Optimal Design of Experiments on Social Networks

Steven Gilmour[1], Ben Parker[1], John Schormans[2]

1 University of Southampton; 2 Queen Mary, University of London

MoN Meeting 2013

Motivating Application

Healthy eating campaigns are to be delivered by text message and a study is being conducted to compare a number of different candidate messages.

Motivating Application

Healthy eating campaigns are to be delivered by text message and a study is being conducted to compare a number of different candidate messages.

The study is conducted at local authority workplaces.

Motivating Application

Healthy eating campaigns are to be delivered by text message and a study is being conducted to compare a number of different candidate messages.

The study is conducted at local authority workplaces.

The subjects in the study work together and might eat lunch together.

Motivating Application

Healthy eating campaigns are to be delivered by text message and a study is being conducted to compare a number of different candidate messages.

The study is conducted at local authority workplaces.

The subjects in the study work together and might eat lunch together.

The text messages could have viral effects, i.e. they could affect those subjects connected to those to whom it they are sent.

Motivating Application

Healthy eating campaigns are to be delivered by text message and a study is being conducted to compare a number of different candidate messages.

The study is conducted at local authority workplaces.

The subjects in the study work together and might eat lunch together.

The text messages could have viral effects, i.e. they could affect those subjects connected to those to whom it they are sent.

If we know the "friendship" networks, how should this influence our allocation of candidate messages to subjects?

Design analysis of experiments

Experiments: Planned and controlled interventions. Allow causal inferences to be drawn.

Design analysis of experiments

Experiments: Planned and controlled interventions. Allow causal inferences to be drawn.

Statistical design of experiments (DOE) is about how to ensure the experiment produces the best data possible, given the practical, financial and ethical constraints.

Design analysis of experiments

Experiments: Planned and controlled interventions. Allow causal inferences to be drawn.

Statistical design of experiments (DOE) is about how to ensure the experiment produces the best data possible, given the practical, financial and ethical constraints.

All experiments have statistical properties in their design (even if we choose not to think about them). There is no non-statistical approach to designing experiments.

Some terminology

The specific interventions (e.g. types of text message) to be made are treatments. Comparing the effects of the treatments should meet the objectives of the experiment.

Some terminology

The specific interventions (e.g. types of text message) to be made are treatments. Comparing the effects of the treatments should meet the objectives of the experiment.

The things (e.g. subjects) to which the treatments are applied are called (experimental) units. One unit can get only one treatment, but different units can get different treatments.

Some terminology

The specific interventions (e.g. types of text message) to be made are treatments. Comparing the effects of the treatments should meet the objectives of the experiment.

The things (e.g. subjects) to which the treatments are applied are called (experimental) units. One unit can get only one treatment, but different units can get different treatments.

One or more responses are measured from each unit. Responses from units with different treatments allow these treatments to be compared.

Classical design: Randomization theory

We have n experimental units.

Classical design: Randomization theory

We have n experimental units.

We will compare n_t treatments.

Classical design: Randomization theory

We have n experimental units.

We will compare n_t treatments.

Assumed model for a response variable y is (on unit i with treatment k applied)

$$y_{i(k)} = u_i + t_k, \tag{1}$$

where u_i would have been the response for a treatment with 0 effect, t_k is the effect of treatment k and the only assumption needed is additivity of unit and treatment effects.

Randomization theory

The form of randomization used determines the appropriate analysis, e.g. for completely randomized design with $\frac{n}{n_t}$ replicates of each treatment, considering all possible outcomes of the randomization, model (1) becomes

$$Y_{i(k)} = \sum_{j=1}^{n} \delta_{ij} u_j + t_k, \qquad (2)$$

where $\delta_{ij} = \begin{cases} 1 & \text{if unit label } i \text{ applies to unit } j; \\ 0 & \text{otherwise.} \end{cases}$

Randomization theory

This implies that the usual linear models analysis gives:

- minimum variance linear unbiased estimators of t_k (and linear functions of t_ks);

Randomization theory

This implies that the usual linear models analysis gives:

- minimum variance linear unbiased estimators of t_k (and linear functions of t_ks);
- minimum variance unbiased estimators of σ^2, the inter-unit variance, and any other variance components;

Randomization theory

This implies that the usual linear models analysis gives:

- minimum variance linear unbiased estimators of t_k (and linear functions of t_ks);
- minimum variance unbiased estimators of σ^2, the inter-unit variance, and any other variance components;
- unbiased estimators of $V(\hat{t}_k)$ (and linear functions);

Randomization theory

This implies that the usual linear models analysis gives:

- minimum variance linear unbiased estimators of t_k (and linear functions of t_ks);
- minimum variance unbiased estimators of σ^2, the inter-unit variance, and any other variance components;
- unbiased estimators of $V(\hat{t}_k)$ (and linear functions);
- $E(MS_{Treat}) = \sigma^2$ if $t_k = 0 \ \forall k$.

