from the tails of the normal distribution. The Z scores were truncated below 3.79 so that half of the generated coordinates were censored on average. This censoring simulates a large effect [26] to avoid false negatives as the purpose here is to demonstrate the clustering method, rather than the sensitivity of the ClusterZ. Each study included an average, after censoring, of 10 randomly (uniformly from a GM mask) distributed coordinates, and up to 5 (depending on the censoring) coordinates spatially distributed about known clustering points; this spatial distribution was Gaussian, using a standard deviation of the distances from the centre points of 4.5mm corresponding to a higher limit (least dense) of standard deviations measured in clusters detected in CBMAs [26]. Three experiments were generated with 20, 50 and 100 studies representing small to large meta-analyses. Clustering was performed with $\phi(\Delta)=0.5$, $\phi(\Delta)=0.1$, and $\phi(\Delta)=0.05$ for each number of studies. It was expected that for smaller clustering distances the clusters would fail to form. Clustering was also performed for fixed clustering distance, with Δ deduced such that $\phi(\Delta)=0.5$ for the 20 study experiment. For this fixed Δ it was expected that the clusters would grow by erroneously recruiting the random coordinates as the number of simulated studies was increased.

Erroneous inclusion of coordinates to clusters, or erroneous exclusion of coordinates from clusters, may make interpretation more difficult, and at least some such errors are inevitable at the edges of clusters where the coordinates truly belonging to the clusters are closest to those that do not belong. As well as being an issue for interpretation, these errors can bias the estimated effect sizes. To consider the bias an experiment involving two clusters was devised, each with standard deviation of coordinate distance from the cluster centre of 4.5mm as above; low density clusters with standard deviation of 8mm are also considered. The Z scores are drawn from a normal distribution with mean magnitude of 3.79 and standard deviation of 1.0, one cluster having negative mean and one positive; with 20 subjects this is an effect size magnitude of $3.79/\sqrt{20}=0.85$. Between the clusters 20 random coordinates per experiment were generated, half of which had a positive mean Z of 3.79, and half with negative 3.79. As above all Z scores were censored below a magnitude of 3.79 and capped at magnitude 6.0. The clusters were placed at three different distances relative to each other, ranging from being placed at distance in different hemispheres to touching. The number of studies considered was 20, 50, and 100. The effect size in each cluster was estimated as a function of the overlap

fraction. It was expected that for small $\phi(\Delta)$ the clusters would fail to form and the effect size would be underestimated. At the limit of large overlap fraction it was expected that random coordinates would be recruited to the clusters and reduce the estimated effect size because of the different effect signs. With the clusters close together it was expected that large overlap fraction (large clustering distance) would cause the clusters to merge resulting in an underestimated mean effect size because of the different effect signs.

4.2.1 *Validation of clustering distance Results*

Figure (2) shows the effect of the clustering distance; note that the ClusterZ depicts clusters as + markers for each member coordinate, rather than as a cluster of voxels, making them appear somewhat blocky. As expected when the overlap fraction is $\phi(\Delta)=0.05$ the clusters can fail to form because Δ is too small for coordinates to overlap, and where they do form they can be fractured. For a wide range of $\phi(\Delta)$ (0.1 to 0.5) the clusters successfully form so, at least for this data, the method is not overly sensitive to the clustering distance. This is because the coordinates in the true clusters are generally considerably more densely packed than the between cluster coordinates.

The important feature of the adaptive clustering scheme is that the clusters are generally independent of the number of studies in the analysis. This does not hold when the clustering distance is held fixed, as seen on the bottom row of figure (2). Fixing the clustering distance causes recruitment of the between-cluster coordinates into the true clusters as the density increases with the number of studies; this is most evident from an increased inclusion of coordinates outside of the cluster boundary for 100 studies and clustering distance 12.94mm. This leads to a paradoxically increasing number of false positives for larger meta-analyses. It is important to note that this is not a feature specific to ClusterZ, as it will affect any method employing a fixed FWHM; the clustering distance equivalent. Indeed this has previously been shown to happen in the ALE algorithm [17, 26].

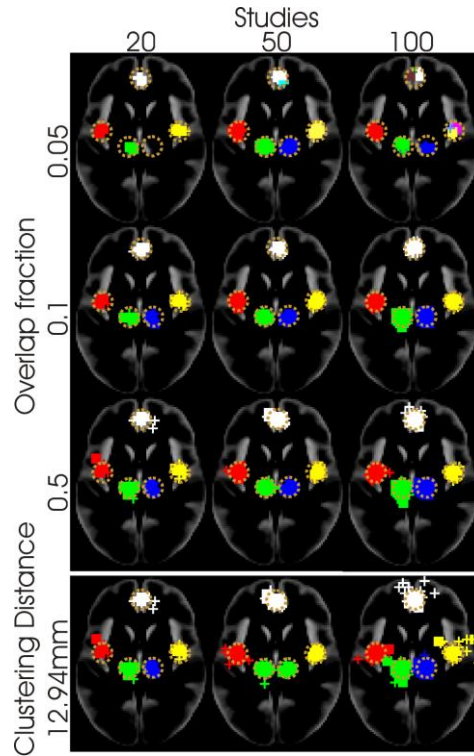


Figure 2. Showing the effect of clustering distance on simulated data with 5 clusters. Significant clusters are represented by + markers for each coordinate in the cluster, and different clusters are indicated by different colour markers; the boundaries of the true clusters are indicated by dashed-line circles. For small overlap fraction (top), clusters can fail to form or be fragmented as indicated by multiple colours within-cluster. For fixed clustering distance (bottom), the clusters grow to erroneously include study specific coordinates. Importantly, allowing the clustering distance to vary with the number of studies produces similar results for large (100 studies) and small (20 studies) meta-analyses.

In figure (3) the overlap fraction is plotted against estimated effect size to show the impact of varying clustering distance. When the two opposite effect size clusters are well spatially separated (figure (3a)) the plots are quite independent of the number of studies; despite the clustering distances for the different numbers of studies being different. For very small $\phi(\Delta)$ the clusters fail to form and the effect sizes are underestimated; this is the same feature seen in figure (2). At the other extreme large $\phi(\Delta)$ results in an underestimated effect size due to recruitment of between cluster coordinates; again this is seen in figure (2). Once the coordinates are fully formed the effect size estimate is accurate, but with increasing bias as the overlap fraction increases, and at $\phi(\Delta)=0.5$ there is some underestimation. This bias increases ever more rapidly as $\phi(\Delta)$ increases, as does the standard error of the estimate. It would appear that for smaller bias a smaller value of cluster fraction is preferable, but this comes at the expense of an increased risk of lower density clusters not forming. This can be seen in figure (3a) where estimates are also given for a cluster of low density that is not fully formed until around $\phi(\Delta)=0.5$.

