
Balancing Interference and Correlation in Spatial Experimental Designs: A Causal Graph Cut Approach

Jin Zhu ^{*1} Jingyi Li ^{*2} Hongyi Zhou ³ Yinan Lin ⁴ Zhenhua Lin ² Chengchun Shi ¹

Abstract

This paper focuses on the design of spatial experiments to optimize the amount of information derived from the experimental data and enhance the accuracy of the resulting causal effect estimator. We propose a surrogate function for the mean squared error (MSE) of the estimator, which facilitates the use of classical graph cut algorithms to learn the optimal design. Our proposal offers three key advances: (1) it accommodates moderate to large spatial interference effects; (2) it adapts to different spatial covariance functions; (3) it is computationally efficient. Theoretical results and numerical experiments based on synthetic environments and a dispatch simulator that models a city-scale ridesharing market, further validate the effectiveness of our design. A python implementation of our method is available at <https://github.com/Mamba413/CausalGraphCut>.

1. Introduction

Background. Before deploying any policy in practice, it is essential to evaluate its impact. This renders randomized control trials or online experiments extensively used for accurate policy evaluation. In numerous applications, the experiments involve multiple units that are either distributed across different spatial regions or connected through a network. The spatial or network dependencies inherent in these experiments pose two unique challenges for policy evaluation: (i) the existence of **interference effects** where

interventions applied to one experimental unit may influence others, leading to the violation of the fundamental stable unit treatment value assumption (SUTVA) for causal inference (Imbens & Rubin, 2015); (ii) the strong **outcome correlations** among units, which increase the variance of the causal effect estimator. We present the following examples to illustrate.

Example 1: A/B testing in marketplaces. A/B testing is frequently used in two-sided markets such as Uber, Lyft, Airbnb and eBay to assess the impact of a newly developed product relative to a standard control. It has become the gold standard for companies in conducting data-driven decision making (Johari et al., 2022). Spatial interference is ubiquitous in marketplaces. For instance, ridesharing companies constantly offer subsidizing policies to drivers or passengers across different regions in a city. Applying a subsidizing policy in one location can attract drivers from neighboring regions to this area, potentially decreasing the driver supply in other regions. Thus, a subsidizing policy in one location can affect outcomes at other locations, inducing interference over space (Shi et al., 2023). Meanwhile, market features like online driver numbers and call order numbers form interconnected spatial networks, leading to considerable spatial correlation (Luo et al., 2024).

Example 2: Environmental and epidemiological applications. Environmental and epidemiological studies frequently utilize spatial data to examine causal effects of certain interventions on health outcomes in different geographical regions (Reich et al., 2021). Spatial correlation and interference are likely to arise in these applications. For instance, in evaluating the impact of vaccination campaigns on disease incidence, infection rates exhibit strong correlations within regions. Furthermore, the vaccination coverage in a given region not only influences its own infection rate but also those in the neighboring regions (VanderWeele et al., 2012; Perez-Heydrich et al., 2014).

Example 3: Experimentation in social networks. Social network sites such as Facebook and LinkedIn extensively conduct online experiments to evaluate user engagement or satisfaction from a new service or feature. In these experiments, the treatment applied to a given user may spill over to other users via underlying social connections

^{*}Equal contribution ¹Department of Statistics, London School of Economics and Political Science, London, United Kingdom ²Department of Statistics and Data Science National University of Singapore, Singapore ³Department of Mathematics, Tsinghua University, Beijing, China ⁴National Center for Applied Mathematics in Chongqing, Chongqing Normal University, Chongqing, China. Correspondence to: Zhenhua Lin <linz@nus.edu.sg>, Chengchun Shi <c.shi7@lse.ac.uk>.

(like friends), leading to network interference. Additionally, since connected users often exhibit similar behaviors, their outcomes are likely correlated (Gui et al., 2015).

Contributions. This paper primarily focuses on the first application, aiming to identify the optimal design that minimizes the mean squared error (MSE) of the resulting average treatment effect (ATE) estimator. Nonetheless, the theories and methods developed herein can be adapted to the other two examples.

Methodologically, our contribution lies in the development of a causal graph cut algorithm for designing experiments with spatial interference and correlation. The advances of our algorithm are summarized as follows:

- **Flexibility:** Our proposal is more flexible than existing works that assume weak interference (Viviano et al., 2023). *It accommodates moderate to large interference effects.*
- **Adaptivity:** Instead of employing existing minimax approaches (Leung, 2022; Viviano et al., 2023; Zhao, 2024b) that minimize the same worst-case MSE across applications with different covariance structures and can be overly conservative, *our algorithm adapts to different correlation structures*, providing a tailored solution to each individual application.
- **Computational efficiency:** The design problem requires to allocate treatment for each region and is inherently NP-hard, as the number of treatment allocation rules grows exponentially fast with respect to the number of regions. Unlike existing solutions that rely on computationally intensive integer programming (Zhao, 2024b), *we propose a new surrogate function for the MSE, which enables the use of spectral clustering-based graph cut algorithms for minimization*, thus substantially enhancing computational efficiency.

Theoretically, our analysis reveals that *the two pivotal factors — interference and correlation — drive the optimal design in opposite directions*. Specifically, experiments with large interference effects are best designed by assigning same policies to neighboring regions whereas those with strong spatial correlations benefit from allocating different policies. By incorporating the effects of the two factors into our surrogate function, the proposed algorithm effectively address both challenges, achieving desirable properties.

Empirically, based on synthetic environments and a city-level simulator — constructed using physical models and a real dataset from a ridesharing company to realistically simulate driver and passenger behaviors, we demonstrate the superior performance of our design compared to existing state-of-the-art methods. Notably, in the simulator, *the proposed ATE estimator achieves an MSE that is 3.5 times smaller than those of the benchmark approaches*.

1.1. Related works

This section reviews related works on spatial causal inference and experimental design. Additional related works on off-policy evaluation are discussed in Appendix A.

Spatial causal inference. Our proposal is closely related to a growing line of research on causal inference in the presence of spatial interference. Depending on the underlying assumptions, the methodologies developed within this area can be broadly grouped into three categories:

1. The first category assumes **partial interference** where units are partitioned into clusters and interference is restricted to within these clusters, not between them (Halloran & Struchiner, 1995; Sobel, 2006; Hudgens & Halloran, 2008; Tchetgen & VanderWeele, 2012; Crépon et al., 2013; Liu et al., 2016; Zigler & Papadogeorgou, 2021).
2. The second category assumes **neighborhood interference** where interference is limited to a unit’s neighbors or a predefined local region (Verbitsky-Savitz & Raudenbush, 2012; Bhattacharya et al., 2020; Fatemi & Zheleva, 2020; Ma & Tresp, 2021; Gao & Ding, 2023; Dai et al., 2024; Jiang et al., 2024; Yang et al., 2024).
3. The last category considers more **general interference** structures (Aronow & Samii, 2017; Forastiere et al., 2021; Puelz et al., 2022; Song & Papadogeorgou, 2024; Zhang et al., 2024; Zhan et al., 2024).

In the machine learning literature, spatial causal inference is also related to recent advances on bandits with interference (Jia et al., 2024; Agarwal et al., 2024; Xu et al., 2024). However, despite the development of various methods to handle interference, experimental designs have been less considered in these papers.

Experimental design under interference. The design of experiments is a classical problem in statistics, motivated by applications ranging from biology and agriculture to psychology and engineering (Fisher et al., 1966). Recently, it has gained considerable attention in the field of machine learning (see e.g., Foster et al., 2021; Blau et al., 2022; Weltz et al., 2023; Fiez et al., 2024; Kato et al., 2024).

Our paper contributes to a growing line of research focusing on identifying optimal treatment assignment strategies in the presence of interference. Several designs have been proposed to guide treatment allocation in settings with interference over time (Hanna et al., 2017; Hu & Wager, 2022; Zhong et al., 2022; Bojinov et al., 2023; Li et al., 2023; Xiong et al., 2024; Sun et al., 2024; Wen et al., 2025) or among entities in marketplaces (Wager & Xu, 2019; Johari et al., 2020; Bajari et al., 2021; Munro et al., 2021; Li et al., 2022; Zhu et al., 2024).

This work studies designs under spatial or network interference. Methodologies developed within this field can be broadly categorized into three types:

1. The first type of methods adopts the principle of balance in the design of experiments. These methods focus on balancing covariates (Kallus, 2018; Liu et al., 2024; Harshaw et al., 2024) and cluster sizes in cluster-randomized designs (Gui et al., 2015; Saveski et al., 2017; Rolnick et al., 2019).
2. The second type of methods analytically calculates the MSE of the ATE estimator and uses it either to heuristically guide the design of cluster-randomized experiments (Ugander et al., 2013) or to determine the optimal number of clusters by minimizing the MSE’s order of magnitude (Leung, 2022; Jia et al., 2023).
3. The last type of methods identifies the optimal design through an optimization perspective (Zhao, 2024a). They apply optimization tools to directly minimize the MSE of the treatment effect estimator or its proxy (Baird et al., 2018; Ni et al., 2023; Jiang & Wang, 2023; Viviano et al., 2023; Chen et al., 2024). In particular, Viviano et al. (2023) markedly advanced the design of cluster-randomized experiments under network interference by proposing a causal clustering algorithm, which offers a rigorous framework to numerically solve the optimal design via graph clustering algorithms. Subsequent developments are explored in several recent studies (see e.g., Eichhorn et al., 2024; Zhao, 2024b).

Our paper falls into the third category: it utilizes classical graph cut algorithms for optimization. Unlike Viviano et al. (2023) that employs a minimax formulation along with a weak interference assumption to simplify the MSE — potentially at the cost of being conservative and restrictive — our proposed algorithm does not rely on minimax formulations or limit itself to weak interference. Despite the more complicated form of the MSE, we identify a suitable surrogate to optimize, backed by theoretical guarantees. Empirically, our proposal substantially outperforms both causal clustering and other state-of-the-art in our application.

2. Preliminaries

In this section, we first formulate the A/B testing problem in spatial experiments, presenting our model and assumptions. We next detail the ATE estimator. Finally, we introduce cluster-randomized designs, which are extensively used in the field (see e.g., Ugander et al., 2013; Karrer et al., 2021; Leung, 2022).

