

ORIGINAL ARTICLE

Optimal block designs for experiments on networks

Vasiliki Koutra¹ | Steven G. Gilmour¹ | Ben M. Parker²

¹Department of Mathematics, King's College London, London, UK

²Department of Mathematics, Brunel University London, London, UK

Correspondence

Vasiliki Koutra, Department of Mathematics, King's College London, London WC2R 2LS, UK.

Email: vasiliki.koutra@kcl.ac.uk

Funding information

EPSRC, Grant/Award Number: EP/T021624/1

Abstract

We propose a method for constructing optimal block designs for experiments on networks. The response model for a given network interference structure extends the linear network effects model to incorporate blocks. The optimality criteria are chosen to reflect the experimental objectives and an exchange algorithm is used to search across the design space for obtaining an efficient design when an exhaustive search is not possible. Our interest lies in estimating the direct comparisons among treatments, in the presence of nuisance network effects that stem from the underlying network interference structure governing the experimental units, or in the network effects themselves. Comparisons of optimal designs under different models, including the standard treatment models, are examined by comparing the variance and bias of treatment effect estimators. We also suggest a way of defining blocks, while taking into account the interrelations of groups of experimental units within a network, using spectral clustering techniques to achieve optimal modularity. We expect connected units within closed-form communities to behave similarly to an external stimulus. We provide evidence that our approach can lead to efficiency gains over conventional designs such as randomised designs that ignore the network structure and we illustrate its usefulness for experiments on networks.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* published by John Wiley & Sons Ltd on behalf of Royal Statistical Society

KEYWORDS

clustering, connected experimental units, graphs, linear network models, treatment interference

1 | INTRODUCTION

Designing experiments on networks is a growing area of research mainly due to the rise and popularity of online social networks and viral marketing. The work presented here is motivated by the need to develop a practical methodology for obtaining efficient designs which control for variation among the experimental units from two sources: blocks and network interference, so that the true effects of the treatments can be detected.

Consider a commercial experiment on a social networking site to compare the effectiveness of different advertisements concerning a product. The network members are connected via virtual friendships. The responses can be the quantities purchased during the week immediately following the advertising campaign. The company's goal is to maximise the appeal of the product by using the advertisement as a tool to affect the purchasing decisions of potential customers and the experiment will be used as a means for comparing different advertisements. The advertisements used might have an effect not only on the recipient but also on their (virtual) friends. Moreover, there may be blocking structures related to different age groups of customers or cliques of close friends who engage in similar behaviours, for example, making similar decisions as to what quantity of product to purchase. We might expect subjects in the same block to have similar responses and subjects in different (non-overlapping) blocks to have dissimilar responses, irrespective of the presence or absence of viral effects of specific advertisements. By allowing for this in the design, we ensure more precise comparisons of the effects of advertisements.

The methodology demonstrated in this paper is straightforward for moderate sizes of networks of hundreds of vertices. For the motivational example below, the optimal design was found within a few minutes. Applications of such networks are plentiful and examples include agriculture, biology, engineering, marketing, pharmaceuticals and other areas. As motivation, consider the network illustrated in Figure 1, which is a small subset of the Facebook network as obtained from the Stanford network data set collection (snap.stanford.edu/data/egonets-Facebook.html), comprising 324 vertices (Facebook members) and 2514 undirected edges (mutual virtual friendships) forming an ego network. The 'centre' vertex of an ego network (the 'ego') is not included in the graph, which consists of only the ego's friends (its contacts).

A noteworthy work which accommodated interference is Besag and Kempton (1986), which discussed spatial techniques and provided appropriate models to adjust for response or treatment interference in the context of agricultural field experiments. With our interest in the case of treatment interference, we focus on the model of Pearce (1957) presented by Besag and Kempton (1986), which jointly takes into account 'local' (i.e. direct) and 'remote' (i.e. indirect or neighbour) effects in the model. A wide variety of candidate models has been suggested for the case of treatment interference thereafter. Important examples include the work of Druilhet (1999), Kunert and Martin (2000) and Kunert and Mersmann (2011), who provided models and efficient designs that concern experiments where units are arranged in a circle or a line and in which neighbours one unit apart can interfere with each other. More recently, Parker et al. (2016) adopted the conceptual approach of Pearce (1957) and introduced a model, called the linear network effects model (LNM). The LNM is similar to the model of Kunert and Martin (2000). However, it differs by relaxing the assumption of neighbour effects existing in only one direction and allows for a network setting in which units can be linked in a way




FIGURE 1 A Facebook ego network with 24 non-overlapping blocks indicated by colours [Colour figure can be viewed at wileyonlinelibrary.com]

that does not form a regular layout. Here, we extend the LNM of Parker et al. (2016) to include the notion of blocks.

Other previous work has focused on representing block designs as graphs and combining graph theoretic and optimality properties (Bailey, 2007; Bailey & Cameron, 2009, 2013; Wit et al., 2005). In contrast, in our current work the edges are not defined by the block structure and instead correspond to existing connections between experimental units.

Typically the blocking structure in an experiment will be pre-defined, for example similarly situated plots in an agricultural experiment or defined via measured covariates in a medical study. The analysis is adjusted appropriately to control for the block-specific effects (sometimes referred to as nuisance effects) to improve precision. Our methodology directly addresses this case where blocks are known prior to decisions about treatment allocation being made. However, in a setting of a social network it may be desirable or necessary to define the blocks based on the connectivity structure. We provide a general method of finding blocks in the context of a network where similarity of the experimental units is based on the physical or logical densely connected clusters on the respective graph. For example people with similar socioeconomic status might form communities within a contact network which can define the blocks of the experimental design.

After we introduce the block network model (BNM) (Section 2), we define two specific design criteria for estimating with minimum variance the direct treatment comparisons in the presence of nuisance network effects or the network effects themselves (Section 3). We describe a systematic exchange algorithm used to obtain our designs (Section 4). We then make some comparisons among optimal designs under different models and investigate issues associated with the design efficiency and bias arising in the analysis of dependent data (Section 5). When blocks arise as cliques in the network, we suggest a method to define the blocks using spectral clustering and the concept of modularity. We give an example illustrating the recommended step-by-step methodology for the special case of two unstructured treatments (Section 6). Finally, we discuss some practical concerns emerging from the suggested methodology and further issues of interest (Section 7).

2 | STATISTICAL MODEL

Suppose that n experimental units are available for experimentation and that they form a network which is represented by means of a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with vertex set \mathcal{V} (of size n) and edge set \mathcal{E} (of size l). The vertices might represent individuals connected through edges which might indicate some form of connection such as friendship, collaboration or communication. The *adjacency matrix* of a graph is an $n \times n$ matrix $A = [A_{jh}]$ with $j, h \in V$, which is a compact way to represent the collection of edges. We focus on connected, undirected and unweighted graphs with $A_{jh} = A_{hj} \in \{0, 1\}$ representing the absence or presence of an edge between the vertices j and h (mutual connections). Thus A is symmetric and binary. The diagonal elements represent the self-links, which by convention are set to zero. In an undirected network, the *degree*, d_j , of the vertex $v_j \in \mathcal{V}$ denotes the number of vertices that vertex v_j is connected to (number of neighbours), that is, $d_j = \sum_{h=1}^n A_{jh}$.

