
Uplift Modeling from Separate Labels

Ikko Yamane^{1,2} Florian Yger^{3,2} Jamal Atif³ Masashi Sugiyama^{2,1}

¹ The University of Tokyo, CHIBA, JAPAN

² RIKEN Center for Advanced Intelligence Project (AIP), TOKYO, JAPAN

³ LAMSADE, CNRS, Université Paris-Dauphine, Université PSL, PARIS, FRANCE

{yamane@ms., sugi@}k.u-tokyo.ac.jp, {florian.yger@, jamal.atif@}dauphine.fr

Abstract

Uplift modeling is aimed at estimating the incremental impact of an action on an individual's behavior, which is useful in various application domains such as targeted marketing (advertisement campaigns) and personalized medicine (medical treatments). Conventional methods of uplift modeling require every instance to be *jointly* equipped with two types of labels: the taken action and its outcome. However, obtaining two labels for each instance at the same time is difficult or expensive in many real-world problems. In this paper, we propose a novel method of uplift modeling that is applicable to a more practical setting where only one type of labels is available for each instance. We show a mean squared error bound for the proposed estimator and demonstrate its effectiveness through experiments.

1 Introduction

In many real-world problems, a central objective is to optimally choose a right action to maximize the profit of interest. For example, in marketing, an advertising campaign is designed to promote people to purchase a product [29]. A marketer can choose whether to deliver an advertisement to each individual or not, and the outcome is the number of purchases of the product. Another example is personalized medicine, where a treatment is chosen depending on each patient to maximize the medical effect and minimize the risk of adverse events or harmful side effects [1, 13]. In this case, giving or not giving a medical treatment to each individual are the possible actions to choose, and the outcome is the rate of recovery or survival from the disease. Hereafter, we use the word *treatment* for taking an action, following the personalized medicine example.

A/B testing [14] is a standard method for such tasks, where two groups of people, A and B, are randomly chosen. The outcomes are measured separately from the two groups after treating all the members of Group A but none of Group B. By comparing the outcomes between the two groups by a statistical test, one can examine whether the treatment positively or negatively affected the outcome. However, *A/B testing* only compares the two extreme options: treating everyone or no one. These two options can be both far from optimal when the treatment has positive effect on some individuals but negative effect on others.

To overcome the drawback of *A/B testing*, *uplift modeling* has been investigated recently [11, 28, 32]. *Uplift modeling* is the problem of estimating the *individual uplift*, the incremental profit brought by the treatment conditioned on features of each individual. Uplift modeling enables us to design a refined decision rule for optimally determining whether to treat each individual or not, depending on his/her features. Such a treatment rule allows us to only target those who positively respond to the treatment and avoid treating negative responders.

In the standard uplift modeling setup, there are two types of labels [11, 28, 32]: One is whether the treatment has been given to the individual and the other is its outcome. Existing uplift modeling methods require each individual to be *jointly* given these two labels for analyzing the association

between outcomes and the treatment [11, 28, 32]. However, joint labels are expensive or hard (or even impossible) to obtain in many real-world problems. For example, when distributing an advertisement by email, we can easily record to whom the advertisement has been sent. However, for technical or privacy reasons, it is difficult to keep track of those people until we observe the outcomes on whether they buy the product or not. Alternatively, we can easily obtain information about purchasers of the product at the moment when the purchases are actually made. However, we cannot know whether those who are buying the product have been exposed to the advertisement or not. Thus, every individual always has one missing label. We term such samples *separately labeled samples*.

In this paper, we consider a more practical uplift modeling setup where no jointly labeled samples are available, but only separately labeled samples are given. Theoretically, we first show that the individual uplift is identifiable when we have two sets of separately labeled samples collected under *different* treatment policies. We then propose a novel method that directly estimates the individual uplift only from separately labeled samples. Finally, we demonstrate the effectiveness of the proposed method through experiments.

2 Problem Setting

This paper focuses on estimation of the *individual uplift* $u(\mathbf{x})$, often called *individual treatment effect (ITE)* in the causal inference literature [31], defined as $u(\mathbf{x}) := \mathbf{E}[Y_1 | \mathbf{x}] - \mathbf{E}[Y_{-1} | \mathbf{x}]$, where $\mathbf{E}[\cdot | \cdot]$ denotes the conditional expectation, and \mathbf{x} is a \mathcal{X} -valued random variable ($\mathcal{X} \subseteq \mathbb{R}^d$) representing features of an individual, and Y_1, Y_{-1} are \mathcal{Y} -valued *potential outcome variables* [31] ($\mathcal{Y} \subseteq \mathbb{R}$) representing outcomes that would be observed if the individual was treated and not treated, respectively. Note that only one of either Y_1 or Y_{-1} can be observed for each individual. We denote the $\{1, -1\}$ -valued random variable of the treatment assignment by t , where $t = 1$ means that the individual has been treated and $t = -1$ not treated. We refer to the population for which we want to evaluate $u(\mathbf{x})$ as the *test population*, and denote the density of the test population by $p(Y_1, Y_{-1}, \mathbf{x}, t)$.