Spatial A/B testing. We consider a spatial settings with R non-overlapping regions across a city. The goal of spatial A/B testing is to compare the impact of implementing a

newly developed policy in the whole city against an existing practice. Its procedure can be summarized as follows:

- At the beginning of the experiment, we observe a covariate matrix $\mathbf{O} = (O_1, \dots, O_R)^\top \in \mathbb{R}^{R \times d}$ where each vector $O_i \in \mathbb{R}^d$ measures certain market features from the i th region. For instance, in ridesharing, O_i could represent the number of initial drivers within that region.
- During the experiment, the decision maker assigns policies by specifying a vector $\mathbf{A} = (A_1, A_2, \dots, A_R)^\top$ where each A_i is a binary variable indicating whether the i -th region receives the newly developed policy ($A_i = 1$) or a standard control policy ($A_i = 0$). In the context of ridesharing, A_i could correspond to whether to assign certain driver-side or passenger-side subsidizing policy at the i th region (Shi et al., 2023; Li et al., 2024).
- After the experiment, we collect an outcome vector $\mathbf{Y} = (Y_1, \dots, Y_R)^\top$ with each Y_i measuring the i -th region’s outcome of interest (such as the total driver income). Throughout this paper, we assume $Y_i = g_i(\mathbf{A}, \mathbf{O}) + e_i$ for some unknown function g_i and mean-zero error e_i independent of \mathbf{O} and \mathbf{A} .
- Finally, the experiment is repeated N times, resulting in N independent $(\mathbf{O}, \mathbf{A}, \mathbf{Y})$ triplets. Based on the data, we aim to estimate the ATE, defined as

$$\text{ATE} = \mathbb{E}[\text{CATE}(\mathbf{O})] = \sum_{i=1}^R \mathbb{E}[g_i(\mathbf{1}, \mathbf{O}) - g_i(\mathbf{0}, \mathbf{O})],$$

which measures the difference in outcome between applying the new policy globally to all regions and implementing the control policy. Here, CATE denotes the conditional ATE as a function of the observation matrix.

We make two remarks. First, unlike existing works on the design of spatial/network experiments (see Section 1.1) which analyze data from a single experiment without repeated measures, we consider settings with repeated measurements (Zhang & Wang, 2024). This is motivated by our ridesharing example where the experiment is conducted over two weeks, and each day’s data can be treated as an independent realization since the number of call orders typically wanes between 1 and 5 am (Luo et al., 2024).

Second, our outcome regression model $Y_i = g_i(\mathbf{A}, \mathbf{O}) + e_i$ manifests the two challenges in spatial A/B testing:

1. **Interference:** g_i depends not only on the i th region’s own treatment A_i , but also on the treatments assigned to other regions as well.
2. **Correlation:** The covariance matrix of the residual $\mathbf{e} = (e_1, \dots, e_R)^\top$, denoted by $\Sigma \in \mathbb{R}^{R \times R}$, is typically non-diagonal, indicating that residuals are correlated across regions.

Estimation. To simplify the estimation under spatial interference, we impose the neighborhood interference assumption introduced in Section 1.1, which restricts the interference to neighboring regions. It is a specific yet widely used exposure mapping assumption (Aronow & Samii, 2017).

Assumption 1 (Neighborhood interference). *For each region i , denote \mathcal{N}_i as the set of spatial neighboring regions for the i -th region, including the i -th region itself. Let $\mathbf{a}_{\mathcal{N}_i}$ be the $|\mathcal{N}_i|$ -dimensional subvector of \mathbf{a} that is formed by the treatments assigned to the regions in \mathcal{N}_i . For any $i \in \{1, 2, \dots, R\}$, $\mathbf{a}, \mathbf{a}' \in \{0, 1\}^R$ and $\mathbf{o} \in \mathbb{R}^d$, $g_i(\mathbf{a}, \mathbf{o}) = g_i(\mathbf{a}', \mathbf{o})$ whenever $\mathbf{a}_{\mathcal{N}_i} = \mathbf{a}'_{\mathcal{N}_i}$.*

For any $1 \leq t \leq n$ and $1 \leq i \leq R$, let $(A_{i,t}, Y_{i,t})$ denote the policy-outcome pair associated with the i -th region collected from the t -th experiment. Assumption 1 motivates us to consider the following importance sampling (IS) estimator popularly employed in the literature to estimate the ATE (see e.g., Ugander et al., 2013; Zhao, 2024b),

$$\widehat{\text{ATE}}^{\text{IS}} = \frac{1}{N} \sum_{t=1}^N \sum_{i=1}^R \left(\frac{T_{i,t}(\mathbf{1})}{\mathbb{E}[T_{i,t}(\mathbf{1})]} - \frac{T_{i,t}(\mathbf{0})}{\mathbb{E}[T_{i,t}(\mathbf{0})]} \right) Y_{i,t},$$

where $T_{i,t}(\mathbf{a}) := \prod_{j \in \mathcal{N}_i} \mathbb{I}(A_{j,t} = a_j)$ so that $T_{i,t}(\mathbf{1})$ (or $T_{i,t}(\mathbf{0})$) indicates whether all neighbors of the i -th region (including itself) are treated with the new policy (or the standard control).

A well-known limitation of IS is its large variance (Dudík et al., 2011). To mitigate this issue, we employ the following doubly robust (DR) estimator developed by Yang et al. (2024) for ATE estimation. DR attains theoretically no larger and empirically often smaller MSE than IS, and is defined as

$$\widehat{\text{ATE}}^{\text{DR}} = \frac{1}{N} \sum_{t=1}^N \sum_{i=1}^R [\nu_{i,t}(\mathbf{1}) - \nu_{i,t}(\mathbf{0})], \quad (1)$$

where $\nu_{i,t}(\mathbf{a}) = g_i(\mathbf{a}, \mathbf{O}_t) + \frac{T_{i,t}(\mathbf{a})}{\mathbb{E}[T_{i,t}(\mathbf{a})]} [Y_{i,t} - g_i(\mathbf{a}, \mathbf{O}_t)]$ and \mathbf{O}_t denotes the observation matrix from the t -th experiment. Notice that $\nu_{i,t}$ depends on g_i , which is unknown and must be estimated. In practice, estimation of g_i can be achieved either parametrically, using a specified functional form based on lower-dimensional statistics of actions and observations (Hu et al., 2022), or nonparametrically through neural networks (Leung & Loupos, 2022; Dai et al., 2024; Wang et al., 2024). Thanks to its double robustness property and the fact that the randomization probability $\mathbb{E}[T_{i,t}(\mathbf{a})]$ is known by design, the resulting estimator remains unbiased even when the estimated g_i is inconsistent. Notably, when the estimated g_i is set to zero, the resulting estimator simplifies to IS.

Design. In our context, each design specifies a treatment allocation rule that determines the joint probability distri-




Figure 1. Illustrations of the optimal cluster-randomized designs in three experiments, where different colors represent different clusters. (a) With weakly correlated residuals, the optimal design simplifies to the global design. (b) Without interference, the optimal design reduces to the individual design. (c) With moderate to large levels of interference and correlation, the optimal design falls between the two extremes, resulting in a 3-cluster design.

bution function of \mathbf{A} . Following the existing practice, we focus on the class of *cluster-randomized designs* in this paper. These designs partition the R regions into m disjoint clusters $\mathcal{C}_1, \dots, \mathcal{C}_m$ where regions within the same cluster receive the same treatment. Meanwhile, treatments across different clusters are independent, each following a Bernoulli distribution with a probability p of receiving the new policy. We further fix $p = 0.5$ throughout the paper to ensure a *balanced design*, since the optimal design is indeed balanced (Yang et al., 2024).

Two special cases are worthwhile to mention: (i) When $m = 1$, there is only one cluster. The resulting design is reduced to a *global design* where all regions receive the same policy during each experiment. (ii) When $m = R$, each region forms a single cluster, resulting in an *individual design* where policies are *i.i.d.* across regions.

Our objective lies in determining the optimal cluster-randomized design, for which we propose a causal graph cut algorithm to minimize the MSE of the ATE estimator in (1). Details are provided in the next section.

3. Causal Graph Cut

This section is organized as follows. Section 3.1 analyzes the roles of spatial interference and correlation in designing cluster-randomized experiments. Section 3.2 derives our surrogate function for the MSE of the ATE estimator. Finally, Section 3.3 details our causal graph cut algorithm.

3.1. Cluster-randomized design: the roles of interference and correlation

We will show that interference and correlation are the two driving factors in determining the MSE of ATE estimators under different designs. Interestingly, they influence the design of experiments in completely different directions:

Claim 1. *Experiments with large interference are best designed by applying same policies to neighboring regions.*

Claim 2. *Experiments with strong correlation are best designed by allocating different policies to different regions.*

To illustrate these two claims, we consider three experiments visualized in Figure 1, with the same spatial grid. They differ in the sizes of interference effects and correlation: (i) In the left experiment, there is moderate to large interference, while the residuals e_1, \dots, e_R are *weakly* correlated across regions. (ii) In the middle experiment, SUTVA holds, meaning there is *no interference*, but there is moderate to large correlation. (iii) In the right experiment, both levels of interference and correlation range from *moderate to large*.

Figure 1 further visualizes the optimal cluster-randomized designs for the three experiments, with different colors representing different clusters. It clearly shows that:

- (i) With weakly correlated errors, interference plays the leading role in determining the MSE. Since all regions are connected, according to Claim 1, the optimal allocation rule assigns the same policy to all regions, resulting in a global design.
- (ii) Without interference effects, correlation becomes the sole factor in determining the MSE. According to Claim 2, the optimal allocation rule assigns different treatments across regions, creating an individual design.
- (iii) When both interference and correlation are present, the optimal design typically falls between the two extremes, resulting in a cluster-randomized design. Specifically, in the right experiment, the proposed algorithm identifies a three-cluster design as the optimal design.

For any two regions i and i' , let $\Sigma_{ii'}$ denote the covariance between two residuals e_i and $e_{i'}$. Let $\delta > 0$ represent the size of the spatial correlation so that $|\Sigma_{ii'}| \leq \delta$ for any two non-neighboring regions i, i' . The following two propositions formally summarize the aforementioned findings.

Proposition 1. *When δ is sufficiently small (i.e., spatial correlation is weak among non-neighboring regions), and the covariance $\Sigma_{ii'}$ between any two neighboring regions i, i' is positive, the global design minimizes the MSE of the DR estimator (1).*

Proposition 2. *When SUTVA (i.e., the no-interference assumption) holds and all entries of Σ are non-negative, the individual design minimizes the MSE of the DR estimator.*

To verify the two claims and propositions, we offer a decomposition of the DR estimator’s MSE in the following theorem. Recall that N denotes the number of repeated experiments. For a given cluster \mathcal{C} , let $\partial\mathcal{C}$ denote its boundary.

Theorem 1. *Under Assumption 1, the DR estimator’s MSE under a cluster-randomized design with m clusters*

$\{\mathcal{C}_j\}_{j=1}^m$ *is given by*

$$\begin{aligned} \text{MSE}(\widehat{\text{ATE}}^{\text{DR}}) &= \underbrace{\frac{1}{N} \text{Var} \left\{ \text{CATE}(\mathbf{O}_t) \right\}}_{\text{design-agnostic (DA) term}} + \underbrace{\frac{4}{N} \left[\sum_{j=1}^m \sum_{i, i' \in \mathcal{C}_j} \Sigma_{ii'} \right]}_{\text{spatial correlation (SC) term}} \\ &\quad + \underbrace{\frac{8}{N} \left[\sum_{j \neq k} \sum_{i \in \mathcal{C}_j} \sum_{i' \in \partial\mathcal{C}_k} \Sigma_{ii'} \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset) \right]}_{\text{first-order interference term } I_1} \\ &\quad + \underbrace{O \left[\frac{1}{N} \sum_{j, k=1}^m \sum_{i \in \partial\mathcal{C}_j} \sum_{i' \in \partial\mathcal{C}_k} |\Sigma_{ii'}| \right]}_{\text{second-order interference term } I_2}, \end{aligned}$$

where the event $\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset$ indicates that the i' -th region is adjacent to the j -th cluster.

