

Optimizing Impression Counts for Outdoor Advertising

Yipeng Zhang*
RMIT University
yipeng.zhang@rmit.edu.au

Yuchen Li†
Singapore Management University
yuchenli@smu.edu.sg

Zhifeng Bao*
RMIT University
zhifeng.bao@rmit.edu.au

Songsong Mo‡
Wuhan University
songsong945@whu.edu.cn

Ping Zhang§
Huawei
zhangping62@huawei.com

ABSTRACT

In this paper we propose and study the problem of optimizing the influence of outdoor advertising (ad) when impression counts are taken into consideration. Given a database \mathcal{U} of billboards, each of which has a location and a non-uniform cost, a trajectory database \mathcal{T} and a budget B , it aims to find a set of billboards that has the maximum influence under the budget. In line with the advertising consumer behavior studies, we adopt the logistic function to take into account the impression counts of an ad (placed at different billboards) to a user trajectory when defining the influence measurement. However, this poses two challenges: (1) our problem is NP-hard to approximate within a factor of $O(|\mathcal{T}|^{1-\epsilon})$ for any $\epsilon > 0$ in polynomial time; (2) the influence measurement is non-submodular, which means a straightforward greedy approach is not applicable. Therefore, we propose a tangent line based algorithm to compute a submodular function to estimate the upper bound of influence. Henceforth, we introduce a branch-and-bound framework with a θ -termination condition, achieving $\frac{\theta}{2}(1 - 1/e)$ approximation ratio. However, this framework is time-consuming when $|\mathcal{U}|$ is huge. Thus, we further optimize it with a progressive pruning upper bound estimation approach which achieves $\frac{\theta}{2}(1 - 1/e - \epsilon)$ approximation ratio and significantly decreases the running-time. We conduct the experiments on real-world billboard and trajectory datasets, and show that the proposed approaches outperform the baselines by 95% in effectiveness. Moreover, the optimized approach is around two orders of magnitude faster than the original framework.

CCS CONCEPTS

• **Mathematics of computing** → **Combinatorial optimization**; *Enumeration*; • **Applied computing** → *Marketing*.

KEYWORDS

Outdoor Advertising, Influence Maximization, Moving Trajectory, Non-submodularity, Logistic Function

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '19, August 4–8, 2019, Anchorage, AK, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6201-6/19/08...\$15.00

<https://doi.org/10.1145/3292500.3330829>

ACM Reference Format:

Yipeng Zhang, Yuchen Li, Zhifeng Bao, Songsong Mo, and Ping Zhang. 2019. Optimizing Impression Counts for Outdoor Advertising. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19)*, August 4–8, 2019, Anchorage, AK, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3292500.3330829>

1 INTRODUCTION

Outdoor advertising (ad) has been a market of 29 billion dollars since 2017 and its revenue is expected to grow by 3% to 4% per year to reach 33 billion dollars by 2021¹. 74% of its growth comes from the billboard segment [1]. The main audiences of billboards are people moving along their trips, by vehicles, motorcycles, bikes, etc. More than 80% drivers notice billboards when driving². Enabled by the prevalence of positioning devices, tremendous amounts of trajectories have been generated and recorded [32]. Moreover, the evidence in our experimental study shows that (Figure 4a), more than 50% travellers are impressed by more than five billboards on each trip.

The aforementioned opportunities motivate us to propose and study a novel research problem, namely optimizing Impression Counts for Outdoor Advertising (ICOA). Given a billboard database \mathcal{U} , a trajectory database \mathcal{T} and a budget B , ICOA aims to find a set of billboards that have the maximum influence under the budget. Solving ICOA is imperative as it facilitates the decision making of an advertiser to achieve the highest return on investment. For example, the average cost of renting a billboard is \$14,000/month in New York City (NYC) [12]; the total cost of renting 50 billboards is \$700,000/month. That means we can save about \$70,000/month if we can improve the influence by 10%.

To the best of our knowledge, this is the first problem that draws the inspiration from the intersection of (1) budget constraints, (2) moving trajectories when impression counts are considered, and (3) non-uniform costs of renting a billboard. As a result, the following challenges are important to be addressed.

The first challenge is how to appropriately measure the influence from a billboard to a user. Studies in consumer behavior report that, in the real world, users are unlikely to take a meaningful action when they receive only one message from an ad [7, 9, 13, 24, 28]. Meanwhile, there is evidence showing that the effect of ad repetition should be measured as an S-shaped function [4, 19, 23, 26], which

*Zhifeng Bao is the corresponding author.

‡ This work was done when Songsong Mo was a visiting student at RMIT.

¹ <https://www.marketing-interactive.com/ooh-advertising-spend-to-soar-to-us33-billion-by-2021/>

² <http://www.runningboards.com.au/outdoor/relocatable-billboards>

means the ad effectiveness will increase at low levels of repetition and then decrease as ad repetition increases. The logistic function has been widely adopted to measure the influence of an ad in many domain [10, 16, 25, 27, 29], since it matches the above characteristic of ad effectiveness. More importantly, as aforementioned, most people pass through more than five billboards during one trip. It is feasible for a company to rent multiple billboards to ingrain the ads in the user's mind. Therefore, we employ the logistic function to measure the influence from the ads (i.e., a set of billboards) to a user (i.e., a trajectory recording the user's travel). The influence of billboards is small when a user only saw the ad a few times, while it will increase dramatically upon seeing more. However, when this user has already seen a sufficient number of the same ad displayed on different billboards, the effect of additional impressions will decline as the effect of extra information diminishes.

The second challenge is posed by the property of the logistic function. The influence model based on the logistic function is non-submodular, which means any straightforward greedy-based approach is not applicable to address the ICOA problem (as elaborated in Section 3). Even worse, the non-uniform cost of different billboards makes the optimization problem intricate. We show that ICOA is NP-hard to approximate within any constant factor.

In order to address this algorithmic challenge, we propose an upper bound estimation method that tightly upper bounds the logistic function value, by means of a tangent line that intersects with the logistic S-curve. Based on the upper bound estimation method, we propose a branch-and-bound framework (Section 4.1). However, the efficiency and scalability of the branch-and-bound framework is limited – there is a potential for an exponential search space against the number of billboards and each upper bound estimation may visit a large number of trajectories. Even in one city, there are thousands of billboards and millions of trajectories (as evident in the real-world datasets used in our experimental study). It is thus time-consuming to explore all branches to get the optimal placement plan. Therefore, to further improve the efficiency of our framework, we devise a θ -termination method (Section 4.3) and a progressive upper bound estimation method (Section 5) which can provide a trade-off between efficiency and effectiveness.

The main contributions are summarized as below:

- We propose and study the ICOA problem for the first time, and show that the influence model based on the logistic function is non-submodular. We also prove that ICOA is NP-hard to approximate (Section 3).
- We propose an upper bound estimation method by adaptively solving submodular optimization problems. Based on the upper bound function, we propose a branch-and-bound framework. We further introduce a θ -termination method to achieve a trade-off between the efficiency and effectiveness. It achieves $\frac{\theta}{2}(1 - 1/e)$ approximation ratio (Section 4).
- To further boost efficiency, we optimize this framework with a progressive upper-bound estimation method, which achieves $\frac{\theta}{2}(1 - 1/e - \epsilon)$ approximation ratio and significant reduction in running-time (Section 5).
- We conduct extensive experiments on real-world trajectory and billboard datasets in the two largest cities of USA: NYC and LA. The results validate the effectiveness, efficiency and scalability of our methods (Section 6).

2 RELATED WORK

In the following, we discuss the most relevant literature to this paper: Trajectory-driven Influential Billboard placement (TIP), Site Selection, and Location-aware IM (LIM). The main differences between existing works and ICOA are summarized in Table 1.

TIP [32] is closely related to our problem, which also studies billboard placement to achieve the best advertising outcome. The core difference lies in the influence model. In particular, TIP assumes that a user (i.e., trajectory) can be influenced so long as one billboard is close enough to the trajectory the user travels along. Under such an influence model, when multiple billboards are close to a trajectory, the marginal influence is reduced to capture the property of diminishing returns. Therefore, TIP focuses on identifying and reducing the overlap of the influence among different billboards to the same trajectories, while keeping the budget constraint into consideration. That is, TIP can maximize the number of distinct users by impressing as many people as possible for one time. It does not consider the relationship between the influence effect and counts of impressions on one user because the model assumes one time impression is enough. ICOA is built upon a logistic influence model which has been widely adopted in consumer behavior studies. To maximize the influence to users, we need to control the overlap to some extent by impressing the same users several times. Unfortunately, the logistic influence model is non-submodular. Adapting the greedy approach to ICOA, which effectively solves TIP, could lead to arbitrarily bad solutions due to the non-submodular of the influence function.