We assume that there are groups of experimental units which are expected to give similar responses, for example, because they share properties. We label these groups $1, \dots, \kappa$ and let group g have $n_{(g)}$ experimental units within it. We model the responses of experimental units by the block network model (BNM), which is an extension of the LNM (Parker et al., 2016) and is described by the equation

$$\text{BNM: } y_{ij} = \mu + \tau_{r(ij)} + b_i + \sum_{g=1}^{\kappa} \sum_{h=1}^{n_{(g)}} A_{\{ij,gh\}} \gamma_{r(gh)} + \epsilon_{ij},$$

where $i = 1, 2, \dots, \kappa; j = 1, 2, \dots, n_{(i)}$, y_{ij} is the continuous response from unit j in the i th block receiving the treatment $s = r(ij) \in \{1, \dots, m\}$, μ represents the response for a baseline treatment or (unit) average, b_i is the effect of block i , $\tau_{r(ij)}$ is the (direct) treatment effect, $A_{\{ij,gh\}}/A = A_{\{ij,gh\}}$ with $j, h \in V$ is the adjacency matrix indicating the edge between units j and h belonging to blocks i and $g \in \{1, 2, \dots, \kappa\}$ respectively, $\gamma_{r(gh)}$ is the network effect (neighbour or indirect treatment effect) and ϵ_{ij} are the errors, which we assume to be independent and identically distributed with mean 0 and constant variance σ^2 .

To overcome the model overparametrisation requires imposing some constraints, otherwise the normal equations have an infinite number of solutions and our parameters cannot be uniquely estimated. In our examples, without loss of generality, we assume the set-last-to-zero linear constraints with the m th treatment effect τ_m and κ th block effect b_{κ} to be set equal to zero. A wider set of network structures can be tackled by introducing a constraint on the model parameters corresponding to the network effects, that is, $\gamma_m = 0$. For example, in practical terms that would allow regular graphs to be considered.

The expectation of the response for this model in matrix form is

$$\mathbb{E}[\mathbf{y}] = (\mathbf{1} \mathbf{u}_1 \dots \mathbf{u}_{m-1} \mathbf{w}_1 \dots \mathbf{w}_{\kappa-1} \mathbf{A} \mathbf{u}_1 \dots \mathbf{A} \mathbf{u}_m) \boldsymbol{\beta},$$

where $\boldsymbol{\beta} = (\mu \ \boldsymbol{\tau}^T \ \mathbf{b}^T \ \boldsymbol{\gamma}^T)^T = (\mu \ \tau_1 \dots \tau_{m-1} \ b_1 \dots b_{\kappa-1} \ \gamma_1 \dots \gamma_m)^T$ is the vector of parameters. There are no columns corresponding to the m th treatment effect and the κ th block effect since we have assumed them to be zero. Following Bailey (2008) Section 2.7), vector \mathbf{u}_s corresponds to an $n \times 1$ vector where for each of the treatments ($s = 1, 2, \dots, m$) all the elements of the vector equal zero, except for those which correspond to the units which receive that treatment and are equal to one. For instance \mathbf{u}_1 is the indicator vector with ones for the unit(s) receiving treatment $s = 1$ and zeros elsewhere. The same holds for the vector \mathbf{w}_i for block i . The symmetric information matrix M is

$$M = \begin{pmatrix} n & \mathbf{1}^T X_\tau^* & \mathbf{1}^T X_b^* & \mathbf{1}^T A X_\tau \\ X_\tau^{*T} \mathbf{1} & X_\tau^{*T} X_\tau^* & X_\tau^{*T} X_b^* & X_\tau^{*T} A X_\tau \\ X_b^{*T} \mathbf{1} & X_b^{*T} X_\tau^* & X_b^{*T} X_b^* & X_b^{*T} A X_\tau \\ X_\tau^T A \mathbf{1} & X_\tau^T A X_\tau^* & X_\tau^T A X_b^* & X_\tau^T A^2 X_\tau \end{pmatrix},$$

where X_τ and X_b are the incidence matrices for treatment and block effects respectively and X_τ^* and X_b^* are the same matrices with their last columns deleted. We should note at this point that only certain types of network structures will enable us to estimate the network effects. With regular graphs the estimation of network effects is not possible due to the singular information matrix and therefore imposing additional restrictions on the values of the γ parameters may be appropriate for such cases.

3 | DESIGNS WITH BLOCK AND NETWORK EFFECTS

In order to determine the optimal block designs we use the L -optimality criterion that minimises the average variance of the estimators of a pre-specified set of linear functions of the parameters $S^T \boldsymbol{\beta}$, where S corresponds to multiple vectors of known constants (Atkinson et al., 2007, Ch. 10). The optimality criterion function to be minimised is thus a scalar function of $\text{var}(S^T \hat{\boldsymbol{\beta}})$, where $\hat{\boldsymbol{\beta}}$ is the least squares estimator of $\boldsymbol{\beta}$ which is proportional to $S^T \{M(\xi)\}^{-1} S$, where $M(\xi) = M$ corresponds to the information matrix and ξ is a design chosen from Ξ the set of all possible designs, where the design is a choice of treatment assignments to the experimental units that correspond to the vertices of the network. The minimisation of the optimality criterion function

$$\phi(\xi) = \text{Tr}(S^T \{M(\xi)\}^{-1} S) = \text{Tr}(\{M(\xi)\}^{-1} L),$$

where $L = SS^T$, leads to an L -optimal design. For the L -optimal design ξ^* , $\phi^* = \phi(\xi^*) = \min_{\xi \in \Xi} \phi(\xi)$, is the *optimal function value*.

To obtain these pairwise differences we have to explicitly define the s vectors that compose S .

Definition Let $s(\alpha_1, \alpha_2)$ be a contrast vector formed in the following way:

1. Form a vector of $2m + \kappa + 1$ zeroes, corresponding to the intercept, m treatment effects, $\boldsymbol{\tau}$, followed by κ block effects, \mathbf{b} , followed by m network effects, $\boldsymbol{\gamma}$.
2. Let the α_1 th and α_2 th elements be 1 and -1 respectively, corresponding to the particular effects we wish to estimate in a given contrast.
3. Delete the $(m + 1)$ th and $(m + \kappa + 1)$ th elements, corresponding to our restriction that $\tau_m = 0$ and $b_\kappa = 0$ to estimate the treatment and block effects uniquely.

The resulting vectors of length $(2m + \kappa - 1)$ are pre- and post-multiplied by the $(2m + \kappa - 1) \times (2m + \kappa - 1)$ matrix M^{-1} in the summation that defines the optimality criteria ϕ_1 and ϕ_2 which we now define.

We seek to minimise the average variance of all pairwise differences of treatment effects

$$\frac{2}{m(m-1)} \sum_{s=1}^{m-1} \sum_{s'=s+1}^m \text{var}(\widehat{\tau_s - \tau_{s'}}),$$

which is proportional to

$$\phi_1 = \sum_{v=2}^m \sum_{q=v+1}^{m+1} s^T(v, q)M^{-1}s(v, q).$$

Alternatively we seek to minimise the average variance of all pairwise differences of network effects

$$\phi_2 = \sum_{v=m+\kappa+2}^{2m+\kappa} \sum_{q=v+1}^{2m+\kappa+1} s^T(v, q)M^{-1}s(v, q).$$

For instance, for the case of $m = 2$ treatments and $\kappa = 2$ blocks the criteria become

$$\begin{aligned} \phi_1 &= s^T(2, 3)M^{-1}s(2, 3) \\ \phi_2 &= s^T(6, 7)M^{-1}s(6, 7), \end{aligned}$$

where $s(2, 3) = (0 \ 1 \ 0 \ 0 \ 0)^T$ and $s(6, 7) = (0 \ 0 \ 0 \ 1 \ -1)^T$. The first corresponds to the estimate of the treatment effect difference, and the second to the estimate of the difference between network effects.

