

Grab the Reins of Crowds: Estimating the Effects of Crowd Movement Guidance Using Causal Inference

Koh Takeuchi
Kyoto University
takeuchi@i.kyoto-u.ac.jp

Hisashi Kashima
Kyoto University
kashima@i.kyoto-u.ac.jp

Ryo Nishida
Tohoku University, AIST
ryo.nishida@aist.go.jp

Masaki Onishi
AIST
onishi@ni.aist.go.jp

ABSTRACT

Crowd movement guidance has been a fascinating problem in various fields, such as easing traffic congestion in unusual events and evacuating people from an emergency-affected area. To grab the reins of crowds, there has been considerable demand for a decision support system that can answer a typical question: “*what will be the outcomes of each of the possible options in the current situation?*”. In this paper, we consider the problem of estimating the effects of crowd movement guidance from past data. To cope with limited amount of available data biased by past decision-makers, we leverage two recent techniques in deep representation learning for spatial data analysis and causal inference. We use a spatial convolutional operator to extract effective spatial features of crowds from a small amount of data and use balanced representation learning based on the integral probability metrics to mitigate the selection bias and missing counterfactual outcomes. To evaluate the performance on estimating the treatment effects of possible guidance, we use a multi-agent simulator to generate realistic data on evacuation scenarios in a crowded theater, since there are no available datasets recording outcomes of all possible crowd movement guidance. The results of three experiments demonstrate that our proposed method reduces the estimation error by at most 56% from state-of-the-art methods.

KEYWORDS

Deep Learning, Causal Inference, Multi-Agent Simulator

ACM Reference Format:

Koh Takeuchi, Ryo Nishida, Hisashi Kashima, and Masaki Onishi. 2021. Grab the Reins of Crowds: Estimating the Effects of Crowd Movement Guidance Using Causal Inference. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 9 pages.

1 INTRODUCTION

Crowd movement guidance has been a fascinating problem in various fields, such as easing traffic congestion in normal or abnormal events [25, 26, 33] and evacuating people from an emergency-affected areas such as a crowded building [7, 17, 32, 35, 43]. To grab the reins of crowds, given a situation, decision-makers need to devote extensive efforts to quickly select guidance that consists

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online. © 2021 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.




Figure 1: Example of crowd movement guidance: Evacuation events from an opera house. The green and red dots represent people who are moving or stopping, respectively.

of a complex compound of actions. For example, a traffic network manager needs to control traffic signals, road capacities, and directions of crowd flow together while considering the interactions among them. We show an example of crowd movement guidance in Figure 1. Let us assume a situation in which theater managers need to decide guidance required for a crowd evacuation, and there are two guide plans: *A* and *B*. If they choose plan *A*, congestion of crowds does not occur, and thus people can quickly escape from the theater. In contrast, plan *B* disturbs the crowd movements, which results in longer evacuation times than in plan *A*. The choice of plan has a significant impact on the outcome, and therefore, there has been considerable demand for a decision support system that can answer a typical question posed by decision-makers: “*what will be the outcomes of each of the possible options in the current situation?*”.

However, due to the limited number of such events or evacuation situations, there are not enough data records for estimating such effects with sufficient accuracy. To make matters worse, these records contain biases introduced by the plans taken by past decision-makers, because they try to make the best decisions for the situations at the time, and we cannot expect randomized control trials (RCTs) were performed to obtain unbiased data due to ethical or other reasons. In addition, due to the counterfactual nature of the data acquisition process, we cannot observe guidance outcomes other than those of the actual ones, and therefore, we cannot directly assess the effects of guidance from such datasets.




Figure 2: Framework of crowd movement guidance based on causal inferences.

In this paper, we focus on the problem of estimating the effects of crowd movement guidance from a relatively small amount of biased data. We propose a *spatial convolutional counterfactual regression* (SC-CFR), which takes advantage of two recent developments in deep representation learning in spatial data analysis and causal inference. Because the locational distribution of crowds has a significant impact on the effectiveness of movement guidance, SC-CFR aggregates crowd distribution information by means of the spatial convolutional layer [19, 41] and predicts the outcomes of movement guidance with high accuracy even with a limited number of data. On the other hand, to cancel the data biases made by past decision-makers, we resort to the causal inference framework that has been successfully applied to answer the what-if problems in various domains, such as healthcare [24], economics [13], and education [42]. In particular, SC-CFR employs balanced representation learning based on integral probability metrics, which mitigates selection bias and missing counterfactual outcomes [24].

Evaluating the performance of estimating the possible guidance’s treatment effects requires datasets recording outcomes of all possible crowd movement guidance. Since there are no datasets available for our purpose, we used a multi-agent simulator [36] to generate realistic data on evacuation scenarios in a crowded theater. From these datasets, we conducted experiments for estimating the treatment effects on outcomes, such as the maximum, average, and standard deviation of evacuation times required by crowds. We demonstrate that our proposed method reduced the estimation error by at most 56% from state-of-the-art baseline methods that do not consider spatial attributes.

Our contributions are summarized as follows:

- We first address the what-if problem in crowd movement guidance as causal effect estimation from a relatively small amount of biased data.
- We propose SC-CFR, which takes advantage of two recent developments in deep representation learning in spatial data analysis and causal inference.
- We demonstrate that our method reduces the estimation error by at most 56% from state-of-the-art baselines by experiments with a multi-agent simulator.

2 PRELIMINARIES

We use an uppercase letter to denote a random variable (i.e., X), a lowercase letter to indicate a realization x drawn from a particular distribution (i.e., $x \sim P(X)$), and a calligraphy letter \mathcal{X} to represent a sample space (i.e., $x \in \mathcal{X}$).

2.1 Crowd Movement Guidance

Let us consider a place or area with crowds of freely moving people. When the place is crowded, a manager should want to carry out a crowd guidance plan to reduce the congestion caused by crowds moving toward other places. For example, if a park hosts a big festival with a sell-out crowd, a park manager has to make a decision on guidance to avoid traffic jams. Another example is of a manager working in a stadium who needs to make a decision on guidance to safely and quickly evacuate crowds during a fire emergency. A crowd guidance plan can comprise various types of actions, such as route recommendations for crowds or a modification of road width, which could have complicated interactions.