The site selection problem has received lots of attention, due to its importance in a wide spectrum of applications. For example, the supply chain management problem involves a set of spatially distributed customers and a set of facilities to serve customer demands [2, 5, 15, 20, 21]. The potential locations of facilities and customers are inputted as a set of fixed locations. The given metrics are used to measure the distances, travel times or costs between customers and facilities. Despite the different metrics used, their goal is to minimize the objective function, e.g., the sum of the distance, time or cost. Another example is the Maximized Bichromatic Reverse k Nearest Neighbor (MaxBR k NN) problem. It aims to find an optimal location, where this location is a k NN of the maximum number of users based on the spatial distance between this location and users' locations [6, 18, 30, 31, 33]. Although the supply chain management problem, MaxBR k NN and ICOA fall under the general category of the site location problem, they are different in the following ways. ICOA aims to maximize the ad influence, whereas the rest seek to minimize the objective function. In addition to the differences in influence models, the supply chain management problem and MaxBR k NN assume that each user remains at a fixed location, whereas ICOA leverages moving trajectories to model the influence between users and billboards when a user travels along a trajectory and meets a number of billboards.

The Location-aware IM (LIM) problem [14, 17] is extended from Influence Maximization (IM) problem, which aims to select a size- k subset from a given social network. The difference is that LIM only measures the spread of influence on users who are located in the given search region. Although IM/LIM and ICOA share the same ultimate goal, which is to maximize influence, they are different

Table 1: Related work

	ICOA	TIP	Site Selection
Audience	Moving trajectory	Moving trajectory	Fixed location
Influence	Logistic	One-time impression	N.A.
S	Budget Constrained	Budget Constrained	Predefined k -size




Figure 1: The upper-bound influence function




Figure 2: An example of ICOA

in two aspects. First, influence models of IM/LIM are submodular, whereas the logistic model is not. This implies that the simple greedy approach is not suitable for ICOA. Second, in the IM problem, the influence can be spread from a user to others. In our model, a user can be only impressed by a billboard through which this user passed under a range constraint.

3 PROBLEM FORMULATION

In this section, we introduce how billboards influence trajectories and the formulation of the ICOA problem. The frequently used notations are presented in Table 2.

3.1 Preliminary

A trajectory database is denoted as $\mathcal{T} = \{t_1, t_2, \dots, t_{|\mathcal{T}|}\}$, where each trajectory $t = \{p_1, p_2, \dots, p_{|t|}\}$ is a set of points generated from the trajectory of a user. Point p consists of the latitude lat and longitude lng . Given a billboard database $\mathcal{U} = \{o_1, o_2, \dots, o_{|\mathcal{U}|}\}$, each billboard o is a tuple $\{loc, w\}$, where loc is also a coordinate and w is the cost of billboard o .

Although different ad companies use different strategies to set the renting cost, the cost is usually proportional to the real impression. How to measure the impression depends on the application needs, such as the panel size, the exposure frequency, the travel speed or the travel direction. Here we generate the cost of a billboard based on the number of trajectories impressed by this billboard. The detail is shown in Section 6.2. We also show that various strategies of setting the cost do not affect the performance of our methods in Appendix A.2.

We assume there are only two states of whether a user meets a billboard. When the user meets a billboard, we say this billboard impresses the user; otherwise, no impression is delivered. Therefore, we use the Bernoulli random variable $I(o, t)$ denoting the states whether o impresses t , where $I(o, t) = 1$ denotes that o delivers an impression to t , otherwise $I(o, t) = 0$.

DEFINITION 3.1. We define that o impresses t , denoted as $I(o, t) = 1$, if $\exists t.p_i$, such that $dist(t.p_i, o.loc) \leq \lambda$, where $dist(\cdot)$ computes the Euclidean distance between p_i and $o.loc$, λ is a given distance threshold.

Our influence model is based on the logistic function. We use the following equation to compute the *effective* influence of an ad

Table 2: Frequently used notations

Symbol	Description
$p(S, t)$ ($p^\dagger(S, t)$)	The (upper bound) influence of S to t
$I(S)$ ($I^\dagger(S)$)	The (upper bound) influence of S to \mathcal{T}
$\Delta(S_2 S_1)$	The marginal influence of adding S_2 into S_1

placed at a billboard set S which can impress a trajectory t :

$$p(S, t) = \begin{cases} \frac{1}{1 + \exp\{\alpha - \beta \cdot \sum_{o_i \in S} I(o_i, t)\}} & \text{if } \exists o_i \in S I(o_i, t) = 1 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

α and β are the parameters that control t 's turning point for being influenced, where α controls the overall influence of S to t and β controls the incremental influence of o to t . The S-curve in Figure 1 shows the shape of the logistic influence model. When β is fixed, with the increasing of α , it is hard to have any significant influence with a small number of billboard impressions. In contrast, when α is fixed, with the increasing of β , each billboard impression triggers a noticeable influence impact.

Next, we define the influence of S to a trajectory database \mathcal{T} as follows:

$$I(S) = \sum_{t \in \mathcal{T}} p(S, t) \quad (2)$$

Example 3.1. Let $S = \{o_3, o_4\}$ be a set of billboards chosen from all billboards in Figure 2, and trajectories t_1 , t_2 and t_3 are influenced by at least one billboard in S since there is at least one point p of each trajectory, which is in the red circle of a billboard. The red circle indicates the impression range of a billboard with its radius controlled by λ . Assuming the parameters $\alpha = 3$, $\beta = 1$, based on Equation 1, we have $p(S, t_1) = 0.119$, $p(S, t_2) = p(S, t_3) = 0.269$, and $p(S, t_4) = 0$, respectively. Hence, the overall influence of S is $I(S) = 0.119 + 0.269 + 0.269 = 0.657$.

3.2 Problem Definition

We are now ready to formally introduce the ICOA problem.

DEFINITION 3.2. (ICOA) Given a billboard database \mathcal{U} , a trajectory database \mathcal{T} , a budget constraint B and the influence model $I(S)$, the ICOA problem is to find a subset $S \subseteq \mathcal{U}$ that maximizes the overall influence of S such that the total cost of S does not exceed B . Formally,

$$\hat{S} = \underset{\text{cost}(S) \leq B}{\operatorname{argmax}} I(S) \quad (3)$$

Non-submodularity of ICOA. Given two sets of billboards S_1 and S_2 , the marginal influence of adding S_2 into S_1 is $\Delta(S_2|S_1) = I(S_1 \cup S_2) - I(S_1)$. Then, we define the monotonicity and submodularity of an influence function as follows. $I(S)$ is monotone iff, $I(S_1) \leq I(S_2)$ for all $S_1 \subseteq S_2$. Furthermore, $I(S)$ is submodular iff, given any set of billboards S^* , it satisfies $\Delta(S^*|S_1) \geq \Delta(S^*|S_2)$ for all $S_1 \subseteq S_2$. The following presents a counterexample for the influence function to be submodular.

Example 3.2. Assuming the same settings as in Example 3.1, we choose three billboard sets $S_1 = \{\}$, $S_2 = \{o_3\}$, $S_3 = \{o_4\}$. The respective influence and marginal influence values can be calculated as $I(S_1) = 0$, $I(S_2) = 0.3576$, $I(S_3) = 0.2384$, $\Delta(S_3|S_1) = 0.2384$, $\Delta(S_3|S_2) = 0.657 - 0.357 = 0.2994$. Since $S_1 \subseteq S_2$ and $\Delta(S_3|S_1) < \Delta(S_3|S_2)$. We thus conclude $I(S)$ is not submodular.

Due to the non-submodularity of ICOA, a greedy-based heuristic method cannot guarantee any constant approximation ratio.

