Designing Experiments on Networks

Vasiliki Koutra

18th June 2026 – mODa14

Overview

- Motivating experiments
- Background and treatment interference
- Optimal designs for direct and indirect effects
- Sub-network design via subsampling

- Koutra, V., Gilmour, S.G., and Parker, B.M. (2021). *JRSSC*, 70, 596–618.
- Koutra, V., Gilmour, S.G., Parker, B.M., and Mead, A. (2023). *JABES*, 28, 526–548.
- Ayres, L. and Koutra, V. (2026). Working paper.

Motivation

- A hospital campaign wishes to compare different patient information leaflets to choose which one will minimise infection rates.
- A study is conducted on a contact network of patients, with a survey used to identify interpersonal contacts and in-hospital movement.

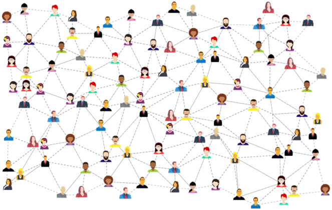

Credit: *pixabay*

- Patients might share the leaflets and therefore messages could have viral effects (indirect treatment effects).
- Patients in the same wards or other locations in hospital settings are likely to exhibit similar responses (block effects).

Other examples

Field of study	Intervention	Connections	Response
Health:	Infection control information	Patient contacts in hospital	Disease incidence
Agriculture:	Pesticide	Geographic proximity	Crop yield
Marketing:	Advertisement	Virtual friendships	Product awareness
Politics:	Direct mailing	Voter interactions	Voting behaviour
Education:	Incentivised food choices	Social links	Snack choice
Ecology:	Reward-based intervention	Animal interactions	Reaction speed
Law enforcement:	Surveillance	Geographical proximity	Crime rate

Treatment interference

What do these experiments have in common? Possible treatment interference (or treatment carryover or spillover).

1. Ignoring substantial treatment interference can lead to biased estimates of differences between the direct treatment effects.
2. Mitigation may be possible (e.g. "wash-out" periods, "guard plots"), but often not feasible or ethical.
3. In many cases the indirect effect of each treatment is itself of interest (e.g. viral effects).

Hence, it is of interest to study designs and models which account for treatment interference.

Objective

Obtain good designs in the presence of interference when interest is in estimating direct and/or indirect treatment effects.

Three core challenges for modern networked experiments:

1. **Scalability**: Networks too large for design search to be computationally tractable.
2. **Practical constraints**: Cost, access, or ethical constraints prevent treatment assignment to the whole network.
3. **Uncertainty**: Networks which are incompletely or inaccurately known.

SUTVA

Standard model for analysing a designed experiment:

$$Y_i = \mu + \beta_{b(i)} + \tau_{z(i)} + \epsilon_i \, ,$$

where $\beta_{b(i)}$ is the effect of the block containing unit i, $\tau_{z(i)}$ is the effect of treatment $z(i)$ assigned to unit i, and ϵ_i is a an error term with constant variance.

This model makes the stable unit treatment value assumption, which states that the response from any particular unit is unaffected by the assignment of treatments to other units (Cox, 1958).

Modelling connections between units as a graph provides a mechanism for incorporating interference into the response model.

Unit connections as a graph

Vertices and Edges Graph Adjacency Matrix Blocks

Suppose that n units are formed into a network where connections represent potential treatment interference.

- Vertex set $\mathcal{V}$ represents the units.
- Edge set $\mathcal{E}$ (of size l) represents the connections.

Unit connections as a graph

Vertices and Edges **Graph** Adjacency Matrix Blocks

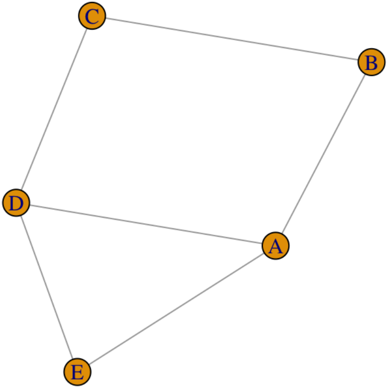

Figure: Example graph $\mathcal{G}$ with $|\mathcal{V}| = 5$ vertices and $|\mathcal{E}| = 6$ undirected edges.

- The adjacency matrix A encodes the connections.
- $A_{jh} = A_{hj}$ for undirected graphs. If no edge between j and h, $A_{jh} = A_{hj} = 0$.

