

Design of Networked Experiments

Ben M Parker

Queen Mary, University of London

September 16th 2011



Background

Often we assume when designing experiments that subjects are independent.

Often we assume when designing experiments that subjects are independent.

Particularly, if we apply a treatment to one subject, we generally assume that the treatment does not affect other subjects.

Often we assume when designing experiments that subjects are independent.

Particularly, if we apply a treatment to one subject, we generally assume that the treatment does not affect other subjects.

In this work, we investigate how the structure of relationships between subjects affects the design of experiments on these subjects.

Introduction

- We consider a graph, as a collection of nodes N and edges E .

Introduction

- We consider a graph, as a collection of nodes N and edges E .
- The nodes represent subjects on which we apply some treatment. We have $|N| = n$ subjects.

Introduction

- We consider a graph, as a collection of nodes N and edges E .
- The nodes represent subjects on which we apply some treatment. We have $|N| = n$ subjects.
- The edges represent some relationship between the subjects.

Introduction

- We consider a graph, as a collection of nodes N and edges E .
- The nodes represent subjects on which we apply some treatment. We have $|N| = n$ subjects.
- The edges represent some relationship between the subjects.
- We allow the relationship between our subjects to be specified by the adjacency matrix A where $A_{ij} = 1$ if i and j are related and $A_{ij} = 0$ otherwise. By convention, $A_{ii} = 0$.

Introduction


- We consider a graph, as a collection of nodes N and edges E .
- The nodes represent subjects on which we apply some treatment. We have $|N| = n$ subjects.
- The edges represent some relationship between the subjects.
- We allow the relationship between our subjects to be specified by the adjacency matrix A where $A_{ij} = 1$ if i and j are related and $A_{ij} = 0$ otherwise. By convention, $A_{ii} = 0$.
- We assume here that links are non-directional, such that $A_{ij} = A_{ji}$, i.e. A is symmetrical.

Introduction

- We consider a graph, as a collection of nodes N and edges E .
- The nodes represent subjects on which we apply some treatment. We have $|N| = n$ subjects.
- The edges represent some relationship between the subjects.
- We allow the relationship between our subjects to be specified by the adjacency matrix A where $A_{ij} = 1$ if i and j are related and $A_{ij} = 0$ otherwise. By convention, $A_{ii} = 0$.
- We assume here that links are non-directional, such that $A_{ij} = A_{ji}$, i.e. A is symmetrical.


Introduction

- We consider a graph, as a collection of nodes N and edges E .
- The nodes represent subjects on which we apply some treatment. We have $|N| = n$ subjects.
- The edges represent some relationship between the subjects.
- We allow the relationship between our subjects to be specified by the adjacency matrix A where $A_{ij} = 1$ if i and j are related and $A_{ij} = 0$ otherwise. By convention, $A_{ii} = 0$.
- We assume here that links are non-directional, such that $A_{ij} = A_{ji}$, i.e. A is symmetrical.



Introduction

- We consider a graph, as a collection of nodes N and edges E .
- The nodes represent subjects on which we apply some treatment. We have $|N| = n$ subjects.
- The edges represent some relationship between the subjects.
- We allow the relationship between our subjects to be specified by the adjacency matrix A where $A_{ij} = 1$ if i and j are related and $A_{ij} = 0$ otherwise. By convention, $A_{ii} = 0$.
- We assume here that links are non-directional, such that $A_{ij} = A_{ji}$, i.e. A is symmetrical.



$$A = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}$$

Examples of problems we might be interested in

- The nodes may be subjects in a marketing experiment.

Examples of problems we might be interested in

- The nodes may be subjects in a marketing experiment.
- The links may be a social relationship, such as being “friends” on

Facebook.



Examples of problems we might be interested in

- The nodes may be subjects in a marketing experiment.
- The links may be a social relationship, such as being “friends” on



Facebook.

- The treatments may be different advertising campaigns, such as “Buy

Coke”



or “Buy Pepsi”



Examples of problems we might be interested in

- The nodes may be subjects in a marketing experiment.
- The links may be a social relationship, such as being “friends” on



Facebook.