Algorithm 1: Branch-and-Bound

Input: $\mathcal{U}, \mathcal{T}, B$
Output: \hat{S}

```

1.1  $\hat{S} \leftarrow \phi, S \leftarrow \phi, \bar{S} \leftarrow \mathcal{U}$ 
1.2  $L_G \leftarrow 0, U_G \leftarrow \infty$ 
1.3 Initialize max heap  $H \leftarrow (S, \bar{S}, U)$ 
1.4 while  $L_G < U_G$  do
1.5    $(S, \bar{S}, U) \leftarrow \text{top of } H$ 
1.6   Select  $o \in \bar{S}$ 
1.7   if  $\text{cost}\{S\} + o.w \leq B$  then
1.8      $\bar{S} \leftarrow \bar{S} \setminus \{o\}$ 
1.9      $S^a \leftarrow S \cup \{o\}$ 
1.10     $S^b \leftarrow S$ 
1.11     $(S^c, L^a, U^a) \leftarrow \text{ComputeBound}(S^a, \bar{S})$ 
1.12    if  $L^a > L_G$  then
1.13       $L_G \leftarrow L^a, \hat{S} \leftarrow S^c$ 
1.14    if  $U^a > L_G$  then
1.15       $H \leftarrow H \cup (S^a, \bar{S}, U^a)$ 
1.16  Repeat Lines 1.11-1.15 for  $S_b$ 

```

THEOREM 3.1. *The ICOA problem is NP-hard to approximate within a factor of $O(|\mathcal{T}|^{1-\epsilon})$ for any $\epsilon > 0$ in polynomial time.*

PROOF. Kindly refer to the proof in Appendix A.4. \square

4 OUR FRAMEWORK

According to Theorem 3.1, there does not exist any efficient algorithm with constant approximation ratio to ICOA. One naïve solution is to enumerate all feasible billboard subsets and compute their influence. However, this is not scalable against thousands of billboards and millions of trajectories.

Thereby, we propose a branch-and-bound framework (Section 4.1). It explores branches, which represent respective feasible billboard sets that have not yet exhausted the budget and can be filled with more billboards. In particular, we propose a novel bound estimation technique for each branch under exploration by setting a submodular function to tightly upper bound $p(S, t) \forall t \in \mathcal{T}$ (Section 4.2). The estimation technique will obtain a candidate solution (i.e. the billboard set which cannot be further expanded due to the budget constraint) when calculating the upper bound score of a branch. The exploration terminates when the upper bound of all remaining branches does not exceed the influence value of the best candidate solution. Theoretically, the branch-and-bound framework achieves an approximation ratio of $\frac{1}{2}(1 - 1/e)$.

4.1 Branch-and-Bound

Algorithm 1 depicts the pseudocode of the branch-and-bound framework. We first initialize the global upper bound U_G and global lower bound L_G , and a max heap H with each entry denoted as (S, \bar{S}, U) (Lines 1.1-1.3), where S is the set of billboards that have been selected as a feasible set, \bar{S} is the set of billboards that have not been considered yet, and U is the upper bound influence of the corresponding search space. H is ordered by the upper bound value of each S . While $L_G < U_G$, H will pop the top entry that has the maximum upper bound influence (Lines 1.4-1.5). For each entry, as long as it matches the budget constraint, it will generate two new

Algorithm 2: ComputeBound

Input: S^a, \bar{S}
Output: (S^c, L^a, U^a)

```

2.1  $S^* \leftarrow \phi$ 
2.2 while  $\text{cost}(S^a) + \text{cost}(S^*) \leq B$  and  $|\bar{S}| \neq 0$  do
2.3   Select  $o \in \bar{S}$  that maximize  $\frac{\Delta^\uparrow(\{o\}|S^a \cup S^*)}{o.w}$ 
2.4   if  $\text{cost}(S^a \cup S^*) + o.w \leq B$  then
2.5      $S^* \leftarrow S^* \cup \{o\}$ 
2.6    $\bar{S} \leftarrow \bar{S} \setminus \{o\}$ 
2.7  $S^c \leftarrow (S^a \cup S^*), L^a \leftarrow I(S^a \cup S^*), U^a \leftarrow I^\uparrow(S^a \cup S^*)$ 

```

branches (S^a and S^b) based on S respectively (Lines 1.6 - 1.10). S^a denotes a feasible set where a billboard $o \in \bar{S}$ can be further added into S , and S^b denotes a feasible set excluding o .

Based on S^a (or S^b) and the corresponding \bar{S} , $\text{ComputeBound}(\cdot)$ will return a triple, i.e., (S^c, L^a, U^a) or (S^c, L^b, U^b) , where S^c is a candidate solution set returned by $\text{ComputeBound}(\cdot)$. L^a and U^a are the lower-bound influence and upper bound influence of S^c respectively (Line 1.11). If $L^a > L_G$, which means S^c is better than the current best feasible solution \hat{S} , then \hat{S} will be replaced by S^c , and the global L_G is updated (Lines 1.12-1.13). If $U^a > L_G$, it is possible that S^a is a subset of the optimal solution. Therefore, (S^a, \bar{S}, U^a) will be pushed into H (Lines 1.14-1.15). We repeat the search loop for all branches until $L_G \geq U_G$.

4.2 ComputeBound

To estimate the upper bound of a branch w.r.t. a feasible set S^a , we devise a submodular function (i.e., $p^\uparrow(S, t)$ and $S = S^a \cup S^*$) which tightly upper bounds the non-submodular influence function $p(S, t)$ (Equation 1). Let $x(S)$ denote the number of effective impressions to t obtained by placing ads in billboard set S (i.e., $x(S) = \sum_{o_i \in S} I(o_i, t)$) and $f(x(S)) = 1/(1 + \exp\{\alpha - \beta \cdot x(S)\}) = p(S, t)$. We draw a tangent line $l(x)$ to upper bound $f(x)$. $l(x)$ intersects $f(x)$ at two points: $(x(S^a), f(x(S^a)))$ and $(x_t^{S^a}, f(x_t^{S^a}))$ where the latter denotes the tangential point (see Figure 1). Formally, we define the upper bound function $p^\uparrow(S, t)$ for $S = S^a \cup S^*$ as follows

If $l(x)$ exists:

$$p^\uparrow(S, t) = \begin{cases} l(x) & \text{if } x(S^a) \leq x \leq x_t^{S^a} \\ f(x) & \text{if } x_t^{S^a} < x \end{cases} \quad (4)$$

Otherwise:

$$p^\uparrow(S, t) = f(x) \quad (5)$$

As shown in Figure 1, $p^\uparrow(S, t)$ is submodular as it concatenates two submodular functions: $l(x)$ and $f(x)$ for different domains of x . Furthermore, we define the following submodular function that upper bounds the influence function $I(S)$ for $S = S^a \cup S^*$.

$$I^\uparrow(S) = \sum_{t \in \mathcal{T}} p^\uparrow(S, t) \quad (6)$$

It is easy to see that $I^\uparrow(S) \geq I(S)$ as $p^\uparrow(S, t) \geq p(S, t)$ for all $S = S^a \cup S^*$. Furthermore, $I^\uparrow(S)$ is submodular because it is a sum of submodular functions. To ease our presentation, we define the marginal influence $I^\uparrow(\cdot)$ of adding S_2 into S_1 as below:

$$\Delta^\uparrow(S_2|S_1) = I^\uparrow(S_1 \cup S_2) - I^\uparrow(S_1) \quad (7)$$

Given the upper bound function $I^\uparrow(\cdot)$, we introduce Algorithm 2 to estimate the upper bound of a branch. S^a is one of the two

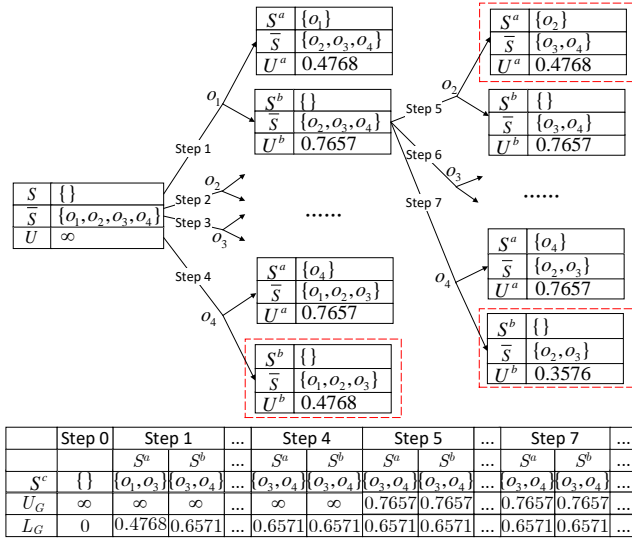


Figure 3: A running example of branch-and-bound framework

branches under exploration in Algorithm 1. \bar{S} is the set of billboards that have not yet been considered. S^* is the set of billboards that have been chosen from \bar{S} by using the cost-effective greedy heuristic under the budget constraint. The algorithm terminates until adding any remaining billboards will exceed the budget (Lines 2.3-2.6). Then, a triple (S^c, L^a, U^a) is returned.

Since $I^\uparrow(\cdot)$ is submodular, S^c achieves a constant approximation ratio w.r.t. $I^\uparrow(\cdot)$ for all candidate solutions containing S^a . It is noted that the tangent line of $p^\uparrow(S, t)$ for $S = S^a \cup S^*$ will be adjusted according to different S^a before the launch of Algorithm 2. The adjustment of the tangent line ensures a tighter upper bound as depicted in Figure 1. While the cardinality of S^a increases, the point $(x(S^a), f(x(S^a)))$ will move to the right. Therefore, the tangent line will shift from Line 1 to Line 2. Since the distance from $l(x)$ to $f(x)$ decreases, it achieves a tighter upper bound.