Randomization theory

Exact inference requires the use of permutation tests.

Randomization theory

Exact inference requires the use of permutation tests.

We usually use a Normal approximation, i.e. model (2) becomes

$$Y_{i(k)} \quad = \quad \mu_i + t_k; \ \mu_i \overset{\text{iid}}{\sim} N(\mu, \sigma^2)$$

Randomization theory

Exact inference requires the use of permutation tests.

We usually use a Normal approximation, i.e. model (2) becomes

$$\begin{aligned} Y_{i(k)} &= \mu_i + t_k;\ \mu_i \overset{\text{iid}}{\sim} N(\mu, \sigma^2) \\ \text{i.e. } Y_{i(k)} &= \mu + t_k + \epsilon_i;\ \epsilon_i \overset{\text{iid}}{\sim} N(0, \sigma^2). \end{aligned}$$

Randomization theory

Exact inference requires the use of permutation tests.

We usually use a Normal approximation, i.e. model (2) becomes

$$
\begin{aligned}
Y_{i(k)} &= \mu_i + t_k;\ \mu_i \overset{\text{iid}}{\sim} N(\mu, \sigma^2) \\
\text{i.e. } Y_{i(k)} &= \mu + t_k + \epsilon_i;\ \epsilon_i \overset{\text{iid}}{\sim} N(0, \sigma^2).
\end{aligned}
$$

Extensions of this normal approximation often used almost automatically.

Randomization theory

Exact inference requires the use of permutation tests.

We usually use a Normal approximation, i.e. model (2) becomes

$$
\begin{aligned}
Y_{i(k)} &= \mu_i + t_k;\ \mu_i \stackrel{\text{iid}}{\sim} N(\mu, \sigma^2) \\
\text{i.e. } Y_{i(k)} &= \mu + t_k + \epsilon_i;\ \epsilon_i \stackrel{\text{iid}}{\sim} N(0, \sigma^2).
\end{aligned}
$$

Extensions of this normal approximation often used almost automatically.

They are the basis for mixed models analysis of nonorthogonal split-plot and multi-stratum designs, etc.

Nonadditivity of Treatment and Unit Effects

Randomization-based inference fails if model (1) is not true, but we might still be interested in treatment effects, especially if they affect neighbouring units, e.g.

Nonadditivity of Treatment and Unit Effects

Randomization-based inference fails if model (1) is not true, but we might still be interested in treatment effects, especially if they affect neighbouring units, e.g.

- carry-over effects in cross-over designs, e.g. in sensory evaluation;

Nonadditivity of Treatment and Unit Effects

Randomization-based inference fails if model (1) is not true, but we might still be interested in treatment effects, especially if they affect neighbouring units, e.g.

- carry-over effects in cross-over designs, e.g. in sensory evaluation;
- interference or competition effects in agricultural field experiments;

Nonadditivity of Treatment and Unit Effects

Randomization-based inference fails if model (1) is not true, but we might still be interested in treatment effects, especially if they affect neighbouring units, e.g.

- carry-over effects in cross-over designs, e.g. in sensory evaluation;
- interference or competition effects in agricultural field experiments;
- viral effects in experiments on social networks.

Nonadditivity of Treatment and Unit Effects

Randomization-based inference fails if model (1) is not true, but we might still be interested in treatment effects, especially if they affect neighbouring units, e.g.

- carry-over effects in cross-over designs, e.g. in sensory evaluation;
- interference or competition effects in agricultural field experiments;
- viral effects in experiments on social networks.

Nonadditivity of Treatment and Unit Effects

Randomization-based inference fails if model (1) is not true, but we might still be interested in treatment effects, especially if they affect neighbouring units, e.g.

- carry-over effects in cross-over designs, e.g. in sensory evaluation;
- interference or competition effects in agricultural field experiments;
- viral effects in experiments on social networks.

Two approaches:

- design in such a way that a standard analysis will be acceptable;

Nonadditivity of Treatment and Unit Effects

Randomization-based inference fails if model (1) is not true, but we might still be interested in treatment effects, especially if they affect neighbouring units, e.g.

- carry-over effects in cross-over designs, e.g. in sensory evaluation;
- interference or competition effects in agricultural field experiments;
- viral effects in experiments on social networks.

Two approaches:

- design in such a way that a standard analysis will be acceptable;
- model the interference and optimise the design for such a model.

Nonadditivity of Treatment and Unit Effects

Randomization-based inference fails if model (1) is not true, but we might still be interested in treatment effects, especially if they affect neighbouring units, e.g.

- carry-over effects in cross-over designs, e.g. in sensory evaluation;
- interference or competition effects in agricultural field experiments;
- viral effects in experiments on social networks.