We will often compare the performance of two designs using their *relative efficiency*, which, with respect to the objective function ϕ of a design ξ_2 compared with a design ξ_1 , is given by $\text{Eff}(\xi_1, \xi_2) = \phi(\xi_1)/\phi(\xi_2)$. We can also define the *L-efficiency* of a design ξ as $\text{Eff}(\xi) = \text{Eff}(\xi^*, \xi)$, where ξ^* is the *L-optimal* design. We return to these definitions when assessing the performance of different optimal designs based on different models for a number of given networks.

4 | EXCHANGE ALGORITHM

When exhaustive search to find an optimal design is not possible due to the large number of units and/or treatments, we implement an approximate method. In the literature, there exist a number of computationally efficient algorithms for finding near-optimal designs in a practical amount of time using iterative methods, for example, Fedorov (1972) and Cook and Nachtsheim (1980). The main steps involved in the majority of those algorithms are the following: (i) initialisation of the search, for example, random generation of a non-singular design (i.e. the matrix M must be non-singular to ensure that the parameters in ϕ are estimable); (ii) modification of the current solution, for example, make exchanges in the treatment set; (iii) assessment of new solution, that is, design is assessed with respect to an objective function. Steps (ii)–(iii) are repeated until no change improves the design value; (iv) termination of the search process and return of the final design (which is assumed to be optimal).

We adapt the ‘Modified Fedorov Exchange Algorithm’ of Cook and Nachtsheim (1980) to obtain an exchange algorithm for finding near-optimal designs for unstructured treatments on networks. The algorithm can be described as follows:

Point Exchange on Networks (PEN) Algorithm

1. Step 1: Generate a random initial design (i.e. random assignment of a treatment label to each unit). Units are labelled from 1 to n .
2. Step 2: Calculate the optimality criterion function ϕ (ϕ_1, ϕ_2 or any other defined criterion) for the arbitrary design of step 1.

3. Step 3: Iterate from $j = 1$ to n to find a better design (where j corresponds to the j -th unit). In the j th iteration:
 - a. Exchange the treatment applied to unit j with another treatment.
 - b. Calculate the information matrix. If it is non-singular, compute the new value of the chosen criterion. Otherwise go to d.
 - c. If the new criterion function value is better than the current one, keep the exchange of treatments. Otherwise, undo the exchange.
 - d. If $j < n$, set $j \leftarrow j + 1$.
4. Step 4: Repeat step 3 until a complete pass of all the vertices yields no further improvement and then go to step 5.
5. Step 5: Rerun all the above steps for several randomly generated initial designs and return the design with the lowest criterion function value ϕ^* . This is the near-optimal design, which will usually be globally optimal, although this cannot be guaranteed. End.

The last step is required to obtain an efficient final design (and overcome the problem of becoming stuck in a local optimum) by using multiple random initialisations. For the particular problems of interest in this work, the PEN algorithm appears to be powerful enough and simple to implement.

5 | COMPARISON OF OPTIMAL DESIGNS UNDER DIFFERENT MODELS

We aim to construct experimental designs on networks that are efficient for estimating treatment or network effects. However, we want to investigate the bias of the estimated model parameters introduced by possible model misspecification. In this section we provide comparisons of optimal designs for estimating the treatment and network effects under traditional models that ignore network interference and models that account for it. In doing so we compare the optimal designs with network effects to the randomised balanced designs (equal replication), providing evidence that depending on the outcome of the randomisation, the balanced design is typically not very good. In the majority of cases designs that account for the network structure have higher efficiency than the standard designs.

We also obtain the bias introduced in the parameter estimators due to the model misspecification as a function of the unknown model parameters. The models we consider are the model, CRM, derived from the completely randomised design, the model, RBM, derived from the generalised randomised block design (which allows blocks potentially to be different sizes), the linear network effects model, LNM, and the block linear network model, BNM, defined as:

$$\text{CRM: } y_j = \mu + \tau_{r(j)} + \epsilon_j$$

$$\text{RBM: } y_{ij} = \mu + \tau_{r(ij)} + b_i + \epsilon_{ij}$$

$$\text{LNM: } y_j = \mu + \tau_{r(j)} + \sum_{h=1}^n A_{\{j,h\}} \gamma_{r(h)} + \epsilon_j$$

$$\text{BNM: } y_{ij} = \mu + \tau_{r(ij)} + b_i + \sum_{g=1}^{\kappa} \sum_{h=1}^{n(g)} A_{\{ij,gh\}} \gamma_{r(gh)} + \epsilon_{ij}$$

In all cases, we assume that the errors are independent and randomly distributed with zero mean and constant variance. The CRM and RBM are the standard treatment models that are derived from randomisation schemes without and with blocks respectively. The designs for which these models are

appropriate are the simplest forms of designs to compare different treatments by randomly assigning them to experimental units (which for RBM are additionally arranged in κ blocks). The LNM and BNM are extensions of CRM and RBM respectively, including a network term for capturing the connections between units. Figure 2 illustrates the hierarchy among the models by means of a Hasse diagram (see, e.g. Bailey, 2008). Here, we use the Hasse diagram to describe the collection of models considered, with dots representing models and lines representing nesting relationships between models. If the true model is a submodel of the assumed model, the bias in estimating treatment effects will be zero. This will be further explained in the following examples. Designs which are optimal for each of these models are labelled CRD (for the Completely Randomised Design), RBD (for the Randomised Block Design), LND (for the Linear Network Design) and BND (for the Block Network Design) respectively.

We now compare the properties of each design for each model. In doing so we obtain the optimal function values for the two optimality criteria considering each model. Under the assumption of independent errors that have a common variance σ^2 , the variance–covariance matrix of the least squares estimator $\hat{\beta}$ is $\text{var}(\hat{\beta}) = \sigma^2 M^{-1}$. Given that interest lies in the comparisons of the designs, the value σ^2 is not relevant since the value is the same if the model is identical for all proposed designs for a particular experiment.

5.1 | Efficiencies of randomised designs

If we ignore the network structure in our design, we would generally assume that all experimental units in the network (for the CRD) or within a block (for the RBD) were exchangeable. There are therefore a large number of designs which we could choose as optimal under the CRD or RBD. In general, we would choose one of these equivalent designs at random, and we refer to this as a randomised design.

In this section we compare the optimal designs with network effects to the randomised balanced designs (while ignoring the network structure), providing evidence that when under randomisation, the balanced design is typically not very good. In the majority of cases designs that account for the network structure have higher efficiency than the standard designs. The optimal randomised designs, under the CRM and the RBM, have near-equal replication of each treatment.