Example 4.1. Figure 3 shows the running example of the branch-and-bound framework. The top part is the search tree. Each node is an entry (S, \bar{S}, U) . The table shows the values of S^c , U_G and L_U for each search step. Let $B = 50$, $o_1.w = 10$, $o_2.w = 30$, $o_3.w = 30$, $o_4.w = 20$. The rest of the parameters are the same as Example 3.1. We first initialize H with an entry $(\{\}, \{o_1, o_2, o_3, o_4\}, \infty)$ as the root node of the searching tree. Next, at step 1, we pop this entry, and invoke Algorithm 2 to generate two branches based on this entry, and push them into H . Since $U^b > U^a > L_G$ and $L^b > L^a > L_G$, we update \hat{S} with $S^c = \{o_3, o_4\}$ and set $L_G = L^b = 0.6571$. Similarly, at step 4, the branch S^a is pushed into H since $U^a > L_G$. Note that, the branch S^b will be pruned since $U^b < L_G$. For now, all branches based on the root node have been generated, and those with $U > L_G$ have been pushed into H . Subsequently, the branch with $S^b = \{\}$ and $\bar{S} = \{o_2, o_3, o_4\}$ is popped because it has the highest upper bound in H . U_G will be replaced by U of this branch, which means $U_G = U = 0.7657$. Six branches are generated with or without o_2 , o_3 and o_4 respectively.

The branches which $U > L_G$ will be added into H , such as the branch S^b of step 5 and the branch S^a of step 7. When L_G is larger

than U_G , it means that, in the rest of branches in H , the one with the highest U is worse than the optimal solution. The algorithm thus terminates.

The upper bounding techniques lead to a constant approximation ratio for the solution returned by the branch-and-bound framework.

THEOREM 4.1. *The branch-and-bound framework invoking Algorithm 2 achieves an approximation factor of $\frac{1}{2}(1 - \frac{1}{e})$ for the ICOA problem.*

PROOF. Let OPT denote the optimal solution. Let \hat{S} denote the solution set from Algorithm 1. $S^* = \{o_1, \dots, o_n\}$. Let S^* denote the feasible set returned from Algorithm 2. In each iteration, o will be added into S^* . Let $S_n^* = \{o_1, \dots, o_n\}$, where o_n is the last one added into S_{n-1}^* within the budget constraint. o_{n+1} is the first billboard considered but not added to S_n^* , because its addition would violate the budget B .

It has been proven in the BMC work [11] that, based on the cost-effective greedy method, $\Delta^\uparrow(S_n^*|S^a) + \Delta^\uparrow(\{o_{n+1}\}|S^a) \geq \Delta^\uparrow(S_n^* \cup \{o_{n+1}\}|S^a) \geq (1 - \frac{1}{e}) \cdot \Delta^\uparrow(OPT|S^a)$ holds. Therefore, at least one of $\Delta^\uparrow(S_n^*|S^a)$ and $\Delta^\uparrow(\{o_{n+1}\}|S^a)$ is not smaller than $\frac{1}{2}(1 - \frac{1}{e}) \cdot \Delta^\uparrow(OPT|S^a)$. While one of $(S_n^* \cup S^a)$ and $\{o_{n+1} \cup S^a\}$ will be \hat{S} , we have $\Delta^\uparrow(\hat{S}) \geq \frac{1}{2}(1 - \frac{1}{e}) \cdot \Delta^\uparrow(OPT|S^a)$ based on all explored S^a . Hence, $I^\uparrow(\hat{S}) \geq \frac{1}{2}(1 - \frac{1}{e}) \cdot I^\uparrow(OPT|S^a)$. For any branch that has not been searched, under the termination condition $L < U$, we have $I(\hat{S}) \geq I^\uparrow(\hat{S})$. Therefore, Algorithm 1 achieves $I(\hat{S}) \geq \frac{1}{2}(1 - \frac{1}{e}) \cdot I(OPT)$. \square

4.3 Branch-and-bound with θ -termination

It is noted that the search space of the branch-and-bound framework is exponential in the worst case. Although the upper bound estimation technique can prune a large number of branches, it inherently overestimates the influence value. As a result, there could be cases where the optimal solution has already been obtained but the search cannot terminate, because there exists an unexplored branch containing a solution which is also near optimal. In this case, the unexplored branch will have a higher upper bound value than the optimal influence value and the branch thereby cannot be pruned. It is unnecessarily expensive for exploring all remaining branches to reach other near-optimal solutions.

Therefore, we introduce a tunable early termination method. In Algorithm 1 (Line 1.4), we utilize a parameter θ to control the termination condition – instead of using $L_G < U_G$, we use $L_G < \theta U_G$ as the termination condition, where $\theta \in (0, 1]$. When L and U are close enough, the search process terminates. It can be easily shown that the early termination technique achieves $\frac{\theta}{2}(1 - 1/e)$ approximation ratio.

THEOREM 4.2. *The branch-and-bound framework invoking Algorithm 2 with θ -termination achieves an approximation ratio of $\frac{\theta}{2}(1 - \frac{1}{e})$ for the ICOA problem.*

5 PROGRESSIVE BRANCH-AND-BOUND

The branch-and-bound framework heavily invokes Algorithm 2 for bound estimations. For every search iteration in Algorithm 2 (Lines 2.3-2.6), the greedy selection needs to recalculate the marginal gain $\Delta^\uparrow(\{o\}|S^a \cup S^*)/o.w$ for all $o \in \bar{S}$, and chooses the maximum.

Algorithm 3: Progressive-ComputeBound

Input: S^a, \bar{S}, ϵ
Output: (S^c, L^a, U^a)

```

3.1 Reorder  $o \in \bar{S}$  by  $\frac{\Delta^\uparrow(\{o\}|S^a)}{o.w}$ 
3.2  $h \leftarrow \max_{o \in \bar{S}} \frac{\Delta^\uparrow(\{o\}|S^a)}{o.w}$ ,  $r = B - \text{cost}(S^a)$ 
3.3  $S^* \leftarrow \phi$ 
3.4 while  $\text{cost}(S) + \text{cost}(S^*) \leq B$  do
3.5   for  $o \in \bar{S}$  do
3.6     if  $\text{cost}(S^a \cup S^*) + o.\text{cost} \leq B$  then
3.7        $\delta(\{o\}, S^*) = \Delta^\uparrow(S^* \cup \{o\}|S^a) - \Delta^\uparrow(S^*|S^a)$ 
3.8       if  $\frac{\delta(\{o\}, S^*)}{o.w} \geq h$  then
3.9          $\bar{S} \leftarrow \bar{S} \setminus \{o\}$ 
3.10         $S^* \leftarrow S^* \cup \{o\}$ 
3.11       if  $\frac{\delta(\{o\}, S^*)}{o.w} < h$  then
3.12         Break
3.13    $h \leftarrow \frac{h}{1+\epsilon}$ 
3.14   if  $h \leq \frac{\Delta^\uparrow(S^*|S^a)}{r} \cdot \frac{e^{-1}}{1-e^{-1}}$  then
3.15     Break
3.16  $S^c \leftarrow (S^a \cup S^*)$ ,  $L^a \leftarrow I(S^a \cup S^*)$ ,  $U^a \leftarrow I^\uparrow(S^* \cup S^a)$ 

```

When $|\bar{S}|$ is huge, such bound estimation approach incurs significant computation overhead. Motivated by this, we propose a progressive upper bound estimation method without traversing all billboards to estimate the bound with an approximation ratio of $\frac{\theta}{2}(1 - 1/e - \epsilon)$, where ϵ is a tunable parameter that provides a trade-off between efficiency with accuracy.

In particular, instead of exploring all billboards to find the one with maximal $\delta^\uparrow(\{o\}|S^a)/o.w$ in each iteration, we sort them by $\delta^\uparrow(\{o\}|S^a)/o.w$ first (Line 3.1). Then, we set a threshold h as the maximal $\delta^\uparrow(\{o\}|S^*)/o.w$ (Line 3.2). We progressively decrease h by a factor of $(1 + \epsilon)$ and add more $o \in \bar{S}$ into S^* (Lines 3.5-3.13). The algorithm terminates when there is no billboard whose $\Delta^\uparrow(\{o\}|S^a)/o.w \geq h$. As a result, we do not need to explore all billboards in order to find the best one, and when h is small enough, the algorithm can terminate early.

In the rest of this section, we analyze the approximation ratio of Algorithm 3.