Two approaches:

- design in such a way that a standard analysis will be acceptable;
- model the interference and optimise the design for such a model.

Nonadditivity of Treatment and Unit Effects

Randomization-based inference fails if model (1) is not true, but we might still be interested in treatment effects, especially if they affect neighbouring units, e.g.

- carry-over effects in cross-over designs, e.g. in sensory evaluation;
- interference or competition effects in agricultural field experiments;
- viral effects in experiments on social networks.

Two approaches:

- design in such a way that a standard analysis will be acceptable;
- model the interference and optimise the design for such a model.

In social networks, only the second approach is viable.

A Model

A network G is an undirected graph, a collection of nodes N and edges E where a node represents a unit on which we apply some treatment and an edge represents a relationship between units.

A Model

A network G is an undirected graph, a collection of nodes N and edges E where a node represents a unit on which we apply some treatment and an edge represents a relationship between units.

We have $|N| = n$ units, and m treatments.

A Model

A network G is an undirected graph, a collection of nodes N and edges E where a node represents a unit on which we apply some treatment and an edge represents a relationship between units.

We have $|N| = n$ units, and m treatments. Relationship between units is specified by the adjacency matrix A, where $A_{ik} = 1$ iff i and k are related and $A_{ik} = 0$ otherwise ($A_{ii} = 0$) and $A_{ik} = A_{ki}$.

A Model

A network G is an undirected graph, a collection of nodes N and edges E where a node represents a unit on which we apply some treatment and an edge represents a relationship between units.

We have $|N| = n$ units, and m treatments. Relationship between units is specified by the adjacency matrix A, where $A_{ik} = 1$ iff i and k are related and $A_{ik} = 0$ otherwise $(A_{ii} = 0)$ and $A_{ik} = A_{ki}$.

Assume a "subject effect" from unit i of β_j when treatment $j = j(i)$ is given to that unit, and a "network effect" of γ_l if treatment $l = l(k)$ is given to a connected unit k if $A_{ik} = 1$.

A Model

A network G is an undirected graph, a collection of nodes N and edges E where a node represents a unit on which we apply some treatment and an edge represents a relationship between units.

We have $|N| = n$ units, and m treatments. Relationship between units is specified by the adjacency matrix A, where $A_{ik} = 1$ iff i and k are related and $A_{ik} = 0$ otherwise $(A_{ii} = 0)$ and $A_{ik} = A_{ki}$.

Assume a "subject effect" from unit i of β_j when treatment $j = j(i)$ is given to that unit, and a "network effect" of γ_l if treatment $l = l(k)$ is given to a connected unit k if $A_{ik} = 1$.

We measure the response Y_i for unit i.

A model

This gives the linear network effects model,

$$Y_i = \mu + \tau_{j(i)} + \sum_{k=1}^{n} A_{ik}\gamma_{l(k)} + \epsilon_i. \qquad (3)$$

A model

This gives the linear network effects model,

$$Y_i = \mu + \tau_{j(i)} + \sum_{k=1}^{n} A_{ik}\gamma_{l(k)} + \epsilon_i. \quad (3)$$

Assume ϵ_i are independent and identically distributed with mean 0 and variance σ^2.

A model

This gives the linear network effects model,

$$Y_i = \mu + \tau_{j(i)} + \sum_{k=1}^{n} A_{ik}\gamma_{l(k)} + \epsilon_i. \tag{3}$$

Assume ϵ_i are independent and identically distributed with mean 0 and variance σ^2. Assume without loss of generality that $\tau_m = 0$.

Optimality Criteria

We might seek to minimise:

1. the average variance of all pairwise differences of treatment effects,

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

Optimality Criteria

We might seek to minimise:

1. the average variance of all pairwise differences of treatment effects,

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

2. the average variance of all pairwise differences of network effects,

$$\frac{2}{m(m-1)} \sum_{j=1}^{m-1} \sum_{l=j+1}^{m} Var(\widehat{\gamma_j - \gamma_l}).$$

Optimality Criteria

We might seek to minimise:

1. the average variance of all pairwise differences of treatment effects,
$$\frac{2}{m(m-1)} \sum_{j=1}^{m-1} \sum_{l=j+1}^{m} Var(\widehat{\tau_j - \tau_l}).$$
2. the average variance of all pairwise differences of network effects,
$$\frac{2}{m(m-1)} \sum_{j=1}^{m-1} \sum_{l=j+1}^{m} Var(\widehat{\gamma_j - \gamma_l}).$$
3. the average variance in estimating the total effects of a treatment. The total treatment effect for treatment i is
$$\theta_j = n_j \tau_j + \gamma_j \left(\sum_{l=1}^{m} n_{lj} \right),$$
where n_j is the number of units given treatment j and n_{lj} is the number of times units given treatment j and l are connected.