Theorem 1 decomposes the MSE into four terms. We elaborate each of them below.

1. The **design-agnostic term** (DA) accounts for the variation in covariates across N experiments. It is independent of the design and hence *design-agnostic*.
2. The **spatial correlation term** (SC) captures the residual covariances between regions within the same cluster, as illustrated in Figure 2(a). When the covariance function is non-negative, it becomes evident that minimizing this term is equivalent to assigning different policies to different regions, thus demonstrating Claim 2. Notably, it attains the minimum value under the individual design, which in turn proves Proposition 2.
3. The **interference terms** I_1 and I_2 equal zero in the absence of spatial interference. Unlike the SC which captures only within-cluster correlation, both I_1 and I_2 account for between-cluster correlation, as illustrated in Figure 2(b) and (c). Under neighborhood interference, the first-order term I_1 requires at least one region to lie on the boundaries between the two clusters, representing a *first-order* boundary effect. The second-order term I_2 further requires that both regions be located on the boundary, thus characterizing a *second-order* boundary effect. For any two neighboring regions, assigning them different policies causes them to belong to connected yet different clusters, which creates boundaries. Thus, minimizing these terms is equivalent to assigning the same policy to neighboring regions, which demonstrates Claim 1. Apparently, there is no boundary effect under the global design. This proves Proposition 1.

In summary, interference and correlation drive the optimal design in opposite directions. Interference induces estimation errors at cluster boundaries. Thus, assigning the same policy to neighboring regions eliminates the boundary effect and minimizes these errors. On the contrary, allocating




Figure 2. An illustration of each term in the MSE decomposition from Theorem 1 with $m = 2$ clusters distinguished by different colors. (a) SC captures the residual covariance between any i, i' within the same cluster. (b) I_1 captures the residual covariance between any i, i' that belong to different clusters, with one region lying on the boundary between the two clusters. (c) I_2 captures residual covariance between any i, i' from different clusters, with both regions located on the boundary.

different policies to different regions effectively negates the positive correlation between their residuals, thus reducing the estimation errors caused by spatial correlation.

3.2. Balancing interference and correlation: a new surrogate function for the MSE

In this section, we design a surrogate function aimed at achieving the following properties:

Property 1 (Flexibility). *The surrogate function shall accommodate moderate to large interference effects.*

Property 2 (Adaptivity). *The surrogate function shall adapt to different correlation structures.*

Property 3 (Computational efficiency). *The surrogate function can be optimized efficiently.*

Flexibility is crucial, since limiting the methodology to weak interference would be overly restrictive. Additionally, computational efficiency is essential for the practical implementation of the design. Below, we provide a numerical example to highlight the importance of adaptivity.

A numerical example. Consider a square spatial grid with $R = 144$ regions visualized in Figure 3(a) and an exponential spatial correlation structure: $\Sigma_{ij} = \rho^{|i-j|}$ for some $0 < \rho \leq 1$ that characterizes the strength of spatial correlation. We numerically compute MSEs of DR estimators under different cluster-randomized designs, each dividing the spatial grid into m equally sized square rectangles (see Figure 3(a) for an example with $m = 4$). Notice that varying ρ induces different covariance functions whereas varying m yields different designs. We report the MSE for each combination of ρ and m in Figure 3(b), where the MSE values on the y -axis are presented using a logarithmic scale for clearer visualization and comparison. These results emphasize two key motivations, which we discuss below.

1. We first note that the MSE varies considerably with cluster size. Notably, when $\rho = 0.9$, the MSE reaches its minimum value of 20 with 4 clusters, but increases




Figure 3. (a) A 12×12 grid divided into four square clusters. (b) A scatter plot visualizing the DR estimator's MSE as a function of cluster size. ρ is selected from $\{0.1, 0.3, 0.5, 0.7, 0.9\}$.

to over 40 under the global design. *This highlights the motivation for addressing the design problem itself.*

2. We also observe that the optimal cluster size varies with ρ . *This highlights the motivation for developing an adaptive method capable of identifying the optimal design tailored to each individual covariance function.*

A new surrogate function. It remains challenging to design a surrogate function for the MSE that satisfies the aforementioned three properties. Under neighborhood interference, the primary challenge arises from the presence of the two interference terms I_1 and I_2 (see Theorem 1), the calculation of which requires to specify cluster boundaries. Since clusters under different designs have different boundaries, optimizing these terms becomes an NP-hard combinatorial optimization problem, conflicting with Property 3. One solution is to impose the weak interference assumption as in Viviano et al. (2023), under which the boundary effect becomes negligible (see their proof of Lemma 3.2). However, this approach violates Property 1.

For illustration purposes, we will now focus on designs with two clusters to discuss how we address these challenges and detail our surrogate function. Its more general form, applicable to general numbers of clusters, is presented in Appendix B. Let \mathcal{C}_1 and \mathcal{C}_2 denote a partition of all R regions, our surrogate function is defined as follows:

$$\frac{8R}{N} \sum_{i \in \mathcal{C}_1} \sum_{i' \in \mathcal{C}_2} W_{ii'} \Sigma_{ii'}^+ - \frac{8}{N} \sum_{i \in \mathcal{C}_1} \sum_{i' \in \mathcal{C}_2} \Sigma_{ii'}, \quad (2)$$

where $\{W_{ii'}\}_{i,i'=1,\dots,R}$ denotes the binary adjacency matrix indicating whether two regions are adjacent, and $\Sigma_{ii'}^+ = \max(\Sigma_{ii'}, 0)$. This objective function is built upon Theorem 1 and is motivated by the following observations:

1. The DA term does not need to be incorporated in the surrogate, since it is design-agnostic.
2. Minimizing the SC term, i.e., the within-cluster covariance, is equivalent to maximizing the between-cluster covariance, represented by the second term in (2).
3. Instead of directly minimizing I_1 , we use a surrogate upper bound given by the first term in (2) that does not

involve cluster boundaries and can be more efficiently optimized.

4. Applying a similar upper bound to I_2 will result in a quartic objective function, rather than quadratic, rendering spectral clustering-based cut algorithms inapplicable. However, as indicated by Figure 2(c), only a small fraction of the regions lie at the boundary, making I_2 a high-order and negligible term when compared with I_1 . This justifies omitting I_2 from the objective function.

Proposition 3 below formally establishes that the first term in (2) serves as a valid upper bound for I_1 , under the following assumption which is automatically satisfied when the covariance function decays with distance — a condition commonly imposed in spatial statistics (Cressie, 2015).

Assumption 2 (Decaying covariance). *For any three disjoint regions i_1, i_2, i_3 such that only i_1 and i_2 are neighbors, it holds that $\Sigma_{i_1, i_2} \geq \Sigma_{i_1, i_3}$.*

Proposition 3. *Under Assumption 2, the first term in (2) upper bounds I_1 .*

Meanwhile, when the covariance function is bounded away from zero, i.e., $\min_{i, i'} \Sigma_{ii'} > 0$, the following Proposition shows that this upper bound remains tight in the sense that it is of the same order of magnitude as I_1 .

Proposition 4. *When $\epsilon > 0$, the first term in (2) is upper bounded by $d_{\max} \sigma I_1$ where d_{\max} denotes the maximum number of neighbors across regions and σ denotes the ratio $\max_{i \neq i'} \Sigma_{ii'} / \min_{i, i'} \Sigma_{ii'}$.*

It is also worthwhile to mention that our surrogate function explicitly relies on the covariance function Σ , allowing the resulting design to be dependent upon Σ . Alternatively, one could employ a minimax approach similar to Viviano et al. (2023) and Zhao (2024b) that derives a worst-case surrogate among all possible covariance functions. However, the resulting design is no longer adaptive to Σ .

In summary, the two terms in (2) measure the effects of interference and correlation, respectively. Optimizing these terms drives the resulting designs as described in Claims 1 and 2. By incorporating both, our surrogate function balances interference and correlation, successfully achieving the first two desired properties. We will discuss how to efficiently optimize this function in the next section.

3.3. Causal graph cut: the detailed algorithm

We focus on the case with two clusters to illustrate our algorithm. Notice that the surrogate function in (2) is proportional to $\sum_{i \in \mathcal{C}_1} \sum_{i' \in \mathcal{C}_2} \omega_{ii'}$ where $\omega_{ii'} = RW_{ii'} \Sigma_{ii'}^+ - \Sigma_{ii'}$ can be viewed as the weight between any two regions. Assigning two regions to different clusters results in a loss of their weight, and minimizing the surrogate function is equivalent to determining the optimal clustering that minimizes the total lost weight. This formulation allows us




Figure 4. (a) Spatial layout of 4 regions to be partitioned. (b) Each region is represented as a vertex, with edges determined by the weight ω . (c) The cut highlighted in red minimizes the total loss weight. (d) The spatial layout obtained after applying the 2-cut, illustrating how the regions are partitioned based on the graph cut.

to employ classical graph cut algorithms (see e.g., Goldschmidt & Hochbaum, 1994; Stoer & Wagner, 1997) to optimize our surrogate function, as illustrated in Figure 4.

Following Hagen & Kahng (1992), let L represent the Laplacian matrix of $\Omega = \{\omega_{ii'}\}_{i, i'}$ and define the partition vector \mathbf{x} such that $x_i = 1/\sqrt{R}$ if $i \in \mathcal{C}_1$ and $x_i = -1/\sqrt{R}$ if $i \in \mathcal{C}_2$. A key observation is that the cut loss function can be represented using the following quadratic form $\mathbf{x}^\top L \mathbf{x} / 4$. Instead of employing integer programming to optimize this quadratic form, we relax the search space to the entire unit ball to facilitate computation. This simplifies the optimization to identifying the eigenvectors of L . Given that the Laplacian matrix has a trivial eigenvector (a vector of all ones) with an associated eigenvalue of zero, classical graph cut algorithms typically compute the eigenvector associated with the second smallest eigenvalue, known as the Fiedler vector (Fiedler, 1973; 1989).

However, our problem slightly differs from the classical graph cut problem in that the weight matrix is not guaranteed to be positive definite. Consequently, if the smallest eigenvalue equals zero, we search for the Fiedler vector. Otherwise, we employ the first eigenvector. Once we have obtained the appropriate eigenvector, we apply the k -means algorithm to finalize the partition. In more general settings with $m > 2$ clusters, instead of searching for the first or second eigenvector, we search for a few (i.e., $\log_2(m)$ many), and apply the k -means algorithm to determine the m clusters. By varying m from 1 to a predefined maximum value m_{\max} — set to $R^{2/3}$ in our implementation based on Leung (2022)’s recommendation — we obtain m_{\max} many designs. We next evaluate their performance by plugging each design into the MSE formula derived from Theorem 1 (instead of the surrogate function) and choose the number of clusters that minimizes the MSE. Finally, the aforementioned discussion implicitly assumes that Σ is known; however, in practical applications, direct access to Σ is often not feasible. Since the experiment is repeated N times, we can estimate Σ to construct the MSE estimator and update the design in an iterative manner. We summarize the procedure in Algorithm 1.