LEMMA 5.1. *Let $r = B - \text{cost}(S^a)$, which is the remaining budget. In Algorithm 3, at the i th iteration of search loop, after o_i has been added into S_{i-1}^* , the following holds:*

$$\Delta^\uparrow(S_i^*|S^a) - \Delta^\uparrow(S_{i-1}^*|S^a) \geq \frac{o_i.w}{(1+\epsilon)r} \cdot (\Delta^\uparrow(\text{OPT}|S^a) - \Delta^\uparrow(S_{i-1}^*|S^a))$$

PROOF. Let OPT denote the optimal S^* returned from Algorithm 3. Let o_i denote the billboard to be added into S_{i-1}^* at a given threshold h . Since $I^\uparrow(\cdot)$ is submodular, it holds that:

$$\frac{\delta(\{o\}, S^*)}{o.w} \begin{cases} \geq h & \text{if } o = o_i \\ \leq h(1+\epsilon) & \text{if } o \in \text{OPT} \setminus (S^* \cup \{o_i\}) \end{cases} \quad (8)$$

where S^* is the set of candidate billboards and $\delta(\{o\}, S^*)$ is defined as Line 3.7 of Algorithm 3. From Equation 8 we can find that, for any $o \in \text{OPT} \setminus S^*$, $\delta(\{o_i\}, S^*)/o_i.w \geq \delta(\{o\}, S^*)/o.w(1+\epsilon)$. Therefore, we have $\delta(\{o_i\}, S^*) \geq \frac{o_i.w}{(1+\epsilon)(\text{cost}(\text{OPT} \setminus S^*))} \sum_{o \in \text{OPT} \setminus S^*} \delta(\{o\}, S^*) \geq \frac{o_i.w}{(1+\epsilon)r} \sum_{o \in \text{OPT} \setminus S^*} \delta(\{o\}, S^*)$.

Let S_i^* denote the feasible set that o_i has been added into, i.e., $S_i^* = \{o_1, \dots, o_i\}$. Then we have:

$$\begin{aligned} \delta(\{o_i\}, S_{i-1}^*) &\geq \frac{o_i.w}{(1+\epsilon)r} \sum_{o \in \text{OPT} \setminus S_{i-1}^*} \delta(\{o\}, S_{i-1}^*) \\ &\geq \frac{o_i.w}{(1+\epsilon)r} \cdot (\Delta^\uparrow(\text{OPT} \cup S_{i-1}^*|S^a) - \Delta^\uparrow(S_{i-1}^*|S^a)) \\ &\geq \frac{o_i.w}{(1+\epsilon)r} \cdot (\Delta^\uparrow(\text{OPT}|S^a) - \Delta^\uparrow(S_{i-1}^*|S^a)) \end{aligned}$$

Then by the definition of $\delta(\{o\}, S^*)$, we have:

$$\Delta^\uparrow(S_i^*|S^a) - \Delta^\uparrow(S_{i-1}^*|S^a) \geq \frac{o_i.w}{(1+\epsilon)r} \cdot (\Delta^\uparrow(\text{OPT}|S^a) - \Delta^\uparrow(S_{i-1}^*|S^a))$$

Thus, the lemma is proved. \square

THEOREM 5.2. *The branch-and-bound framework invoking Algorithm 3 achieves an approximation ratio of $\frac{\theta}{2}(1 - 1/e - \epsilon)$.*

PROOF. Based on Lemma 5.1, it is easy to get:

$$\begin{aligned} \Delta^\uparrow(S_{n'}^*|S^a) &\geq \left(1 - \prod_{k=1}^{n+1} \left(1 - \frac{o_k.w}{(1+\epsilon)r}\right)\right) \cdot \Delta^\uparrow(\text{OPT}|S^a) \\ &\geq \left(1 - e^{-\frac{(n+1)}{(1+\epsilon)(n'+1)}}\right) \cdot \Delta^\uparrow(\text{OPT}|S^a) \end{aligned}$$

Therefore, we have:

$$\Delta^\uparrow(S_{n'}^*|S^a) \geq \frac{1}{2} \left(1 - e^{-\frac{(n+1)}{(1+\epsilon)(n'+1)}}\right) \cdot \Delta^\uparrow(\text{OPT}|S^a)$$

where $n = |S_n^*|$, which is the number of billboards returned from Algorithm 3 when the budget is exhausted, n' is the number of billboards when the budget is not exhausted. Because Algorithm 3 will be terminated when $h \leq \frac{\Delta^\uparrow(S^*|S^a)}{r} \cdot \frac{e^{-1}}{1-e^{-1}}$, we have $n' \leq n$.

Case 1: when $n' = n$, we have:

$$\Delta^\uparrow(S^*|S^a) \geq \frac{1}{2} (1 - e^{-1} - \epsilon) \cdot \Delta^\uparrow(\text{OPT}|S^a)$$

Case 2: when $n' < n$, $n' \cdot o.w < r$, we have:

$$\begin{aligned} \Delta^\uparrow(\text{OPT}|S^a) &\leq \Delta^\uparrow(\text{OPT} \cup S^*|S^a) \\ &\leq \frac{\Delta^\uparrow(S^*|S^a)}{r} \cdot \frac{e^{-1}}{1-e^{-1}} \cdot \sum_{o \in \text{OPT} \setminus S^*} o.w + \Delta^\uparrow(S^*|S^a) \\ &\leq \frac{1}{1-e^{-1}} \Delta^\uparrow(S^*|S^a) \end{aligned}$$

Therefore, $\Delta^\uparrow(S^*|S^a) \geq \frac{1}{2} (1 - \frac{1}{e} - \epsilon) \cdot \Delta^\uparrow(\text{OPT}|S^a)$. Similar to the proof of Theorem 4.2, we have: $I^\uparrow(\bar{S}) \geq \frac{\theta}{2} (1 - \frac{1}{e} - \epsilon) \cdot I^\uparrow(\text{OPT})$. Theorem 5.2 is proved. \square

6 EXPERIMENT

6.1 Experimental Setup

Datasets. The real-world billboard datasets for the two largest cities in the US (NYC and LA) are crawled from LAMAR [12], one of the largest outdoor advertising companies worldwide. The real-world trajectory datasets are obtained as follows. For NYC, we collect five hundred thousand taxi trips from TLC trip record³. Each trip record includes the pick-up and drop-off locations, time and trip distances. We use Google Maps API⁴ to generate the trajectories. Similar to [32], we only keep the trajectory if (1) the distance of the generated

³http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml


⁴<https://developers.google.com/maps/>

Table 3: Statistics of datasets

	$ \mathcal{T} $	$ \mathcal{U} $	AvgDistance	AvgTravelTime	AvgPoint
NYC	600k	1500	2.9km	569s	159
LA	250k	2500	2.7km	511s	138

Table 4: Parameter settings

Parameter	Values
B	100k, 200k, 300k , 400k, 500k
$ \mathcal{T} $ (NYC)	100k, 200k, 300k , 400k, 500k, (1m, ..., 5m)
$ \mathcal{T} $ (LA)	50k, 100k, 150k , 200k, 250k
β/α	3/7, 3/8, 3/9, 3/10, 3/11
ϵ	10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , 1
θ	0.86, 0.88, 0.90 , 0.92, 0.94
λ	25m, 50m , 75m, 100m

**Figure 4: Distribution of Datasets in NYC**

trajectory is close to that of the trip record, and (2) the travel time of the generated trajectory is close to that of the trip record (within 5% error rate). For LA, as there is no public taxi record, we collect the Foursquare check-in data and generate the trajectories using Google Maps API by randomly selecting the pick-up and drop-out locations from the check-ins.

Figure 4a shows the distribution of trajectories according to the number of billboards that a trajectory can pass over in NYC. We find more than 50% trajectories can pass over more than 5 billboards, which validates the motivation of this work as well as our use of the logistic function for influence modelling. Figure 4b shows the distribution of billboard cost.

6.2 Experiment Setting

Performance Measurement. For each method we evaluate the runtime and the influence value of the selected billboards. Each experiment is repeated ten times, and the average result is reported. **Billboard Costs.** There is no exact leasing cost of billboard provided by any advertising companies. To the best of our knowledge, they only provide a range of costs based on different areas. For example, according to the LAMAR company, the cost of the most expensive billboard is about thirty times higher than the cheapest one. Therefore, we generate the costs of billboard based on how many trajectories this billboard can influence, $o.w = \lceil k \sum_{t_i \in \mathcal{T}} I(o, t_i) \rceil$, where k is a factor chosen from 0.5 to 2 randomly to simulate various ratio of influence to cost. Then we map the costs of billboards to an interval from \$2,000 to \$60,000.