We assume that t is *unconfounded* with either of Y_1 and Y_{-1} conditioned on \mathbf{x} , i.e. $p(Y_1 | \mathbf{x}, t) = p(Y_1 | \mathbf{x})$ and $p(Y_{-1} | \mathbf{x}, t) = p(Y_{-1} | \mathbf{x})$. Unconfoundedness is an assumption commonly made in observational studies [5, 33]. For notational convenience, we denote by $y := Y_t$ the outcome of the treatment assignment t . Furthermore, we refer to any conditional density of t given \mathbf{x} as a *treatment policy*.

In addition to the test population, we suppose that there are two *training populations* $k = 1, 2$, whose joint probability density $p_k(Y_1, Y_{-1}, \mathbf{x}, t)$ satisfy

$$p_k(Y_{t_0} = y_0 | \mathbf{x} = \mathbf{x}_0) = p(Y_{t_0} = y_0 | \mathbf{x} = \mathbf{x}_0) \quad (\text{for } k = 1, 2), \quad (1)$$

$$p_1(t = t_0 | \mathbf{x} = \mathbf{x}_0) \neq p_2(t = t_0 | \mathbf{x} = \mathbf{x}_0), \quad (2)$$

for all possible realizations $\mathbf{x}_0 \in \mathcal{X}$, $t_0 \in \{-1, 1\}$, and $y_0 \in \mathcal{Y}$. Intuitively, Eq. (1) means that potential outcomes depend on \mathbf{x} in the same way as those in the test population, and Eq. (2) states that those two policies give a treatment with different probabilities for every $\mathbf{x} = \mathbf{x}_0$.

We suppose that the following four training data sets, which we call *separately labeled samples*, are given:

$$\{(\mathbf{x}_i^{(k)}, y_i^{(k)})\}_{i=1}^{n_k} \stackrel{\text{i.i.d.}}{\sim} p_k(\mathbf{x}, y), \quad \{(\tilde{\mathbf{x}}_i^{(k)}, t_i^{(k)})\}_{i=1}^{\tilde{n}_k} \stackrel{\text{i.i.d.}}{\sim} p_k(\mathbf{x}, t) \quad (\text{for } k = 1, 2),$$

where n_k and \tilde{n}_k , $k = 1, 2$, are positive integers. Under Assumptions (1), (2), and the unconfoundedness, we have $p_k(Y_t | \mathbf{x}, t = t_0) = p(Y_{t_0} | \mathbf{x}, t = t_0) = p(Y_{t_0} | \mathbf{x})$ for $t_0 \in \{-1, 1\}$ and $k \in \{1, 2\}$. Note that we can safely denote $p(y | \mathbf{x}, t) := p_k(y | \mathbf{x}, t)$. Moreover, we have $\mathbf{E}[Y_{t_0} | \mathbf{x}] = \mathbf{E}[y | \mathbf{x}, t = t_0]$ for $t_0 = 1, -1$, and thus our goal boils down to the estimation of

$$u(\mathbf{x}) = \mathbf{E}[y | \mathbf{x}, t = 1] - \mathbf{E}[y | \mathbf{x}, t = -1] \quad (3)$$

from the separately labeled samples, where the conditional expectation is taken over $p(y | \mathbf{x}, t)$.

Estimation of the individual uplift is important for the following reasons.

It enables the estimation of the average uplift. The *average uplift* $U(\pi)$ of the treatment policy $\pi(t | \mathbf{x})$ is the average outcome of π , subtracted by that of the policy π_- , which constantly assigns

the treatment as $t = -1$, i.e., $\pi_-(t = \tau | \mathbf{x}) := 1[\tau = -1]$, where $1[\cdot]$ denotes the indicator function:

$$\begin{aligned} U(\pi) &:= \iint \sum_{t=-1,1} yp(y | \mathbf{x}, t)\pi(t | \mathbf{x})p(\mathbf{x})dyd\mathbf{x} - \iint \sum_{t=-1,1} yp(y | \mathbf{x}, t)\pi_-(t | \mathbf{x})p(\mathbf{x})dyd\mathbf{x} \\ &= \int u(\mathbf{x})\pi(t = 1 | \mathbf{x})p(\mathbf{x})d\mathbf{x}. \end{aligned} \quad (4)$$

This quantity can be estimated from samples of \mathbf{x} once we obtain an estimate of $u(\mathbf{x})$.

It provides the optimal treatment policy. The treatment policy given by $\pi(t = 1 | \mathbf{x}) = 1[0 \leq u(\mathbf{x})]$ is the optimal treatment that maximizes the average uplift $U(\pi)$ and equivalently the average outcome $\iint \sum_{t=-1,1} yp(y | \mathbf{x}, t)\pi(t | \mathbf{x})p(\mathbf{x})dyd\mathbf{x}$ (see Eq. (4)) [32].