Example 1 The social network in Figure 3 is a real-world co-authorship network, comprising 22 subjects (PhD students, supervisors and co-authors) and 27 edges, which indicate the ties between

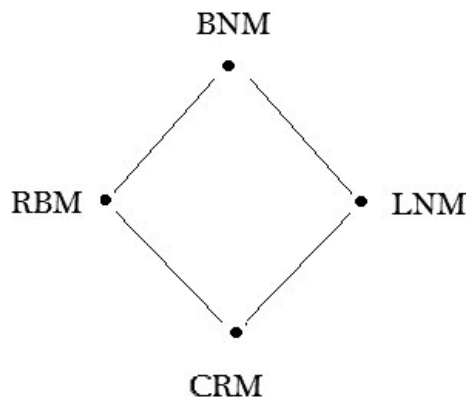


FIGURE 2 Hasse diagram for the collection of models




FIGURE 3 Social network with three blocks indicated by colours [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

individuals with common publications, or ties between PhD students and their supervisors. Let us assume that we want to conduct an experiment on this network in order to compare two distinct treatments. Consider a small hypothetical experiment on these work colleagues, for example, to compare different text messages with the object of raising charitable funds. The responses of the subjects could be measured by means of their donations. The subjects work together and might hold discussions with each other. Therefore, the text messages could have viral effects, as such affecting not only the person receiving the text message but also that person's collaborators. The three blocks have been defined following the method we suggest in Section 6, but any other means of defining blocks could be addressed in the same way. The resulting blocks, which emerged from the connectivity structure, nicely captured collaboration communities within a similar research area, who are likely to spend lots of time working on common projects.

We obtain the L -optimal designs (using exhaustive search) for this network which are illustrated in Figure 4. Overall the optimal design for ϕ_1 has all treatments equally allocated to the subjects, that is, 11 subjects receive each of the two treatments. Subjects allocated within each treatment have similar first and second order degrees (numbers of connections between units of distance one or two). The design for ϕ_1 also has equal replication within each block. For ϕ_2 , the design is highly dependent on the particular network structure. For instance subjects located at the ends of the network tend to receive the same treatment which is different from their better connected immediate neighbours.

Tables 1 and 2 illustrate the optimality function values for each design (columns) under the different models (rows). We can obtain the efficiencies of the designs with respect to the optimal designs. The criterion function values of the optimal designs are highlighted in bold. Recall that the smaller the criterion function value the better the design is.

A key observation derived from these tables is that a randomised design which ignores the network effects is on average highly inefficient, especially with respect to ϕ_2 . Note the values for CRD and RBD are the mean values over all 705,432 and 105,840 possible randomisations respectively.




FIGURE 4 Optimal designs (under BNM), ϕ_1 (left) and ϕ_2 (right); different colours indicate different blocks and different vertex shapes different treatments [Colour figure can be viewed at wileyonlinelibrary.com]

TABLE 1 Comparisons of the designs for ϕ_1 under different models

Models	Optimal designs			
	CRD	RBD	LND	BND
CRM	0.1818	0.1818	0.1818	0.1818
RBM	0.2034	0.1818	0.2685	0.1818
LNM	0.2126	0.2191	0.1822	0.1827
BNM	0.2500	0.2270	0.3621	0.1828

TABLE 2 Comparisons of the designs for ϕ_2 under different models

Models	Optimal designs			
	CRD	RBD	LND	BND
LNM	0.1447	0.1927	0.0237	0.0366
BNM	0.2354	0.2503	0.0998	0.0388

In Table 1, we can observe that when assuming the LNM to be true, the BND, which has equal replication within blocks, is almost as efficient as the LND (with efficiency $0.1822/0.1827 = 99.7\%$). Under BNM the LND, which has equal replication overall but not within blocks, is 50.4% efficient. The poor performance of the LND relative to the BND under the BNM results from having a completely different structure and different allocation to unequal and relatively small-sized blocks. BND is balanced overall but also balanced within blocks. On the other hand, the optimisation process for the LNM drives the design away from having an equal treatment allocation within blocks, resulting in many neighbouring vertices receiving the same treatment. Moreover, under the BNM, a RBD which has equal replication within blocks performs better than LND on average. This means that, if we account for the blocks, we almost certainly do better. Another interesting observation is that CRD performs better than LND on average under BNM. This is because we consider all possible randomisations implying that in general randomisation will be beneficial if blocks are not known or included. It is worth noting that the optimal function values, which are 0.1822 (under the LNM)




FIGURE 5 Boxplots of efficiencies for ϕ_1 (assuming the BNM)

and 0.1828 (under the BNM) are approximately equal to the minimum average variance possible for the unstructured case when having independent units (i.e. $\sigma^2/n_1 + \sigma^2/n_2 = 2\sigma^2m/n = 0.1818\sigma^2$ with $n_1 = n_2 = n/2$).

In Table 2 the CRD and RBD perform poorly under both network models with their efficiencies being on average 16.4% and 16.5% under the LNM and BNM respectively when ignoring blocks (for CRD), and 12.3% and 15.5% under the LNM and BNM respectively when accounting for blocks (for RBD). These low efficiencies can be explained by the fact that these standard designs ignore the indirect treatment effects resulting from the network structure among units. We are able to estimate the network effects only under the LNM and the BNM. Moreover, when assuming the BNM to be true, the LND is only 39% efficient.

The boxplots in Figure 5 show the L -efficiencies of the designs based on ϕ_1 for all CRDs and RBDs. For ϕ_1 , the CRD and RBD have median efficiencies 0.78 and 0.83 respectively, with minimum efficiencies 0.08 and 0.38 respectively. On the other hand for ϕ_2 in Figure 6, CRD and RBD perform similarly with median efficiencies being 0.17 and 0.18, the lower quartiles being 0.05 and 0.10 and the upper quartiles being 0.62 and 0.55. This results from ignoring the network effects. The LND is approximately 50% and 40% as efficient as the BND for ϕ_1 and ϕ_2 respectively.

Example 2 The second example network (see Figure 1) is a subset of the Facebook network where we wish to compare two treatments, which could correspond, for instance, to advertisements as discussed earlier.

The comparisons of the (near-)optimal designs under ϕ_1 and under ϕ_2 are given in Tables 3 and 4 respectively. Note the values for CRD and RBD are the mean values over a large number of possible randomisations (we have evaluated 50,000 designs). Figures 7 and 8 depict the boxplots of efficiencies under L -optimality of the design based on ϕ_1 and ϕ_2 respectively for a number of possible designs (chosen at random) for CRD and RBD which do not take into account network effects, although they exist.

A key observation derived from these tables is that a randomised design, which ignores the network effects is on average highly inefficient, at least with respect to ϕ_2 . We can observe that when




FIGURE 6 Boxplots of efficiencies for ϕ_2 (assuming the BNM)

TABLE 3 Comparisons of the designs for ϕ_1 under different models

Models	Optimal designs for ϕ_1 ($\times 10^2$)			
	CRD	RBD	LND	BND
CRM	1.2346	1.2347	1.2346	1.2346
RBM	1.3298	1.2432	1.3042	1.2432
LNM	1.2481	1.2747	1.2346	1.2348
BNM	1.3749	1.2907	1.3177	1.2432

TABLE 4 Comparisons of the designs for ϕ_2 under different models

Models	Optimal designs for ϕ_2 ($\times 10^2$)			
	CRD	RBD	LND	BND
LNM	0.1121	0.1474	0.0119	0.0149
BNM	0.1553	0.1647	0.0312	0.0230

assuming LNM to be true, the BND, which is completely balanced within blocks, is almost as efficient as the LND. Under BNM the LND, which is balanced overall but unbalanced with respect to blocks, is 93.9% efficient. Moreover, under BNM, a random balanced RBD which is balanced within blocks performs better than LND on average, with the latter being as efficient as a random balanced CRD on average. In Table 4 the random balanced designs CRD and RBD perform poorly under both models with their efficiencies ranging on average between 10.6% and 14.8% respectively when ignoring blocks and 8.1% and 13.9% respectively when accounting for blocks. These low efficiencies arise because these standard designs ignore the spillover effects resulting from the network structure among units.