Experiment Environment. All codes are implemented in Java. Experiments are conducted on a laptop with Intel Core i7-8550U CPU and 16GB memory running Windows 10.

Parameter Settings. Table 4 shows the settings of all parameters, and the default one is highlighted in bold. In all experiments, we

only vary one parameter and keep the rest by default. α and β are the parameters in the logistic function that control t 's turning point for influence. We set $\beta = 3$, and vary α from 7 to 11. ϵ is used in Algorithm 3 to trade efficiency with accuracy, and θ is used in Algorithm 1 to control the termination condition. The methodology on selecting the default settings is shown in Appendix A.

Metrics & Methods for Comparison. In particular, we would like to evaluate the efficiency, effectiveness and scalability of our methods. To the best of our knowledge, this is the first work studying how to optimize outdoor ad influence by considering the impression counts over moving trajectories. Despite that, we compare the following baselines.

- Greedy: A basic greedy algorithm. In each iteration, it adds o with the maximum ratio of marginal influence to cost (i.e., $\frac{\Delta^I(\{o\}|\mathcal{S}^*)}{o.w}$) into \mathcal{S}^* until reaching the budget constraint.
- Top- k : In each iteration, it chooses o which can influence the maximum number of trajectories until reaching the budget constraint.
- BBS: The branch-and-bound framework with θ -termination and Algorithm 2 for bound estimations.
- PBBS: The branch-and-bound framework with θ -termination and Algorithm 3 for bound estimations.
- LazyProbe: The best-performing method in the most recent trajectory-driven billboard placement study [32]. Recall Section 2, our work and [32] are developed based on different influence models; LazyProbe can only work with a submodular influence function. Although it is not fair for our methods to be compared with a submodular influence model, we still compare our method with LazyProbe in Section 6.6, while we neglect it in other experiments.

6.3 Varying the Budget B

Figure 5 shows the effectiveness and efficiency of all algorithms when varying the budget B in NYC and LA, respectively.

Effectiveness. From Figure 5a and Figure 5c, we make the following observations. First, when the budget raises from 100k to 500k, both BBS and PBBS outperform Greedy, from 10% to 95%. Second, PBBS is slightly worse than BBS by up to 8%. It is because Algorithm 3 may terminate early for some branches and miss some ideal selections. Third, Top- k has the worst performance, as it gives preference to the billboard with the highest influence in each iteration, which is usually the most expensive billboard in the real world. Hence, it can only choose a few of billboards when the budget is fewer. A few billboards are unlikely to overlap, which makes our solutions better than it by at most 25 times in NYC, and 15 times in LA, respectively. When the budget is big enough, the growing effectiveness is mainly contributed from a growing number of billboards, which makes the advantage of our solutions dwindle to about 3 times in NYC, and 1 time in LA, respectively. Last, the advantage of BBS and PBBS in LA is less than that in NYC. One possible reason is that the distribution of trajectories in LA is comparatively even, making it more possible to have influence overlaps among billboards.

Efficiency. From Figure 5b and Figure 5d, we have three observations. First, the running time of BBS increases significantly w.r.t. the budget B . This is because in every branch, BBS has to invoke Algorithm 2, which needs to calculate the unit marginal influence for all $o \in \bar{\mathcal{S}}$ in each iteration. When B increases, more billboards



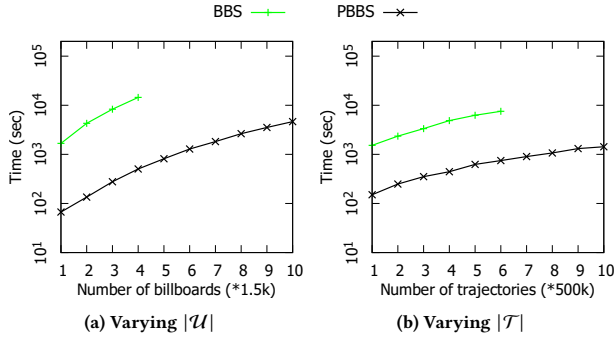
Figure 5: Varying the budget B Figure 6: Varying the number of trajectories $|\mathcal{T}|$ 

Figure 7: Scalability test in NYC

can be selected into S^* , and thus more iterations are needed. Second, the running time of PBBS increases by 100% when B increases. The reason is that, Algorithm 3 does not need to recalculate the unit marginal influence for every $o \in \bar{S}$. Last, Top- k is the fastest one since it only scans all billboards once.

6.4 Varying the Number of Trajectories $|\mathcal{T}|$

Effectiveness. As shown in Figure 6a and Figure 6c, we have three observations. First, the influence of all methods increase because more trajectories can be influenced. Second, the effectiveness of BBS and PBBS consistently outperform that of Greedy and Top- k by up to 60% and 300%, respectively. Last, the advantage of efficiency of BBS and PBBS in LA is less than that in NYC for similar reason mentioned in Section 6.3.

Efficiency. From Figure 6b and Figure 6d, we have the following observations. First, PBBS is about one order of magnitude faster than BBS. Second, the running time of all methods increase almost

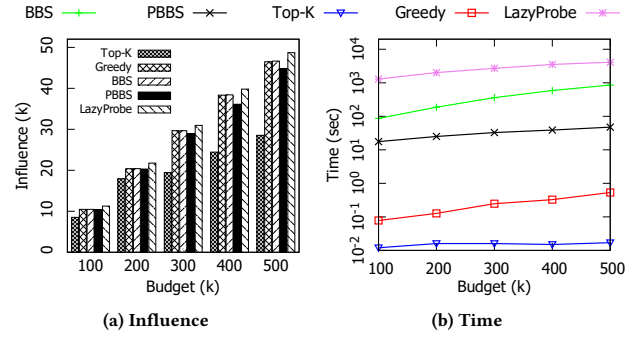


Figure 8: Comparison with LazyProbe

linearly w.r.t. $|\mathcal{T}|$, except for Top- k , because it chooses the billboards that can influence the most number of trajectories, which can be calculated off-line.

6.5 Scalability Test

To evaluate the scalability of our methods BBS and PBBS, we vary $|\mathcal{T}|$ from 500k to 5M, and $|\mathcal{U}|$ from 1,500 to 15,000. As shown in Figure 7a, the efficiency of BBS is more sensitive than that of PBBS when varying $|\mathcal{U}|$. With the increase of $|\mathcal{U}|$, the growth rate of running time of BBS is larger than that of PBBS. In particular, PBBS is 15 times faster than BBS at $|\mathcal{U}| = 1,500$, while such performance gap increases to almost two orders of magnitude at $|\mathcal{U}| = 6,000$. Since BBS takes more than 10^4 seconds to complete when the number of billboards $|\mathcal{U}|$ exceeds 6,000, its result is omitted in Figure 7a. Varying $|\mathcal{T}|$ leads to a similar result. As shown in Figure 7b, the results are omitted for BBS since the approach cannot terminate within 10^4 seconds when $|\mathcal{T}|$ exceeds 3.5M.

6.6 Our methods vs. LazyProbe [32]

In this experiment, we try to compare our work with the related work [32], although its goal is to maximize the influence while reducing the influence overlap, which is different from ours. In particular, we choose the best-performing method in [32], LazyProbe, and use the default setting as specified in [32]. Since LazyProbe can only work with a submodular influence model while ours is non-submodular, we have to adapt our influence model to be submodular by adjusting α and β , although it is actually not fair for our methods. It means if a trajectory is influenced by a billboard, i.e., if $\exists o_i \in S, I(o_i, t) = 1$, then $p(S, t) = 1$.

Figure 8 shows the experiment result. LazyProbe has the best effectiveness, which outperforms BBS, Greedy and PBBS by up to 3%, 3% and 6% respectively. It is because LazyProbe enumerates all feasible billboard sets whose cardinality is no larger than 3, and invokes a greedy-based dynamic computation function to get the result, which makes it achieve $(1 - 1/e)$ approximation ratio. In contrast, BBS and PBBS are much faster than LazyProbe, around one order of magnitude and two orders of magnitude respectively. The reason lies in the influence function. As aforementioned, $p(S, t) = 1$ if $\exists o_i \in S$, otherwise $p(S, t) = 0$. Therefore, the upper bound is equal to the lower bound. The branch-and-bound framework does not search for further, deeper branches. It will try to complete S^c with or without a billboard and terminates after finishing the first iteration. Hence, BBS is slightly better than Greedy, but worse than LazyProbe on effectiveness. Since PBBS may miss some ideal selections, it is slightly worse than the rest except Top- k .