It is the optimal ranking scoring function. From a practical viewpoint, it may be useful to prioritize individuals to be treated according to some ranking scores especially when the treatment is costly and only a limited number of individuals can be treated due to some budget constraint. In fact, $u(\mathbf{x})$ serves as the optimal ranking scores for this purpose [36]. More specifically, we define a family of treatment policies $\{\pi_{f,\alpha}\}_{\alpha \in \mathbb{R}}$ associated with scoring function f by $\pi_{f,\alpha}(t = 1 | \mathbf{x}) = 1[\alpha \leq f(\mathbf{x})]$. Then, under some technical condition, $f = u$ maximizes the *area under the uplift curve (AUUC)* defined as

$$\begin{aligned} \text{AUUC}(f) &:= \int_0^1 U(\pi_{f,\alpha})dC_\alpha \\ &= \int_0^1 \int u(\mathbf{x})1[\alpha \leq f(\mathbf{x})]p(\mathbf{x})d\mathbf{x}dC_\alpha \\ &= \mathbf{E}[1[f(\mathbf{x}) \leq f(\mathbf{x}')]u(\mathbf{x}')], \end{aligned}$$

where $C_\alpha := \Pr[f(\mathbf{x}) < \alpha]$, $\mathbf{x}, \mathbf{x}' \stackrel{\text{i.i.d.}}{\sim} p(\mathbf{x})$, and \mathbf{E} denotes the expectation with respect to these variables. AUUC is a standard performance measure for uplift modeling methods [11, 25, 28, 32]. For more details, see Appendix B in the supplementary material.

Remark on the problem setting: Uplift modeling is often referred to as individual treatment effect estimation or heterogeneous treatment effect estimation and has been extensively studied especially in the causal inference literature [5, 7, 9, 12, 16, 24, 31, 37]. In particular, recent research has investigated the problem under the setting of *observational studies*, inference using data obtained from *uncontrolled experiments* because of its practical importance [33]. Here, experiments are said to be uncontrolled when some of treatment variables are not controlled to have designed values.

Given that treatment policies are unknown, our problem setting is also of observational studies but poses an additional challenge that stems from missing labels. What makes our problem feasible is that we have two kinds of data sets following different treatment policies.

It is also important to note that our setting generalizes the standard setting for observational studies since the former is reduced to the latter when one of the treatment policies always assigns individuals to the treatment group, and the other to the control group.

Our problem is also closely related to individual treatment effect estimation via instrumental variables [2, 6, 10, 19].¹

3 Naive Estimators

A naive approach is first estimating the conditional density $p_k(y | \mathbf{x})$ and $p_k(t | \mathbf{x})$ from training samples by some conditional density estimator [4, 34], and then solving the following linear system for $p(y | \mathbf{x}, t = 1)$ and $p(y | \mathbf{x}, t = -1)$:

$$\underbrace{p_k(y | \mathbf{x})}_{\text{Estimated from } \{(\mathbf{x}_i^{(k)}, y_i^{(k)})\}_{i=1}^n} = \sum_{t=-1,1} p(y | \mathbf{x}, t) \underbrace{p_k(t | \mathbf{x})}_{\text{Estimated from } \{(\tilde{\mathbf{x}}_i^{(k)}, t_i^{(k)})\}_{i=1}^{\tilde{n}}} \quad (\text{for } k = 1, 2). \quad (5)$$

¹Among the related papers mentioned above, the most relevant one is Lewis and Syrgkanis [19], which is concurrent work with ours.

After that, the conditional expectations of y over $p(y | \mathbf{x}, t = 1)$ and $p(y | \mathbf{x}, t = -1)$ are calculated by numerical integration, and finally their difference is calculated to obtain another estimate of $u(\mathbf{x})$.

However, this may not yield a good estimate due to the difficulty of conditional density estimation and the instability of numerical integration. This issue may be alleviated by working on the following linear system implied by Eq. (5) instead: $\mathbf{E}_k[y | \mathbf{x}] = \sum_{t=-1,1} \mathbf{E}[y | \mathbf{x}, t] p_k(t | \mathbf{x})$, $k = 1, 2$, where $\mathbf{E}_k[y | \mathbf{x}]$ and $p_k(t | \mathbf{x})$ can be estimated from our samples. Solving this new system for $\mathbf{E}[y | \mathbf{x}, t = 1]$ and $\mathbf{E}[y | \mathbf{x}, t = -1]$ and taking their difference gives an estimate of $u(\mathbf{x})$. A method called *two-stage least-squares* for instrumental variable regression takes such an approach [10].

The second approach of estimation $E_k[y|x]$ and $p_k(t|x)$ avoids both conditional density estimation and numerical integration, but it still involves post processing of solving the linear system and subtraction, being a potential cause of performance deterioration.