	A	B	C	D	E
A	0	1	0	1	1
B	1	0	1	0	0
C	0	1	0	1	0
D	1	0	1	0	1
E	1	0	0	1	0

Unit connections as a graph

Vertices and Edges Graph Adjacency Matrix **Blocks**

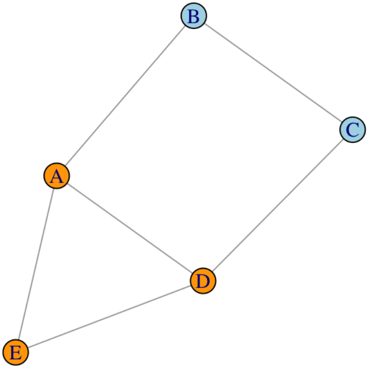

Figure: Example graph with $|\mathcal{V}| = 5$ vertices, $|\mathcal{E}| = 6$ undirected edges. The colours represent blocks identified through spectral clustering.

- Blocking can remove sources of variability that would otherwise contribute to noise.
- In networks, blocks may be based on connectivity patterns, e.g. using **spectral clustering** with optimal modularity (Koutra et al., 2021).

Network Model

The modelling framework underpinning this work extends the Generalised Additive Network Effects (GANE) model (Bui et al., 2023) to incorporate multiple treatments and blocking structure.

Network Model

$$Y_i = \mu + \tau_{z(i)} + \sum_{k=1}^{t} f_{ik}(G, \boldsymbol{\gamma}) + \sum_{q=1}^{Q} \zeta^{(q)}_{s_q(i)} + \epsilon_i, \qquad i = 1, \ldots, n,$$

- $z(i) \in \{1, \ldots, t\}$: treatment assigned to node i; $\tau_{z(i)}$ is the direct treatment effect.
- $f_{ik}(\cdot, \cdot)$: spillover onto node i from neighbours receiving treatment k, depending on $G = (\mathcal{V}, \mathcal{E})$ and $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_t)^T$.
- $s_q(i)$: level of the qth (fixed) stratum containing node i, with effect vector $\boldsymbol{\zeta}^{(q)}$, $q = 1, \ldots, Q$.

The parameters $\boldsymbol{\tau}$ and $\boldsymbol{\gamma}$ are the primary inferential targets; the $\boldsymbol{\zeta}^{(q)}$ are nuisance stratum effects.

Special Cases of the Network Model

Linear Network Model (LNM) (Parker et al., 2017): no strata ($Q = 0$), with $f_{ik}(G, \boldsymbol{\gamma}) = \gamma_k \sum_{j=1}^{n} A_{ij}\, \mathbb{1}\{z(j) = k\}$, so $\sum_{k=1}^{t} f_{ik}(G, \boldsymbol{\gamma}) = \sum_{j=1}^{n} A_{ij}\, \gamma_{z(j)}$:

$$Y_i = \mu + \tau_{z(i)} + \sum_{j=1}^{n} A_{ij}\, \gamma_{z(j)} + \epsilon_i.$$

Block Network Model (BNM) (Koutra et al., 2021): one fixed stratum (blocks, $b(i)$):

$$Y_i = \mu + \beta_{b(i)} + \tau_{z(i)} + \sum_{j=1}^{n} A_{ij}\, \gamma_{z(j)} + \epsilon_i.$$

Block Row–Column Network Model (BRCNM) (Koutra et al., 2023): three fixed strata (block $b(i)$, row $r(i)$, column $c(i)$):

$$Y_i = \mu + \beta_{b(i)} + \rho_{r(i)} + \kappa_{c(i)} + \tau_{z(i)} + \sum_{j=1}^{n} A_{ij}\, \gamma_{z(j)} + \epsilon_i.$$

Optimal Designs

The design is the treatment allocation vector $\mathbf{z} = (z(1), \ldots, z(n))^T$.

An optimal design minimises an objective function $\phi(\mathbf{z})$ that depends on the Fisher information matrix $M(\mathbf{z}, G)$.

For efficient estimation of **direct treatment differences**:

$$\phi_\tau = \frac{2}{t(t-1)} \sum_{s=1}^{t-1} \sum_{s'=s+1}^{t} \mathrm{var}(\widehat{\tau_s - \tau_{s'}}).$$

For efficient estimation of **indirect (network) effects**:

$$\phi_\gamma = \frac{2}{t(t-1)} \sum_{s=1}^{t-1} \sum_{s'=s+1}^{t} \mathrm{var}(\widehat{\gamma_s - \gamma_{s'}}).$$

Designs found using standard optimisation algorithms, such as exchange (-interchange) algorithms.

Example 1: Co-authorship Network

Links between academics within a university research group (Koutra et al., 2021).

- 22 nodes, split into three blocks, and 27 edges.
- Focus on estimating the effects of two treatments.