Algorithm 1 Causal graph cut (CGC)

Input: A batch sample size B and an initial clustering \mathcal{C} .

- 1: **for** $l = 1, \dots, N/B$ **do**
- 2: Implement the design based on \mathcal{C} and collect the dataset $\mathcal{D}^{(l)} = \{(\mathbf{Y}_t, \mathbf{A}_t, \mathbf{O}_t)\}_{t=1}^B$.
- 3: Use datasets $\mathcal{D}^{(1)}, \dots, \mathcal{D}^{(l)}$ to estimate g_i (denoted by \hat{g}_i) and compute $\hat{e}_{it} = Y_{it} - \hat{g}_i(\mathbf{A}_{\mathcal{N}_i,t}, \mathbf{O}_{\mathcal{N}_i,t})$.
- 4: Compute the estimated covariance matrix $\hat{\Sigma}$ whose (i, i') -th entry is given by $\hat{\Sigma}_{ii'} = \sum_{t=1}^B \hat{e}_{it} \hat{e}_{i't} / B$.
- 5: **for** $m = 1, \dots, m_{\max}$ **do**
- 6: Use $\hat{\Sigma}$ to estimate the MSE and apply graph cut to learn the optimal design with m clusters.
- 7: **end for**
- 8: Select the number of clusters whose clustering minimizes the estimated MSE and update \mathcal{C} .
- 9: Calculate the DR estimator $\widehat{\text{ATE}}_l$ according to (1) using the data $\mathcal{D}^{(l)}$.
- 10: **end for**

Output: $\widehat{\text{ATE}} = (B/N) \sum_{l=1}^{N/B} \widehat{\text{ATE}}_l$

4. Experiments

In this section, we conduct numerical experiments to compare the ATE estimator computed via the proposed causal graph cut (CGC) algorithm against those constructed based on the following benchmark methods: (i) spectral clustering (SC, [Leung, 2022](#)); (ii) the three-net algorithm (TNET, [Ugander et al., 2013](#)); (iii) causal clustering (CC, [Viviano et al., 2023](#)). we also compare against the DR estimators constructed based on (iv) the global design (GD) and (v) the individual design (ID). We also compare against (vi) an oracle version of the proposed CGC algorithm (denoted by OCGC) which works as well as if the covariance matrix Σ is known in advance.

Synthetic environments. We design four spatial settings visualized in Figure 5(a), covering four classical urban morphology ([Kropf, 2017](#)). We set function g_i to be a sinusoidal function that incorporates both spatial covariates and treatments, with correlation parameter ρ indicating degree of spatial dependence, under three different settings; see Appendix D.1 for further details. We study the MSE of estimators as ρ increases from 0.1 to 0.9. The numerical results are shown in Figure 6, with values on the y -axis presented on a logarithmic scale for clearer visualization. It can be seen that our algorithms CGC and OCGC achieve MSEs that are substantially smaller than CC, TNET and SC. Besides, regardless of the spatial setting and covariance structure, our design consistently surpasses both the naive global design and individual design, with 54% and 72% improvements, which reflects the advantage of our proposed algorithms. Finally, it is worth mentioning that the CGC closely matches the performance of OCGC, indicating the

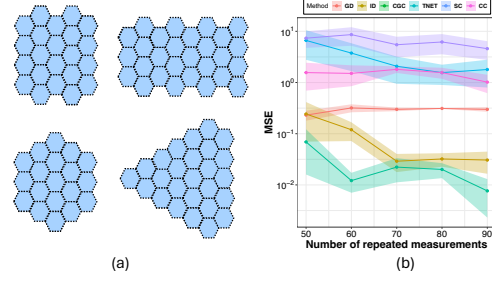


Figure 5. (a) Four types of spatial grids in synthetic environments: square (top-left), rectangle (top-right), circle (bottom-left), fans (bottom-right). (b) MSEs with different numbers of repetitions N in the real-data-based simulator. The shaded area visualizes the confidence interval.

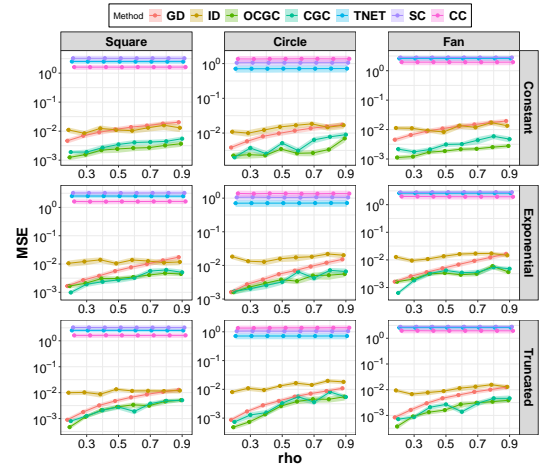


Figure 6. The plot of ρ versus MSE, with each panel representing a different spatial arrangement setting.

effectiveness of our iterative framework.

Real-data-based simulator. We develop a simulator using a five-day historical dataset from a ridesharing company to evaluate the proposed approach. This simulator replicates key dynamics of a city-scale ridesharing market with $R = 85$ hexagonal regions. Every 2 seconds, the simulator updates, among others: (1) drivers' decisions to accept assigned orders, taking into account historical driver and order characteristics; (2) driver movements to pick-up locations or idle movement patterns based on historical trajectories; (3) the driver pool, accounting for new active drivers and those going offline; and (4) new orders based on historical data, incorporating factors like passenger subsidies and processing unassigned and new requests. Particularly, drivers' decisions and reposition leverage the prediction of machine learning models trained with datasets that includes about 1 million observations. Furthermore, to match with the real world setting, when each simulation starts, drivers are distributed across the city based on their empirical distribution from the offline dataset that has more than 70,000 observations. See Appendix D.2 for more details.

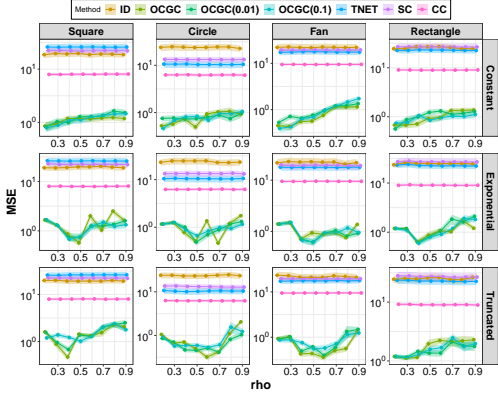


Figure 7. The plot of ρ versus MSE on single experiment setting, with each panel representing a different spatial arrangement setting. OCGC (0.01) and OCGC (0.1) represent our method using a noised covariance matrix, where the noise follows a normal variable with mean 0 and variance 0.01 and 0.1, respectively.

In the experiments, the covariate is the number of initial drivers available at the start of each time period. The outcome, measured by gross merchandise value (GMV), will be used to evaluate the effectiveness of the policy. We exclude OCGC since we cannot assess the spatial correlation of the real-world environment. The results are presented in Figure 5(b).

First, we observe that our proposed method achieves the lowest MSE in most cases, with its MSE consistently decreasing as the number of repeated experiments increases, demonstrating the learning capability of our iterative framework. The improvement over the benchmarks is substantial, with the MSE being approximately 3.5 times lower. Second, due to the strong spatial correlation in real-world environments, the global design is significantly less effective than the individual design, further supporting Proposition 1. The individual design, in turn, is substantially less effective than the proposed method, highlighting the advantages of our approach in balancing spatial correlation and interference.

Single-experiment setting. Our core methodology does not rely on the assumption of repeated experiments. It remains applicable to single-experiment settings when we have certain prior knowledge regarding the underlying covariance matrix, either from a pilot study or based on historical data.

To satisfy the practical requirement, we utilized a proxy covariance matrix, obtained by inserting noises into the true covariance matrix, and the results are shown in Figure 7. The results demonstrate that our estimator: (i) maintains optimality against existing methods, (ii) achieves near-oracle performance (comparable to the oracle method with the true covariance matrix), (iii) remains robust to the approximation errors.

We also remark that while cross-fitting helps simplify the theory (by avoiding the need for imposing VC-class conditions on to establish the asymptotics of the ATE estimator), our method remains effective without it - numerical results above were obtained without cross-fitting.

Acknowledgements

Zhu’s and Shi’s research were supported by the EPSRC grant EP/W014971/1. Hongyi’s research was partially supported by the China Scholarship Council. Lin’s research was partially supported by MOE AcRF Tier 1 grant. The authors express their gratitude for the constructive comments provided by the referees, which enhanced the initial version of the paper.

Impact statement

The broader impact of this work lies in its potential to enhance policy evaluation and decision-making across diverse domains, including urban planning, transportation, healthcare, and social sciences. Its scalability and adaptability provide opportunities to address complex societal challenges where spatial data plays a pivotal role. However, the application of our adaptive experimental design algorithm requires careful consideration of data privacy, transparency, and informed consent—especially in sensitive contexts involving personal or sensitive information—to ensure its responsible and ethical deployment.

References

- Agarwal, A., Agarwal, A., Masoero, L., and Whitehouse, J. Multi-armed bandits with network interference. *arXiv preprint arXiv:2405.18621*, 2024.
- Aronow, P. M. and Samii, C. Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics*, 11(4), December 2017. ISSN 1932-6157. doi: 10.1214/16-aos1005. URL <http://dx.doi.org/10.1214/16-AOS1005>.
- Baird, S., Bohren, J. A., McIntosh, C., and Özler, B. Optimal design of experiments in the presence of interference. *The Review of Economics and Statistics*, 100(5): 844–860, 12 2018. ISSN 0034-6535. doi: 10.1162/rest_a.00716. URL https://doi.org/10.1162/rest_a_00716.
- Bajari, P., Burdick, B., Imbens, G. W., Masoero, L., McQueen, J., Richardson, T., and Rosen, I. M. Multiple randomization designs. *arXiv preprint arXiv:2112.13495*, 2021.
- Bhattacharya, R., Malinsky, D., and Shpitser, I. Causal