As the clusters are placed closer together, another phenomenon becomes apparent. For large clustering distance it is not possible for the clustering algorithm to resolve two nearby clusters. In figure (3b) this impacts the 20 study experiment mostly as the larger clustering distances cause coordinates in the two clusters to overlap. This is even more obvious when the clusters slightly overlap (row c), where only the experiment with 100 studies is able to properly resolve the two clusters with an overlap fraction of 0.5. When clusters of opposite effect size are not resolved the effect estimate due to the combined coordinates may be near zero, and consequently no significant cluster would be declared. However, this can be detected by inspection of forest plots (see figure (9)), where the mixing of positive and negative effects would be highlighted. A solution employed by other algorithms is to perform separate positive and negative reported effect meta-analyses, but this is not ideal as the effects are not independent. The solution offered by ClusterZ is to modify the clustering algorithm such that only coordinates of the same effect sign can overlap, producing positive effect clusters and negative effect clusters with no sign mixing. The application of this to the touching clusters is shown in figure (3d), which shows the bias has been eliminated and the clusters are resolved. There is a caveat to this modification, however, in that it can mask true heterogeneity in reported effects. This modified algorithm should therefore only be used if the forest plots indicate merged clusters with opposite sign.

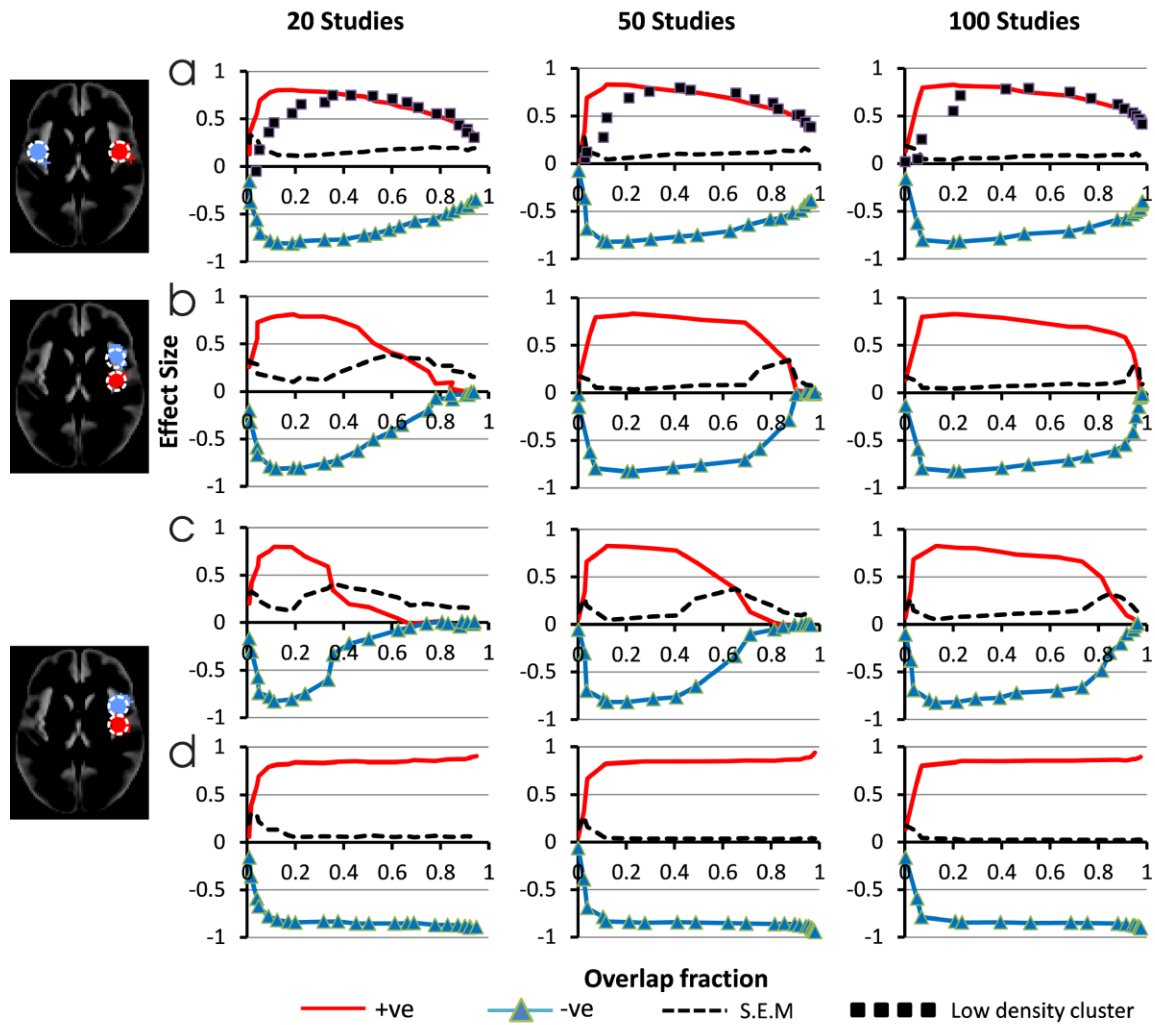


Figure 3. Showing the impact of overlap fraction (and therefore clustering distance) on effect size estimates. Two clusters are simulated, as depicted on the left; the clusters are placed at distance, near, or overlapping. True effect sizes are $+0.85$ and -0.85 , and between cluster coordinates have the same magnitude effect. Here S.E.M is the standard error on the effect size estimate. Coordinates are placed with a standard deviation of 4.5mm from the cluster centre, or 8mm for the low density example.

4.3 Type 1 error control

To confirm that FCDR controls the FWER under the null hypothesis, pseudo experiments were generated by randomly placing coordinates, independently and with uniform probability, within a GM mask. This simulation used 40 studies, and the Z scores were sampled at random from a Gaussian distribution with mean 3.79 , standard deviation 1.0 , and truncated such that $3.79 \leq |Z| \leq 6$; this simulates Z scores from studies reporting significant and censored activation (for example). Therefore, incidental clusters formed by the random coordinates may be expected to have a significant positive mean, yet there should be few significant results with FWER control. Five hundred such experiments were analysed and the number that produced significant results counted while controlling the FCDR at 0.01 , 0.05 , and 0.1 representing conservative, typical, and liberal settings respectively. It was

expected that the pseudo experiments would produce significant results around 1%, 5%, and 10%, respectively, of the time.

Control of the FWER is important, but does not then place a quantifiable limit on the number of clusters, from those declared significant, that might be expected under the null hypothesis. This is an aim of FCDR. To test this, experiments involving known numbers of true clusters (1, 3, and 5) were simulated; with 500 simulations per experiment. Since the point of the experiment is control of false positive results, forty studies were simulated, providing sufficient statistical power to avoid many false negative results. As above, all Z scores were sampled at random from a Gaussian distribution with mean 3.79, standard deviation 1.0, and truncated such that $3.79 \leq |Z| \leq 6$. Each study had twenty random coordinates, which was reduced to 10 on average after censoring. Added to these were a set of true clusters distributed about fixed Talairach [27] coordinates; the spatial distribution was Gaussian, using a standard deviation of the distances from the centre points of 4.5mm corresponding to a higher limit (least dense) of standard deviations measured in clusters detected in CBMA [26].