Assume that a situation of crowds, such as their current locations, is observed as covariates. We denote a situation of crowds (covariates) as $X \in \mathbb{R}^d$, where d is the number of people. A guide (treatment) is $Z \in \mathcal{Z} \subset \{0, 1\}^t$ that consists of t possible actions, where \mathcal{Z} is a set of possible guides. A guide can select single or multiple actions (i.e., a combination of actions); therefore, $|\mathcal{Z}|$ grows exponentially as t increases. We denote a potential outcome as $Y \in \mathbb{R}$, such as an average moving cost or the maximum movement cost consumed by crowds. For example, let us suppose an evacuation scenario from a place. Given an initial situation of crowds $x \sim P(X)$, a decision-maker, who needs to guide crowds to a safe place, selects a treatment $z \sim P(Z|X)$, after which crowds start to move. We can observe an outcome $y \sim P(Y|X, Z)$, such as evacuation times, after finishing the scenario. This procedure is built upon by the Rubin-Neyman potential outcomes framework [22]. We show a framework for crowd movement guidance in Figure 2.

2.2 Causal Effect Estimation

Suppose that there is a healthcare record (covariate) of a patient, and a doctor needs to choose whether to prescribe medicine, which we will call a treatment. We can observe a future health condition of the patient as a subsequent outcome. Causal effect estimation is a task for estimating outcomes given not actual treatments and to assess the difference in outcomes between treatments. In this paper, we consider the situation of crowds as covariates, guidance as a treatment, and costs of crowd movements as outcomes. Effects of a guide should differ from situation to situation since it depends on the crowd conditions and types of events. Thus, decision-making on guide planning is one of the most difficult tasks for managers, and they need a useful tool to help their decision-making. Surprisingly, the effect of guidance has not been considered in the existing literature.

We suppose that two arbitrary guide plans A and B correspond to guides z and z' , respectively. Given covariates x and a guide z , we denote the cost consumed by crowds as $Y_x(z)$. The conditional average treatment effect (CATE) for covariates x can be defined as $\mathbb{E}[Y_x(z)] - \mathbb{E}[Y_x(z')]$, where we use \mathbb{E} to denote the expectation over Y . Intuitively, this quantity indicates how much of an advantage (or disadvantage) plan A has in guiding the crowd movements compared with plan B . Knowing this quantity helps a decision-maker compare the difference between treatments for a particular situation x ; however, we cannot directly calculate CATE from observed outcomes because we can only observe the outcome of one of the treatments.

2.3 Selection Bias and Missing Counterfactuals

Estimating the outcome of guidance requires a deep understanding of the complex relationships within covariates and guidance. One of the difficulties in estimating treatment effects on crowd movement is that an observed set of guidance is biased in the given dataset, because the manager usually selects guidance according to their policy. Thus, even if the observed covariates are random, we cannot observe the outcome with a pair of covariates and guidance that managers do not prefer. Another obstacle is that we can only observe a factual outcome, but cannot observe counterfactual outcomes. Usually, we can only deploy a single guide in a situation, and thus the outcomes for the other possible guides remain missing.

Formally, in an observed study, we can obtain only one factual outcome y^f depending on a selected treatment z , and thus never see counterfactual outcomes y^{cf} corresponding to the other treatments that are not selected by the decision-maker. More formally, when a decision-maker chooses a treatment z , a factual outcome y^f is equal to the outcome $y_x(z)$, and the counterfactual outcomes y^{cf} are $y_x(z')$, $\forall z' \in \mathcal{Z} \setminus z$. We suppose that a decision-maker has a policy for choosing a treatment z based on a situation x , and thus, observed treatments are not at random and always biased. This bias, which is called a selection bias in the Rubin-Neyman framework, leads to a simple model for estimating outcomes, not accurate and biased results. Thus, we need to develop a method that learns a model for predicting outcomes and can simultaneously cancel the effects of the selection bias.

In this paper, we make the common assumption called *strong ignorability* in the Rubin-Neyman framework.

Assumption 1 (Stable Unit Treatment Value Assumption). *The potential outcomes for any situations do not vary with the treatment assigned to other situations, and, for each situation, there are no different forms or versions of each treatment level, which lead to different potential outcomes.*

This assumption emphasizes the independence of each situation and that there are no interactions between situations.

Assumption 2 (Ignorability). *Given the covariate, X , treatment assignment Z is independent to the potential outcomes, i.e., $Z \perp Y_X(z) \mid X, \forall z \in \mathcal{Z}$.*

The ignorability assumption is also referred to the unconfoundedness assumption. With this unconfoundedness assumption, the situations with the same covariates X , their treatment assignment can be viewed as random.

Assumption 3 (Positivity). *For any value of X , treatment assignment is not deterministic: $P(Z = z \mid X = x) > 0, \forall z$ and x .*

This assumption implies that all the factors determining the outcome of each treatment are observed. From the above assumptions, we follow strong ignorability assumption that is referred to the no-hidden confounders assumption. We formalize this assumption by using the standard strong ignorability condition:

Assumption 4 (Strong Ignorability). *$Z \perp Y_X(z) \mid X, \forall z \in [0, 1]^t$ and $P(Z = z \mid X = x) > 0, \forall z$ and x .*

The strong ignorability assumption is a sufficient condition for CATE function f to be identifiable [13].




Figure 3: Model architecture of spatial convolutional counterfactual regression (SC-CFR).

3 SPATIAL CONVOLUTIONAL COUNTERFACTUAL REGRESSION

To predict what the outcomes of each of the possible options in the current situation will be, we develop a model $y = f(x, z)$ that predicts the outcome for any treatments z . Our idea is twofold: a spatial convolutional operator extracting effective spatial features of crowds from a small amount of data, and balanced representation learning to mitigate the selection bias and missing counterfactuals.