7 CONCLUSION

We first introduced a non-submodular influence model, which is widely adopted in many areas such as consumer behaviour and advertising marketing, etc. Based on this influence model, we studied the ICOA problem and proved that it is NP-hard to approximate. More importantly, a simple cost-effective greedy method cannot work well since the logistic-based influence model is not submodular. Then, we proposed a method to compute the upper bound by using a dynamic tangent line to tightly bound the real influence. By utilizing this upper bound computation method, we built a branch-and-bound framework to solve ICOA problem. To reduce the computational cost, we further proposed a θ -termination method and a progressive bound computation algorithm. Lastly, we conducted experiments on real-world datasets to verify the efficiency, effectiveness adaptability, and scalability of our methods.

ACKNOWLEDGMENTS

Zhifeng Bao was partially supported by ARC DP170102726, DP180102050, and NSFC 61728204, 91646204, and Google Faculty Award. This research was supported by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier I grant MSS18C001.

REFERENCES

- [1] Penneco Outdoor Advertising. 2016. Billboard Statistics. <https://www.pennecooutdoor.com/billboard-statistics>
- [2] Sara Ahmadian, Zachary Friggstad, and Chaitanya Swamy. 2013. Local-Search based Approximation Algorithms for Mobile Facility Location Problems. In *SODA*. SIAM, 1607–1621.
- [3] Christoph Ambühl, Monaldo Mastrolilli, and Ola Svensson. 2011. Inapproximability Results for Maximum Edge Biclique, Minimum Linear Arrangement, and Sparsest Cut. *SIAM J. Comput.* 40, 2 (2011), 567–596.
- [4] Margaret C Campbell and Kevin Lane Keller. 2003. Brand familiarity and advertising repetition effects. *Journal of consumer research* 30, 2 (2003), 292–304.
- [5] Can Chen, Junming Liu, Qiao Li, Yijun Wang, Hui Xiong, and Shanshan Wu. 2017. Warehouse Site Selection for Online Retailers in Inter-Connected Warehouse Networks. In *ICDM*. IEEE, 805–810.
- [6] Farhana Murtaza Choudhury, J. Shane Culpepper, Zhifeng Bao, and Timos Sellis. 2018. Finding the optimal location and keywords in obstructed and unobstructed space. *Vldb J.* 27, 4 (2018), 445–470.
- [7] Gershon Feder, Richard E Just, and David Zilberman. 1985. Adoption of agricultural innovations in developing countries: A survey. *Economic development and cultural change* 33, 2 (1985), 255–298.
- [8] Gerald J Gorn and Marvin E Goldberg. 1980. Children's responses to repetitive television commercials. *Journal of Consumer Research* 6, 4 (1980), 421–424.
- [9] William H Greene. 2003. *Econometric analysis*. Pearson Education India.
- [10] Johny K Johansson. 1979. Advertising and the S-curve: A new approach. *Journal of Marketing Research* (1979), 346–354.
- [11] Samir Khuller, Anna Moss, and Joseph Naor. 1999. The Budgeted Maximum Coverage Problem. *Inf. Process. Lett.* 70, 1 (1999), 39–45.
- [12] LAMAR. 2017. National Rate Card. <http://apps.lamar.com/demographicrates/content/salesdocuments/nationalratecard.xlsx>
- [13] Sang Yup Lee. 2014. Examining the factors that influence early adopters' smartphone adoption: The case of college students. *Telematics and Informatics* 31, 2 (2014), 308–318.
- [14] Guoliang Li, Shuo Chen, Jianhua Feng, Kian-Lee Tan, and Wen-Syan Li. 2014. Efficient location-aware influence maximization. In *SIGMOD*. ACM, 87–98.
- [15] Shi Li. 2019. On Facility Location with General Lower Bounds. In *SODA*. SIAM, 2279–2290.
- [16] John DC Little. 1979. Aggregate advertising models: The state of the art. *Operations research* 27, 4 (1979), 629–667.
- [17] Dongyu Liu, Di Weng, Yuhong Li, Jie Bao, Yu Zheng, Huamin Qu, and Yingcai Wu. 2017. SmartAdP: Visual Analytics of Large-scale Taxi Trajectories for Selecting Billboard Locations. *IEEE Trans. Vis. Comput. Graph.* 23, 1 (2017), 1–10.
- [18] Yubao Liu, Raymond Chi-Wing Wong, Ke Wang, Zhijie Li, Cheng Chen, and Zitong Chen. 2013. A new approach for maximizing bichromatic reverse nearest neighbor search. *Knowl. Inf. Syst.* 36, 1 (2013), 23–58.
- [19] Prashant Malaviya. 2007. The moderating influence of advertising context on ad repetition effects: The role of amount and type of elaboration. *Journal of Consumer Research* 34, 1 (2007), 32–40.
- [20] M. Teresa Melo, Stefan Nickel, and Francisco Saldanha-da-Gama. 2006. Dynamic multi-commodity capacitated facility location: a mathematical modeling framework for strategic supply chain planning. *Computers & OR* 33 (2006), 181–208.
- [21] M. Teresa Melo, Stefan Nickel, and Francisco Saldanha-da-Gama. 2009. Facility location and supply chain management - A review. *European Journal of Operational Research* 196, 2 (2009), 401–412.
- [22] George A Miller. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review* 63, 2 (1956), 81.
- [23] Kristian S Palda. 1965. The measurement of cumulative advertising effects. *The Journal of Business* 38, 2 (1965), 162–179.
- [24] William Sierzechula, Sjoerd Bakker, Kees Maat, and Bert Van Wee. 2014. The influence of financial incentives and other socio-economic factors on electric vehicle adoption. *Energy Policy* 68 (2014), 183–194.
- [25] Julian L Simon and Johan Arndt. 1980. The shape of the advertising response function. *Journal of Advertising Research* (1980).
- [26] Jennifer Taylor, Rachel Kennedy, and Byron Sharp. 2009. Is once really enough? Making generalizations about advertising's convex sales response function. *Journal of Advertising Research* 49, 2 (2009), 198–200.
- [27] Gerard J Tellis. 1988. Advertising exposure, loyalty, and brand purchase: A two-stage model of choice. *Journal of marketing research* (1988), 134–144.
- [28] Kenneth E Train. 2009. *Discrete choice methods with simulation*. Cambridge university press.
- [29] Demetrios Vakratsas, Fred M Feinberg, Frank M Bass, and Gurumurthy Kalyanaram. 2004. The shape of advertising response functions revisited: A model of dynamic probabilistic thresholds. *Marketing Science* 23, 1 (2004), 109–119.
- [30] Sheng Wang, Zhifeng Bao, J. Shane Culpepper, Timos Sellis, and Gao Cong. 2018. Reverse k Nearest Neighbor Search over Trajectories. *IEEE Trans. Knowl. Data Eng.* 30, 4 (2018), 757–771.
- [31] Raymond Chi-Wing Wong, M. Tamer Özsu, Philip S. Yu, Ada Wai-Chee Fu, and Lian Liu. 2009. Efficient Method for Maximizing Bichromatic Reverse Nearest Neighbor. *PVLDB* 2, 1 (2009), 1126–1137.
- [32] Ping Zhang, Zhifeng Bao, Yuchen Li, Guoliang Li, Yipeng Zhang, and Zhiyong Peng. 2018. Trajectory-driven Influential Billboard Placement. In *SIGKDD*. ACM, 2748–2757.
- [33] Zenan Zhou, Wei Wu, Xiaohui Li, Mong-Li Lee, and Wynne Hsu. 2011. MaxFirst for MaxBRkNN. In *ICDE*. IEEE, 828–839.

A APPENDIX

A.1 Parameter Sensitivity Test

In this part, we conduct the sensitivity experiments of our methods by varying parameters. We only vary one parameter and keep the rest by default setting as shown in Table 4. The experimental results will be explained, and we will describe how we choose the default settings.

A.1.1 Varying β/α

As we mentioned in Section 3.2, α and β are the parameters that control the turning point for the influence. The bigger the β/α is, the more times of impression are needed to change a user's adoption. With the increase of β , the influence from one billboard to one user becomes larger. A huge β leads to a sudden change of a user's adoption, whereas a small β makes the changing process smooth.

Moreover, studies show that there is a high chance that a person will change her adoption action after noticing a message more than four times [4, 8, 19, 22]. Therefore, we make the parameter setting satisfy the above assumption, which means that, when a trajectory has been influenced more than four times, the influence is larger than 0.95 (i.e., $p(S, t) > 0.95$ when $\sum_{o_i \in S} I(o_i, t) \geq 4$). Hence, we fix $\beta = 3$ and vary α from 7 to 11. Based on Equation 1, β/α rises when α drops, which increases the influence from o to t . It also leads to higher overall influence.