4 Proposed Method

In this section, we develop a method that can overcome the aforementioned problems by directly estimating the individual uplift.

4.1 Direct Least-Square Estimation of the Individual Uplift

First, we will show an important lemma that directly relates the marginal distributions of separately labeled samples to the individual uplift $u(\mathbf{x})$.

Lemma 1. *For every \mathbf{x} such that $p_1(\mathbf{x}) \neq p_2(\mathbf{x})$, $u(\mathbf{x})$ can be expressed as*

$$u(\mathbf{x}) = 2 \times \frac{\mathbf{E}_{y \sim p_1(y|\mathbf{x})}[y] - \mathbf{E}_{y \sim p_2(y|\mathbf{x})}[y]}{\mathbf{E}_{t \sim p_1(t|\mathbf{x})}[t] - \mathbf{E}_{t \sim p_2(t|\mathbf{x})}[t]}. \quad (6)$$

For a proof, refer to Appendix C in the supplementary material.

Using Eq. (6), we can re-interpret the naive methods described in Section 3 as estimating the conditional expectations on the right-hand side by separately performing regression on $\{(\mathbf{x}_i^{(1)}, y_i^{(1)})\}_{i=1}^{n_1}$, $\{(\mathbf{x}_i^{(2)}, y_i^{(2)})\}_{i=1}^{n_2}$, $\{(\tilde{\mathbf{x}}_i^{(1)}, t_i^{(1)})\}_{i=1}^{\tilde{n}_1}$, and $\{(\tilde{\mathbf{x}}_i^{(2)}, t_i^{(2)})\}_{i=1}^{\tilde{n}_2}$. This approach may result in unreliable performance when the denominator is close to zero, i.e., $p_1(t | \mathbf{x}) \simeq p_2(t | \mathbf{x})$.

Lemma 1 can be simplified by introducing auxiliary variables z and w , which are \mathcal{Z} -valued and $\{-1, 1\}$ -valued random variables whose conditional probability density and mass are defined by

$$\begin{aligned} p(z = z_0 | \mathbf{x}) &= \frac{1}{2} p_1(y = z_0 | \mathbf{x}) + \frac{1}{2} p_2(y = -z_0 | \mathbf{x}), \\ p(w = w_0 | \mathbf{x}) &= \frac{1}{2} p_1(t = w_0 | \mathbf{x}) + \frac{1}{2} p_2(t = -w_0 | \mathbf{x}), \end{aligned}$$

for any $z_0 \in \mathcal{Z}$ and any $w_0 \in \{-1, 1\}$, where $\mathcal{Z} := \{s_0 y_0 | y_0 \in \mathcal{Y}, s_0 \in \{1, -1\}\}$.

Lemma 2. *For every \mathbf{x} such that $p_1(\mathbf{x}) \neq p_2(\mathbf{x})$, $u(\mathbf{x})$ can be expressed as*

$$u(\mathbf{x}) = 2 \times \frac{\mathbf{E}[z | \mathbf{x}]}{\mathbf{E}[w | \mathbf{x}]},$$

where $\mathbf{E}[z | \mathbf{x}]$ and $\mathbf{E}[w | \mathbf{x}]$ are the conditional expectations of z given \mathbf{x} over $p(z | \mathbf{x})$ and w given \mathbf{x} over $p(w | \mathbf{x})$, respectively.

A proof can be found in Appendix D in the supplementary material.

Let $w_i^{(k)} := (-1)^{k-1} t_i^{(k)}$ and $z_i^{(k)} := (-1)^{k-1} y_i^{(k)}$. Assuming that $p_1(\mathbf{x}) = p_2(\mathbf{x}) =: p(\mathbf{x})$, $n_1 = n_2$, and $\tilde{n}_1 = \tilde{n}_2$ for simplicity, $\{(\tilde{\mathbf{x}}_i, w_i)\}_{i=1}^{\tilde{n}} := \{(\tilde{\mathbf{x}}_i^{(k)}, w_i^{(k)})\}_{k=1,2; i=1, \dots, \tilde{n}_k}$ and $\{(\mathbf{x}_i, z_i)\}_{i=1}^n := \{(\mathbf{x}_i^{(k)}, z_i^{(k)})\}_{k=1,2; i=1, \dots, n_k}$ can be seen as samples drawn from $p(\mathbf{x}, z) := p(z | \mathbf{x}) p(\mathbf{x})$ and $p(\mathbf{x}, w) := p(w | \mathbf{x}) p(\mathbf{x})$, respectively, where $n = n_1 + n_2$ and $\tilde{n} = \tilde{n}_1 + \tilde{n}_2$. The more general cases where $p_1(\mathbf{x}) \neq p_2(\mathbf{x})$, $n_1 \neq n_2$, or $\tilde{n}_1 \neq \tilde{n}_2$ are discussed in Appendix I in the supplementary material.