We assume that the biased observed dataset $\mathcal{D} = \{x_n, z_n, y_n^f\}_{n=1}^N$ is available for learning our model, where N is the number of situations. Note that, our dataset contains only factual outcomes for assigned treatments; thus, we do not have access to the unobserved counterfactual outcomes y^{cf} for unassigned treatments.

3.1 Model Architecture

We first give an overview of our proposed model called spatial convolutional counterfactual regression (SC-CFR) shown in Figure 3. SC-CFR takes two inputs: covariates x and a treatment z , and outputs the outcome estimate y . SC-CFR first extracts effective spatial features Φ_x from covariates x with the spatial feature extractor g . Then, a multi-layer perceptron h makes a prediction $y = h(\Phi_x, z) = h(g(x), z)$.

In the training phase of SC-CFR, to reduce the biases made by past decision-makers, we utilize the IPM regularizer that makes the distributions of the extracted representations $\Phi_x = g(x)$ as similar as possible between different treatments.

3.2 Spatial Feature Extraction

Now, we explain the details of learning spatial features Φ_x . We utilize spatial information, such as GPS locations of crowds for extracting spatial features. We split spatial attributes vertically and horizontally to make a spatial discrete grid. Then, from covariates x , we construct a two-dimensional array $a \in \mathbb{R}^{d_1 \times d_2}$ filled with values such as the aggregated numbers of people in each block.

We employ a spatial convolution operator and an average pooling operator to extract spatial features Φ_x from the two-dimensional array a . These operators have been widely adopted in spatial data analysis, and have achieved significant improvements in various applications such as traffic forecasting [18, 19, 41]. These operators enable us to extract spatial distributions of the crowds by aggregating information of nearby grids, and the hierarchical structure of the operators captures both local and global interactions of the crowds. Our two-dimensional spatial convolution operation [8]

takes the inner product or measures the cross-correlation between an input array \mathbf{a} and a kernel $\mathbf{k} \in \mathbb{R}^{k_1 \times k_2}$, defined as

$$(\mathbf{k} * \mathbf{a})(i, j) = \sum_{m=1}^{k_1} \sum_{n=1}^{k_2} \mathbf{a}_{i+m, j+n} \mathbf{k}_{m,n}. \quad (1)$$

The average pooling operator is the convolution operator whose kernel values are set to 1. We employ two repetitions of a convolution operator and an average pooling operator as $g(\mathbf{x})$ in this paper (See Figure 3).

3.3 Balanced Representation

If we do nothing to address the observation biases in data, the distribution of the extracted spatial feature $\Phi_{\mathbf{x}}$ is biased depending on \mathbf{z} , which could made the prediction accuracy of the model deteriorate. To mitigate the bias and make $\Phi_{\mathbf{x}}$ a balanced representation, we utilize the integral probability metrics (IPM) [20, 27] in this paper. Let us consider a binary treatment case $\mathcal{Z} = \{0, 1\}$: IPM can measure distance between two distributions. Given two probability density functions $P(\Phi_{\mathbf{x}}|z_i)$ and $P(\Phi_{\mathbf{x}}|z_j)$, for $\mathcal{S} \subset \mathbb{R}^d$ and G , a family of functions from \mathcal{S} to \mathbb{R} , IPM is defined as

$$\text{IPM}_G(p, q) = \sup_{g \in G} \left| \int_{\mathcal{S}} g(s) (p(s) - q(s)) ds \right|. \quad (2)$$

Intuitively, minimizing IPM makes the distributions of the extracted spatial features for the factual and counterfactual treatments hard to distinguish, and thus the distribution of the training data becomes close to those of RCTs [10, 24].

In this paper, we consider a case where a guiding plan $\mathbf{z} \in \mathcal{Z} = \{0, 1\}^t$ consists of multiple actions, such as route recommendations for crowds and a modification of the road width. However, the IPM compares only two distributions, and therefore, is not directly applicable to multiple treatments [39]. One of the most natural extensions of IPM to multiple treatments is to measure IPM between all possible pairs of treatments; however, this requires $|\mathcal{Z}|^2$ IPM terms and is not tractable for a large number of treatments. To cope with this problem, we utilize another simple extension of IPM that considers the most frequent treatment in a mini-batch as the control treatment and only use the IPMs between the control and the others.

$$\sum_{z' \in \mathcal{Z} \setminus z} \text{IPM}_G(p(\Phi_{\mathbf{x}}|z), q(\Phi_{\mathbf{x}}|z')). \quad (3)$$

3.4 Outcome Prediction

Finally, we utilize the spatial representations balanced for all the treatments to construct a predictive model $h(\Phi_{\mathbf{x}}, \mathbf{z})$ to infer an outcome \hat{y} . The original counterfactual regression [24] and its variants [23] employ multi-head networks where each multi-layer network corresponds to a treatment z_i and shares common base layers to avoid the risk of losing treatment information. However, since the total number of treatments $|\mathcal{Z}|$ rapidly grows with the number of actions, and training all of the heads requires a considerable amount of observations that are not available in our setting, we employ a single-head network instead built on a multi-layer perceptron (MLP) that takes a concatenation of $\Phi_{\mathbf{x}}$ and \mathbf{z} as the input.

$$h(\Phi_{\mathbf{x}}, \mathbf{z}) = \text{MLP}(\Phi_{\mathbf{x}}, \mathbf{z}). \quad (4)$$

The single-head network architecture has showed better performance when the number of treatments is more than a dozen [30].

To train the SC-CFR model, we find the optimal parameter that minimizes the loss function.

$$\sum_{i=1}^N \|y_n - \text{MLP}(\Phi_{\mathbf{x}_n}, \mathbf{z}_n)\|_2^2 + \lambda \sum_{z' \in \mathcal{Z} \setminus z} \text{IPM}_G(p(\Phi_{\mathbf{x}}|z_n), q(\Phi_{\mathbf{x}}|z'_n)), \quad (5)$$

where $\lambda > 0$ is a hyper parameter for controlling the effect of IPM. We employ the ADAM optimiser to optimize parameters in a mini-batch manner [16]. Among several candidates for the IPM, we employ the maximum mean discrepancy (MMD) [9, 27] for its simplicity in this paper.