- inference under interference and network uncertainty. In *Uncertainty in Artificial Intelligence*, pp. 1028–1038. PMLR, 2020.
- Blau, T., Bonilla, E. V., Chades, I., and Dezfouli, A. Optimizing sequential experimental design with deep reinforcement learning. In *International Conference on Machine Learning*, pp. 2107–2128. PMLR, 2022.
- Bojinov, I., Simchi-Levi, D., and Zhao, J. Design and analysis of switchback experiments. *Management Science*, 69(7):3759–3777, 2023.
- Chen, Q., Li, B., Deng, L., and Wang, Y. Optimized covariance design for a/b test on social network under interference. *Advances in Neural Information Processing Systems*, 36, 2024.
- Chen, X. and Qi, Z. On well-posedness and minimax optimal rates of nonparametric q-function estimation in off-policy evaluation. In *International Conference on Machine Learning*, pp. 3558–3582. PMLR, 2022.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018. ISSN 1368-4221. doi: 10.1111/ectj.12097. URL <https://doi.org/10.1111/ectj.12097>.
- Chernozhukov, V., Newey, W. K., and Singh, R. Debiased machine learning of global and local parameters using regularized riesz representers. *The Econometrics Journal*, 25(3):576–601, 2022.
- Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P. Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *The quarterly journal of economics*, 128(2):531–580, 2013.
- Cressie, N. *Statistics for spatial data*. John Wiley & Sons, 2015.
- Dai, B., Nachum, O., Chow, Y., Li, L., Szepesvári, C., and Schuurmans, D. Coindice: Off-policy confidence interval estimation. *Advances in Neural Information Processing Systems*, 33:9398–9411, 2020.
- Dai, R., Wang, J., Zhou, F., Luo, S., Qin, Z., Shi, C., and Zhu, H. Causal deepsets for off-policy evaluation under spatial or spatio-temporal interferences. *arXiv preprint arXiv:2407.17910*, 2024.
- Dudík, M., Langford, J., and Li, L. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pp. 1097–1104, 2011.
- Eichhorn, M., Khan, S., Ugander, J., and Yu, C. L. Low-order outcomes and clustered designs: combining design and analysis for causal inference under network interference. *arXiv preprint arXiv:2405.07979*, 2024.
- Fatemi, Z. and Zheleva, E. Minimizing interference and selection bias in network experiment design. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pp. 176–186, 2020.
- Feng, Y., Ren, T., Tang, Z., and Liu, Q. Accountable off-policy evaluation with kernel bellman statistics. In *International Conference on Machine Learning*, pp. 3102–3111. PMLR, 2020.
- Fiedler, M. Algebraic connectivity of graphs. *Czechoslovak mathematical journal*, 23(2):298–305, 1973.
- Fiedler, M. Laplacian of graphs and algebraic connectivity. *Banach Center Publications*, 1(25):57–70, 1989.
- Fiez, T., Nassif, H., Chen, Y.-C., Gamez, S., and Jain, L. Best of three worlds: Adaptive experimentation for digital marketing in practice. In *Proceedings of the ACM on Web Conference 2024*, pp. 3586–3597, 2024.
- Fisher, R. A., Fisher, R. A., Genetiker, S., Fisher, R. A., Geneticien, S., Britain, G., Fisher, R. A., and Généticien, S. *The design of experiments*, volume 21. Springer, 1966.
- Forastiere, L., Airolidi, E. M., and Mealli, F. Identification and estimation of treatment and interference effects in observational studies on networks. *Journal of the American Statistical Association*, 116(534):901–918, 2021.
- Foster, A., Ivanova, D. R., Malik, I., and Rainforth, T. Deep adaptive design: Amortizing sequential bayesian experimental design. In *International Conference on Machine Learning*, pp. 3384–3395. PMLR, 2021.
- Gao, M. and Ding, P. Causal inference in network experiments: regression-based analysis and design-based properties. *arXiv preprint arXiv:2309.07476*, 2023.
- Goldschmidt, O. and Hochbaum, D. S. A polynomial algorithm for the k-cut problem for fixed k. *Mathematics of operations research*, 19(1):24–37, 1994.
- Gui, H., Xu, Y., Bhasin, A., and Han, J. Network a/b testing: From sampling to estimation. In *Proceedings of the 24th International Conference on World Wide Web*, pp. 399–409, 2015.
- Hagen, L. and Kahng, A. B. New spectral methods for ratio cut partitioning and clustering. *IEEE transactions on computer-aided design of integrated circuits and systems*, 11(9):1074–1085, 1992.

- Hahn, J. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, pp. 315–331, 1998.
- Halloran, M. E. and Struchiner, C. J. Causal inference in infectious diseases. *Epidemiology*, 6:142–151, 1995.
- Hanna, J. P., Thomas, P. S., Stone, P., and Niekum, S. Data-efficient policy evaluation through behavior policy search. In *International Conference on Machine Learning*, pp. 1394–1403. PMLR, 2017.
- Hao, B., Ji, X., Duan, Y., Lu, H., Szepesvári, C., and Wang, M. Bootstrapping fitted q-evaluation for off-policy inference. In *International Conference on Machine Learning*, pp. 4074–4084. PMLR, 2021.
- Harshaw, C., Sävje, F., Spielman, D. A., and Zhang, P. Balancing covariates in randomized experiments with the gram–schmidt walk design. *Journal of the American Statistical Association*, 119(548):2934–2946, 2024.
- Heckman, J. J., Ichimura, H., and Todd, P. Matching as an econometric evaluation estimator. *The review of economic studies*, 65(2):261–294, 1998.
- Hu, Y. and Wager, S. Switchback experiments under geometric mixing. *arXiv preprint arXiv:2209.00197*, 2022.
- Hu, Y. and Wager, S. Off-policy evaluation in partially observed markov decision processes under sequential ignorability. *The Annals of Statistics*, 51(4):1561–1585, 2023.
- Hu, Y., Li, S., and Wager, S. Average direct and indirect causal effects under interference. *Biometrika*, 109(4): 1165–1172, 2022.
- Hudgens, M. G. and Halloran, M. E. Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842, 2008.
- Imbens, G. W. and Rubin, D. B. *Causal inference in statistics, social, and biomedical sciences*. Cambridge university press, 2015.
- Jia, S., Kallus, N., and Yu, C. L. Clustered switchback experiments: Near-optimal rates under spatiotemporal interference. *arXiv preprint arXiv:2312.15574*, 2023.
- Jia, S., Frazier, P., and Kallus, N. Multi-armed bandits with interference. *arXiv preprint arXiv:2402.01845*, 2024.
- Jiang, N. and Li, L. Doubly robust off-policy value evaluation for reinforcement learning. In *International conference on machine learning*, pp. 652–661. PMLR, 2016.
- Jiang, Y. and Wang, H. Causal inference under network interference using a mixture of randomized experiments. *arXiv preprint arXiv:2309.00141*, 2023.
- Jiang, Y., Deng, L., Wang, Y., and Wang, H. Causal inference in social platforms under approximate interference networks. *arXiv preprint arXiv:2408.04441*, 2024.
- Johari, R., Li, H., and Weintraub, G. Experimental design in two-sided platforms: An analysis of bias. *arXiv preprint arXiv:2002.05670*, 2020. "To appear in ACM Economics and Computation 2020".
- Johari, R., Koomen, P., Pekelis, L., and Walsh, D. Always valid inference: Continuous monitoring of a/b tests. *Operations Research*, 70(3):1806–1821, 2022.
- Kallus, N. Optimal a priori balance in the design of controlled experiments. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 80(1):85–112, 2018.
- Kallus, N. and Uehara, M. Efficiently breaking the curse of horizon in off-policy evaluation with double reinforcement learning. *Operations Research*, 70(6):3282–3302, 2022.
- Karrer, B., Shi, L., Bhole, M., Goldman, M., Palmer, T., Gelman, C., Konutgan, M., and Sun, F. Network experimentation at scale. In *Proceedings of the 27th acm sigkdd conference on knowledge discovery & data mining*, pp. 3106–3116, 2021.
- Kato, M., Oga, A., Komatsubara, W., and Inokuchi, R. Active adaptive experimental design for treatment effect estimation with covariate choices. *arXiv preprint arXiv:2403.03589*, 2024.
- Kropf, K. *The handbook of urban morphology*. John Wiley & Sons, 2017.
- Kuzborskij, I., Vernade, C., Gyorgy, A., and Szepesvári, C. Confident off-policy evaluation and selection through self-normalized importance weighting. In *International Conference on Artificial Intelligence and Statistics*, pp. 640–648. PMLR, 2021.
- Le, H., Voloshin, C., and Yue, Y. Batch policy learning under constraints. In *International Conference on Machine Learning*, pp. 3703–3712. PMLR, 2019.
- Leung, M. P. Rate-optimal cluster-randomized designs for spatial interference. *The Annals of Statistics*, 50(5): 3064–3087, 2022.
- Leung, M. P. and Loupos, P. Graph neural networks for causal inference under network confounding. *arXiv preprint arXiv:2211.07823*, 2022.
- Li, H., Zhao, G., Johari, R., and Weintraub, G. Y. Interference, bias, and variance in two-sided marketplace experimentation: Guidance for platforms. In *Proceedings of the ACM Web Conference 2022*, pp. 182–192, 2022.

- Li, T., Shi, C., Wang, J., Zhou, F., et al. Optimal treatment allocation for efficient policy evaluation in sequential decision making. *Advances in Neural Information Processing Systems*, 36, 2023.
- Li, T., Shi, C., Lu, Z., Li, Y., and Zhu, H. Evaluating dynamic conditional quantile treatment effects with applications in ridesharing. *Journal of the American Statistical Association*, pp. 1–15, 2024.
- Liu, L., Hudgens, M. G., and Becker-Dreps, S. On inverse probability-weighted estimators in the presence of interference. *Biometrika*, 103(4):829–842, 2016.
- Liu, L., Hudgens, M. G., Saul, B., Clemens, J. D., Ali, M., and Emch, M. E. Doubly robust estimation in observational studies with partial interference. *Stat*, 8(1):e214, 2019.
- Liu, Q., Li, L., Tang, Z., and Zhou, D. Breaking the curse of horizon: Infinite-horizon off-policy estimation. *Advances in Neural Information Processing Systems*, 31, 2018.
- Liu, Y., Zhou, Y., Li, P., and Hu, F. Cluster-adaptive network a/b testing: From randomization to estimation. *Journal of Machine Learning Research*, 25(170):1–48, 2024.
- Luckett, D. J., Laber, E. B., Kahkoska, A. R., Maahs, D. M., Mayer-Davis, E., and Kosorok, M. R. Estimating dynamic treatment regimes in mobile health using v-learning. *Journal of the American Statistical Association*, 115(530):692–706, 2020.
- Luo, S., Yang, Y., Shi, C., Yao, F., Ye, J., and Zhu, H. Policy evaluation for temporal and/or spatial dependent experiments. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, pp. qkad136, 2024.
- Ma, Y. and Tresp, V. Causal inference under networked interference and intervention policy enhancement. In *International Conference on Artificial Intelligence and Statistics*, pp. 3700–3708. PMLR, 2021.
- Munro, E., Wager, S., and Xu, K. Treatment effects in market equilibrium. *arXiv preprint arXiv:2109.11647*, 2021.
- Ni, T., Bojinov, I., and Zhao, J. Design of panel experiments with spatial and temporal interference. *Available at SSRN 4466598*, 2023.
- Perez-Heydrich, C., Hudgens, M. G., Halloran, M. E., Clemens, J. D., Ali, M., and Emch, M. E. Assessing effects of cholera vaccination in the presence of interference. *Biometrics*, 70(3):731–741, 2014.
- Puelz, D., Basse, G., Feller, A., and Toulis, P. A graph-theoretic approach to randomization tests of causal effects under general interference. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):174–204, 2022.
- Reich, B. J., Yang, S., Guan, Y., Giffin, A. B., Miller, M. J., and Rappold, A. A review of spatial causal inference methods for environmental and epidemiological applications. *International Statistical Review*, 89(3):605–634, 2021.
- Rolnick, D., Aydin, K., Pouget-Abadie, J., Kamali, S., Mirrokni, V., and Najmi, A. Randomized experimental design via geographic clustering. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2745–2753, 2019.
- Saveski, M., Pouget-Abadie, J., Saint-Jacques, G., Duan, W., Ghosh, S., Xu, Y., and Airolidi, E. M. Detecting network effects: Randomizing over randomized experiments. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1027–1035, 2017.
- Shi, C., Zhang, S., Lu, W., and Song, R. Statistical inference of the value function for reinforcement learning in infinite-horizon settings. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(3):765–793, 2022.
- Shi, C., Wan, R., Song, G., Luo, S., Zhu, H., and Song, R. A multiagent reinforcement learning framework for off-policy evaluation in two-sided markets. *The Annals of Applied Statistics*, 17(4):2701–2722, 2023.
- Sobel, M. E. What do randomized studies of housing mobility demonstrate? causal inference in the face of interference. *Journal of the American Statistical Association*, 101(476):1398–1407, 2006.
- Song, Z. and Papadogeorgou, G. Bipartite causal inference with interference, time series data, and a random network. *arXiv preprint arXiv:2404.04775*, 2024.
- Stoer, M. and Wagner, F. A simple min-cut algorithm. *Journal of the ACM (JACM)*, 44(4):585–591, 1997.
- Sun, K., Kong, L., Zhu, H., and Shi, C. Arma-design: Optimal treatment allocation strategies for a/b testing in partially observable time series experiments. *arXiv preprint arXiv:2408.05342*, 2024.
- Swaminathan, A. and Joachims, T. The self-normalized estimator for counterfactual learning. *advances in neural information processing systems*, 28, 2015.