Example

Comparing two treatments on 20 connected experimental units.

Example

Comparing two treatments on 20 connected experimental units.

Interest is in the direct effects and the viral effects.

Example

Comparing two treatments on 20 connected experimental units.

Interest is in the direct effects and the viral effects.

We compare designs by exhaustive search.

Example

Comparing two treatments on 20 connected experimental units.

Interest is in the direct effects and the viral effects.

We compare designs by exhaustive search.

Different designs are optimal for different criteria.

Optimal designs for direct and network effects

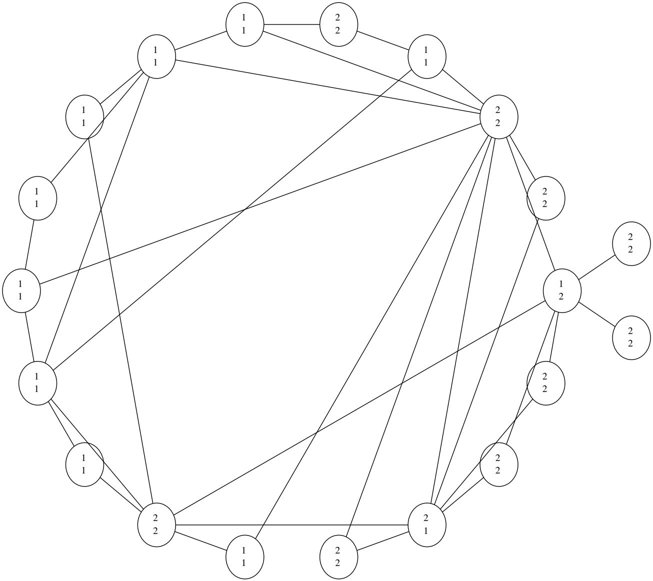

Conclusions from several examples

So far, we have studied the $MA(1)$ model in a known network graph:

- ignoring network effects in analysis can lead to severe bias;

Conclusions from several examples

So far, we have studied the $MA(1)$ model in a known network graph:

- ignoring network effects in analysis can lead to severe bias;
- ignoring network effects in the design can lead to very low efficiencies, even singularity;

Conclusions from several examples

So far, we have studied the $MA(1)$ model in a known network graph:

- ignoring network effects in analysis can lead to severe bias;
- ignoring network effects in the design can lead to very low efficiencies, even singularity;
- optimal design depends on the local properties of the network, so requires heavy computing and/or smart algorithms.

Conclusions from several examples

So far, we have studied the $MA(1)$ model in a known network graph:

- ignoring network effects in analysis can lead to severe bias;
- ignoring network effects in the design can lead to very low efficiencies, even singularity;
- optimal design depends on the local properties of the network, so requires heavy computing and/or smart algorithms.

Conclusions from several examples

So far, we have studied the $MA(1)$ model in a known network graph:

- ignoring network effects in analysis can lead to severe bias;
- ignoring network effects in the design can lead to very low efficiencies, even singularity;
- optimal design depends on the local properties of the network, so requires heavy computing and/or smart algorithms.

Extensions:

- larger networks;

Conclusions from several examples

So far, we have studied the $MA(1)$ model in a known network graph:

- ignoring network effects in analysis can lead to severe bias;
- ignoring network effects in the design can lead to very low efficiencies, even singularity;
- optimal design depends on the local properties of the network, so requires heavy computing and/or smart algorithms.

Extensions:

- larger networks;
- factorial treatment structures;

Conclusions from several examples

So far, we have studied the $MA(1)$ model in a known network graph:

- ignoring network effects in analysis can lead to severe bias;
- ignoring network effects in the design can lead to very low efficiencies, even singularity;
- optimal design depends on the local properties of the network, so requires heavy computing and/or smart algorithms.

Extensions:

- larger networks;
- factorial treatment structures;
- discrete responses;

Conclusions from several examples

So far, we have studied the $MA(1)$ model in a known network graph:

- ignoring network effects in analysis can lead to severe bias;
- ignoring network effects in the design can lead to very low efficiencies, even singularity;
- optimal design depends on the local properties of the network, so requires heavy computing and/or smart algorithms.

Extensions:

- larger networks;
- factorial treatment structures;
- discrete responses;
- $AR(1)$ models;

Conclusions from several examples

So far, we have studied the $MA(1)$ model in a known network graph:

- ignoring network effects in analysis can lead to severe bias;
- ignoring network effects in the design can lead to very low efficiencies, even singularity;
- optimal design depends on the local properties of the network, so requires heavy computing and/or smart algorithms.

Extensions:

- larger networks;
- factorial treatment structures;
- discrete responses;
- $AR(1)$ models;
- partially known networks.

Questions / Comments?

Thank you for your attention.

Contact details: s.gilmour@southampton.ac.uk