4 RELATED WORK

Modeling and simulating large-scale crowd movements [3], including traffic network flows in cities [25, 33] and at sea [26], and evacuation from buildings [7, 17, 32, 43] have been actively studied in the field of multi-agent systems. Researchers have exploited multi-agent simulators to find optimal decision-making policies for crowd management. Comparisons of the properties of various types of simulators are found in [21]. Most of the existing research focuses on bottom-up decision-making by individual crowds, i.e., each agent making decisions based on individual or partial information to achieve one's goals. This study focuses on the indirect control of crowd movement through top-down decision-making by managers and the estimation of its effects.

A convolutional operator is a feature extraction method developed in computer vision for capturing low and high level image features [8]. Researchers in spatial data analysis have recently applied convolutional operators to handle the spatial correlations in urban areas [19]. They have extracted local and city-wide global features by stacking spatial convolutional operators and utilized it to solve their domain-specific problems [18]. With such rich features, many researchers have reported significant improvements in the predictive accuracy on forecasting future amounts and speeds of transportation networks [40, 41] and rides of on-demand ride services [15], to name a few. Based on their results, we utilize a spatial convolutional operator for extracting effective local spatial features of crowds from a small amount of data.

In the field of causal inference, novel methods for estimating the causal effect on individual and population levels have been proposed [12, 37]. The bayesian additive regression tree (BART) [11] is one of the most popular methods built on a simple decision tree model [2, 34] and it shows state-of-the-art performance. Recently, in the field of machine learning, novel methods based on deep neural networks to remove biases in data introduced by past decision-makers have been proposed [1, 5, 14, 24, 38]; they use representation learning techniques to learn balanced representations of confounders. In particular, the counterfactual regression (CFR) [24] and its various extensions [10, 23, 39] have received widespread interest due to its significant performance improvement in estimating treatment effects compared to existing methods, such as BART. However, most of the existing studies have not considered the use of spatial information to address data deficiencies and observation biases. In this paper, we use spatial convolution to extract effective spatial features for estimating treatment effects from the



Figure 4: An evacuation.




Figure 5: Route guidance.



Figure 6: A jam of crowds.




Figure 7: Building structures, seat layouts, and actions; a route guide for people (blue) and doors at six exits (red). Four blocks A, B, C, and D for data generation.

distribution of crowds given as images [19, 41]. The problem of estimating treatment effects is closely related to bandit feedback [4, 28] and off-policy reinforce learning [29]. Some of them assume that the policies of the agents used to collect data are known [10, 12]. Meanwhile, we do not assume prior information about the data collection mechanisms.

5 EXPERIMENTS

5.1 Experimental Setup

Evaluating the performance of estimating possible guidance treatment effects requires datasets recording the outcomes of all possible crowd movement guidance. Since such datasets are not available, we used an open-source multi-agent simulator [36] to generate realistic data on evacuation situations in a crowded theater. We utilized the actual structure of an opera house, namely New National Theatre, Tokyo in our simulator, whose seating capacity is $d = 868$. We show the building structures in Figure 7, in which the yellow blocks correspond to the seats. The blue lines and red texts represent actions of treatments explained, as in the following paragraph. Crowds evacuate from the auditorium to the outside using one of the six exits. We conducted a real evacuation drill at the opera house and measured the evacuation movements of 570 people using eight stereo cameras. Then, we employed the simulator [36] to generate crowd movement. We used the default setting parameters and set agents at the same initial locations as in the drill. We confirmed that the absolute error of the maximum evacuation time between measured and simulated data was less

than five percent. We show images of an evacuation drill held in the opera house in Figures 4, 5, and 6.

In our experiments, we considered two types of actions. The first type of action was *route guidance* to guide the audience to the evacuation exits. If no route guide is placed, each audience evacuates from the nearest exit. We employed a route guide whose maximum evacuation time was the smallest than other guides when all seats were occupied and all doors were open. The second type of the actions was *door operations* that determine whether the door of an exit to be fully open or half-open. In our scenario, we considered a risk, such as spreading fires, and allowed only two of six doors to be fully open. We show snapshots of simulations having the same covariates but with and without route guidance in Figures 9 and 8, where the green and red dots represent agents that are moving and stopping, respectively. We can observe the positive effect of route guidance by comparing the snapshots at 160 sec; without route guidance, congestion occurred (Figure 8), while it is dissolved if we give route guidance (Figure 9). Our codes and datasets are publicly available at a repository ¹.

5.2 Data Generations

We generated crowd movement datasets using our simulator with the following settings. We first split the auditorium into four blocks (Blocks A, B, C, and D in Figure 7) and set the occupancy rate of each block from $\{0.1, 0.5, 0.9\}$. We randomly set the occupation of each seat with the occupancy rate and constructed covariates x by setting x_i to 1 if the i -th seat was occupied and set to 0 otherwise. We ran the seat selection ten times with different random seeds. Then, we put agents on the occupied seats and ran evacuation simulations for all possible treatments $z \in \mathcal{Z}$, where $|\mathcal{Z}| = 2 * (6 * 5) / 2 = 30$. We omitted simulations where the total number of people was less than 400, since the guidance was not effective in such situations. The total number of data N was 423 and each data has 30 possible treatments and their corresponding outcomes. We observed three outcomes for each run: the maximum, average, and standard deviation of evacuation times required by each person. Then, we added observation noises to the outcomes that were randomly sampled from a normal distribution $\mathcal{N}(0, \sigma^2)$ where $\sigma = 2$. We show the distributions of the outcomes in Figures 10, 11, and 12. The red and blue lines correspond to the outcomes with and without route guidance, respectively.

To select a factual treatment z for guidance, we employed two independent probabilistic policies (the propensity scores). We set the first dimension of treatment $z \in \{0, 1\}^7$ to represent the action

¹<https://github.com/koh-t/SC-CFR>




Figure 8: Snapshots of a multi-agent simulation without route guidance.