Theorem 1. Assume that $\mu_w, \mu_z \in L^2(p)$ and $\mu_w(\mathbf{x}) \neq 0$ for every \mathbf{x} such that $p(\mathbf{x}) > 0$, where $L^2(p) := \{f : \mathcal{X} \rightarrow \mathbb{R} \mid \mathbf{E}_{\mathbf{x} \sim p(\mathbf{x})}[f(\mathbf{x})^2] < \infty\}$. The individual uplift $u(\mathbf{x})$ equals the solution to the following least-squares problem:

$$u(\mathbf{x}) = \operatorname{argmin}_{f \in L^2(p)} \mathbf{E}[(\mu_w(\mathbf{x})f(\mathbf{x}) - 2\mu_z(\mathbf{x}))^2], \quad (7)$$

where \mathbf{E} denotes the expectation over $p(\mathbf{x})$, $\mu_w(\mathbf{x}) := \mathbf{E}[w \mid \mathbf{x}]$, and $\mu_z(\mathbf{x}) := \mathbf{E}[z \mid \mathbf{x}]$.

Theorem 1 follows from Lemma 2. Note that $p_1(\mathbf{x}) \neq p_2(\mathbf{x})$ in Eq. (2) implies $\mu_w(\mathbf{x}) \neq 0$.

In what follows, we develop a method that directly estimates $u(\mathbf{x})$ by solving Eq. (7). A challenge here is that it is not straightforward to evaluate the objective functional since it involves unknown functions, μ_w and μ_z .

4.2 Disentanglement of z and w

Our idea is to transform the objective functional in Eq. (7) into another form in which $\mu_w(\mathbf{x})$ and $\mu_z(\mathbf{x})$ appear separately and linearly inside the expectation operator so that we can approximate them using our separately labeled samples.

For any function $g \in L^2(p)$ and any $\mathbf{x} \in \mathcal{X}$, expanding the left-hand side of the inequality $\mathbf{E}[(\mu_w(\mathbf{x})f(\mathbf{x}) - 2\mu_z(\mathbf{x}) - g(\mathbf{x}))^2] \geq 0$, we have

$$\mathbf{E}[(\mu_w(\mathbf{x})f(\mathbf{x}) - 2\mu_z(\mathbf{x}))^2] \geq 2\mathbf{E}[(\mu_w(\mathbf{x})f(\mathbf{x}) - 2\mu_z(\mathbf{x}))g(\mathbf{x})] - \mathbf{E}[g(\mathbf{x})^2] =: J(f, g). \quad (8)$$

The equality is attained when $g(\mathbf{x}) = \mu_w(\mathbf{x})f(\mathbf{x}) - 2\mu_z(\mathbf{x})$ for any fixed f . This means that the objective functional of Eq. (7) can be calculated by maximizing $J(f, g)$ with respect to g . Hence,

$$u(\mathbf{x}) = \operatorname{argmin}_{f \in L^2(p)} \max_{g \in L^2(p)} J(f, g). \quad (9)$$

Furthermore, μ_w and μ_z are separately and linearly included in $J(f, g)$, which makes it possible to write it in terms of z and w as

$$J(f, g) = 2\mathbf{E}[wf(\mathbf{x})g(\mathbf{x})] - 4\mathbf{E}[zg(\mathbf{x})] - \mathbf{E}[g(\mathbf{x})^2]. \quad (10)$$

Unlike the original objective functional in Eq. (7), $J(f, g)$ can be easily estimated using sample averages by

$$\widehat{J}(f, g) = \frac{2}{\tilde{n}} \sum_{i=1}^{\tilde{n}} w_i f(\tilde{\mathbf{x}}_i) g(\tilde{\mathbf{x}}_i) - \frac{4}{n} \sum_{i=1}^n z_i g(\mathbf{x}_i) - \frac{1}{2n} \sum_{i=1}^n g(\mathbf{x}_i)^2 - \frac{1}{2\tilde{n}} \sum_{i=1}^{\tilde{n}} g(\tilde{\mathbf{x}}_i)^2. \quad (11)$$

In practice, we solve the following regularized empirical optimization problem:

$$\min_{f \in F} \max_{g \in G} \widehat{J}(f, g) + \Omega(f, g), \quad (12)$$

where F, G are models for f, g respectively, and $\Omega(f, g)$ is some regularizer.

An advantage of the proposed framework is that it is model-independent, and any models can be trained by optimizing the above objective.