To compare with the ALE and ES-SDM algorithms, 50 experiments were generated (as detailed above) for each of 0, 1, 3, and 5 fixed clusters and saved in the ALE and ES-SDM formats; these files are included as supplemental material. The distributions of clusters detected using ALE and ES-SDM were plotted as histograms, along with equivalent results from ClusterZ using the experiments detailed above with FCDR 0.05. FDR of 0.05 was used with a minimum 200mm³ cluster size for ALE [26], and the recommended settings (uncorrected p-value 0.005, minimum cluster size 10 voxels, and a minimum Z threshold of 1) were used with ES-SDM [14]; in addition the default anisotropic smoothing kernels were employed [28]. Ideally the recommended cluster based threshold method would have been employed with the ALE algorithm, but the execution time (around one day per experiment) prevented its use for the 200 experiments performed. Nevertheless, the limitations of FDR in the context of CBMA are well understood [26] and can be considered in the comparison.

4.3.1 Type 1 error control Results

Five hundred pseudo experiments, with zero fixed clusters, were processed. A histogram of the number of clusters detected are shown in figure (4). At FCDR of 0.01 the total number of experiments declaring significant clusters was 9/500, which is an estimated FWER of 0.018. Similarly at FCDR of 0.05 and 0.1 the total number of experiments declaring significant clusters was 24/500 and 54/500 respectively, representing family wise error rates of 0.048 and 0.11. This experiment suggests that FCDR is able to control the FWER under the null hypothesis, just as the FDR scheme it is based on. The ability of FCDR to control the type 1 error rate was also tested in experiments with up to 5 true significant clusters, and also shown in figure (4). The number of clusters declared significant was correct in at least 80% of experiments, and as expected the number of false negatives was highest for

FCDR of 0.01, the number of false positives was highest for 0.1, while the typical setting of 0.05 fell between the two. This experiment demonstrates that, at least for the simulated data, FCDR does control both the FWE rate and the cluster-wise type 1 error rate.

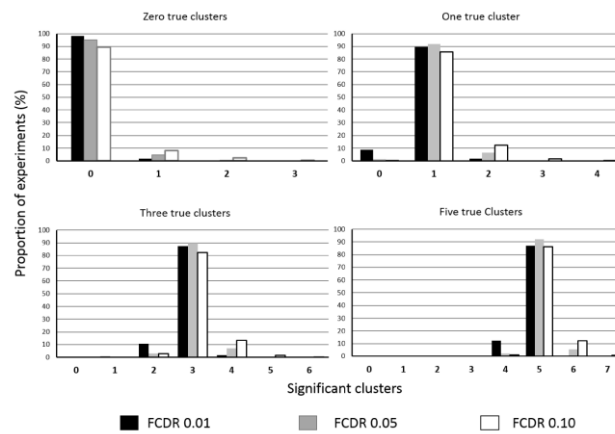


Figure 4. The number of clusters declared significant by ClusterZ for known numbers of clusters (0, 1, 3, and 5) and FCDR (0.01, 0.05, and 0.1).

The results of the comparison between ClusterZ and the ALE and ES-SDM algorithms is shown in figures (5) & (6). The ALE algorithm controls the FWER when there are zero true clusters since FDR was used, but the rate of false positive clusters begins to increase for 3 and 5 clusters. However, these extra clusters were small and are a known consequence of voxel-wise FDR in ALE. Taking into account the limitations of voxel-wise FDR, the ALE algorithm performs similarly with ClusterZ for this data, as can be seen in figure (6), where the small false clusters are highlighted; note the smoothness of the ALE algorithm is due to the Gaussian kernel used by the method, while ClusterZ depicts significant clusters by + markers for each coordinate contributing to the cluster. The ES-SDM algorithm, on the other hand, has been unable to demonstrate control of either the FWE rate, or the number of false clusters. This is probably a result of the uncorrected p-value threshold employed. Figure (6) shows the extent of the issue, with multiple clusters detected under the null hypothesis (zero true clusters) and results that are quite different to both ALE (despite the same coordinates being used) and ClusterZ, and also not constrained to the regions where the true clusters are placed.

The contrast in clusters produced by ES-SDM, on the one hand, and ClusterZ and ALE on the other, is expected to be due to the parameters being set for different aims. The ALE algorithm attempts to model the spatial probability distribution of activation foci on a voxel-wise basis [7]. A statistical threshold applied to this probability reveals voxel clusters that might then be loosely interpreted as the distribution of reported activation foci. ClusterZ also attempts to find the distribution of reported coordinates (activation and deactivation), but by clustering coordinates rather than voxels. ES-SDM has a different aim of attempting to estimate the whole brain activation pattern in order, where

possible, to incorporate statistical parametric maps into the analysis [14]. As a consequence the clusters are much larger than those of ALE and ClusterZ.

To consider whether the differences in clusters between ES-SDM and ClusterZ/ALE might be completely explained by the differences in parameters, a set of equivalent (as far as is possible) parameters values were established to reflect those used by the ALE algorithm: p-value threshold due to FDR, the FWHM, and the minimum cluster volume. These parameters were then used with ES-SDM, and the analysis shown in figure (6) repeated. For zero, one, three, and five clusters, the threshold p-values were 0.0, 4.4×10^{-6} , 1.7×10^{-5} , and 2.9×10^{-5} respectively in the respective ALE analyses; however ES-SDM allows a minimum p-value threshold of 1.0×10^{-5} , which was therefore used where necessary. The FWHM was set to 10mm, and the minimum cluster size 200mm^3 ; these are similar to those typically used in ALE. There is no equivalent of the anisotropic kernels or Z value threshold in ALE, so they were set to isotropic and the default $Z > 1$. The results shown in figure (6) (bottom row) demonstrate that using similar parameter settings modifies the output of ES-SDM to be similar to ALE, and ClusterZ, for this experiment. However, this is at the expense of changing the original aim of the method.