Figure 9: Snapshots of a multi-agent simulation with route guidance.

of the route guide, and others to indicate the actions of six doors. We randomly determined whether to use the route guide using a distribution $p(z_{\text{guide}} = 1) = 1/(1 + \exp(-\sum_{i=1}^d x_i/d + 1))$. We set z_{guide} to 1 when the route guidance is given. To decide which exit doors to be fully open, we calculated the populations of occupied seats within 8 meters $\sum_{i \in \mathcal{D}_j} x_i/|\mathcal{D}_j|$ from each door, where \mathcal{D}_j is the number of the seats within 8 meters from the j -th door. Finally, we randomly sampled two doors to be open by using a multinomial distribution whose parameters were set to those populations. We treated all the other treatments as counterfactual treatments.

5.3 Model Settings

We randomly sampled 90% of the seat assignments as a training dataset and used the rest as a test dataset. As the baseline methods, we employed nine supervised learning methods [6]: Lasso [31], the ridge regression (Ridge), the support vector regression (SVR), the random forest (RF), the gradient boosting regressor (GBR), the bayesian additive regression tree (BART) [11], and the multi-layer perceptron (MLP) consisting of five layers. Since these methods were not applicable to identify differences between covariates and a treatment, we simply utilized a concatenation $[\mathbf{x}, \mathbf{z}]$ as an input. The treatment-agnostic representation network (TARNET) and the counterfactual regression (CFR) are recently proposed neural network-based causal effect estimation methods in machine learning fields [24]. For TARNET and CFR, we used a two-layer MLP to learn the representation of \mathbf{x} and a three-layer MLP whose inputs were a concatenation of the extracted representation and treatments to predict the outcome. We employed two variants of our proposed methods. SC-CFR is our proposed method using IPM, and a spatial convolutional treatment-agnostic representation network (SC-TARNET) is the proposed method without IPM ($\lambda = 0$). We constructed two-dimensional arrays $\mathbf{a} \in \mathbb{R}^{22 \times 42}$ from covariates based on seat layouts. We show examples of the arrays in Figure 13. For our methods, we utilized two repetitions of a block that consisted of a convolutional operator and an average pooling operator, whose kernel size k and paddings were set to 3 to extract the

spatial feature $\Phi_{\mathbf{x}}$. Then, we used a three-layer MLP to predict the outcome. We selected model architectures and hyperparameters by minimizing the training errors.

5.4 Metrics

To compare the performance of models, we employed several metrics that could measure the accuracy of predicted outcomes \hat{y} under the potential outcome framework. We evaluated models in two different settings; the *within-sample* setting, where the task is to estimate outcomes for all scenarios in the training dataset, evaluates what happens if a decision-maker chose other treatments in the past evacuation events. On the other hand, the *out-of-sample* setting, where the task is to estimate outcomes for scenarios in the test dataset, corresponds to the case of a new evacuation event, and the goal is to estimate the outcomes of all possible treatments.

Root mean squared error (RMSE). RMSE is a widely used metric in supervised learning problems. We employed RMSE to measure the accuracy of each predicted outcome depending on the covariates and all possible treatments. We define ϵ_{RMSE} as

$$\left[\frac{1}{|\mathcal{Z}'|} \frac{1}{N} \sum_{z_i \in \mathcal{Z}'} \sum_{n=1}^N (\hat{y}_{\mathbf{x}_n}(z_i) - \mathbb{E}[Y_{\mathbf{x}_n}(z_i)])^2 \right]^{1/2}, \quad (6)$$

where \mathcal{Z}' is a set of whole pairs of possible guiding plans.

Precision in estimation of heterogeneous effect (PEHE). In the binary setting, PEHE measures the ability of a predictive model to estimate the difference between true and predicted outcomes between two treatments z_0 and z_1 for covariates \mathbf{x} . We define PEHE as the RMSE between the true difference of the outcomes of two treatments z_i and z_j , that is, $\mathbb{E}[Y_{\mathbf{x}_n}(z_i)] - \mathbb{E}[Y_{\mathbf{x}_n}(z_j)]$, and the predicted difference, $\hat{y}_{\mathbf{x}_n}(z_i) - \hat{y}_{\mathbf{x}_n}(z_j)$. Obtaining a better PEHE requires accurate estimations of both the factual and counterfactual outcomes. We set ϵ_{PEHE} as

$$\left[\frac{1}{N} \sum_{n=1}^N (\hat{y}_{\mathbf{x}_n}(z_i) - \hat{y}_{\mathbf{x}_n}(z_j) - (\mathbb{E}[Y_{\mathbf{x}_n}(z_i)] - \mathbb{E}[Y_{\mathbf{x}_n}(z_j)]))^2 \right]^{1/2}. \quad (7)$$




Figure 10: Maximum evacuation times.




Figure 11: Average evacuation times.




Figure 12: Standard deviation of the evacuation times.




Figure 13: Examples of two-dimensional covariate arrays.

Average treatment Effect (ATE). We employed the error on estimating a conditional ATE. In contrast to PEHE, ATE considers only the average difference of treatment effects over the population. ATE is not as important as PEHE for models optimised for CATE estimation, but it can be a useful indicator of how well an estimator performs at comparing two treatments across the entire population. We define ϵ_{ATE} as

$$\frac{1}{N} \sum_{n=1}^N (\hat{y}_{x_n}(z_i) - \hat{y}_{x_n}(z_j)) - \frac{1}{N} \sum_{n=1}^N (\mathbb{E}[Y_{x_n}(z_i)] - \mathbb{E}[Y_{x_n}(z_j)]) \quad (8)$$

Multiple Treatments. Let us consider a general case where combinations of multiple actions are available for making a guiding plan. By following [23], we extended PEHE and ATE to be applicable to multiple treatments by considering the average PEHE and ATE between every possible pair of guiding plans defined as

$$\epsilon_{mPEHE} = \frac{1}{\binom{t}{2}} \sum_{(z_i, z_j) \in \mathcal{Z}} \epsilon_{PEHE}, \quad \epsilon_{mATE} = \frac{1}{\binom{t}{2}} \sum_{(z_i, z_j) \in \mathcal{Z}} \epsilon_{ATE}. \quad (9)$$

5.5 Results

Tables 1, 2, and 3 show the results of prediction of the maximum, average, and standard deviation of evacuation times, respectively. Each column represents the average and the standard deviation of each metric. A bold letter indicates a method that achieved the smallest (best) error in each column. To assess the effect on bias cancellation, we employed the t-test on two related samples of the scores between TARNET and CFR, and SC-TARNET and SC-CFR. The values with a single star or double stars indicate that

the p-values between the two models were $p < 0.10$ and $p < 0.05$, respectively.