- The treatments may be different advertising campaigns, such as “Buy



Coke”

or “Buy Pepsi”



- The responses may be the amount of a products bought by a subject, or how much the subject likes the product.

Examples of problems we might be interested in

- The nodes may be subjects in a marketing experiment.
- The links may be a social relationship, such as being “friends” on



Facebook.

- The treatments may be different advertising campaigns, such as “Buy



Coke”

or “Buy Pepsi”



- The responses may be the amount of a products bought by a subject, or how much the subject likes the product.

Examples of problems we might be interested in

- The nodes may be subjects in a marketing experiment.
- The links may be a social relationship, such as being “friends” on



Facebook.

- The treatments may be different advertising campaigns, such as “Buy

Coke”



or “Buy Pepsi”



- The responses may be the amount of a products bought by a subject, or how much the subject likes the product.

If a subject is exposed to marketing, their friend may see this marketing and their response may be altered (positively or negatively). We assume that if a relationship exists between two individuals, the response of the first subject is dependent on the treatment applied to the second.

A less trivial example

The nodes may be subjects in a public health experiment. The links may be some friendship relationship, some geographical relationship, or some familial relationship. The treatments may be different interventions in public health (e.g. “Eat five fruit or vegetables a day”) and the response the change in the number of fruits or vegetables eaten after the campaign.

A less trivial example

The nodes may be subjects in a public health experiment. The links may be some friendship relationship, some geographical relationship, or some familial relationship. The treatments may be different interventions in public health (e.g. “Eat five fruit or vegetables a day”) and the response the change in the number of fruits or vegetables eaten after the campaign. For example, a parent may receive information on public health, and whether or not they alter their own behaviour, may bring about a change in their child’s response.

Optimal design without any network effect

Model

$$Y_i = \mu + \tau_j + \epsilon_i$$

where ϵ_i are i.i.d normally with mean 0, and variance σ^2 , and τ_j are the treatment effects for applying treatment j to subject i , $j = 1, \dots, m$. Let $\tau_m = 0$ for uniqueness.

Optimal design without any network effect

Model

$$Y_i = \mu + \tau_j + \epsilon_i$$

where ϵ_i are i.i.d normally with mean 0, and variance σ^2 , and τ_j are the treatment effects for applying treatment j to subject i , $j = 1, \dots, m$. Let $\tau_m = 0$ for uniqueness.

Optimality Criteria: We assume we will generally wish to measure the difference between some treatment effects, for example $\tau_1 - \tau_2$.

Optimal design without any network effect

Model

$$Y_i = \mu + \tau_j + \epsilon_i$$

where ϵ_i are i.i.d normally with mean 0, and variance σ^2 , and τ_j are the treatment effects for applying treatment j to subject i , $j = 1, \dots, m$. Let $\tau_m = 0$ for uniqueness.

Optimality Criteria: We assume we will generally wish to measure the difference between some treatment effects, for example $\tau_1 - \tau_2$. For example, we minimise the average variance of all estimates of pairwise differences of treatment effects,

$$\frac{2}{m(m-1)} \sum_{j=1}^m \sum_{l>j}^m \text{Var}(\widehat{\tau_j - \tau_l}).$$

Optimal design without any network effect

Model

$$Y_i = \mu + \tau_j + \epsilon_i$$

where ϵ_i are i.i.d normally with mean 0, and variance σ^2 , and τ_j are the treatment effects for applying treatment j to subject i , $j = 1, \dots, m$. Let $\tau_m = 0$ for uniqueness.

Optimality Criteria: We assume we will generally wish to measure the difference between some treatment effects, for example $\tau_1 - \tau_2$. For example, we minimise the average variance of all estimates of pairwise differences of treatment effects,

$$\frac{2}{m(m-1)} \sum_{j=1}^m \sum_{l>j}^m \text{Var}(\widehat{\tau_j - \tau_l}).$$

This is A-optimality for estimating the difference in treatment effects.