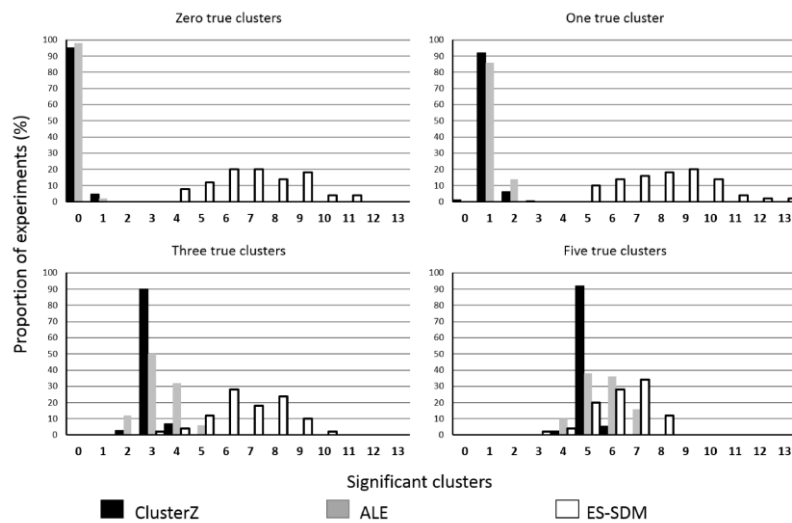


Figure 5. Comparison of ClusterZ, ALE, and ES-SDM using simulated data with known numbers of clusters. For ClusterZ an FCDR of 0.05 was employed, and FDR of 0.05 with a minimum cluster size of 200mm^3 was used for the ALE algorithm, and the default $p < 0.005$ and cluster extent of 10 voxels used for ES-SDM.

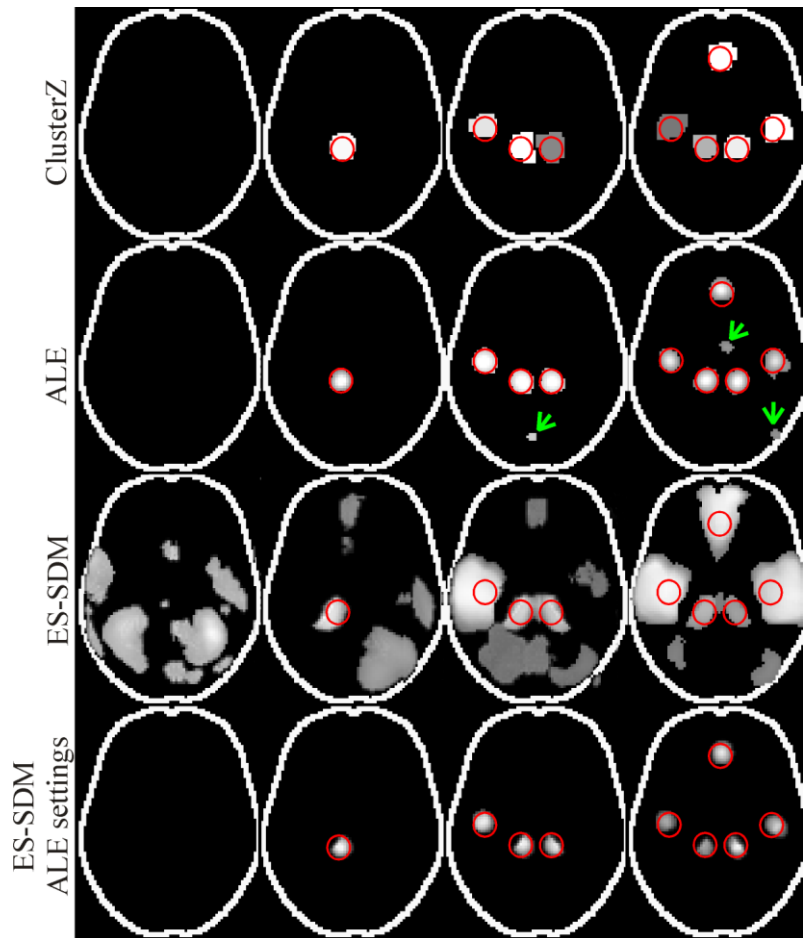


Figure 6. Typical results of numerical experiments with fixed numbers of clusters superimposed onto an axial outline of the brain. The red circles indicate the placement of coordinates belonging to the fixed clusters. The resulting significant clusters are shown as maximum intensity projections. For the ALE results, the arrows indicate small false clusters due to the use of voxel-wise FDR. The results for ES-SDM are shown with the default parameter settings and with settings equivalent to those used by ALE (bottom).

4.4 Coordinate Based Random Effect analysis of real data

Numerical experiments verify the functionality of the algorithms and demonstrate their features, but it remains to be shown that ClusterZ performs on real data. Full meta-analyses are a study in themselves and beyond this demonstration, so data adapted from two previously published analyses [29, 30] are used. The first is a meta-analysis of VBM studies of clinically isolated syndrome (CIS) and multiple sclerosis (MS); subjects diagnosed with CIS are at risk of developing clinically definite MS, which is known to result in grey and white matter atrophy. The studies compared grey matter density or volume of patients to healthy controls. In total there are 21 studies reporting 29 experiments comparing patients to healthy controls, and of these 4 were CIS, 16 RRMS, 3 BMS, 2 SPMS, 2 PPMS, 1 classical MS, and 1 cortical MS; details are in supplement 2. The second is a meta-analysis of fMRI studies of mechanically induced pain in healthy volunteers. These studies compared functional activation under painful stimulus to activation at rest or innocuous stimulus. In total this includes 24 studies involving 318 volunteers; details are in supplement 3. Each analysis was

performed using ClusterZ (FCDR 0.05 & FWE 0.05), ALE (default cluster forming threshold $p < 0.001$, cluster-level FWE 0.05, 1000 iterations), and ES-SDM (default uncorrected $p < 0.005$, minimum cluster size 10 voxels, $Z > 1$, and anisotropic smoothing kernels).

ClusterZ reports its results in several ways. Cluster images use colour coding to show each significant cluster, and the reported foci are indicated by + markers. Details of each cluster are tabulated, including effect size estimates, cluster centroid, and Talairach labels as detailed in [31] to locate the anatomy implicated. ClusterZ does not work voxel-wise, so the cluster volumes typically reported by both the ES-SDM and ALE methods are not directly computable. However, an equivalent voxel-based volume is estimated by filling an image array with overlapping cube markers for each reported coordinate to form a solid cluster of voxels with computable volume that encompasses the distribution of reported foci; these volumes are also reported. An important aspect of MA is data checking and forest plots are the standard way to visualise the data in context, which is invaluable for quickly identifying problems. Data checking is particularly important for ClusterZ, where multiple meta-analyses (one per cluster formed) are performed. The use of random effects meta-analysis in clusters means that forest plots become viable, and code (written for the R statistical package [32]) to draw the plots is automatically generated for each cluster detected regardless of significance level. Forest plots for the most significant clusters found are presented.