From Tables 1, 2, and 3, we confirmed that our proposed SC-CFR achieved the best performance on ϵ_{RMSE} and ϵ_{mPEHE} . SC-TARNET came in the second place. GBR or BART placed third. In Table 1, SC-CFR reduced the estimation error of ϵ_{RMSE} for within-sample and out-of-sample by 33% and 38% from GBR, and 55% and 56% from Ridge, respectively.

The improvements of SC-CFR on ϵ_{RMSE} and ϵ_{mPEHE} for both within-sample and out-of-sample settings indicated that our method successfully extracted spatial features and learned balanced representations, thus providing more accurate predictions of outcomes than the state-of-the-art baselines. Our methods achieved the best estimation results of ϵ_{mATE} in Tables 1 and 3, while BART showed the best performance in Table 2. Since the distribution of the standard deviations had a relatively simpler shape than the others, BART, which is a simpler non-linear method than ours, was sufficient to estimate ATE.

By comparing MLP and TARNET, TARNET showed better scores than MLP because of its representation learning mechanism, which MLP does not have. The performance of SC-CFR was significantly better than that of CFR. With these results, we could argue that the spatial convolutional operators worked correctly and greatly contributed to this improvement. As regarding the effect of bias cancellation, we observed that IPM contributed to the performance improvements of CFR from TARNET and SC-CFR from SC-TARNET. This result indicates that a balanced representation was effective in precisely estimating CATE of crowd movement guidance, which is consistent with existing literature [24].

6 CONCLUSION

In this paper, we focused on the problem of estimating the effects of crowd movement guidance from a relatively small amount of biased data to answer a typical question posed by decision-makers: “*what will be the outcomes of each of the possible options in the current situation?*”. We proposed a Spatial Convolutional Counterfactual Regression (SC-CFR), which takes advantage of two recent developments in deep representation learning in spatial data analysis and causal inference. We utilized spatial convolutional operators to extract effective spatial features of crowds from a small amount of data, and applied balanced representation learning using IPM to mitigate the selection bias and missing counterfactual outcomes. With

Table 1: Results on estimating the maximum evacuation time.

Method	Within-sample			Out-of-sample		
	ϵ_{RMSE}	ϵ_{mPEHE}	ϵ_{mATE}	ϵ_{RMSE}	ϵ_{mPEHE}	ϵ_{mATE}
Lasso	19.047 _(0.452)	17.334 _(0.282)	5.604 _(0.327)	19.290 _(0.777)	18.524 _(1.216)	5.947 _(0.674)
Ridge	18.396 _(0.317)	16.886 _(0.576)	7.216 _(0.642)	18.575 _(0.841)	24.369 _(1.962)	7.276 _(0.841)
SVR	19.125 _(0.159)	14.895 _(0.179)	5.534 _(0.138)	19.365 _(0.892)	15.519 _(0.714)	5.860 _(0.469)
RF	16.117 _(0.613)	13.595 _(0.444)	4.952 _(0.362)	16.406 _(0.906)	14.824 _(0.779)	5.234 _(0.667)
GBR	13.933 _(0.577)	11.617 _(0.503)	3.492 _(0.675)	14.290 _(1.221)	12.841 _(0.890)	3.966 _(1.079)
BART	14.429 _(1.601)	11.620 _(1.299)	4.541 _(1.324)	14.796 _(1.840)	12.508 _(1.375)	4.863 _(1.559)
MLP	19.125 _(0.286)	17.911 _(0.310)	5.381 _(0.235)	19.342 _(0.781)	21.243 _(1.628)	5.656 _(0.511)
TARNET	15.115 _(0.281)	12.992 _(0.227)	3.782 _(0.576)	15.433 _(0.797)	13.981 _(0.863)	4.147 _(0.923)
CFR	*14.399 _(0.713)	**12.660 _(0.576)	3.658 _(0.701)	*14.716 _(1.259)	**13.511 _(1.052)	4.025 _(1.057)
SC-TARNET	9.767 _(1.029)	7.748 _(0.746)	3.017 _(0.506)	9.932 _(0.938)	7.920 _(0.595)	3.288 _(0.738)
SC-CFR	*9.636 _(1.032)	**7.658 _(0.787)	*2.933 _(0.512)	*9.802 _(0.902)	**7.816 _(0.579)	**3.209 _(0.761)

Table 2: Results on estimating the average evacuation time.

Method	Within-sample			Out-of-sample		
	ϵ_{RMSE}	ϵ_{mPEHE}	ϵ_{mATE}	ϵ_{RMSE}	ϵ_{mPEHE}	ϵ_{mATE}
Lasso	3.717 _(0.075)	6.497 _(0.167)	2.133 _(0.106)	3.743 _(0.155)	6.889 _(0.604)	2.168 _(0.140)
Ridge	2.726 _(0.108)	3.075 _(0.096)	0.758 _(0.236)	2.708 _(0.176)	5.034 _(0.448)	0.720 _(0.257)
SVR	2.747 _(0.074)	2.860 _(0.093)	1.188 _(0.098)	2.765 _(0.161)	2.859 _(0.166)	1.185 _(0.165)
RF	3.293 _(0.148)	3.185 _(0.105)	1.819 _(0.118)	3.321 _(0.178)	3.627 _(0.203)	1.848 _(0.162)
GBR	2.255 _(0.172)	2.336 _(0.145)	0.810 _(0.163)	2.277 _(0.161)	2.793 _(0.209)	0.827 _(0.155)
BART	2.296 _(0.128)	2.298 _(0.133)	0.507 _(0.076)	2.328 _(0.160)	2.503 _(0.215)	0.523 _(0.074)
MLP	3.323 _(0.098)	4.132 _(0.078)	1.580 _(0.112)	3.361 _(0.139)	5.688 _(0.543)	1.611 _(0.171)
TARNET	2.470 _(0.042)	2.589 _(0.077)	0.761 _(0.159)	2.494 _(0.145)	3.030 _(0.301)	0.811 _(0.146)
CFR	**2.337 _(0.065)	**2.537 _(0.074)	**0.702 _(0.110)	2.357 _(0.148)	**3.014 _(0.329)	0.731 _(0.115)
SC-TARNET	1.947 _(0.221)	1.849 _(0.240)	0.569 _(0.105)	1.976 _(0.224)	1.929 _(0.252)	0.602 _(0.106)
SC-CFR	1.922 _(0.250)	1.816 _(0.260)	0.542 _(0.106)	1.955 _(0.208)	1.869 _(0.237)	0.572 _(0.085)