Optimal design without any network effect

Model

$$Y_i = \mu + \tau_j + \epsilon_i$$

where ϵ_i are i.i.d normally with mean 0, and variance σ^2 , and τ_j are the treatment effects for applying treatment j to subject i , $j = 1, \dots, m$. Let $\tau_m = 0$ for uniqueness.

Optimality Criteria: We assume we will generally wish to measure the difference between some treatment effects, for example $\tau_1 - \tau_2$.

For example, we minimise the average variance of all estimates of pairwise differences of treatment effects,

$$\frac{2}{m(m-1)} \sum_{j=1}^m \sum_{l>j}^m \text{Var}(\widehat{\tau_j - \tau_l}).$$

This is A-optimality for estimating the difference in treatment effects.


For $m = 2$ treatments, it is clear that the balanced design with an equal number of subjects given each treatment is optimal.

The idea

We assume that if we apply a treatment I to a subject, there will be a **network effect** of γ_I to all neighbours of that subject.


The idea

We assume that if we apply a treatment I to a subject, there will be a **network effect** of γ_I to all neighbours of that subject.



The idea


We assume that if we apply a treatment I to a subject, there will be a **network effect** of γ_I to all neighbours of that subject.



For example, if we apply treatment 1 to subject 2, there will be a **network effect** of γ_1 on subject 1

The idea

We assume that if we apply a treatment I to a subject, there will be a **network effect** of γ_I to all neighbours of that subject.



For example, if we apply treatment 1 to subject 2, there will be a **network effect** of γ_1 on subject 1 as well as the standard **subject effect** on subject 1 from its own treatment.

Optimal design with network effect

We introduce:

We introduce:

Linear network effect model

$$Y_i = \mu + \tau_j + \sum_{k=1} A_{ik} \gamma_{I(k)} + \epsilon_i$$

where $I = I(k)$ is the treatment given to subject k , and γ_I is the corresponding **network effect**, which is the change in the behaviour on a subject due to giving a connected subject a particular treatment.

We introduce:


Linear network effect model

$$Y_i = \mu + \tau_j + \sum_{k=1} A_{ik} \gamma_{I(k)} + \epsilon_i$$


where $I = I(k)$ is the treatment given to subject k , and γ_I is the corresponding **network effect**, which is the change in the behaviour on a subject due to giving a connected subject a particular treatment.

It is no longer clear that balanced designs are A-optimal for estimating the difference in treatment effects.


Example 1



Example 1




Example 1




A-optimal criterion: 0.418

Example 1




Example 1




By performing exhaustive search over all possible designs, we find that the optimal designs for the $m = 2$ two treatment case is :

- $\{1, 1, 2, 2, 2, 1, 2, 1, 1, 1\}$ for optimality in estimating the difference in the subject effects (i.e we give treatment 1 to subjects 1,2,6,8,9, and 10 and treatment 2 to the other subjects).
- $\{2, 2, 1, 1, 1, 1, 2, 2, 2, 1\}$ for optimality in estimating the difference in the network effects γ_i .

Example 2



Example 2



A-optimal criterion: 0.406

Neglecting the network effect

Let us imagine our linear network effect model

$$Y_i = \mu + \tau_j + \sum_{k=1} A_{ik} \gamma_l(k) + \epsilon_i$$

is true, but we neglect it in the design.

Neglecting the network effect

Let us imagine our linear network effect model

$$Y_i = \mu + \tau_j + \sum_{k=1} A_{ik} \gamma_l(k) + \epsilon_i$$

is true, but we neglect it in the design.

- We would incorrectly choose a (randomised) balanced design.

Neglecting the network effect

Let us imagine our linear network effect model

$$Y_i = \mu + \tau_j + \sum_{k=1} A_{ik} \gamma_l(k) + \epsilon_i$$

is true, but we neglect it in the design.

- We would incorrectly choose a (randomised) balanced design.
- If there is no network effect, then all permutations of a design are equivalent.

Neglecting the network effect

Let us imagine our linear network effect model

$$Y_i = \mu + \tau_j + \sum_{k=1} A_{ik} \gamma_l(k) + \epsilon_i$$

is true, but we neglect it in the design.