- Tan, Z. Bounded, efficient and doubly robust estimation with inverse weighting. *Biometrika*, 97(3):661–682, 2010.
- Tchetgen, E. J. T. and VanderWeele, T. J. On causal inference in the presence of interference. *Statistical methods in medical research*, 21(1):55–75, 2012.
- Thomas, P. and Brunskill, E. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pp. 2139–2148. PMLR, 2016.
- Thomas, P., Theocharous, G., and Ghavamzadeh, M. High-confidence off-policy evaluation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- Ugander, J., Karrer, B., Backstrom, L., and Kleinberg, J. Graph cluster randomization: Network exposure to multiple universes. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 329–337, 2013.
- VanderWeele, T. J., Vandenbroucke, J. P., Tchetgen, E. J. T., and Robins, J. M. A mapping between interactions and interference: implications for vaccine trials. *Epidemiology*, 23(2):285–292, 2012.
- Verbitsky-Savitz, N. and Raudenbush, S. W. Causal inference under interference in spatial settings: a case study evaluating community policing program in chicago. *Epidemiologic Methods*, 1(1):107–130, 2012.
- Viviano, D., Lei, L., Imbens, G., Karrer, B., Schrijvers, O., and Shi, L. Causal clustering: design of cluster experiments under network interference. *arXiv preprint arXiv:2310.14983*, 2023.
- Wager, S. and Xu, K. Experimenting in equilibrium. *arXiv preprint arXiv:1903.02124*, 2019.
- Wan, R., Kveton, B., and Song, R. Safe exploration for efficient policy evaluation and comparison. In *International Conference on Machine Learning*, pp. 22491–22511. PMLR, 2022.
- Wang, J., Qi, Z., and Wong, R. K. Projected state-action balancing weights for offline reinforcement learning. *The Annals of Statistics*, 51(4):1639–1665, 2023.
- Wang, Y., Gu, C., and Otsu, T. Graph neural networks: Theory for estimation with application on network heterogeneity. *arXiv preprint arXiv:2401.16275*, 2024.
- Weltz, J., Fiez, T., Volfovsky, A., Laber, E., Mason, B., Jain, L., et al. Experimental designs for heteroskedastic variance. *Advances in Neural Information Processing Systems*, 36, 2023.
- Wen, Q., Shi, C., Yang, Y., Tang, N., and Zhu, H. Unraveling the interplay between carryover effects and reward autocorrelations in switchback experiments. In *International Conference on Machine Learning*. PMLR, 2025.
- Xiong, R., Chin, A., and Taylor, S. J. Data-driven switchback experiments: Theoretical tradeoffs and empirical bayes designs. *arXiv preprint arXiv:2406.06768*, 2024.
- Xu, Y., Lu, W., and Song, R. Linear contextual bandits with interference. *arXiv preprint arXiv:2409.15682*, 2024.
- Yang, Y., Shi, C., Yao, F., Wang, S., and Zhu, H. Spatially randomized designs can enhance policy evaluation. *arXiv preprint arXiv:2403.11400*, 2024.
- Zhan, R., Han, S., Hu, Y., and Jiang, Z. Estimating treatment effects under recommender interference: A structured neural networks approach. *arXiv preprint arXiv:2406.14380*, 2024.
- Zhang, B., Tsiatis, A. A., Laber, E. B., and Davidian, M. A robust method for estimating optimal treatment regimes. *Biometrics*, 68(4):1010–1018, 2012.
- Zhang, W., Yang, Y., and Yao, F. Spatial interference detection in treatment effect model. *arXiv preprint arXiv:2409.04836*, 2024.
- Zhang, Z. and Wang, Z. Online experimental design with estimation-regret trade-off under network interference. *arXiv preprint arXiv:2412.03727*, 2024.
- Zhao, J. Experimental design for causal inference through an optimization lens. In *Tutorials in Operations Research: Smarter Decisions for a Better World*, pp. 146–188. INFORMS, 2024a.
- Zhao, J. A simple formulation for causal clustering. *Available at SSRN 5008213*, 2024b.
- Zhao, Y., Zeng, D., Rush, A. J., and Kosorok, M. R. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118, 2012.
- Zhong, R., Zhang, D., Schäfer, L., Albrecht, S., and Hanna, J. Robust on-policy sampling for data-efficient policy evaluation in reinforcement learning. *Advances in Neural Information Processing Systems*, 35:37376–37388, 2022.
- Zhu, Z., Cai, Z., Zheng, L., and Si, N. Seller-side experiments under interference induced by feedback loops in two-sided platforms. *arXiv preprint arXiv:2401.15811*, 2024.

Zigler, C. M. and Papadogeorgou, G. Bipartite causal inference with interference. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 36(1): 109, 2021.

A. Additional related works

Off-policy evaluation (OPE). Our work is also closely related to OPE in contextual bandits and reinforcement learning. OPE aims to learn the expected return under a target policy using an offline dataset collected from a possibly different behavior policy. There are three commonly-used OPE methods:

1. Direct method derives the policy value estimator by learning a value function that measures the (cumulative) reward under the target policy (Hahn, 1998; Le et al., 2019; Feng et al., 2020; Luckett et al., 2020; Hao et al., 2021; Chen & Qi, 2022; Shi et al., 2022).
2. Importance sampling (IS) method reweights the observed reward using IS ratios to address the distributional shift between target and behavior policies (Heckman et al., 1998; Zhao et al., 2012; Swaminathan & Joachims, 2015; Thomas et al., 2015; Liu et al., 2018; Dai et al., 2020; Kuzborskij et al., 2021; Wang et al., 2023; Hu & Wager, 2023).
3. Doubly robust (DR) method combines the direct and IS estimators to achieve more accurate and robust estimation (Tan, 2010; Dudík et al., 2011; Zhang et al., 2012; Jiang & Li, 2016; Thomas & Brunskill, 2016; Liu et al., 2019; Chernozhukov et al., 2022; Kallus & Uehara, 2022).

However, none of these aforementioned studies address interference or focus on experimental design. While works such as Wan et al. (2022) and Hanna et al. (2017) consider the design problem to optimize the efficiency of the OPE estimator, they do not address spatial/network interference either.

B. Surrogate function for general m

For the general case where $m > 2$, similarly, to minimize MSE that is affected by the clustering configuration, it is sufficient to minimize the upper bound of $SC + I_1$ by addressing the following objective function:

$$\sum_{j < k} \sum_{i \in \mathcal{C}_j} \sum_{i' \in \mathcal{C}_k} \frac{2R}{m} \Sigma_{ii'}^+ W_{ii'} - \Sigma_{ii'}, \quad (3)$$

where the first term comes from I_1 and the later one comes from SC, collectively forming a graph-cut problem.

C. Proofs

C.1. Proof of Theorem 1

Proof. According to Theorem 2 in Yang et al. (2024) and assuming functions g_i are known, we have

$$\text{MSE}(\widehat{\text{ATE}}^{\text{DR}}) = \sigma_1^2 + \text{DA},$$

where

$$\begin{aligned} \sigma_1^2 &= \frac{1}{N} \sum_{i, i'} \left(\frac{1}{p^{m_{ii'}}} + \frac{1}{(1-p)^{m_{ii'}}} \right) \Sigma_{ii'} \mathbb{I}(m_{ii'} > 0), \\ \text{DA} &= \frac{1}{N} \text{Var} \left\{ \sum_{i=1}^R [g_i(\mathbf{1}, \mathbf{O}_t) - g_i(\mathbf{0}, \mathbf{O}_t)] \right\}, \end{aligned}$$

and $m_{ii'} = \sum_{j=1}^m \mathbb{I}(\mathcal{N}_i \cap \mathcal{C}_j \neq \emptyset) \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset)$ represents the number of common clusters that the i -th region with its neighbor and the i' -th region with its neighbor belong to. We now elaborate that $\sigma_1^2 = \text{SC} + I_1 + I_2$.

Note that we set $p = 0.5$ throughout this paper for balanced design. Let \mathcal{C}_j^0 denote the interior region of cluster \mathcal{C}_j . We first rewrite σ_1^2 in the following form:

$$\sigma_1^2 = \frac{1}{N} \text{Var} \left[\sum_{j=1}^m \left(\sum_{i \in \mathcal{C}_j^0} \frac{2A_i - 1}{p} e_i + \sum_{i \in \partial \mathcal{C}_j} \frac{\mathbb{I}(A_{\mathcal{N}_i} = 1) - \mathbb{I}(A_{\mathcal{N}_i} = 0)}{p^{m_i}} e_i \right) \right],$$

which can be decomposed into the sum of a within-cluster term σ_w^2 and a between-cluster term σ_b^2 . Specifically,