The boxplots in Figure 7 depict the efficiencies under L-optimality of the design based on ϕ_1 for a number of possible balanced designs (as chosen at random) for CRD and RBD which do not take




FIGURE 7 Boxplots of efficiencies for ϕ_1




FIGURE 8 Boxplots of efficiencies for ϕ_2

into account network effects, although they exist. The outliers are detected by setting the upper/lower ends of the whiskers at three standard deviations. For ϕ_1 CRD and RBD have median efficiencies 0.92 and 0.97 respectively, with minimum efficiencies 0.68 and 0.76 respectively. On the other hand for ϕ_2 in Figure 8, CRD and RBD perform similarly with median efficiencies between 0.10 and 0.25, lower quartiles between 0.15 and 0.17 and upper quartiles between 0.27 and 0.29. This results from not taking into account the network effects. For ϕ_1 , LND performs on average better than most of the balanced CRDs, but worse than the majority of balanced RBDs, while for ϕ_2 LND is approximately 70% as efficient as BND.

5.1.1 | Recommendations

These and other design examples led us to make the following general guidelines:

1. On average the randomisation process assuming no network structure leads to fairly poor designs, that is high variance (this is especially true under ϕ_2). For Example 2, all designs for ϕ_2 were less than 35% efficient which indicates that randomised designs are highly inefficient.
2. When ignoring both block and network effects, the design can perform poorly. When taking into account the network effects but not the block effects, the design performs satisfactorily for ϕ_2 but can perform poorly for ϕ_1 when there is strong evidence of an underlying block structure.
3. By using blocks and ignoring network effects, with approximately a 25% chance (upper quartile) one can do just as well as when taking into account the network effects, but on average one does worse (for ϕ_1).
4. Evidence suggests that we will be better off by using blocks instead of ignoring them, independently of whether we are taking into account network effects (for ϕ_1).

The optimal designs under the second optimality criterion are generally unbalanced, that is, the treatments have unequal replication in the design. Therefore, it is not surprising that the standard designs perform poorly with median efficiencies lower than 35%. Thus in practice, when there is a strong belief that the units are governed by a network structure and/or when blocking structure exists, then it is of importance to account for those in order to obtain an efficient design.

5.2 | Bias due to model misspecification

In the previous section, we showed that randomised designs that ignore network effects are not efficient when used to fit the BNM. However, practitioners might use standard randomised designs and analyse their data using models that ignore network effects such as CRM and RBM. We now show that when we randomise and fit the corresponding model, the design can result in biased estimators of the quantities of interest. Although the bias has not been taken into account in the design optimality criteria, we can obtain it as a function of the model parameters and show that, under the false assumption that there are no network effects, the bias can be practically important.

We employ Example 1 to investigate the potential impact of model bias. We consider wrongly fitting a simpler reduced or nested model when a more complicated model is true. For example, when fitting a model with no network effect (e.g. assuming RBM) we calculate the expectation of the bias introduced in the parameter estimates for all possible balanced designs (which are optimal when there is no network effect) in terms of the true unknown parameters (for true BNM). In the Appendix, we derive the design bias in the least squares estimators of the model parameters for the RBM, \mathcal{W} , due to model misspecification, that is,

$$\mathcal{W} = \begin{pmatrix} \mathbf{0} B^{-1} \Gamma \\ \mathbf{0} -I_m \end{pmatrix} \beta,$$

where

$$B = \begin{pmatrix} \mathbf{1}^T \mathbf{1} & \mathbf{1}^T X_\tau^* & \mathbf{1}^T X_b^* \\ X_\tau^{*T} \mathbf{1} & X_\tau^{*T} X_\tau^* & X_\tau^{*T} X_b^* \\ X_b^{*T} \mathbf{1} & X_b^{*T} X_\tau^* & X_b^{*T} X_b^* \end{pmatrix}, \quad \Gamma = \begin{pmatrix} \mathbf{1}^T A X_\tau \\ X_\tau^{*T} A X_\tau \\ X_b^{*T} A X_\tau \end{pmatrix},$$

and I_m is the $m \times m$ identity matrix.

Note that any parameters absent from the assumed model are estimated as zero, leading directly to a bias equal to the negative of the true parameter value. We obtain the biases under the different models with respect to the model parameters in Table 5. Each bias results from fitting an incorrect model and the fitted model is the same as that assumed when finding the design. Thus we can perform an experiment where we wrongly assume that there are either no network effects or no blocking effects, while the true model is of the form BNM. Note that if the true model is a submodel of the assumed model, the bias in estimating treatment effects will be zero. Moreover, the expressions resulted from averaging the biases across all possible balanced randomised designs. We see that if the network effects are substantial compared to the (direct) treatment effects, then ignoring them can potentially lead to over- or under-estimated treatment effects. Thus, by not taking into account network effects in our design, we produce an experiment which can have higher variance than necessary if the correct model is fitted and biased estimators if a simpler model is fitted.

We can extract the bias of the treatment effects estimates (the second line in each matrix). Assuming that we wrongly ignore network effects, we have the following two cases of the bias for the direct treatment effects estimator, where the expectation is taken over all possible randomisations

$$\text{CRD under LNM: } \mathbb{E} [\hat{\tau}_1] - \tau_1 = -0.48\gamma_1 - 0.24\gamma_2, \tag{1}$$

$$\text{RBD under BNM: } \mathbb{E} [\hat{\tau}_1] - \tau_1 = -0.72\gamma_1 - 0.05\gamma_2. \tag{2}$$

If the true values of the network effects are zero (so that $\gamma_1 = \gamma_2 = 0$), the balanced design will produce unbiased estimators. We should note that the network effects are not in general of the same magnitude. However, for our investigation we assume them to be equal ($\gamma_1 = \gamma_2 = \gamma$), and without loss of generality we set $\gamma = 1$. Under this assumption the bias in (2) is slightly larger than

TABLE 5 Bias under model misspecification

True model	Optimal designs			
	CRD	RBD	LND	BND
CRM	$\mathbf{0}_{2 \times 2}$	$\mathbf{0}_{2 \times 2}$	$\mathbf{0}_{2 \times 2}$	$\mathbf{0}_{2 \times 2}$
RBM	$\begin{bmatrix} 0 & 0 & 0.17 & 0.34 \\ 0 & 0 & 0.02 & 0.03 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \begin{pmatrix} \mu \\ \tau_1 \\ b_1 \\ b_2 \end{pmatrix}$	$\mathbf{0}_{4 \times 4}$	Non-nested model	$\mathbf{0}_{4 \times 4}$
LNM	$\begin{bmatrix} 0 & 0 & 1.28 & 1.53 \\ 0 & 0 & -0.48 & -0.24 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \begin{pmatrix} \mu \\ \tau_1 \\ \gamma_1 \\ \gamma_2 \end{pmatrix}$	Non-nested model	$\mathbf{0}_{4 \times 4}$	$\mathbf{0}_{4 \times 4}$
BNM	$\begin{bmatrix} 0 & 0 & 0.17 & 0.34 & 1.28 & 1.53 \\ 0 & 0 & 0.02 & 0.03 & -0.48 & -0.24 \\ 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 \end{bmatrix} \begin{pmatrix} \mu \\ \tau_1 \\ b_1 \\ b_2 \\ \gamma_1 \\ \gamma_2 \end{pmatrix}$	$\begin{bmatrix} 0 & 0 & 0 & 1.18 & 1.60 \\ 0 & 0 & 0 & -0.72 & -0.05 \\ 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0.05 & -0.45 \\ 0 & 0 & 0 & 0.56 & -0.21 \end{bmatrix} \begin{pmatrix} \mu \\ \tau_1 \\ b_1 \\ b_2 \\ \gamma_1 \\ \gamma_2 \end{pmatrix}$	$\begin{bmatrix} 0 & 0 & 0.32 & 0.08 & 0 & 0 \\ 0 & 0 & -0.18 & 0.52 & 0 & 0 \\ 0 & 0 & -0.01 & 0.15 & 0 & 0 \\ 0 & 0 & -0.02 & -0.13 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 \end{bmatrix} \begin{pmatrix} \mu \\ \tau_1 \\ b_1 \\ b_2 \\ \gamma_1 \\ \gamma_2 \end{pmatrix}$	$\mathbf{0}_{6 \times 6}$