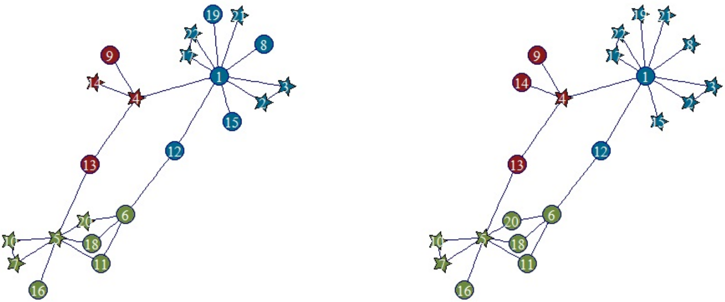

Figure: Block designs for a co-authorship network. Colours represent blocks (identified via spectral clustering), and plot symbols indicate treatment allocation (1 or 2). Left: Optimal design for estimating direct effects. Right: Optimal design for estimating indirect effects.

Example 1: Co-authorship Network

Efficiencies for Estimating Direct Effects. Efficiencies are calculated **within each row**.

Model	**Designs**			
	CRD	RBD	LND	BND
CRM	1.00	1.00	1.00	1.00
RBM	0.89	1.00	0.68	1.00
LNM	0.86	0.83	1.00	1.00
BNM	0.73	0.81	0.50	1.00

Table: Efficiencies for estimating direct treatment effects across different designs, considering blocking and network effects under various model assumptions.

- The BND is always 100% efficient across all models.
- Ignoring network effects (CRD, RBD) loses 15–27% efficiency.
- Ignoring blocks (LND) can lose 50% efficiency.

Example 1: Co-authorship Network

Efficiencies for Estimating Indirect Effects. Efficiencies are calculated **within each row**.

Model	**Designs**			
	CRD	RBD	LND	BND
LNM	0.16	0.12	1.00	0.64
BNM	0.16	0.16	0.39	1.00

Table: Efficiencies for estimating network effects under various design models.

Note: The loss in efficiency for designs that ignore network structure is now substantial ($> 80\%$).

Example 2: Subset of the Facebook Network

Mutual virtual friendships between users within the Facebook network.

- 324 nodes, divided into 24 blocks, with 2514 edges.
- Interest lies in estimating the effects of two treatments.

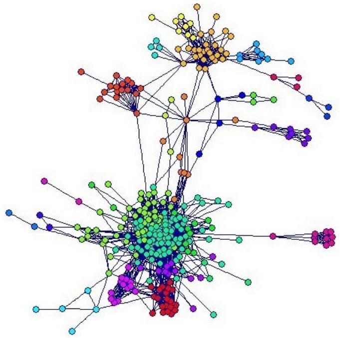

Figure: Subset of the Facebook network from the Stanford Data Collection. Colours indicate 24 distinct blocks, or communities of users.

Table: Treatment Effects

Designs	BNM
CRD	0.904
RBD	0.969
LND	0.948
BND	**1.000**

Table: Network Effects

Designs	BNM
CRD	0.173
RBD	0.168
LND	0.761
BND	**1.000**

Mutual virtual friendships between users within the Facebook network.

- 324 nodes, divided into 24 blocks, with 2514 edges.
- Interest lies in estimating the effects of two treatments

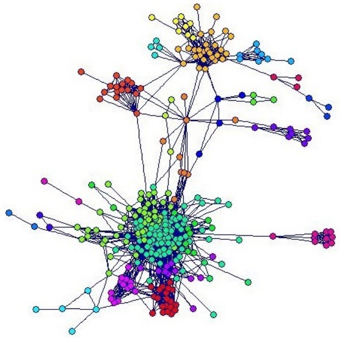

Figure: Subset of the Facebook network from the Stanford Data Collection. Colours indicate 24 distinct blocks, or communities of users.

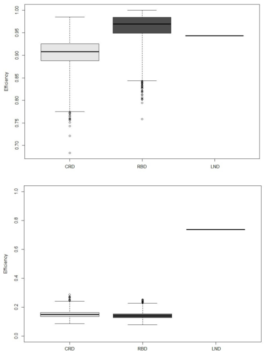

Example 3: Agricultural Experiment

The impact of neighbouring plots in field trials has been widely studied, particularly regarding indirect effects (Besag and Kempton, 1986).


Figure: Left: A typical field trial layout, clearly illustrating the proximity of neighbouring plots. Right: Cereal aphid

This experiment at Rothamsted focused on natural cereal aphid colonisation in 21 varieties of wheat (Koutra et al., 2023).