4.4.1 Real data CBRES Results

ClusterZ detected 6 significant clusters in the MS data using FCDR, which are listed in table 1. Most covered more than one Talairach structure as indicated in the table. Using ClusterZ FWE method reduced the significant clusters to 3, but it is worth noting that they are identical to the respective three from the FCDR analysis; clusters, formed by the clustering algorithm, are independent of the threshold level in ClusterZ unlike voxel-wise analyses. The ALE algorithm has declared 5 significant clusters in the MS data, which are very similarly located to those detected by ClusterZ. The similarity between the ALE and ClusterZ results extends somewhat to the characteristic cluster shapes, as can be seen by comparing rows a & b with row c in both figure (7) and figure (8), however the ALE clusters tend to be smaller. This is a result of the different way that ALE represents clusters (voxel-wise activation likelihood measure) compared to ClusterZ. Comparison with the ES-SDM algorithm is not straight forward. ES-SDM declares clusters that are comparatively large, and having different centroids; this is evident from both figure (7) and (8). The relatively large clusters reflect the different aim of ES-SDM compared to ALE and ClusterZ. Where possible the reported cluster centroids were matched to those detected by ClusterZ and ALE and reported in table 1.

ClusterZ declared 9 significant clusters with the pain data using FCDR; table (2) and figure (8). A useful feature of ClusterZ is that the FCDR is estimated for every cluster, and post-hoc analysis revealed a tenth cluster just beyond significant at FCDR 0.06; ClusterZ automatically estimates, and reports, FCDR and FWE rates for every cluster detected. This is reduced to 7 clusters when FWE is employed. ALE agreed on 8 of these. Again the ES-SDM results were somewhat different, and matching cluster centres in table 2 was not possible for most.

Figure (9) shows forest plots from the most significant clusters from the MS and pain analyses. These indicate the reported effect sizes and a range deduced from the within-study standard deviation, an overall estimate of grand mean and standard deviation, and the p-value for the cluster. Censored studies are also depicted in the plots by empty circle markers and dashed line range indicators derived from the study thresholds. For censored studies the markers do not indicate the contribution to the estimate of the model parameters, but the dashed lines indicate the range over which the contribution is computed using equations (9-11).

Talairach	Grand mean (μ), between study σ	ClusterZ FCDR {x,y,z}mm, volume mm ³	ClusterZ FEW {x,y,z}mm, volume mm ³	ALE {x,y,z}mm, volume mm ³	ES-SDM {x,y,z}mm, volume mm ³
Right Thalamus	-1.21, 0.30	10.4 -26.6 7.7 9248	10.4 -26.6 7.7 9248	10.4 -26.9 7.6 3872	6 -16 12 55584
Left Thalamus	-1.17, 0.28	-12.3 -27.5 5.3 10264	-12.3 -27.5 5.3 10264	-11.6 -25.4 6.5 4504	-
Left BA 3/4/40	-0.74, 0.34	-48.7 -21.3 40.4 8216	-48.7 -21.3 40.4 8216	-47.5 -24.2 41.4 2160	-36 -16 43 14808
Left BA 9	-0.65, 0.26	-44.2 14.6 32.1 4216	-	-44.7 14.0 32.3 632	-
Right Thalamus/Caudate	-0.61, 0.45	5.1 -3.2 17.4 3688	-	-	-
Right BA 3/4	-0.80, 0.19	44.0 -16.3 43.1 5208	44.0 -16.3 43.1 5208	40.5 -17.7 41.6 624	50 -20 50 12336
Right Insula	-	-	-	-	32 -1 15 19160
Corpus Callosum	-	-	-	-	2 18 12 19976

Table 1. Significant clusters with their estimated grand mean effect size and between-study standard deviation, volume, and centre Talairach location, for the MS data. The results are shown to compare ClusterZ, ALE, and ES-SDM methods.

Talairach	Grand mean (μ), between study σ	ClusterZ FCDR {x,y,z}mm, volume mm ³	ClusterZ FEW {x,y,z}mm, volume mm ³	ALE {x,y,z}mm, volume mm ³	ES-SDM {x,y,z}mm, volume mm ³
Left BA 13/40	1.20, 0.56	-52.3 -28.9 19.9 13512	-52.3 -28.9 19.9 13512	-52.2 -27.5 20 5760	-
Right BA 13/44	1.18, 0.21	41.3 8.4 10.1 21920	41.3 8.4 10.1 21920	41.3 9.7 6.8 6472	-
Right BA 40	1.17, 0.27	53.2 -25.1 26.0 12048	53.2 -25.1 26.0 12048	52.3 -24.4 20.5 3832	-
Left BA 13	1.1, 0.27	-38.2 5.3 11.3 14520	-38.2 5.3 11.3 14520	-38.2 1.9 10.4 4432	-
Right BA 9/10/46	0.75, 0.54	36.7 37.3 23.7 6264	36.7 37.3 23.7 6264	36.3 39.2 24 1144	-
Right Thalamus	0.82, 0.58	13.4 -12.6 9.6 8208	13.4 -12.6 9.6 8208	12.9 -12.7 9.7 1672	-
Right Caudate/putamen	0.67, 0.62	13.4 6.5 5.4 4832	-	14.6 6.9 5.2 1160	-
Right BA 6/8/24	0.91, 0.37	-2.0 17.6 40.1 15248	-2.0 17.6 40.1 15248	-1.6 17.8 38.8 5072	-3 22 42 28120
Left Thalamus	0.65, 0.15	-10.9 -17.7 5.7 4136	-	-	-
Right BA 7/40	*0.56, 0.56	*38.5 -55.7 47.4 3336	-	-	-
Right BA 44/45/48	-	-	-	-	51 -5 13 110088
Left BA 40/41/48	-	-	-	-	-38 -13 8 76352
Undefined	-	-	-	-	11 -1 22 2296
Left Caudate Nucleus	-	-	-	-	-11 6 18 792

Table 2. Significant clusters with their estimated grand mean effect size and between-study standard deviation, volume, and centre Talairach location, for the pain data. The results are shown to compare ClusterZ, ALE, and ES-SDM methods. Clusters marginally beyond significant indicated by * (ClusterZ only).

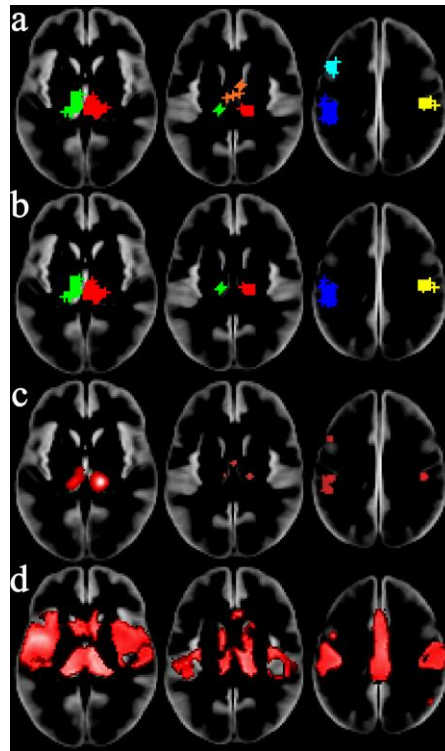


Figure 7. Significant clusters detected in the MS coordinates. a) is ClusterZ using FCDR 0.05, b) is ClusterZ using FWE 0.05, c) is ALE algorithm employing $p < 0.001$ cluster forming threshold and cluster threshold of 0.05 (FWE corrected), and d) is ES-SDM using the recommended $p < 0.005$ threshold and cluster extent of 10 voxels. For the ClusterZ results coordinates contributing to clusters are indicated by a + marker, and different clusters are indicated by different colours.