The function g can be interpreted as a *critic* of f as follows. Minimizing Eq. (10) with respect to f is equivalent to minimizing $J(f, g) = \mathbf{E}[g(\mathbf{x})\{\mu_w(\mathbf{x})f(\mathbf{x}) - 2\mu_z(\mathbf{x})\}]$. $g(\mathbf{x})$ serves as a good critic of $f(\mathbf{x})$ when it makes the cost $g(\mathbf{x})\{\mu_w(\mathbf{x})f(\mathbf{x}) - 2\mu_z(\mathbf{x})\}$ larger for \mathbf{x} at which f makes a larger error $|\mu_w(\mathbf{x})f(\mathbf{x}) - 2\mu_z(\mathbf{x})|$. In particular, g maximizes the objective above when $g(\mathbf{x}) = \mu_w(\mathbf{x})f(\mathbf{x}) - 2\mu_z(\mathbf{x})$ for any f , and the maximum coincides with the least-squares objective in Eq. (7).

Suppose that F and G are linear-in-parameter models: $F = \{f_\alpha : \mathbf{x} \mapsto \alpha^\top \phi(\mathbf{x}) \mid \alpha \in \mathbb{R}^{b_f}\}$ and $G = \{g_\beta : \mathbf{x} \mapsto \beta^\top \psi(\mathbf{x}) \mid \beta \in \mathbb{R}^{b_g}\}$, where ϕ and ψ are b_f -dimensional and b_g -dimensional vectors of basis functions in $L^2(p)$. Then, $\widehat{J}(f_\alpha, g_\beta) = 2\alpha^\top \mathbf{A}\beta - 4\mathbf{b}^\top \beta - \beta^\top \mathbf{C}\beta$, where

$$\begin{aligned} \mathbf{A} &:= \frac{1}{\tilde{n}} \sum_{i=1}^{\tilde{n}} w_i \phi(\tilde{\mathbf{x}}_i) \psi(\tilde{\mathbf{x}}_i)^\top, & \mathbf{b} &:= \frac{1}{n} \sum_{i=1}^n z_i \psi(\mathbf{x}_i), \\ \mathbf{C} &:= \frac{1}{2n} \sum_{i=1}^n \psi(\mathbf{x}_i) \psi(\mathbf{x}_i)^\top + \frac{1}{2\tilde{n}} \sum_{i=1}^{\tilde{n}} \psi(\tilde{\mathbf{x}}_i) \psi(\tilde{\mathbf{x}}_i)^\top. \end{aligned}$$

Using ℓ_2 -regularizers, $\Omega(f, g) = \lambda_f \alpha^\top \alpha - \lambda_g \beta^\top \beta$ with some positive constants λ_f and λ_g , the solution to the inner maximization problem can be obtained in the following analytical form:

$$\hat{\beta}_\alpha := \operatorname{argmax}_\beta \hat{J}(f_\alpha, g_\beta) = \tilde{\mathbf{C}}^{-1}(\mathbf{A}^\top \alpha - 2\mathbf{b}),$$

where $\tilde{\mathbf{C}} = \mathbf{C} + \lambda_g \mathbf{I}_{b_g}$ and \mathbf{I}_{b_g} is the b_g -by- b_g identity matrix. Then, we can obtain the solution to Eq. (12) analytically as

$$\hat{\alpha} := \operatorname{argmin}_\alpha \hat{J}(f_\alpha, g_{\hat{\beta}_\alpha}) = 2(\mathbf{A} \tilde{\mathbf{C}}^{-1} \mathbf{A}^\top + \lambda_f \mathbf{I}_{b_g})^{-1} \mathbf{A} \tilde{\mathbf{C}}^{-1} \mathbf{b}.$$

Finally, from Eq. (7), our estimate of $u(\mathbf{x})$ is given as $\hat{\alpha}^\top \phi(\mathbf{x})$.

Remark on model selection: Model selection for F and G is not straightforward since the test performance measure cannot be directly evaluated with (held out) training data of our problem. Instead, we may evaluate the value of $J(\hat{f}, \hat{g})$, where $(\hat{f}, \hat{g}) \in F \times G$ is the optimal solution pair to $\min_{f \in F} \max_{g \in G} \hat{J}(f, g)$. However, it is still nontrivial to tell if the objective value is small because the solution is good in terms of the outer minimization, or because it is poor in terms of the inner maximization. We leave this issue for future work.

5 Theoretical Analysis

A theoretically appealing property of the proposed method is that its objective consists of simple sample averages. This enables us to establish a generalization error bound in terms of the Rademacher complexity [15, 22].

Denote $\varepsilon_G(f) := \sup_{g \in L^2(p)} J(f, g) - \sup_{g \in G} J(f, g)$. Also, let $\mathfrak{R}_q^N(H)$ denote the *Rademacher complexity* of a set of functions H over N random variables following probability density q (refer to Appendix E for the definition). Proofs of the following theorems and corollary can be found in Appendix E, Appendix F, and Appendix G in the supplementary material.