Example 3: Agricultural Experiment

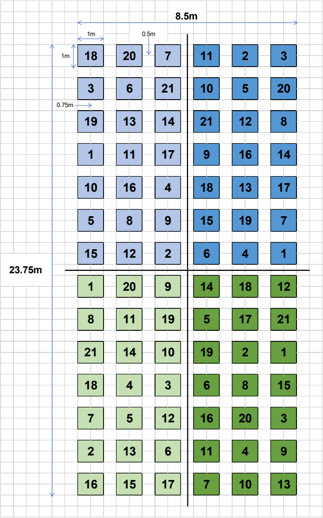

Figure: Plot layout for the agricultural experiment, with original design obtained from CycDesigN.

- 14×6 grid of $1m \times 1m$ plots.
- Each treatment is replicated four times.
- Response is square root of total number of aphids in a plot.
- Treatment interference due to the varying susceptibility of different wheat varieties and the likelihood of aphids moving between plots.
- Network describes the interference pattern.

Designs	Efficiency
CRD	0.495
RBD	0.539
BRCD	0.637
RCND	0.621
BRCND	**1**
α-RCD (implemented)	0.619

Table: Efficiencies under the full BRCNM.

Example 3: Agricultural Experiment

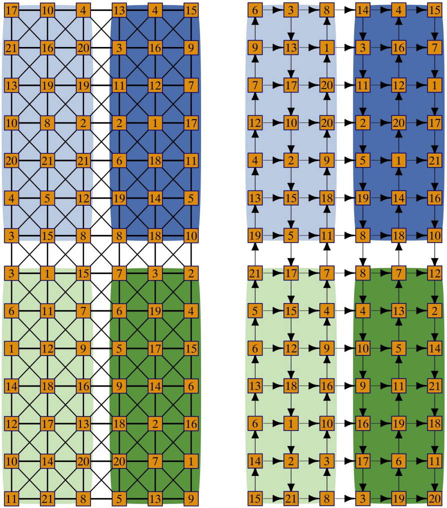

Figure: Left: Accounts for field variation due to the adjacency of plants (weighted network, with weights based on the inverse spatial distance between plots' centroids). Right: Accounts for variation introduced by cultivation and farming techniques (directed network).

Why Subsample?

In many experiments on large networks, not all nodes can be used:

- **Cost**: Each impression on a social media platform requires payment.
- **Access**: Only a proportion of agricultural plots may be available.
- **Computation**: Design search on the full network is intractable.

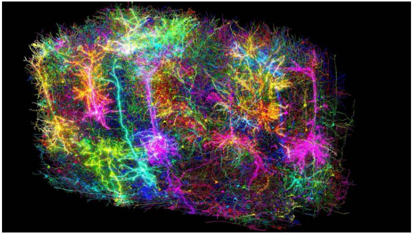

Figure: 76k neurons, >500M synapses in a 1 mm³ of mouse cortex (Nature, 2024)

Key insight: Under the Network Model, the parameters τ and γ are properties of the treatment mechanism, not the specific network. Sub-network and full-network designs target the same estimand; only efficiency depends on the sub-network choice.

Network Properties and Design Quality

Under the LNM with $t = 2$ treatments:

- $\phi_\tau \approx 4/n$ regardless of network structure $\Rightarrow$ depends only on size.
- ϕ_γ depends strongly on network structure.

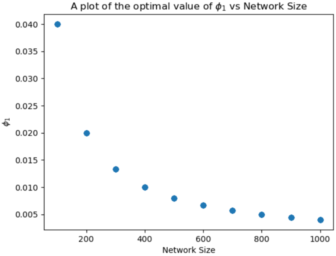

(a) ϕ_τ vs. network size

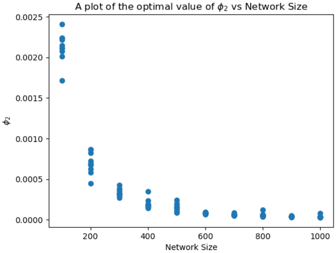

(b) ϕ_γ vs. network size

Figure: Optimal values for preferential attachment networks.

Edge Density is Key for ϕ_γ

Networks with edge density close to 0.5 give the best optimal designs for estimating indirect effects.

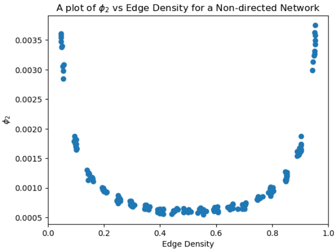

(a) Erdős–Rényi model

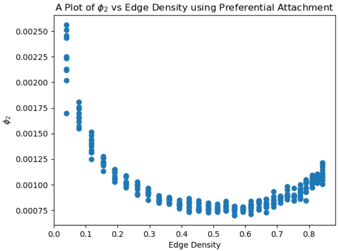

(b) Preferential attachment model

Figure: Optimal ϕ_γ vs. edge density for networks of size 100.