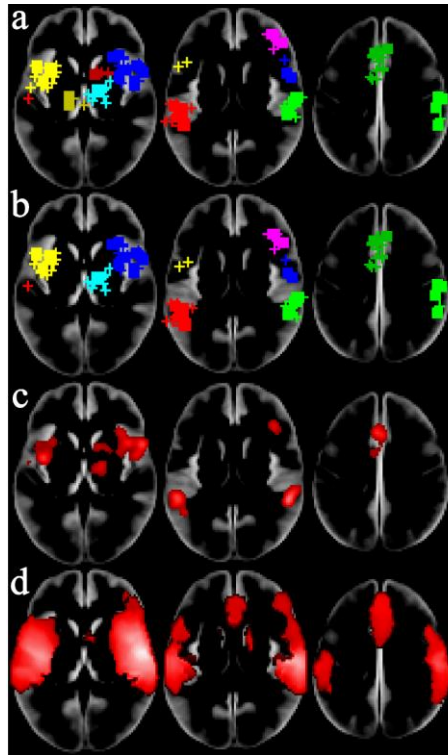


Figure 8. Significant clusters detected in the pain coordinates. a) is ClusterZ using FCDR 0.05, b) is ClusterZ using FWE 0.05, c) is ALE algorithm employing $p < 0.001$ cluster forming threshold and cluster threshold of 0.05 (FWE corrected), and d) is ES-SDM using the recommended $p < 0.005$ threshold and cluster extent of 10 voxels. For the ClusterZ results coordinates contributing to clusters are indicated by a + marker, and different clusters are indicated by different colours.

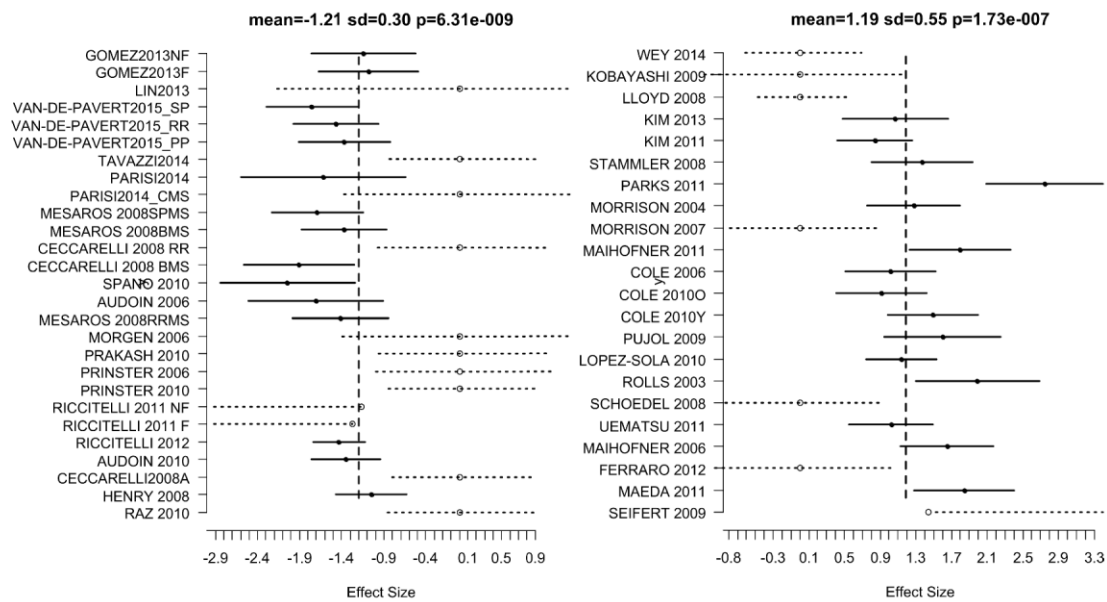


Figure 9. Forest plots of the effect sizes in the most significant cluster from the MS (left) and pain (right) meta-analyses. Solid circle markers indicate the effect size reported by the study in the cluster, while the confidence intervals are depicted as solid horizontal lines spanning ± 1.96 times the within study standard deviation of the effect size. Censored values are indicated by open circle markers and the intervals by dashed lines (---o---); these are determined by the study thresholds and indicate regions where the likelihood contributions are computed using equations 9, 10, or 11.

5 Discussion

Here a method of performing coordinate based random effect meta-analysis and meta-regression was presented. ClusterZ utilises the Z scores (or t statistics) and coordinates, typically reported by functional MRI or VBM studies, to detect where studies report effects consistently. The reported statistics are not suitable effect measures for meta-analysis directly, but they can be transformed (equations 1-5) into an approximate effect size that is suitable for a random effects meta-analysis. The advantages of ClusterZ include estimates of effect size, the possibility of meta-regression, and that cluster-wise statistical significance is determined by the reported effect rather than the density of clustering, which is not so interpretable. Furthermore, analysis at the cluster-, rather than voxel-, level makes type 1 error control using the easy to interpret false cluster discovery rate feasible.

ClusterZ relies on multiple established methods to perform its analysis. Firstly, clusters are formed based on the DBSCAN algorithm, which aims to differentiate true clusters from study specific effects. The clustering algorithm has a parameter that is analogous to the FWHM parameter in CBMA algorithms, but in ClusterZ this automatically adapts to the experiment to avoid false negatives when there are few studies and avoid false positives when there are many studies; the fixed FWHM used in other algorithms paradoxically increases the false positives for increasing numbers of studies. Estimation of model parameters is based on the standard statistical technique of maximum likelihood estimation, which also allows for censoring. The subsequent inference on the parameters is based on the generalised likelihood ratio test, which has well known statistical properties in the asymptotic limit. Finally, principled control of the false positive clusters is based on the popular FDR method, which controls the FWER under the null hypothesis, and then limits the proportion of clusters expected under the null hypothesis to a user specified level; by comparison the ES-SDM method employs no disciplined control of the false positive clusters. This is perhaps the most important feature of ClusterZ, since false positives results might propagate and even amplify the very issues that motivated the meta-analysis in the first place.

Maximum likelihood estimation is used to estimate model parameters in the presence of censoring. MLE is one of multiple methods of performing meta-analysis, and has been shown to perform well in comparison [33]. Numerically MLE is a nonlinear problem, with no guarantee of convergence to the correct solution. However, because ClusterZ computes estimates only at clusters, rather than voxels, the algorithm is very efficient (~1minute for typical analyses such as the MS and pain studies) so considerable computations effort is dedicated to finding the global maximum solutions. Numerical experiments confirmed its validity, being able to accurately estimate the parameters of the models represented by equation (6).