$$\begin{aligned}
 \sigma_w^2 &= \frac{1}{N} \left[\sum_{j=1}^m \sum_{i, i' \in \mathcal{C}_j^0} \frac{\Sigma_{ii'}}{p^2} + 2 \sum_{j=1}^m \sum_{i \in \mathcal{C}_j^0} \sum_{i' \in \partial \mathcal{C}_j} \frac{\Sigma_{ii'}}{p^2} + \sum_{j=1}^m \sum_{i, i' \in \partial \mathcal{C}_j} \frac{\Sigma_{ii'}}{p^{1+m_{ii'}}} \right] \\
 &= \frac{1}{N} \left[\sum_{j=1}^m \sum_{i, i' \in \mathcal{C}_j} \frac{\Sigma_{ii'}}{p^2} + \sum_{j=1}^m \sum_{i, i' \in \partial \mathcal{C}_j} \frac{\Sigma_{ii'}(1-p^{m_{ii'}-1})}{p^{1+m_{ii'}}} \right] \\
 &= SC + \underbrace{\frac{1}{N} \left[\sum_{j=1}^m \sum_{i, i' \in \partial \mathcal{C}_j} \frac{\Sigma_{ii'}(1-p^{m_{ii'}-1})}{p^{1+m_{ii'}}} \right]}_{J_1:=}, \\
 \sigma_b^2 &= \frac{2}{N} \left[\sum_{j \neq k} \sum_{i \in \mathcal{C}_j^0} \sum_{i' \in \partial \mathcal{C}_k} \frac{\Sigma_{ii'}}{p^2} \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset) \right] + \frac{1}{N} \left[\sum_{j \neq k} \sum_{i \in \partial \mathcal{C}_j} \sum_{i' \in \partial \mathcal{C}_k} \frac{\Sigma_{ii'}}{p^{m_{ii'}+1}} \mathbb{I}(m_{ii'} > 0) \right] \\
 &= \frac{2}{N} \left[\sum_{j \neq k} \sum_{i \in \mathcal{C}_j} \sum_{i' \in \partial \mathcal{C}_k} \frac{\Sigma_{ii'}}{p^2} \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset) \right] + \frac{1}{N} \left[\sum_{j \neq k} \sum_{i \in \partial \mathcal{C}_j} \sum_{i' \in \partial \mathcal{C}_k} \left(\frac{\Sigma_{ii'}}{p^{m_{ii'}+1}} - \frac{2\Sigma_{ii'}}{p^2} \right) \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset) \right] \\
 &\quad + \frac{1}{N} \left[\sum_{j \neq k} \sum_{i \in \partial \mathcal{C}_j} \sum_{i' \in \partial \mathcal{C}_k} \frac{\Sigma_{ii'}}{p^{m_{ii'}+1}} \left(\mathbb{I}(m_{ii'} > 0) - \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset) \right) \right] \\
 &= I_1 + \underbrace{\frac{1}{N} \left[\sum_{j \neq k} \sum_{i \in \partial \mathcal{C}_j} \sum_{i' \in \partial \mathcal{C}_k} \left(\frac{\Sigma_{ii'}}{p^{m_{ii'}+1}} - \frac{2\Sigma_{ii'}}{p^2} \right) \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset) \right]}_{J_2:=} \\
 &\quad + \underbrace{\frac{1}{N} \left[\sum_{j \neq k} \sum_{i \in \partial \mathcal{C}_j} \sum_{i' \in \partial \mathcal{C}_k} \frac{\Sigma_{ii'}}{p^{m_{ii'}+1}} \left(\mathbb{I}(m_{ii'} > 0) - \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset) \right) \right]}_{J_3:=}.
 \end{aligned}$$

We denote $I_2 = J_1 + J_2 + J_3$. Then we have $\sigma_1^2 = SC + I_1 + I_2$. \square

C.2. Proof of Proposition 1

Proof. Because the correlation between non-neighboring regions is sufficiently small and the correlation between neighboring regions is positive, we can simplify our proof by setting $\Sigma_{ii'} = 0$ for any two non-neighboring regions. Under such condition, we have

$$\begin{aligned}
 \frac{N}{4} SC &= \sum_{ii'} \Sigma_{ii'} - \sum_{j \neq k} \sum_{i \in \mathcal{C}_j} \sum_{i' \in \mathcal{C}_k} \Sigma_{ii'} \\
 &= \sum_{ii'} \Sigma_{ii'} - \sum_{j \neq k} \sum_{i \in \partial \mathcal{C}_j} \sum_{i' \in \partial \mathcal{C}_k} \Sigma_{ii'} W_{ii'}, \\
 \frac{N}{4} I_1 &= 2 \sum_{j \neq k} \sum_{i \in \partial \mathcal{C}_j} \sum_{i' \in \partial \mathcal{C}_k} \Sigma_{ii'} W_{ii'}.
 \end{aligned}$$

Thus minimizing $SC + I_1$ is equivalent to minimizing

$$\sum_{j \neq k} \sum_{i \in \partial \mathcal{C}_j} \sum_{i' \in \partial \mathcal{C}_k} \Sigma_{ii'} W_{ii'},$$

which attains its minimum under the global design.

For the second-order interference I_2 , we analyze its three components separately to demonstrate that the global design is optimal:

- For J_1 , when the i, i' -th regions are within the same cluster \mathcal{C}_j , we have $m_{ii'} \geq 1$, and thus $\frac{\Sigma_{ii'}(1-p^{m_{ii'}-1})}{p^{1+m_{ii'}}} \geq 0$. Therefore, J_1 is minimized under the global design where there are no boundaries.

- For J_2 , we have

$$NJ_2 = \sum_{j \neq k} \sum_{i \in \partial \mathcal{C}_j} \sum_{i' \in \partial \mathcal{C}_k} \left(\frac{\Sigma_{ii'}}{p^{m_{ii'}+1}} - \frac{2\Sigma_{ii'}}{p^2} \right) \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset) \mathbb{I}(i, i' \text{ are two neighboring regions}).$$

Since $m_{ii'} \geq 2$ when i, i' are adjacent, together with the fact that $p = 0.5$, $\frac{\Sigma_{ii'}}{p^{m_{ii'}+1}} - \frac{2\Sigma_{ii'}}{p^2}$ must be a non-negative value. The results ensures a global design minimizes J_2 .

- For J_3 , because $\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset$ implies $m_{ii'} \geq 1$, then we have $\mathbb{I}(m_{ii'} > 0) \geq \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset)$, which ensures $\frac{\Sigma_{ii'}}{p^{m_{ii'}+1}} (\mathbb{I}(m_{ii'} > 0) - \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_j \neq \emptyset))$ is always non-negative. This fact implies J_3 is minimized under the global design.

Since $\text{SC} + I_1$ and I_2 are each minimized under the global design, their sum is also minimized under the global design. \square

C.3. Proof of Proposition 2

Proof. When SUTVA holds, the interference terms I_1 and I_2 equal zero. For SC, we have

$$\text{SC} = \frac{4}{N} \left[\sum_{ii'} \Sigma_{ii'} - \sum_{j \neq k} \sum_{i \in \mathcal{C}_j} \sum_{i' \in \mathcal{C}_k} \Sigma_{ii'} \right]. \quad (4)$$

Thus minimizing MSE is equivalent to maximizing $\sum_{j \neq k} \sum_{i \in \mathcal{C}_j} \sum_{i' \in \mathcal{C}_k} \Sigma_{ii'}$, which equals the summation over correlations between different clusters. Due to the non-negativeness of $\Sigma_{ii'}$, it is maximized when $m = R$, where each region belongs to a unique cluster, i.e., as in the individual design. \square

C.4. Proof of Proposition 3

Proof. As proved in Section C.1, we have

$$\begin{aligned} I_1 &= \frac{8}{N} \left[\sum_{i \in \mathcal{C}_1} \sum_{i' \in \partial \mathcal{C}_2} \Sigma_{ii'} \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_1 \neq \emptyset) + \sum_{i \in \mathcal{C}_2} \sum_{i' \in \mathcal{C}_1} \Sigma_{ii'} \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_2 \neq \emptyset) \right] \\ &\leq \frac{8}{N} \left[\sum_{i \in \mathcal{C}_1} \sum_{i' \in \mathcal{C}_2} \Sigma_{ii'}^+ \sum_{\ell \in \mathcal{C}_1} W_{\ell i'} + \sum_{i \in \mathcal{C}_2} \sum_{i' \in \mathcal{C}_1} \Sigma_{ii'}^+ \sum_{\ell \in \mathcal{C}_2} W_{\ell i'} \right] \\ &\leq \frac{8}{N} \left[\sum_{i \in \mathcal{C}_1} \sum_{i' \in \mathcal{C}_2} \sum_{\ell \in \mathcal{C}_1} \Sigma_{\ell i'}^+ W_{\ell i'} + \sum_{i \in \mathcal{C}_2} \sum_{i' \in \mathcal{C}_1} \sum_{\ell \in \mathcal{C}_2} \Sigma_{\ell i'}^+ W_{\ell i'} \right] \\ &= \frac{8}{N} \left[|\mathcal{C}_1| \sum_{i' \in \mathcal{C}_2} \sum_{\ell \in \mathcal{C}_1} \Sigma_{\ell i'}^+ W_{\ell i'} + |\mathcal{C}_2| \sum_{i' \in \mathcal{C}_1} \sum_{\ell \in \mathcal{C}_2} \Sigma_{\ell i'}^+ W_{\ell i'} \right], \end{aligned} \quad (5)$$

where the second inequality holds under Assumption 2. The last term (5) equals

$$\frac{8}{N} (|\mathcal{C}_1| + |\mathcal{C}_2|) \sum_{i' \in \mathcal{C}_1} \sum_{\ell \in \mathcal{C}_2} \Sigma_{\ell i'}^+ W_{\ell i'} = \frac{8R}{N} \sum_{i' \in \mathcal{C}_1} \sum_{\ell \in \mathcal{C}_2} \Sigma_{\ell i'}^+ W_{\ell i'},$$

which, together with Equation (4), gives us the objective function:

$$\frac{8}{N} \sum_{i \in \mathcal{C}_1} \sum_{i' \in \mathcal{C}_2} \left[R \Sigma_{ii'}^+ W_{ii'} - \Sigma_{ii'} \right].$$

\square

C.5. Proof of Proposition 4

Proof. As stated in Section C.3, we have

$$\begin{aligned}
 \frac{8R}{N} \sum_{i' \in \mathcal{C}_1} \sum_{\ell \in \mathcal{C}_2} \Sigma_{\ell i'}^+ W_{\ell i'} &= \frac{8}{N} \left[\sum_{i \in \mathcal{C}_1} \sum_{i' \in \mathcal{C}_2} \sum_{\ell \in \mathcal{C}_1} \Sigma_{\ell i'}^+ W_{\ell i'} + \sum_{i \in \mathcal{C}_2} \sum_{i' \in \mathcal{C}_1} \sum_{\ell \in \mathcal{C}_2} \Sigma_{\ell i'}^+ W_{\ell i'} \right] \\
 &= \frac{8}{N} \left[\sum_{i \in \mathcal{C}_1} \sum_{i' \in \partial \mathcal{C}_2} \sum_{\ell \in \mathcal{C}_1} \Sigma_{\ell i'}^+ W_{\ell i'} + \sum_{i \in \mathcal{C}_2} \sum_{i' \in \partial \mathcal{C}_1} \sum_{\ell \in \mathcal{C}_2} \Sigma_{\ell i'}^+ W_{\ell i'} \right] \\
 &\leq \frac{8}{N} \max_{i \neq i'} \Sigma_{ii'} \left[\sum_{i \in \mathcal{C}_1} \sum_{i' \in \partial \mathcal{C}_2} \sum_{\ell \in \mathcal{C}_1} W_{\ell i'} + \sum_{i \in \mathcal{C}_2} \sum_{i' \in \partial \mathcal{C}_1} \sum_{\ell \in \mathcal{C}_2} W_{\ell i'} \right] \\
 &\leq \frac{8}{N} d_{\max} \max_{i \neq i'} \Sigma_{ii'} \left[\sum_{i \in \mathcal{C}_1} \sum_{i' \in \partial \mathcal{C}_2} \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_1 \neq \emptyset) + \sum_{i \in \mathcal{C}_2} \sum_{i' \in \partial \mathcal{C}_1} \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_2 \neq \emptyset) \right] \\
 &\leq \frac{8}{N} d_{\max} \sigma \left[\sum_{i \in \mathcal{C}_1} \sum_{i' \in \partial \mathcal{C}_2} \Sigma_{ii'} \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_1 \neq \emptyset) + \sum_{i \in \mathcal{C}_2} \sum_{i' \in \partial \mathcal{C}_1} \Sigma_{ii'} \mathbb{I}(\mathcal{N}_{i'} \cap \mathcal{C}_2 \neq \emptyset) \right] \\
 &= d_{\max} \sigma I_1,
 \end{aligned}$$

where the second equation holds because $W_{\ell i'} = 0$ when i' falls in the interior region of its cluster, and the second inequality holds because each $i' \in \partial \mathcal{C}_k$ (for $k = 1, 2$) can have at most d_{\max} adjacent neighbors that fall in \mathcal{C}_j . \square