Theorem 2. *Assume that $n_1 = n_2$, $\tilde{n}_1 = \tilde{n}_2$, $p_1(\mathbf{x}) = p_2(\mathbf{x})$, $W := \inf_{\mathbf{x} \in \mathcal{X}} |\mu_w(\mathbf{x})| > 0$, $M_Z := \sup_{z \in \mathcal{Z}} |z| < \infty$, $M_F := \sup_{f \in F, \mathbf{x} \in \mathcal{X}} |f(\mathbf{x})| < \infty$, and $M_G := \sup_{g \in G, \mathbf{x} \in \mathcal{X}} |g(\mathbf{x})| < \infty$. Then, the following holds with probability at least $1 - \delta$ for every $f \in F$:*

$$\mathbf{E}_{\mathbf{x} \sim p(\mathbf{x})} [(f(\mathbf{x}) - u(\mathbf{x}))^2] \leq \frac{1}{W^2} \left[\sup_{g \in G} \hat{J}(f, g) + \mathcal{R}_{F, G}^{n, \tilde{n}} + \left(\frac{M_z}{\sqrt{2n}} + \frac{M_w}{\sqrt{2\tilde{n}}} \right) \sqrt{\log \frac{2}{\delta}} + \varepsilon_G(f) \right],$$

where $M_z := 4M_y M_G + M_G^2/2$, $M_w = 2M_F M_G + M_G^2/2$, and $\mathcal{R}_{F, G}^{n, \tilde{n}} := 2(M_F + 4M_Z) \mathfrak{R}_{p(\mathbf{x}, z)}^n(G) + 2(2M_F + M_G) \mathfrak{R}_{p(\mathbf{x}, w)}^{\tilde{n}}(F) + 2(M_F + M_G) \mathfrak{R}_{p(\mathbf{x}, w)}^{\tilde{n}}(G)$.

In particular, the following bound holds for the linear-in-parameter models.

Corollary 1. *Let $F = \{x \mapsto \alpha^\top \phi(\mathbf{x}) \mid \|\alpha\|_2 \leq \Lambda_F\}$, $G = \{x \mapsto \beta^\top \psi(\mathbf{x}) \mid \|\beta\|_2 \leq \Lambda_G\}$. Assume that $r_F := \sup_{\mathbf{x} \in \mathcal{X}} \|\phi(\mathbf{x})\| < \infty$ and $r_G := \sup_{\mathbf{x} \in \mathcal{X}} \|\psi(\mathbf{x})\| < \infty$, where $\|\cdot\|_2$ is the L_2 -norm. Under the assumptions of Theorem 2, it holds with probability at least $1 - \delta$ that for every $f \in F$,*

$$\mathbf{E}_{\mathbf{x} \sim p(\mathbf{x})} [(f(\mathbf{x}) - u(\mathbf{x}))^2] \leq \frac{1}{W^2} \left[\sup_{g \in G} \hat{J}(f, g) + \frac{C_z \sqrt{\log \frac{2}{\delta}} + D_z}{\sqrt{2n}} + \frac{C_w \sqrt{\log \frac{2}{\delta}} + D_w}{\sqrt{2\tilde{n}}} + \varepsilon_G(f) \right],$$

where $C_z := r_G^2 \Lambda_G^2 + 4r_G \Lambda_G M_y$, $C_w := 2r_F^2 \Lambda_F^2 + 2r_F r_G \Lambda_F \Lambda_G + r_G^2 \Lambda_G^2$, $D_z := r_G^2 \Lambda_G^2/2 + 4r_G \Lambda_G M_y$, and $D_w := r_G^2 \Lambda_G^2/2 + 4r_F r_G \Lambda_F \Lambda_G$.

Theorem 2 and Corollary 1 imply that minimizing $\sup_{g \in G} \hat{J}(f, g)$, as the proposed method does, amounts to minimizing an upper bound of the mean squared error. In fact, for the linear-in-parameter models, it can be shown that the mean squared error of the proposed estimator is upper bounded by $O(1/\sqrt{n} + 1/\sqrt{\tilde{n}})$ plus some model mis-specification error with high probability as follows.

Theorem 3 (Informal). Let $\hat{f} \in F$ be any approximate solution to $\inf_{f \in F} \sup_{g \in G} \hat{J}(f, g)$ with sufficient precision. Under the assumptions of Corollary 1, it holds with probability at least $1 - \delta$ that

$$\mathbf{E}_{\mathbf{x} \sim p(\mathbf{x})} [(\hat{f}(\mathbf{x}) - u(\mathbf{x}))^2] \leq O\left(\left(\frac{1}{\sqrt{n}} + \frac{1}{\sqrt{\bar{n}}}\right) \log \frac{1}{\delta}\right) + \frac{2\varepsilon_G^F + \varepsilon_F}{W^2}, \quad (13)$$

where $\varepsilon_G^F := \sup_{f \in F} \varepsilon_G(f)$ and $\varepsilon_F := \inf_{f \in F} J(f)$.

A more formal version of Theorem 3 can be found in Appendix G.