The adaptive clustering algorithm was tested on simulated data having known numbers of true clusters. Properties of true clusters have been empirically established [26], which was used to define a simulated true cluster with coordinates having a standard deviation of 4.5mm about a centre point. The successful formation of the true clusters depends on the clustering distance parameter Δ , which is automatically estimated from the data. The results showed (figure (2)) that such clusters would form for a wide range of clustering distance, implying that ClusterZ does not critically depend on its exact value. Small values of clustering distance did prevent the clusters forming, or caused them to fracture, but this would be true of any other CBMA algorithm. Keeping Δ fixed increases cluster sizes as the number of studies increases indicating a lack of convergence onto a steady-state result for larger analyses. This is not specific to ClusterZ, but is a feature common to algorithms using fixed FWHM (the analogue of Δ) [17, 26]. An important feature of ClusterZ, compared to other algorithms, is therefore that the clusters converge to represent the spatial distribution of reported foci and escalating false positives are avoided for larger analyses.

However, there is some inevitable uncertainty about membership of coordinates at the edge of clusters. Erroneous inclusion or exclusion of coordinates is controlled by the adaptive clustering distance, but not perfectly. When studies erroneously contribute to clusters because they report a coordinate that is density reachable from the cluster edge, interpretation is more difficult and effect estimates can be biased. Figure (3) shows the estimated effect size for two clusters, one simulated with positive effect size and one with negative. The between cluster study specific coordinates were also simulated with positive or negative effects, and erroneous inclusion to the true clusters did produce bias. The bias is reduced by reducing the clustering distance, and hence the chance of recruiting the study specific coordinates, but at the risk of missing lower density clusters. A further phenomenon in this scenario is the ability to resolve the true clusters when they fall very close together. Clusters are better resolved with large numbers of studies providing that the clustering distance is reduced, as might be intuitively expected. Two clusters having different sign effects that are not resolved would not be declared significant by ClusterZ using the clustering algorithm detailed, but would be made evident by inspection of the forest plots. A modified algorithm is provided with ClusterZ that can help with this, but must be used with care as it can also mask true heterogeneity.

The features of FCDR were also demonstrated by numerical simulation, which showed that the number of pseudo (involving randomly generated coordinates) experiments declaring significant results is as required for control of the FWER. Having controlled the FWER the user has a measure of confidence that a significant experiment is not a false positive. Furthermore, FCDR adapts to the data

to quantify a limit on the proportion of significant results that are expected under the null hypothesis. This was demonstrated on experiments containing a known number of true clusters. This is the expected behaviour of FDR, and therefore FCDR, which was designed for experiments where control of the FWER is undesirably conservative.

By comparison the ALE algorithm performed similarly with simulated data, in that it controlled the family wise error rate and correctly detected the true clusters. One limitation of the ALE experiment was the use of voxel-wise FDR, which caused multiple false positive clusters. This is a known issue, and one that is fixed with the new cluster based FWE algorithm that is now recommended [10]; unfortunately the long execution time made it impracticable for this experiment. While ALE and ClusterZ performed similarly for the simulated data presented, this would not have been the case if experiments with both activations and deactivations were simulated since ALE can only process one or the other while ClusterZ can process both. The comparison with ES-SDM is not straight forward. The algorithm uses a large fixed FWHM (20mm) and a liberal uncorrected p-value to declare significant results. This is in order to achieve the objective of extrapolating a voxel-wise effect size estimate across the whole-brain [14], but is at the expense of FWE control and results in false positive clusters. It was demonstrated in figure (6) that employing parameters typically used with the ALE algorithm made the ES-SDM results comparable to both ALE and ClusterZ. However, this represents a substantial shift from the original aim of the method, and cannot be implemented without prior knowledge of the threshold; obtained here by using ALE with the same data. It should be noted that the coordinates used for the 200 ALE experiments were exactly the same as those used for the ES-SDM analyses, and are included as supplementary material for validation.

Analysis of real data cannot be considered validation as the results are unknown, but is important to demonstrate the output of ClusterZ. The analysis of real data also provides an opportunity to qualitatively compare ClusterZ to the popular ALE and the ES-SDM algorithms. The meta-analysis of VBM studies of MS patients revealed some of the clusters that have been demonstrated to be the most consistently atrophied in a recent study comparing software packages for volumetric analysis in relapsing remitting MS [34]. A further example using real data analyses functional studies of healthy controls receiving painful mechanical stimulus. The ALE algorithm produced results similar to ClusterZ, but the data largely include a single effect sign (all GM loss in the MS data, and mostly activation in the pain data). If the data were a mix of activation and deactivation, ALE would necessitate two separate analyses, which statistically is a type of censoring not considered by the algorithm that could lead to bias. As with the simulated data the ES-SDM algorithm produces qualitatively different results, and comparison is not straight forward. The parameter values employed

by ES-SDM to achieve its aim inevitably compromise the type 1 error control, which has been demonstrated previously to produce apparently significant results where there are none [15].

A limitation of ClusterZ, and of the routine reporting of neuroimaging studies, is that the effect sizes (t statistics and Z scores) are not directly relatable to biological effect [20]. Reported grey matter volume loss in VBM, for example, would arguably be more appropriate. Prospective studies could be used to measure a more appropriate effect in a-priori regions, but first the regions need to be identified. ClusterZ, and other CBMA algorithms can do this. However, location of the regions is only one aspect of prospective study design as a sample size calculation is a requisite. ClusterZ estimates a distribution of effect sizes for each of the regions found to be consistently reported across studies, which might then be used for power calculations [21].

ClusterZ is part of NeuRoi, which is available to download and use freely. To help perform a CBRES MA, an annotated spreadsheet is provided to make extraction of relevant parameters, coordinates, and effect sizes straight forward. Subsequent error checking is performed on the coordinate files, and diagnostics detailed previously [16] are available. When reporting the analysis any relevant diagnostics performed should be provided. In common with other CBMA methods, reported results should include the Talairach coordinates for each significant cluster along with an anatomical reference and a cluster volume. Detailed cluster reports are automatically produced by ClusterZ, which should ideally be given as supplementary material as they provide a breakdown of each cluster down to the constituent foci. In addition to the detail reported for other methods, the effect size estimates should be given and forest plots shown. The text files of reported coordinates should also be included as supplementary material to allow further validation as required for modern data science [35].