Table 3: Results on estimating the standard deviation of evacuation time.

Method	Within-sample			Out-of-sample		
	ϵ_{RMSE}	ϵ_{mPEHE}	ϵ_{mATE}	ϵ_{RMSE}	ϵ_{mPEHE}	ϵ_{mATE}
Lasso	4.957 _(0.103)	6.538 _(0.130)	2.861 _(0.108)	5.012 _(0.294)	6.699 _(0.472)	2.950 _(0.231)
Ridge	3.403 _(0.038)	3.571 _(0.109)	0.852 _(0.099)	3.426 _(0.198)	5.487 _(0.384)	0.941 _(0.208)
SVR	3.658 _(0.087)	3.455 _(0.100)	1.687 _(0.099)	3.694 _(0.253)	3.420 _(0.139)	1.730 _(0.225)
RF	3.730 _(0.258)	3.412 _(0.237)	1.885 _(0.231)	3.756 _(0.197)	3.649 _(0.250)	1.905 _(0.257)
GBR	2.693 _(0.147)	2.565 _(0.164)	0.790 _(0.176)	2.720 _(0.117)	2.832 _(0.183)	0.898 _(0.161)
BART	2.732 _(0.180)	2.567 _(0.098)	0.537 _(0.053)	2.764 _(0.168)	2.772 _(0.292)	0.658 _(0.101)
MLP	4.488 _(0.113)	4.427 _(0.127)	2.353 _(0.114)	4.517 _(0.220)	5.157 _(0.314)	2.395 _(0.175)
TARNET	3.053 _(0.150)	2.961 _(0.144)	0.786 _(0.116)	3.087 _(0.200)	3.280 _(0.350)	0.888 _(0.201)
CFR	**2.772 _(0.215)	**2.866 _(0.124)	*0.734 _(0.092)	2.834 _(0.201)	**3.131 _(0.316)	0.838 _(0.174)
SC-TARNET	2.095 _(0.176)	1.889 _(0.159)	0.561 _(0.123)	2.100 _(0.209)	1.933 _(0.239)	0.652 _(0.170)
SC-CFR	2.062 _(0.174)	1.867 _(0.154)	*0.525 _(0.103)	2.081 _(0.221)	1.954 _(0.229)	**0.609 _(0.138)

datasets generated by a multi-agent simulator on evacuation scenarios in a crowded theater, we conducted experiments to estimate treatment effects on three types of outcomes of evacuation times required by crowds. We demonstrated that SC-CFR reduced the estimation error by at most 56% from the state-of-the-art baseline methods that did not consider spatial attributes.

Acknowledgement This work was supported by JST, PRESTO Grant Number JPMJPR20C5, Japan and JSPS KAKENHI Grant Numbers 20H00609, Japan

REFERENCES

- [1] Susan Athey and Guido Imbens. 2015. *Machine Learning for Estimating Heterogeneous Causal Effects*. Research Papers. Stanford University, Graduate School of