Figure 9 shows the experimental results of varying β/α . We find that with the increase of β/α , our solutions outperform baselines by 50% to 100%. Note that, varying β/α does not affect the efficiency of all methods. Therefore, we choose $\alpha = 7$ and $\beta = 3$ as the default setting since our solutions have the smallest advantage of effectiveness for the setting.

A.1.2 Varying ϵ

ϵ is used to adjust the step distance of decreasing threshold h in Algorithm 3. Figure 10 shows the experimental results of varying ϵ . When ϵ increases from 10^{-4} to 1, our solutions decrease by at most 9% in effectiveness, but speed up by two orders of magnitude. We find that when $\epsilon > 10^{-1}$, the changing of effectiveness and efficiency tends to be stable. Therefore, we choose $\epsilon = 10^{-1}$ as the default setting.

A.1.3 Varying θ

Recall Section 4.1, θ is the parameter controlling the termination condition. A smaller θ leads to better efficiency with worse effectiveness. Figure 11 shows the experimental results of varying θ . When ϵ decreases from 0.94 to 0.86, the effectiveness of our solutions decreases by at most 6%, while the performance is improved at around one order of magnitude speed ups. Therefore, we choose the median value 0.9 as the default setting, because it reaches an ideal balance of efficiency and effectiveness.

A.1.4 Varying λ

Figure 12 shows the influence of changing λ on the effectiveness and efficiency. Recall Section 3.1, λ is a given distance threshold. A user is impressed by a billboard only if the distance between them is smaller than λ (i.e., $\text{dist}(t.p_i, o.loc) \leq \lambda$). We have the following observations. First, With the increase of λ , the effectiveness

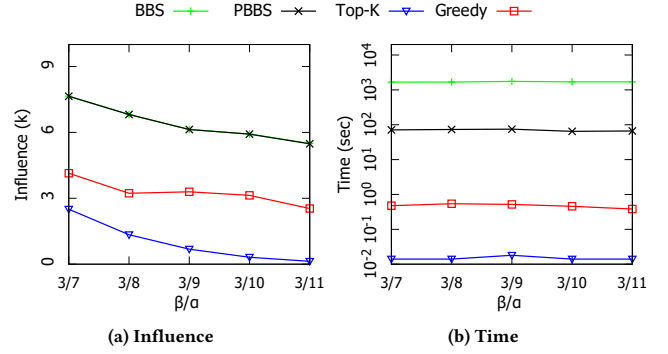


Figure 9: Varying β/α in NYC




Figure 10: Varying ϵ in NYC




Figure 11: Varying θ in NYC




Figure 12: Varying λ in NYC

of all methods improves. The growing effectiveness is mainly contributed from a growing number of trajectories that the billboards can impress. Second, the advantage of our methods decreases, since overlaps become easier with larger λ . Last, except for Top- k , the running-time of all methods increases. The reason for this is similar to that of the first observation.




Figure 13: Test on different cost setting strategies




Figure 14: Test on different cost setting strategies

A.2 Different cost setting strategies

In this part, we study the effect of changing strategies of setting cost. DC is the default strategy as we mentioned in Section 6.2. WI is based on the weekly impression; that is, the impression statistic data from LAMAR [12]. Similar to DC, we map the weekly impression to the cost interval (i.e., \$2,000 to \$60,000). TC1 and TC2 are the trajectory-based strategies of setting cost. Under TC1, we first calculate the square of the number of trajectories impressed by this billboard, and then map this value to the cost interval. TC2 is similar to TC1, the difference is that we use the square root of the number of trajectories impressed by this billboard. As shown in Figure 13a, we make three observations. First, in WI, the effectiveness of all methods decreases. This is because in the LAMAR database, billboards that impress a large number of trajectories have a very high weekly impression. It leads to the consequence that, most billboards have a higher cost compared with WI. Therefore, all methods afford fewer billboards. Second, in TC1, the performance of Top- k is worse compared with itself in other strategies since for all billboards, the higher the cost, the worse the effectiveness. Last, in TC2, Greedy is only worse than BBS by 8%. The reason is that, billboards that have the higher impression is more cost-efficient. Therefore, both Greedy and our methods achieve an ideal result easily. As shown in Figure 13b, the different cost strategies do not affect efficiency.

A.3 Number of Influenced Audiences

In this part, instead of calculating the total influence by Equation 2, we study how many audiences are influenced by S . We define an audience t to be influenced if $p(S, t) \geq \varphi$, where $\varphi \in (0, 1]$. In this case, we assume $\varphi = 0.8$, which means that an audience will be influenced after meeting at least three billboards. Figure 14 shows the experimental result. We have two observations. First, BBS and PBBS have similar effectiveness compared with themselves in Figure

5a. Second, the effectiveness of Greedy and Top- k is worse than themselves in Figure 5a. The reason is that, based on the influence model, the marginal influence decreases when $p(S, t) \geq 0.8$. Before reaching the turning point of the influence model, our methods prefer to impress the same trajectory multiple times. Therefore, the influence areas of our methods are more concentrated, whereas that of Greedy and Top- k are more dispersed.

A.4 Proof of Theorem 3.1

The NP-hardness of ICOA can be easily inferred by reducing the set cover problem to ICOA. Next, we show it is in fact NP-hard to approximate ICOA within any constant factor. To prove the theorem, we reduce the biclique detection (BD) problem to ICOA. Given a bipartite graph $G = (U \cup V, E)$ where U and V denote the vertex sets while E denotes the edge set containing the edges $(u, v) \in E$ such that $u \in U$ and $v \in V$; an instance of BD asks whether there exists vertex subsets $U' \subseteq U$, $V' \subseteq V$ with $|U'| \geq k$, $|V'| \geq k'$ and the induced subgraph of U' and V' is a biclique. The decision version of BD is equivalent to an optimization problem where we fix $|U'| = k$ and ask what is the maximum $|V'|$ s.t. the induced subgraph is a biclique. We reduce the optimization version of BD to ICOA with the following process: (1) $\forall u \in U$ of BD, we create a billboard x_u with a uniform cost of 1; (2) $\forall v \in V$ of BD, we create a trajectory x_v as the trajectory set \mathcal{T} ; (3) $\forall (u, v) \in E$ of BD, we make sure x_u is close to x_v so that $\text{dist}(x_u, x_v) \leq \lambda$; (4) We set the budget of ICOA as k , which is the same as the input of the optimization version of BD; (5) We set $\alpha = k \cdot |\mathcal{T}|$ and $\beta = |\mathcal{T}|$. Clearly, the reduction can be done in polynomial time. We establish the following relationship between BD and ICOA.

LEMMA A.1. *Let OPT_B and OPT_I denote the optimal values of BD and the instance of ICOA that BD is reduced to, respectively. We have $2 \cdot \text{OPT}_I - \frac{1}{|\mathcal{T}|} \leq \text{OPT}_B \leq 2 \cdot \text{OPT}_I$.*

PROOF. We show the second inequality first. Let U' such that $|U'| = k$ represents the vertex set selected to achieve the largest biclique in G w.r.t. the size of $|V'|$. We choose the corresponding $x_u \forall u \in U'$ as the selected billboards. The influence achieved is thus lower bounded by $\text{OPT}_B/2$. We can then conclude $\text{OPT}_I \geq \text{OPT}_B/2$.

We now show the first inequality. Suppose S is selected such that $|S| = k$ as the billboards to place the ad on, we identify a trajectory set $\mathcal{T}_S \in \mathcal{T}$ where all trajectories in \mathcal{T}_S meet all billboards in S . The following inequality holds:

$$I(S) = \sum_{t \in \mathcal{T}_S} p(S, t) + \sum_{t \in \mathcal{T} \setminus \mathcal{T}_S} p(S, t) \leq \frac{\text{OPT}_B}{2} + \frac{|\mathcal{T}| - |\mathcal{T}_S|}{1 + e^{|\mathcal{T}|}} \leq \frac{\text{OPT}_B}{2} + \frac{1}{2|\mathcal{T}|}$$

which completes the proof for the lemma. \square

Finally, we are ready to prove the hardness of ICOA by a gap preserving reduction: let a constant $\theta \geq 1$, if $\text{OPT}_B \geq \theta$ then $\text{OPT}_I \geq \theta/2$, and if $\text{OPT}_I \leq \theta/|\mathcal{T}|^{1-\epsilon}$ for any $\epsilon > 0$, then $\text{OPT}_B \leq \frac{1}{\frac{1}{2}|\mathcal{T}|^{1-\epsilon}} \cdot \frac{\theta}{2}$. Thus, if one can approximate ICOA in polynomial time within $\frac{1}{2}|\mathcal{T}|^{1-\epsilon}$ factor, then BD can be approximated within a factor of $|\mathcal{T}|^{1-\epsilon}$, which leads to a contradiction [3].