in (1), that is, $|0.77| > |0.72|$. A conclusion drawn from this example is that if we use a model that ignores network effects there is no obvious benefit in terms of bias from including blocks. Thus having blocks does not insure us against the bias introduced by wrongly excluding network effects.

In a completely randomised design there is a 50% chance of any specific pair of units receiving both treatments 1 and 2. To investigate if the bias from treatment effects results from network effects passed on from units receiving treatment 1 to those receiving treatment 2, we focus on the number of links connecting the different treatments, which we call l_{12} . The plot in Figure 9 shows the bias when assuming an optimal RBD under the true BNM in treatment effect estimates, as a proportion of the true γ , against the proportion of edges, which connect pairs of units receiving different treatments. Thus the expectation of the bias is over a conditional distribution of randomisations given a fixed proportion of edges. The locations of the plotting symbols are related to the coefficients of the network effects in the bias equation under each design and are dependent on the size of the true γ . As such they correspond to the bias from RBDs for estimating the treatment effects due to the network effects under the assumption that the underlying parameters are equal to one. The expectation of the bias over all designs for the intersection point, which corresponds to approximately half for the proportion of links, is equal to zero. It should be noted that many designs are overlapping and each location of the plotting symbols can represent the bias due to network effects for alternative RBDs. Table 6 illustrates the number of balanced designs for each proportion of edges (connecting units receiving treatment 1 to units receiving treatment 2).





FIGURE 9 Bias in treatment effects due to network effects [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

TABLE 6 Number of balanced designs for each proportion of edges

# designs	38	322	1532	4710	9460	12742	11838	7674	3540	952	112
Prop. of edges	0.37	0.41	0.44	0.48	0.52	0.56	0.59	0.63	0.67	0.7	0.74

FIGURE 10 Bias for treatment effects when $\gamma_1 = \gamma_2$

The boxplots in Figure 10 depict the bias in the estimation of treatment effects over all possible RBDs in the case that we assume that $\gamma_1 = \gamma_2$ for every proportion of links. This expectation is conditional on the proportion of links being fixed, but averaged over the subset of randomisations which respect that. The limits of the boxplots are relative to the size of the true network effect. This plot suggests that in general the complete randomisation while ignoring the network structure does not perform well. In combination with the previous observation, we can conclude that the least expected bias in the treatment effects is achieved when the number of pairs of connected units which receive different treatments equals approximately half the total number of edges of the network (see Figure 10 at the proportion 0.56).

The experimenter should impose restrictions on the allocation of treatment combinations such that the number of connected pairs of units receiving different treatments roughly equals half of the total number of edges in the network. Restrictions on the randomisation in such a way enables us to protect the experimental results against bias in treatment effects stemming from potential network effects.

6 | DEFINING BLOCKS USING SPECTRAL CLUSTERING

In our model, we assume that individuals within the same community exhibit similar responses, which are different to those in other communities, irrespective of the presence or absence of viral effects of specific treatment. Sometimes these communities are defined by features not directly related to the graph structure, such as age group, sex, nationality, etc. At other times they are defined by the structure of the graph itself. Then we have to deduce the community structure by using a clustering algorithm.

Spectral clustering, as the name implies, makes use of the graph spectrum, that is, the set of eigenvalues and eigenvectors of a matrix describing the graph structure, to group vertices into communities. The graph spectrum plays a major role in the understanding of the structure and dynamics in the network since it is linked with numerous graph properties (Chung, 1997; Von Luxburg, 2007). Spectral clustering is a versatile clustering method which can be applied to any network and is considered to give good quality solutions. To identify the number of clusters, we additionally use the modularity of the network, which is a measure which quantifies the strength of the community structure in a network (Newman, 2006).

Apart from the adjacency matrix, a graph \mathcal{G} can be represented by other connectivity matrices such as the Laplacian matrices. The main differences between spectral clustering techniques lie in the choice of the Laplacian matrix. We use the normalised Laplacian graph, $L_{rw} = I - D^{-1}A$, where D is the $n \times n$ diagonal degree matrix (where the entries are the degrees of the vertices). In fact, $D^{-1}A$ is the transition matrix of a standard random walk on the given graph, making it a useful tool for capturing a diffusion process, such as the treatment propagation effects, in a network. Recall that a walk is a sequence of edges connected to a sequence of vertices, where vertices can appear more than once. A random walk is a walk across a network created by taking repeated random steps. The vertices with high degree are more likely to be visited by the random walk because there are more ways of reaching them.

For detecting communities we implement the normalised spectral clustering algorithm, namely the Shi and Malik (SM) algorithm (Shi & Malik, 2000) as described by Von Luxburg (2007). We describe the SM algorithm in the Appendix with appropriate adjustments for the purposes of this work. This clustering algorithm gives no hint about the choice of the number of clusters or ‘goodness’ of each partition. To quantify the quality of partitions and to choose the number of communities based on the ‘best’ partition, we use modularity (Newman & Girvan, 2004), which is defined as the difference between the fraction of the edges that fall within clusters and the fraction of the edges that would be expected to fall within the clusters if the edges were assigned randomly but keeping the degrees of the vertices unchanged. When the communities are not stronger than the random partition or when the network does not exhibit any community structure, the modularity score Q is zero or negative (see Appendix for the detailed definition). As such κ is regarded as the appropriate candidate for indicating the number of intrinsic communities which will define the blocks to be used in the design process, if the κ th partition corresponds to the highest modularity score over all partitions.

Revisiting Example 2 we provide and implement the step-by-step methodology.

Step one: Blocking structure. Expressing the topology of the network through the normalised Laplacian matrix L_{rw} , we produce a number of possible network partitions for different values of κ (dimensionality of the eigenvector space, that is, the number of κ eigenvectors) using the spectral clustering algorithm SM. Then we assess all the partitions produced using the quality function of modularity and choose the number of communities (fixed κ) to be used for the block design, which maximises the value of modularity. For the social network at hand the maximum modularity value over all possible partitions is found for $\kappa = 24$ clusters. Figure 11 illustrates the modularity scores for different numbers




FIGURE 11 Modularity values for partitions of $\kappa = 2, \dots, 162$ [Colour figure can be viewed at wileyonlinelibrary.com]

of communities obtained via the graph partition. This screening stage helps us to choose the number of communities, where the modularity score takes its maximum value. The resulting clustering is illustrated in Figure 1 with different colours indicating the different 24 blocks.