- Business.
- [2] Susan Athey, Julie Tibshirani, Stefan Wager, et al. 2019. Generalized random forests. *The Annals of Statistics* 47, 2 (2019), 1148–1178.
 - [3] M. Balmer, N. Cetin, Kai Nagel, and B. Raney. 2004. Towards truly agent-based traffic and mobility simulations. In *Proceedings of AAMAS*. 60–67.
 - [4] Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. 2011. Contextual Bandit Algorithms with Supervised Learning Guarantees. In *Proceedings of AISTATS*.
 - [5] Ioana Bica, Ahmed M. Alaa, James Jordon, and Mihaela van der Schaar. 2020. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. In *Proceedings of ICLR*.
 - [6] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. 2001. *The Elements of Statistical Learning*. Vol. 1. Springer series in statistics New York.
 - [7] Daniele Gianni, Georgios Loukas, Erol Gelenbe, et al. 2008. A Simulation Framework for the Investigation of Adaptive Behaviours in Largely Populated Building Evacuation Scenarios. In *Proceedings of AAMAS*.
 - [8] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. 2016. *Deep learning*. Vol. 1. MIT press Cambridge.
 - [9] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. 2012. A Kernel Two-Sample Test. *Journal of Machine Learning Research* 13, 25 (2012), 723–773.
 - [10] Negar Hassanpour and Russell Greiner. 2019. Counterfactual Regression with Importance Sampling Weights. In *Proceedings of IJCAI*.
 - [11] Jennifer L. Hill. 2011. Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics* 20, 1 (2011), 217–240.
 - [12] Guido W. Imbens and Donald B. Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.
 - [13] Guido W. Imbens and Jeffrey M. Wooldridge. 2009. Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature* 47, 1 (2009), 5–86.
 - [14] Fredrik D. Johansson, Uri Shalit, and David Sontag. 2016. Learning Representations for Counterfactual Inference. In *Proceedings of ICML*.
 - [15] Jintao Ke, Hongyu Zheng, Hai Yang, Xiqun, and Chen. 2017. Short-Term Forecasting of Passenger Demand under On-Demand Ride Services: A Spatio-Temporal Deep Learning Approach. *Transportation Research Part C: Emerging Technologies* 85 (2017), 591 – 608.
 - [16] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. arXiv:1412.6980
 - [17] Gregor Lammel, Marcel Rieser, and Kai Nagel. 2008. Bottlenecks and Congestion in Evacuation Scenarios: A Microscopic Evacuation Simulation for Large-Scale Disasters. In *Proceedings of AAMAS*.
 - [18] Yuxuan Liang, Kun Ouyang, Lin Jing, Sijie Ruan, Ye Liu, Junbo Zhang, David S. Rosenblum, and Yu Zheng. 2019. UrbanFM: Inferring Fine-Grained Urban Flows. In *Proceedings of SIGKDD*.
 - [19] Xiaolei Ma, Zhuang Dai, Zhengbing He, Jihui Ma, Yong Wang, and Yunpeng Wang. 2017. Learning Traffic as Images: A Deep Convolutional Neural Network for Large-Scale Transportation Network Speed Prediction. *Sensors* 17, 4 (2017), 818.
 - [20] Alfred Müller. 1997. Integral Probability Metrics and Their Generating Classes of Functions. *Advances in Applied Probability* 29, 2 (1997), 429–443.
 - [21] Yara Rizk, Mariette Awad, and Edward W. Tunstel. 2018. Decision Making in Multiagent Systems: A Survey. *IEEE Transactions on Cognitive and Developmental Systems* 10, 3 (2018), 514–529.
 - [22] Donald B. Rubin. 2005. Causal Inference Using Potential Outcomes. *J. Amer. Statist. Assoc.* 100, 469 (2005), 322–331.
 - [23] Patrick Schwab, Lorenz Linhardt, and Walter Karlen. 2018. Perfect Match: A Simple Method for Learning Representations For Counterfactual Inference With Neural Networks. arXiv:1810.00656
 - [24] Uri Shalit, Fredrik D. Johansson, and David Sontag. 2017. Estimating individual treatment effect: generalization bounds and algorithms. In *Proceedings of ICML*.
 - [25] Guni Sharon, Josiah P. Hanna, Tarun Rambha, Michael W. Levin, Michael Albert, Stephen D. Boyles, and Peter Stone. 2017. Real-time adaptive tolling scheme for optimized social welfare in traffic networks. In *Proceedings of AAMAS*.
 - [26] Arambam James Singh, Duc Thien Nguyen, Akshat Kumar, and Hoong Chuin Lau. 2019. Multiagent Decision Making For Maritime Traffic Management. In *Proceedings of AAAI*.
 - [27] Bharath K. Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Bernhard Schölkopf, Gert RG Lanckriet, et al. 2012. On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics* 6 (2012), 1550–1599.
 - [28] Alex Strehl, John Langford, Lihong Li, and Sham M. Kakade. 2010. Learning from Logged Implicit Exploration Data. In *Proceedings of NeurIPS*.
 - [29] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
 - [30] Akira Tanimoto, Tomoya Sakai, Takashi Takenouchi, and Hisashi Kashima. 2021. Regret Minimization for Causal Inference on Large Treatment Space. In *Proceedings of AISTATS*.
 - [31] Robert Tibshirani. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58, 1 (1996), 267–288.
 - [32] Jason Tsai, Natalie Fridman, Emma Bowring, Matthew Brown, Shira Epstein, Gal A. Kaminka, Stacy Marsella, Andrew Ogden, Inbal Rika, Ankur Sheel, et al. 2011. ESCAPES - Evacuation Simulation with Children, Authorities, Parents, Emotions, and Social comparison. In *Proceedings of AAMAS*.
 - [33] Pradeep Varakantham, Hala Mostafa, Na Fu, and Hoong Chuin Lau. 2015. DIRECT: A scalable approach for route guidance in Selfish Orienteering Problems. In *Proceedings of AAMAS*.
 - [34] Stefan Wager and Susan Athey. 2018. Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *J. Amer. Statist. Assoc.* 113, 523 (2018), 1228–1242.
 - [35] Amulya Yadav, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. 2016. sing Social Networks to Aid Homeless Shelters: Dynamic Influence Maximization under Uncertainty. In *Proceedings of AAMAS*.
 - [36] Tomohisa Yamashita, Takashi Okada, and Itsuki Noda. 2013. Implementation of Simulation Environment for Exhaustive Analysis of Huge-Scale Pedestrian Flow. *SICE Journal of Control, Measurement, and System Integration* 6, 2 (2013), 137–146.
 - [37] Liuyi Yao, Zhixuan Chu, Sheng Li, Yaliang Li, Jing Gao, and Aidong Zhang. 2020. A Survey on Causal Inference. (2020). arXiv:2002.02770
 - [38] Liuyi Yao, Sheng Li, Yaliang Li, Mengdi Huai, Jing Gao, and Aidong Zhang. 2018. Representation Learning for Treatment Effect Estimation from Observational Data. In *Proceedings of NeurIPS*.
 - [39] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. 2018. GANITE: Estimation of Individualized Treatment Effects using Generative Adversarial Nets. In *Proceedings of ICLR*.
 - [40] Haiyang Yu, Zhihai Wu, Shuqin Wang, Yunpeng Wang, and Xiaolei Ma. 2017. Spatiotemporal Recurrent Convolutional Networks for Traffic Prediction in Transportation Networks. *Sensors* 17, 7 (2017), 1501.
 - [41] Junbo Zhang, Yu Zheng, and Dekang Qi. 2017. Deep Spatio-Temporal Residual Networks for Citywide Crowd Flows Prediction. In *Proceedings of AAAI*.
 - [42] Siyuan Zhao and Neil Heffernan. 2017. Estimating Individual Treatment Effect from Educational Studies with Residual Counterfactual Networks. In *Proceedings of EDM*.
 - [43] Min Zhou, Hairong Dong, Petros A. Ioannou, Yanbo Zhao, and Fei-Yue Wang. 2019. Guided Crowd Evacuation: Approaches and Challenges. *IEEE/CAA Journal of Automatica Sinica* 6, 5 (2019), 1081–1094.