Since most real-world networks are sparse, sampling algorithms should aim to **maximise edge density** of the sub-network.

Sampling Algorithms

A range of sampling algorithms were assessed:

Algorithm	**Description**
RNS	Random node sampling
RNSR	RNS with replacement of isolated nodes
RNSOBJ	RNS with edge-density objective function
BFS	Breadth first sampling
DFS	Depth first sampling
SS	Snowball sampling (multiple seeds)
FFS	Forest fire sampling
CPS	Contact process sampling
CPSEP	CPS with equal probabilities
RWS	Random walk sampling

Edges between sampled nodes are kept as they appear in the original network, so the sampled network is always a subgraph of the full network.

Subsampling snapshots of Facebook networks (Rossi and Ahmed, 2015) with 769 and 50,515 nodes to 80 and 1,000 nodes respectively:

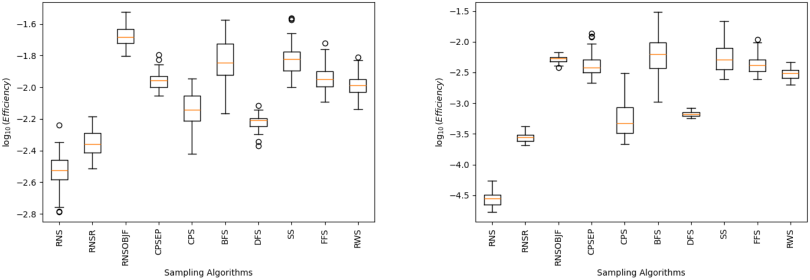

(a) Caltech network (769 nodes)

(b) UMichigan network (50,515 nodes)

Figure: Log efficiencies of ϕ_γ across sampling algorithms.

BFS, **Snowball Sampling** and **RNSOBJ** give the best designs.

Edge Density Drives Efficiency

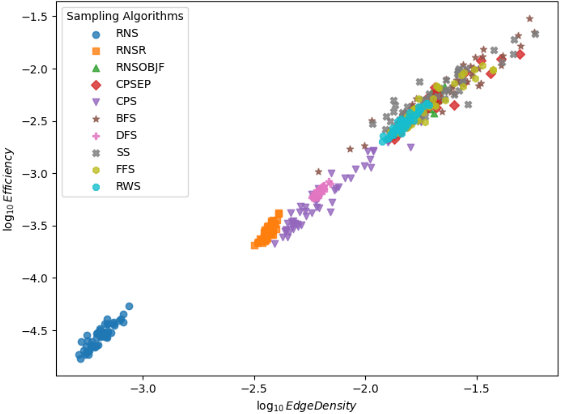

Figure: $\log_{10}$(efficiency) vs. $\log_{10}$(edge density) UMichigan network (50,515 nodes).

This linear relationship on a log-log scale justifies the use of edge density as an objective function in the RNSOBJ algorithm.

Bias from Unsampled Nodes

Subsampling drops units — but a dropped unit is often still a *neighbour* of one we kept, so its spillover leaks in unmodelled.

For each sampled unit i, let $k_i^* - k_i$ count its **missing (unsampled) neighbours**. This boundary leakage is the source of bias:

$$\textbf{Bias}(\hat{\beta}) = \gamma_0 \, (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T(\mathbf{k}^* - \mathbf{k}),$$

where $\mathbf{k}^*$ and $\mathbf{k}$ are neighbour counts (over all nodes vs. sampled nodes only), indexed over the *sampled* units.

- In a full-network experiment there are no missing neighbours ($\mathbf{k}^* = \mathbf{k}$) $\Rightarrow$ no bias.
- Fewer severed edges, or weaker spillover $|\gamma_0|$ $\Rightarrow$ less bias.
- The same dense, well-connected subsamples that give efficient designs also tend to leave fewer neighbours behind.

Discussion

- The network modelling framework provides a flexible foundation for designing and analysing experiments on networks.
- Designs which don't allow for interference effects can be very inefficient ($>80\%$ efficiency loss for indirect effects).
- When full-network experimentation is infeasible, sub-network designs can estimate the *same* parameters. Edge density of the sub-network is the key driver of efficiency.
- BFS, Snowball Sampling, and RNSOBJ are recommended sampling strategies.
- More work to be done on uncertain networks.
  CRAN package: `MOODE` (Koutra, 2024)

Thank you!

vasiliki.koutra@kcl.ac.uk