D. Experiments: details and additional results

D.1. Synthetic environments

Settings. The model for the outcome is represented as follows:

$$Y_{it} = 3O_{it} \sin(l_x + l_y + s \times (A_{it} + 0.5\bar{A}_{it})) + e_{it},$$

where $O_{it} \in \mathbb{R}$ represents the covariates, $\bar{A}_{it} = \frac{1}{|\mathcal{N}_i|-1} \sum_{j \neq i, j \in \mathcal{N}_i} A_{jt}$ represents the average of neighboring treatments, and s denotes the signal strength, which is set to 2.5% for our experiments. For each t , (e_{1t}, \dots, e_{Rt}) are independently drawn from a zero-mean multivariate Gaussian distribution with time-invariant covariance Σ . Under the spatial setting, Σ is set as one of the following correlation structures: (i) constant correlation, $\Sigma_{ij} = \rho^{\mathbb{I}(i \neq j)}$; (ii) truncated constant correlation, $\Sigma_{ij} = \mathbb{I}(i = j) + (\rho - R^{-1}|i - j|)\mathbb{I}(|i - j| \leq \rho R)$; (iii) exponential correlation, $\Sigma_{ij} = \rho^{|i - j|}$. After structuring the regions, the corresponding synthetic environment is created to simulate the conditions that allow us to generate datasets.

We evaluate each method using the relative MSE as criterion. Specifically, we use the relative MSE which is computed as $\mathcal{R}^{-1} \sum_{r=1}^{\mathcal{R}} (\widehat{\text{ATE}}_r - \text{ATE})^2 / (\text{ATE})^2$, where $\widehat{\text{ATE}}_r$ is the estimated ATE returned by one estimator at the r -th independent replication. We set the total number of replications to $\mathcal{R} = 50$.

Additional results. To compare MSE with the correlation parameter ρ , we also conduct simulations using the spatial rectangle structure in Figure 5(a). Results are shown in Figure 8, with the values of y -axis on a logarithmic scale.

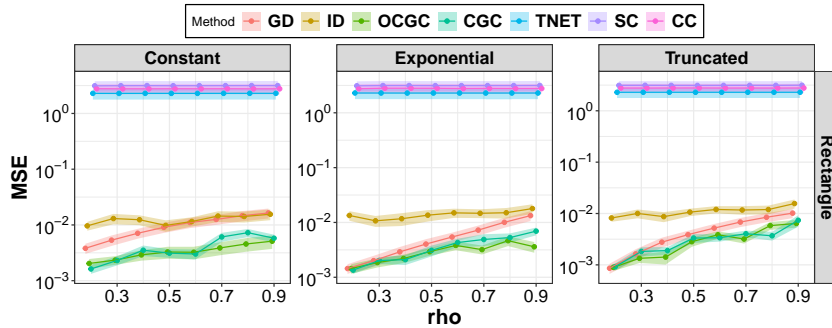


Figure 8. The plot of ρ versus MSE on rectangle spatial setting. Each panel corresponds to a spatial correlation setting.

D.2. Real-data simulator

The update process for the states of drivers and orders is outlined as follows:

1. Drivers who are assigned to orders will decide whether to accept them based on probabilities generated by a pre-trained LightGBM model. This step considers various characteristics of both the drivers and the orders.
2. Drivers who are idle and not currently assigned to any orders are directed to specific areas based on historical data of idle driver movements.
3. Drivers who are subject to repositioning must follow the directives provided by a pre-trained repositioning algorithm implemented by the ridesharing platform.
4. Once drivers accept orders, they proceed to the pickup locations, collect the passengers, and then travel to the specified destinations.
5. The pool of available drivers is continuously updated to reflect new drivers entering the service area and existing drivers who go offline.

As in the synthetic environment where we do not have an explicit expression on outcome and thus cannot obtain the oracle ATE value by mathematical calculation, we use a Monte Carlo (MC) procedure to approximate true ATE value. Specifically, we run the simulator for N independent days under actions 1 and 0, respectively, leading to the reward observations $\{r_{it}^{(1)}\}_{1 \leq i \leq R, 1 \leq t \leq N}$ and $\{r_{it}^{(0)}\}_{1 \leq i \leq R, 1 \leq t \leq N}$. With the two datasets, the MC estimator for ATE is given by:

$$\widehat{\text{ATE}}^{\text{MC}} = \frac{1}{N} \sum_{t=1}^N \sum_{i=1}^R [r_{it}^{(1)} - r_{it}^{(0)}].$$

When $N \rightarrow \infty$, the $\widehat{\text{ATE}}^{\text{MC}}$ converges to ATE. In our experiment, we set N to a large value and the obtained $\widehat{\text{ATE}}^{\text{MC}}$ is set as the surrogate of the true ATE. Specifically, N is set as 1000.

E. Implementation details

Algorithm 1 in the main text skips the estimation procedure of $\{g_i\}_{i=1}^R$ so as to simplify the illustration of the algorithm. Here, we elaborate the details on estimate it. Besides, to circumvent the need for imposing stringent metric entropy conditions on the outcome regression function (Dai et al., 2020), we adopt a recent advanced method that involves data-splitting and cross-fitting (Chernozhukov et al., 2018).

The idea of cross fitting is presented below. Given the repeated observations $\{(Y_{i,t}, A_{i,t}, O_{i,t})\}$, we split the data into K non-overlapped subsets of equal size. Then we estimate the regression function g_i based on the training data $\{(Y_{i,t}, A_{i,t}, O_t) : t \notin \mathcal{I}_k, 1 \leq i \leq R\}$, where \mathcal{I}_k denote the indices of the k th subset. We denote the estimated functions g_i and $\nu_{i,t}(\mathbf{a})$ as $\hat{g}_i^{(k)}$ and $\hat{\nu}_{i,t}^{(k)}(\mathbf{a})$, respectively. Under the cross-fitting framework, we denote the DR estimator as $\widehat{\text{ATE}}^{\text{DR-CF}}$. We summarized the complete algorithm for this in Algorithm 2.

Algorithm 2 Estimating outcome regression function at the l -th iteration.

Input: Datasets $\mathcal{D}^{(j)} = \{(Y_t, \mathbf{A}_t, \mathbf{O}_t)\}_{t=1}^B$ for $j = 1, \dots, l$, as collected in Algorithm 1.

- 1: Split dataset $\mathcal{D}^{(1)} \cup \dots \cup \mathcal{D}^{(l)}$ into K non-overlapped subsets of equal size. Let \mathcal{I}_k denote the set of the indices for the k th subset.
- 2: **for** $k = 1, \dots, K$ **do**
- 3: Compute the estimated outcome regression functions $\{\hat{g}_i^{(k)}\}_{i=1}^R$ based on the data tuples $\{(Y_{i,t}, A_{i,t}, O_{i,t})\} \mid t \notin \mathcal{I}_k, i \in \{1, \dots, R\}$ using random forest.
- 4: Compute $\hat{e}_{it} = Y_{i,t} - \hat{g}_i^{(k)}(A_{i,t}, O_{i,t})$ for all $i \in \{1, \dots, R\}$ and $t \in \mathcal{I}_k$.
- 5: **end for**
- 6: Compute the ATE estimator:

$$\widehat{\text{ATE}}_l^{\text{DR-CF}} = \frac{1}{lB} \sum_{k=1}^K \sum_{t \in \mathcal{I}_k} \sum_{i=1}^R [\hat{\nu}_{i,t}^{(k)}(\mathbf{1}) - \hat{\nu}_{i,t}^{(k)}(\mathbf{0})]$$

Output: $\{\hat{e}_{it}\}_{1 \leq i \leq R, 1 \leq t \leq lB}, \{\widehat{\text{ATE}}_l^{\text{DR-CF}}\}_{1 \leq t \leq lB}$.

We give more details on the Step 3 and Step 5 in Algorithm 2.

- In Step 3, instead of using the observed data at each region to fit regression model, we concatenate the observed data in all region, and augment the data by including the longitude and latitude as two additional features. The advantage of this procedure is that it allows for information sharing across regions, which ensures that even when repeat observations N is small, the estimated outcome regression model can still fit data well. The advantage of this procedure can be reflected by Figure 9, where the logarithmic scale on y -axis values is used for a clearer representation purpose. We can see that, when sample size is small, the procedures (denoted as OCGC, CGC) are much better than fitting g_i individually with dataset $\{(Y_{it}, A_{it}, O_{it})\}_{t=1}^N$ (denote as OCGC-Local, CGC-Local).
- In Step 6, we need to calculate the term $\mathbb{E}(T_{it}(\mathbf{a}))$ for estimation. Because the treatment are randomly assigned and independent to covariate, $\mathbb{E}(T_{it}(\mathbf{a}))$ has a close form expression that is written as:

$$\mathbb{E}[T_{it}(\mathbf{a})] = (p)^a(1-p)^{1-a} \prod_{k \neq j} \left[1 - \prod_{i \in \mathcal{N}_k} \mathbb{I}(i \notin \mathcal{C}_k) (p)^a(1-p)^{1-a} \right] + \prod_{i \in \mathcal{N}_l} \mathbb{I}(i \notin \mathcal{C}_k).$$

Using this expression can simplify the computing on $\widehat{\text{ATE}}$.

- In the algorithm, we use $\mathcal{D}^{(1)}, \dots, \mathcal{D}^{(l)}$ together to refit machine learning model, rather than using the dataset obtained in a single step (i.e., $\mathcal{D}^{(l)}$ generated in the l -th step). This step sufficiently leverages the collected data as reported in Figure 10. We use the logarithmic scale on values of y -axis for a clearer presentation. From Figure 10, leveraging all data (denoted as CGC) generally decreases the MSE when compared to using only one dataset $\mathcal{D}^{(l)}$ (denoted as CGC-ST).




Figure 9. The plot of number of repeated experiments versus MSE on square spatial setting. Each panel corresponds to a spatial correlation setting.




Figure 10. The plot of number of repeated experiments versus MSE on square spatial setting. Each panel corresponds to a spatial correlation setting.