Step two: Optimal design with block and network effects (using PEN). The optimal function values resulting from this search (allowing for multiple initial designs) were found to be $\phi_1^* = 0.0124$ and $\phi_2^* = 0.0002$. Similar to the first example, ϕ_1^* is approximately equal to the estimated variance of the difference of the two treatments when the units are independent and the model does not contain network or block effects (i.e. $0.0123\sigma^2$). While the optimal design for estimating τ_1 has 161 and 163 units receiving treatments 1 and 2 respectively, the optimal design for estimating the difference in network effects has 117 and 207 units receiving treatments 1 and 2 respectively.

7 | DISCUSSION

Many experiments can be regarded as being on networks by appropriate specification of the adjacency matrix. The suggested methods are easily adaptable to a wide class of networked experiments and to different block definitions for enabling the researcher to make comparisons between different treatments both directly and indirectly. However, generalising these methods to large-scale networks with hundreds of thousands of nodes, for example, commercial databases, is an open research problem.

This paper offers a framework for finding optimal block designs with network effects. It provides evidence that conventional designs such as randomised designs while ignoring the network structure

are inefficient compared to block designs with network effects. We saw that a model which wrongly ignores block and/or network effects leads to inefficient designs.

We also suggest a way to define blocks, while taking into account the community structure. We justify our choice of using spectral clustering due to the simplicity of its implementation (standard linear algebra), without requiring any explicit distributional model and without making any kind of assumption on the formation of the clusters. We have three main components: (i) spectral clustering of the given social network to project the vertices of the network onto an eigenvector space; (ii) modularity to choose the ‘best’ partition of the network over all partitions; and (iii) optimal block designs, for the clusters of subjects produced. However, there will be situations where the network is big enough that this method is computationally inefficient and we would have to approximate by using a different graph partitioning algorithm, for example, the Louvain algorithm (Blondel et al., 2008).

Based on a wide variety of examples similar to the ones we presented, we can suggest some general guidelines. Heuristically the experimental design for estimating the treatment effects tends to have equal replication within each block on the units having a similar number of connections, while for estimating the network effects the design is greatly influenced by the network’s first and second order connections.

ACKNOWLEDGEMENTS

We are grateful to the anonymous referees for some insightful comments on an earlier version of the paper. The authors gratefully acknowledge funding for this research. Most of it was supported by a studentship from the ESRC South Coast Doctoral Training Partnership and the research was completed under EPSRC grant EP/T021624/1 Multi-Objective Optimal Design of Experiments.

REFERENCES

- Atkinson, A.C., Donev, A.N. & Tobias, R.D. (2007) *Optimum Experimental Designs, with SAS*. Oxford: Oxford University Press.
- Bailey, R.A. (2007) Designs for two-colour microarray experiments. *Journal of the Royal Statistical Society: Series C*, 56, 365–394.
- Bailey, R.A. (2008) *Design of Comparative Experiments*. Cambridge: Cambridge University Press.
- Bailey, R.A. (2013) Using graphs to find the best block designs. In: Beineke, L.W. & Wilson, R.J. (Eds.), *Topics in Structural Graph Theory*. Cambridge: Cambridge University Press, pp. 282–317.
- Bailey, R.A. & Cameron, P.J. (2009) Combinatorics of optimal designs. *Surveys in Combinatorics*, 365, 19–73.
- Besag, J. & Kempton, R.A. (1986) Statistical analysis of field experiments using neighbouring plots. *Biometrics*, 42, 231–251.
- Blondel, V.D., Guillaume, J.-L., Lambiotte, R. & Lefebvre, E. (2008) Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008, P10008.
- Chung, F.R.K. (1997) *Spectral Graph Theory*. Providence: American Mathematical Society.
- Cook, R.D. & Nachtsheim, C.J. (1980) A comparison of algorithms for constructing exact D-optimal designs. *Technometrics*, 22, 315–324.
- Druilhet, P. (1999) Optimality of neighbor-balanced designs. *Journal of Statistical Planning and Inference*, 81, 141–152.
- Fedorov, V.V. (1972) *Theory of Optimal Experiments*. New York: Academic Press.
- Harville, D.A. (1997) *Matrix Algebra from a Statistician's Perspective*. New York: Springer.
- Kunert, J. & Martin, R.J. (2000) On the determination of optimal designs for an interference model. *The Annals of Statistics*, 28, 1728–1742.
- Kunert, J. & Mersmann, S. (2011) Optimal designs for an interference model. *Journal of Statistical Planning and Inference*, 141, 1623–1632.
- Newman, M.E.J. (2006) Modularity and community structure in networks. *Proceedings of the National Academy of Sciences of the United States of America*, 103, 8577–8582.
- Newman, M.E.J. & Girvan, M. (2004) Finding and evaluating community structure in networks. *Physical Review E*, 69, 026113.

- Parker, B., Gilmour, S.G. & Schormans, J. (2016) Optimal design of experiments on connected units with application to social networks. *Journal of the Royal Statistical Society: Series C*, 66, 455–480.
- Pearce, S.C. (1957) Experimenting with organisms as blocks. *Biometrika*, 44, 141–149.
- Shi, J. & Malik, J. (2000) Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22, 888–905.
- Von Luxburg, U. (2007) A tutorial on spectral clustering. *Statistics and Computing*, 17, 395–416.
- Wit, E., Nobile, A. & Khanin, R. (2005) Near-optimal designs for dual channel microarray studies. *Journal of the Royal Statistical Society: Series C*, 54, 817–830.

How to cite this article: Koutra V, Gilmour SG, Parker BM. Optimal block designs for experiments on networks. *J R Stat Soc Series C*. 2021;70:596–618. <https://doi.org/10.1111/rssc.12473>

APPENDIX

Bias due to model misspecification

We provide the derivation of the bias in the estimators of the model parameters due to model misspecification.

Consider performing an experiment on a network, where the adjacency matrix, A , is given. We assume that the true model is the BNM. However, the postulated model for the experiment is the RBM, which ignores network effects (i.e. $\gamma = \mathbf{0}$). We want to obtain the bias in the parameter estimates due to the model misspecification as a function of the unknown model parameters. This will be accomplished by using generalised inverse matrices (Harville, 1997, Ch.9). Let $\hat{\beta}_C$ and $\hat{\beta}_R$ be estimators of β with $X_C = (\mathbf{1} X_\tau^* X_b^* A X_\tau)^T$ and $X_R = (\mathbf{1} X_\tau^* X_b^* \mathbf{0}_{n \times m})^T$ being the extended design matrices under the BNM and RBM respectively. Note that $\hat{\beta}_C$ is the best linear unbiased estimator for β so that $\mathbb{E}[\hat{\beta}_C] = \beta$ and $\hat{\beta}_R$ is the OLS estimator of β for the RBM (without network effects). With the necessary algebraic calculations on generalised inverse matrices (where M^- is the generalised inverse of a square matrix M), it follows that the bias of the design, \mathcal{W} , under the assumption that there are network effects but we do not take them into account, is