6 Summary

ClusterZ is a new algorithm for performing coordinate based random effect size meta-analysis, or meta-regression, on summary results routinely reported in functional MRI studies and VBM studies. It employs disciplined control of the type 1 error rate, which is most important for a meta-analysis where rigorous statistics is necessary to gauge the strength of evidence for an effect. Advantages over coordinate based meta-analysis are that statistical significance is determined by standardised effects rather than the density of clusters, the censoring present in published summary results can be considered, and meta-regression is feasible. The effect size employed is not ideal and is limited due to the current reporting standards. Nevertheless, ClusterZ detects where the reported summary results are consistent both in terms of statistical effect size and spatial agreement, indicating target regions of

interest for prospective studies of biological effect whilst also providing effect sizes that can be used to power the study.

References

1. Button, K.S., et al., *Power failure: why small sample size undermines the reliability of neuroscience*. Nat Rev Neurosci, 2013. **14**(5): p. 365-76.
2. Bennett, C.M., G.L. Wolford, and M.B. Miller, *The principled control of false positives in neuroimaging*. Soc Cogn Affect Neurosci, 2009. **4**(4): p. 417-22.
3. Eklund, A., T.E. Nichols, and H. Knutsson, *Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates*. Proc Natl Acad Sci U S A, 2016. **113**(28): p. 7900-5.
4. Ioannidis, J.P.A., *Why Most Published Research Findings Are False*. PLoS Med, 2005. **2**(8): p. e124.
5. Lazar, N.A., et al., *Combining brains: a survey of methods for statistical pooling of information*. Neuroimage, 2002. **16**(2): p. 538-50.
6. Salimi-Khorshidi, G., et al., *Meta-analysis of neuroimaging data: a comparison of image-based and coordinate-based pooling of studies*. Neuroimage, 2009. **45**(3): p. 810-23.
7. Turkeltaub, P.E., et al., *Meta-analysis of the functional neuroanatomy of single-word reading: method and validation*. Neuroimage, 2002. **16**(3 Pt 1): p. 765-80.
8. Laird, A.R., et al., *ALE meta-analysis: controlling the false discovery rate and performing statistical contrasts*. Hum Brain Mapp, 2005. **25**(1): p. 155-64.
9. Eickhoff, S.B., et al., *Coordinate-based activation likelihood estimation meta-analysis of neuroimaging data: a random-effects approach based on empirical estimates of spatial uncertainty*. Hum Brain Mapp, 2009. **30**(9): p. 2907-26.
10. Eickhoff, S.B., et al., *Activation likelihood estimation meta-analysis revisited*. Neuroimage, 2012. **59**(3): p. 2349-61.
11. Turkeltaub, P.E., et al., *Minimizing within-experiment and within-group effects in Activation Likelihood Estimation meta-analyses*. Hum Brain Mapp, 2012. **33**(1): p. 1-13.
12. Wager, T.D., M. Lindquist, and L. Kaplan, *Meta-analysis of functional neuroimaging data: current and future directions*. Soc Cogn Affect Neurosci, 2007. **2**(2): p. 150-8.
13. Radua, J., et al., *Meta-analytical comparison of voxel-based morphometry studies in obsessive-compulsive disorder vs other anxiety disorders*. Arch Gen Psychiatry, 2010. **67**(7): p. 701-11.
14. Radua, J., et al., *A new meta-analytic method for neuroimaging studies that combines reported peak coordinates and statistical parametric maps*. Eur Psychiatry, 2012. **27**(8): p. 605-11.
15. Tench, C.R., et al., *Coordinate based meta-analysis does not show grey matter atrophy in narcolepsy*. Neurosci Biobehav Rev, 2016.
16. Tench, C.R., et al., *Coordinate based meta-analysis of functional neuroimaging data; false discovery control and diagnostics*. PLoS One, 2013. **8**(7): p. e70143.
17. Tench, C.R., et al., *Coordinate based meta-analysis of functional neuroimaging data using activation likelihood estimation; full width half max and group comparisons*. PLoS One, 2014. **9**(9): p. e106735.
18. Costafreda, S.G., *Parametric coordinate-based meta-analysis: valid effect size meta-analysis of studies with differing statistical thresholds*. J Neurosci Methods, 2012. **210**(2): p. 291-300.

19. Ester, M., et al., *A density-based algorithm for discovering clusters in large spatial databases with noise*. Kdd, 1996. **96**(34): p. 226-231.
20. Chen, G., P.A. Taylor, and R.W. Cox, *Is the statistic value all we should care about in neuroimaging?* Neuroimage, 2016.
21. Maumet, C. and T.E. Nichols, *Minimal Data Needed for Valid & Accurate Image-Based fMRI Meta-Analysis*. bioRxiv, 2016.
22. Friston, K.J., et al., *Statistical parametric maps in functional imaging: A general linear approach*. Human Brain Mapping, 1994. **2**(4): p. 189-210.
23. Chevillard, S., *The functions erf and erfc computed with arbitrary precision and explicit error bounds*. Information and Computation, 2012. **216**: p. 72-95.
24. Benjamini, Y. and Y. Hochberg, *Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing*. Journal of the Royal Statistical Society. Series B (Methodological), 1995. **57**(1): p. 289-300.
25. Box, G.E.P. and E.M. Mervin, *A note on the Generation of Random Normal Deviates*. The Annals of Mathematical Statistics, 1958. **29**(2): p. 610-611.
26. Eickhoff, S.B., et al., *Behavior, sensitivity, and power of activation likelihood estimation characterized by massive empirical simulation*. Neuroimage, 2016. **137**: p. 70-85.
27. Talairach, J. and P. Tournoux, *Co-planar stereotaxic atlas of the human brain*. 1988, New York: Thieme.
28. Radua, J., et al., *Anisotropic kernels for coordinate-based meta-analyses of neuroimaging studies*. Front Psychiatry, 2014. **5**: p. 13.
29. Lansley, J., et al., *Localized grey matter atrophy in multiple sclerosis: a meta-analysis of voxel-based morphometry studies and associations with functional disability*. Neurosci Biobehav Rev, 2013. **37**(5): p. 819-30.
30. Tanasescu, R., et al., *Functional reorganisation in chronic pain and neural correlates of pain sensitisation: A coordinate based meta-analysis of 266 cutaneous pain fMRI studies*. Neurosci Biobehav Rev, 2016. **68**: p. 120-33.
31. Lancaster, J.L., et al., *Automated Talairach atlas labels for functional brain mapping*. Hum Brain Mapp, 2000. **10**(3): p. 120-31.
32. Team, R.D.C., *R: A Language and Environment for Statistical Computing*. 2008, R Foundation for Statistical Computing.
33. DerSimonian, R. and N. Laird, *Meta-analysis in clinical trials*. Controlled Clinical Trials, 1986. **7**(3): p. 177-188.
34. Popescu, V., et al., *Grey Matter Atrophy in Multiple Sclerosis: Clinical Interpretation Depends on Choice of Analysis Method*. PLoS One, 2016. **11**(1): p. e0143942.
35. Diggle, P.J., *Statistics: a data science for the 21st century*. Journal of the Royal Statistical Society: Series A (Statistics in Society), 2015. **178**(4): p. 793--813.