- We would incorrectly choose a (randomised) balanced design.
- If there is no network effect, then all permutations of a design are equivalent.
- As all subjects are identical, it does not matter which subject we assign a treatment to)

Neglecting the network effect


Let us imagine our linear network effect model

$$Y_i = \mu + \tau_j + \sum_{k=1} A_{ik} \gamma_l(k) + \epsilon_i$$


is true, but we neglect it in the design.

- We would incorrectly choose a (randomised) balanced design.
- If there is no network effect, then all permutations of a design are equivalent.
- As all subjects are identical, it does not matter which subject we assign a treatment to)
- This is not the case for our model.


Example 2



Example 2




Example 2



A-optimal criterion: 0.406

A-criterion for design 1,1,1,1,1,2,2,2,2,2 : 0.430


Example 2



A-optimal criterion: 0.406

A-criterion for design 1,1,1,1,1,2,2,2,2,2 : 0.430 (Efficiency : 94.3%)

Boxplot of efficiencies of balanced designs for 20 random networks



As well as causing problems with a higher variance, a design which fails to take account of a network effect may produce **a biased estimate**.

As well as causing problems with a higher variance, a design which fails to take account of a network effect may produce **a biased estimate**.

Writing R for the (reduced) model without network effects, and C for that with network effects, the bias in the subject and network effects β caused by neglecting a network effect in a model can be shown to be

$$E(\hat{\beta}_R - \hat{\beta}_C) = \left[\begin{pmatrix} (F_R^T F_R)^{-1} F_R^T \\ 0_{m \times n} \end{pmatrix} - (F_C^T F_C)^{-1} F_C^T \right] F_C \beta_C,$$

where F is the extended design matrix for our experiment.

e.g. The bias in the two examples seen so far, for $\beta^T = (\mu \quad \tau_1 \quad \gamma_1 \quad \gamma_2)$ are

$$\begin{pmatrix} 0 & 0.8 & 1.2 & 0.2 \\ 0 & -1.4 & -1 & 0.6 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \beta$$

and

$$\begin{pmatrix} 0 & 0.6 & 2.4 & 2.2 \\ 0 & -1.2 & -0.4 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \beta$$

Properties of the network

- In spectral graph theory, the **spectrum of the graph** is the set of eigenvalues of the adjacency matrix A .

Properties of the network

- In spectral graph theory, the **spectrum of the graph** is the set of eigenvalues of the adjacency matrix A .
- The **degree matrix** D is defined such that $D_{ii} = \sum_j A_{ij}$, $D_{ij} = 0$ for $i \neq j$, which specifies the number of nodes connected to each node.

Properties of the network

- In spectral graph theory, the **spectrum of the graph** is the set of eigenvalues of the adjacency matrix A .
- The **degree matrix** D is defined such that $D_{ii} = \sum_j A_{ij}$, $D_{ij} = 0$ for $i \neq j$, which specifies the number of nodes connected to each node.
- The **Laplacian spectrum**, the set of eigenvalues of the Laplacian matrix, is also often studied, where the Laplacian L is defined as $L = D - A$.

Properties of the network


- In spectral graph theory, the **spectrum of the graph** is the set of eigenvalues of the adjacency matrix A .
- The **degree matrix** D is defined such that $D_{ii} = \sum_j A_{ij}$, $D_{ij} = 0$ for $i \neq j$, which specifies the number of nodes connected to each node.
- The **Laplacian spectrum**, the set of eigenvalues of the Laplacian matrix, is also often studied, where the Laplacian L is defined as $L = D - A$.

Properties of the network


- In spectral graph theory, the **spectrum of the graph** is the set of eigenvalues of the adjacency matrix A .
- The **degree matrix** D is defined such that $D_{ii} = \sum_j A_{ij}$, $D_{ij} = 0$ for $i \neq j$, which specifies the number of nodes connected to each node.
- **The Laplacian spectrum**, the set of eigenvalues of the Laplacian matrix, is also often studied, where the Laplacian L is defined as $L = D - A$.