$$\begin{aligned}
 \mathcal{W} &= \mathbb{E}[\hat{\beta}_R - \hat{\beta}_C] = \mathbb{E}[(X_R^T X_R)^- X_R^T \mathbf{y} - (X_C^T X_C)^- X_C^T \mathbf{y}] \\
 &= ((X_R^T X_R)^- X_R^T - (X_C^T X_C)^- X_C^T) \mathbb{E}[\mathbf{y}] \\
 &= ((X_R^T X_R)^- X_R^T - (X_C^T X_C)^- X_C^T) X_C \beta \\
 &= ((X_R^T X_R)^- X_R^T X_C - I) \beta \\
 &= \left(\left(\begin{array}{c} \mathbf{1} \\ X_\tau^* \\ X_b^* \\ \mathbf{0}_{n \times m} \end{array} \right)^T \left(\begin{array}{c} \mathbf{1} \\ X_\tau^* \\ X_b^* \\ \mathbf{0}_{n \times m} \end{array} \right) \right)^- \left(\begin{array}{c} \mathbf{1} \\ X_\tau^* \\ X_b^* \\ A X_\tau \end{array} \right)^T - I \beta \\
 &= \left(\left(\begin{array}{cccc} \mathbf{1}^T \mathbf{1} & \mathbf{1}^T X_\tau^* & \mathbf{1}^T X_b^* & \mathbf{0} \\ X_\tau^{*T} \mathbf{1} & X_\tau^{*T} X_\tau^* & X_\tau^{*T} X_b^* & \mathbf{0} \\ X_b^{*T} \mathbf{1} & X_b^{*T} X_\tau^* & X_b^{*T} X_b^* & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{array} \right) \left(\begin{array}{cccc} \mathbf{1}^T \mathbf{1} & \mathbf{1}^T X_\tau^* & \mathbf{1}^T X_b^* & \mathbf{1}^T A X_\tau \\ X_\tau^{*T} \mathbf{1} & X_\tau^{*T} X_\tau^* & X_\tau^{*T} X_b^* & X_\tau^{*T} A X_\tau \\ X_b^{*T} \mathbf{1} & X_b^{*T} X_\tau^* & X_b^{*T} X_b^* & X_b^{*T} A X_\tau \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{array} \right) \right)^- - I \beta,
 \end{aligned}$$

where I is the $(2m + \kappa - 1) \times (2m + \kappa - 1)$ identity matrix. Let B and Γ represent the $(m + \kappa - 1) \times (m + \kappa - 1)$ and $(m + \kappa - 1) \times (m - 1)$ matrices respectively

$$B = \begin{pmatrix} \mathbf{1}^T \mathbf{1} & \mathbf{1}^T X_\tau^* & \mathbf{1}^T X_b^* \\ X_\tau^{*T} \mathbf{1} & X_\tau^{*T} X_\tau^* & X_\tau^{*T} X_b^* \\ X_b^{*T} \mathbf{1} & X_b^{*T} X_\tau^* & X_b^{*T} X_b^* \end{pmatrix} \quad \text{and} \quad \Gamma = \begin{pmatrix} \mathbf{1}^T A X_\tau \\ X_\tau^{*T} A X_\tau \\ X_b^{*T} A X_\tau \end{pmatrix}.$$

Then for the $(2m + \kappa - 1) \times (2m + \kappa - 1)$ block-diagonal matrix

$$\Delta = \begin{pmatrix} B & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix},$$

we have that B is a non-singular matrix (it has full rank) and, defining a $(2m + \kappa - 1) \times (2m + \kappa - 1)$ matrix

$$G = \begin{pmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{pmatrix}$$

where G_{11} is of dimension $(m + \kappa - 1) \times (m + \kappa - 1)$, we obtain

$$\Delta G \Delta = \begin{pmatrix} B G_{11} B & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix},$$

implying that G is a generalised inverse of Δ if and only if $B G_{11} B = B$, or if and only if $G_{11} = B^{-1}$. Hence, we have

$$\mathcal{W} = \left(\begin{pmatrix} B^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \begin{pmatrix} B \Gamma \\ \mathbf{0} \mathbf{0} \end{pmatrix} - I \right) \beta = \left(\begin{pmatrix} I B^{-1} \Gamma \\ \mathbf{0} \mathbf{0} \end{pmatrix} - I \right) \beta = \begin{pmatrix} \mathbf{0} & B^{-1} \Gamma \\ \mathbf{0} & -I_m \end{pmatrix} \beta.$$

Thus the bias introduced in the estimates of the parameters, $\hat{\beta}_R$, under the false assumption that there are no network effects is given by the quantity $B^{-1} \Gamma$. This quantity is the result of ignoring the network effects, which is represented by an adjustment of the intercept and the treatment effect estimates. Observe that $\mathbb{E}[\hat{\beta}_R] \neq \mathbb{E}[\hat{\beta}_C]$, unless $B^{-1} \Gamma = \mathbf{0}$ which results from $\Gamma = \mathbf{0}$ or $\gamma = \mathbf{0}$.

Following the same rationale, we can end up in a similar expression of the design bias introduced in the treatment effects, when we perform an experiment where we wrongly assume there are no blocking effects. This means that if we consider the reduced model LNM, while the true model is of the form BNM, the design matrix will be of the form $X_R = (\mathbf{1} X_\tau^* \mathbf{0}_{n \times (\kappa-1)} A X_\tau)^T$.

The Shi and Malik algorithm

The SM algorithm involves three steps: I. *Compute the normalised graph Laplacian* L_{rw} and its spectrum (as based on the known adjacency matrix of the network); II. *Dimensionality reduction*: using the κ first eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_\kappa$ of the graph Laplacian that correspond to the first κ eigenvalues sorted in

ascending order, let $U \in R_{n \times \kappa}$ be the matrix containing $\mathbf{v}_1, \dots, \mathbf{v}_\kappa$ as columns. (Note that the eigenvector space varies based on the chosen dimensionality κ); III. *Clustering step*: treating each row of U as a data point, $(y_i)_{i=1, \dots, n} \in R^\kappa$, group them via the (standard) k-means algorithm into κ (dimensionality of the eigenvector space) clusters, C_1, \dots, C_κ . Therefore, the vertices of the network are projected into a κ -dimensional space, where κ is the number of the first nontrivial eigenvectors of L_{rw} . As a result, each unit is allocated to one cluster.

We perform the clustering step for various numbers of clusters κ , $2 \leq \kappa \leq n/2$, the upper bound chosen so that there should be at least two units within a cluster for achieving treatment comparisons. Note that the partition method can rely on a different clustering step rather than k-means if required. However, we used this standard one, which is the most used partitioning method.

This clustering algorithm gives no hint about the choice of the number of clusters or ‘goodness’ of each partition. To quantify the quality of partitions and to choose the number of communities based on the ‘best’ partition, we use modularity,

$$Q = \frac{1}{2l} \sum_{j,h=1}^n \sum_{i=1}^{\kappa} \left(A_{jh} - \frac{d_j d_h}{2l} \right) s_{ji} s_{hi},$$

where s_{ji} and s_{hi} are binary indicators of whether vertices j and h belong to group i or not (membership vectors) and $2l = \sum_{j=1}^n d_j$ is the total degree of all the vertices. In other words, modularity is defined as the difference between the fraction of the edges that fall within clusters and the fraction of the edges that would be expected to fall within the clusters if the edges were assigned randomly but keeping the degrees of the vertices unchanged. The factor of $1/2$ accounts for the fact that every vertex pair j,h is counted twice. When the communities are not better than the random partition or when the network does not exhibit any community structure, Q is zero or negative.

As such κ is regarded as the appropriate candidate for indicating the number of intrinsic communities which will define the blocks to be used in the design process, if the κ th is the highest modularity score over all produced partitions, that is, $\kappa = \arg \max Q$.