These spectra are a way of summarising some important macroscopic properties of the graph; for example, **the second smallest eigenvalue of the Laplacian matrix is the algebraic connectivity of our network**. This gives a measure of how well connected the overall graph is.

Algebraic connectivity for optimal design





Algebraic connectivity for an arbitrary balanced design





We have found no clear link between the (macro) properties of the graph and those of the design.

We have found no clear link between the (macro) properties of the graph and those of the design.



We have found no clear link between the (macro) properties of the graph and those of the design.



These graphs are clearly very different, although have the same spectrum $(-2,0,0,0,2)$. However, the left hand graph has optimal design for estimating the treatment effects $\{1, 2, 2, 2, 1\}$ with value 2, whereas for the right-hand graph all factors are not estimable; there is no optimal design, or rather all designs are equally ineffective.


Strange shaped fields

Most field trial experiments focus on rectangular fields. Let us assume we have an irregularly shaped field divided into plots as follows:

A field layout

	1	2	3	
	4		5	
6	7	8	9	10
11	12	13	14	15

Example 3: Non-rectangular field trials



Example 3: Non-rectangular field trials

Optimal design

		1	1	2
		2		2
3	3	1	3	2
3	1	1	2	3

- The A-optimal design function value is 0.4055.

Example 3: Non-rectangular field trials

Optimal design

		1	1	2
		2		2
3	3	1	3	2
3	1	1	2	3

- The A-optimal design function value is 0.4055.
- Minimum average variance possible with no restriction on comparisons $0.4\sigma^2$.

Example 3: Non-rectangular field trials

Optimal design

		1	1	2
		2		2
3	3	1	3	2
3	1	1	2	3

- The A-optimal design function value is 0.4055.
- Minimum average variance possible with no restriction on comparisons $0.4\sigma^2$.
- It is interesting to note how close we can get to the variance for the unrestricted case with this unusual shape of field ($0.4/0.4055=98.6\%$ efficiency).

Example 4: Extension to a crossover trial with planned dropouts

Usually in a crossover trial, including a planned drop out in the design be tricky. Let us assume we are performing a crossover trial

- on four subjects (a,b,c, and d).

Example 4: Extension to a crossover trial with planned dropouts

Usually in a crossover trial, including a planned drop out in the design be tricky. Let us assume we are performing a crossover trial

- on four subjects (a,b,c, and d).
- over four periods.

Example 4: Extension to a crossover trial with planned dropouts

Usually in a crossover trial, including a planned drop out in the design be tricky. Let us assume we are performing a crossover trial

- on four subjects (a,b,c, and d).
- over four periods.
- three treatments labelled 1 to 3,

Example 4: Extension to a crossover trial with planned dropouts

Usually in a crossover trial, including a planned drop out in the design be tricky. Let us assume we are performing a crossover trial

- on four subjects (a,b,c, and d).
- over four periods.
- three treatments labelled 1 to 3,
- each treatment is believed to have wash-out time of one period.

Example 4: Extension to a crossover trial with planned dropouts

Usually in a crossover trial, including a planned drop out in the design be tricky. Let us assume we are performing a crossover trial

- on four subjects (a,b,c, and d).
- over four periods.
- three treatments labelled 1 to 3,
- each treatment is believed to have wash-out time of one period.
- and that it becomes clear that participant b will not able to take the treatment in the third period of the trial.

Example 4: Extension to a crossover trial with planned dropouts

Usually in a crossover trial, including a planned drop out in the design be tricky. Let us assume we are performing a crossover trial

- on four subjects (a,b,c, and d).
- over four periods.
- three treatments labelled 1 to 3,
- each treatment is believed to have wash-out time of one period.
- and that it becomes clear that participant b will not able to take the treatment in the third period of the trial.

Example 4: Extension to a crossover trial with planned dropouts

Usually in a crossover trial, including a planned drop out in the design be tricky. Let us assume we are performing a crossover trial

- on four subjects (a,b,c, and d).
- over four periods.
- three treatments labelled 1 to 3,
- each treatment is believed to have wash-out time of one period.
- and that it becomes clear that participant b will not able to take the treatment in the third period of the trial.

We can represent our network as below, with links now unidirectional

Crossover Trial

a: 1-> 2-> 3->4
b: 5-> 6 x 7
c: 8-> 9->10-> 11
d: 12->13->14->15

Our Linear Network Effects Model is still valid, and we can perform exhaustive search as before to minimise our A-optimality criterion for subject effects.

Our Linear Network Effects Model is still valid, and we can perform exhaustive search as before to minimise our A-optimality criterion for subject effects. A is no longer symmetrical.

Our Linear Network Effects Model is still valid, and we can perform exhaustive search as before to minimise our A-optimality criterion for subject effects. A is no longer symmetrical.

Optimal Design

a:	1	1	2	1
b:	1	3	x	2
c:	2	2	3	1
d:	3	3	3	2

Our Linear Network Effects Model is still valid, and we can perform exhaustive search as before to minimise our A-optimality criterion for subject effects. A is no longer symmetrical.

Optimal Design

a:	1	1	2	1
b:	1	3	x	2
c:	2	2	3	1
d:	3	3	3	2

This has optimality criterion 0.4128 (efficiency of 96.9% of the unrestricted design with 15 treatments and no drop out).

Conclusions

- Presented a new method to take account of simple relationships between subjects when constructing a design.

Conclusions

- Presented a new method to take account of simple relationships between subjects when constructing a design.
- Useful for “social networks”, but there are a very wide class of experiments where the methodology might work.

Conclusions

- Presented a new method to take account of simple relationships between subjects when constructing a design.
- Useful for “social networks”, but there are a very wide class of experiments where the methodology might work.
- Interesting challenge to our “standard” assumption of independence of our units.

Conclusions

- Presented a new method to take account of simple relationships between subjects when constructing a design.
- Useful for “social networks”, but there are a very wide class of experiments where the methodology might work.
- Interesting challenge to our “standard” assumption of independence of our units.
- Readily extendible; i.e. our matrix A can represent the strength of a relationship.

Conclusions

- Presented a new method to take account of simple relationships between subjects when constructing a design.
- Useful for “social networks”, but there are a very wide class of experiments where the methodology might work.
- Interesting challenge to our “standard” assumption of independence of our units.
- Readily extendible; i.e. our matrix A can represent the strength of a relationship.
- We can easily introduce further relationship matrices to specify further connections, e.g. blocking effect.

Conclusions

- Presented a new method to take account of simple relationships between subjects when constructing a design.
- Useful for “social networks”, but there are a very wide class of experiments where the methodology might work.
- Interesting challenge to our “standard” assumption of independence of our units.
- Readily extendible; i.e. our matrix A can represent the strength of a relationship.
- We can easily introduce further relationship matrices to specify further connections, e.g. blocking effect.
- We have here considered unstructured treatments, but method also works for, e.g. factorial treatments.

Conclusions

- Presented a new method to take account of simple relationships between subjects when constructing a design.
- Useful for “social networks”, but there are a very wide class of experiments where the methodology might work.
- Interesting challenge to our “standard” assumption of independence of our units.
- Readily extendible; i.e. our matrix A can represent the strength of a relationship.
- We can easily introduce further relationship matrices to specify further connections, e.g. blocking effect.
- We have here considered unstructured treatments, but method also works for, e.g. factorial treatments.
- For more complicated experiments, more work needed on design algorithms.

Conclusions

- Presented a new method to take account of simple relationships between subjects when constructing a design.
- Useful for “social networks”, but there are a very wide class of experiments where the methodology might work.
- Interesting challenge to our “standard” assumption of independence of our units.
- Readily extendible; i.e. our matrix A can represent the strength of a relationship.
- We can easily introduce further relationship matrices to specify further connections, e.g. blocking effect.
- We have here considered unstructured treatments, but method also works for, e.g. factorial treatments.
- For more complicated experiments, more work needed on design algorithms.
- We would very much like to hear of any experimental work in this area.