6 More General Loss Functions

Our framework can be extended to more general loss functions:

$$\inf_{f \in L^2(p)} \mathbf{E}[\ell(\mu_w(\mathbf{x})f(\mathbf{x}), 2\mu_z(\mathbf{x}))], \quad (14)$$

where $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is a loss function that is lower semi-continuous and convex with respect to both the first and the second arguments, where a function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is *lower semi-continuous* if $\liminf_{y \rightarrow y_0} \varphi(y) = \varphi(y_0)$ for every $y_0 \in \mathbb{R}$ [30].² As with the squared loss, a major difficulty in solving this optimization problem is that the operand of the expectation has nonlinear dependency on both $\mu_w(\mathbf{x})$ and $\mu_z(\mathbf{x})$ at the same time. Below, we will show a way to transform the objective functional into a form that can be easily approximated using separately labeled samples.

From the assumptions on ℓ , we have $\ell(y, y') = \sup_{z \in \mathbb{R}} yz - \ell^*(z, y')$, where $\ell^*(\cdot, y')$ is the convex conjugate of the function $y \mapsto \ell(y, y')$ defined for any $y' \in \mathbb{R}$ as $z \mapsto \ell^*(z, y') = \sup_{y \in \mathbb{R}} [yz - \ell(y, y')]$ (see Rockafellar [30]). Hence,

$$\mathbf{E}[\ell(\mu_w(\mathbf{x})f(\mathbf{x}), 2\mu_z(\mathbf{x}))] = \sup_{g \in L^2(p)} \mathbf{E}[\mu_w(\mathbf{x})f(\mathbf{x})g(\mathbf{x}) - \ell^*(g(\mathbf{x}), 2\mu_z(\mathbf{x}))].$$

Similarly, we obtain $\mathbf{E}[\ell^*(g(\mathbf{x}), 2\mu_z(\mathbf{x}))] = \sup_{h \in L^2(p)} 2\mathbf{E}[\mu_z(\mathbf{x})h(\mathbf{x})] - \mathbf{E}[\ell_*^*(g(\mathbf{x}), h(\mathbf{x}))]$, where $\ell_*^*(y, \cdot)$ is the convex conjugate of the function $y' \mapsto \ell^*(y, y')$ defined for any $y, z' \in \mathbb{R}$ by $\ell_*^*(y, z') := \sup_{y' \in \mathbb{R}} [y'z' - \ell^*(y, y')]$. Thus, Eq. (14) can be rewritten as

$$\inf_{f \in L^2(p)} \sup_{g \in L^2(p)} \inf_{h \in L^2(p)} K(f, g, h),$$

where $K(f, g, h) := \mathbf{E}[\mu_w(\mathbf{x})f(\mathbf{x})g(\mathbf{x})] - 2\mathbf{E}[\mu_z(\mathbf{x})h(\mathbf{x})] + \mathbf{E}[\ell_*^*(g(\mathbf{x}), h(\mathbf{x}))]$. Since μ_w and μ_z appear separately and linearly, $K(f, g, h)$ can be approximated by sample averages using separately labeled samples.

7 Experiments

In this section, we test the proposed method and compare it with baselines.

7.1 Data Sets

We use the following data sets for experiments.

Synthetic data: Features \mathbf{x} are drawn from the two-dimensional Gaussian distribution with mean zero and covariance $10\mathbf{I}_2$. We set $p(y \mid \mathbf{x}, t)$ as the following logistic models: $p(y \mid \mathbf{x}, t) = 1/(1 - \exp(-y\mathbf{a}_t^\top \mathbf{x}))$, where $\mathbf{a}_{-1} = (10, 10)^\top$ and $\mathbf{a}_1 = (10, -10)^\top$. We also use the logistic models for $p_k(t \mid \mathbf{x})$: $p_1(t \mid \mathbf{x}) = 1/(1 - \exp(-tx_2))$ and $p_2(t \mid \mathbf{x}) = 1/(1 - \exp(-t\{x_2 + b\}))$, where b is varied over 25 equally spaced points in $[0, 10]$. We investigate how the performance changes when the difference between $p_1(t \mid \mathbf{x})$ and $p_2(t \mid \mathbf{x})$ varies.

Email data: This data set consists of data collected in an email advertisement campaign for promoting customers to visit a website of a store [8, 27]. Outcomes are whether customers visited the website or not. We use 4×5000 and 2000 randomly sub-sampled data points for training and evaluation, respectively.

² $\liminf_{y \rightarrow y_0} \varphi(y) := \lim_{\delta \searrow 0} \inf_{|y - y_0| \leq \delta} \varphi(y)$.

Jobs data: This data set consists of randomized experimental data obtained from a job training program called the National Supported Work Demonstration [17], available at <http://users.nber.org/~rdehejia/data/nswdata2.html>. There are 9 features, and outcomes are income levels after the training program. The sample sizes are 297 for the treatment group and 425 for the control group. We use 4×50 randomly sub-sampled data points for training and 100 for evaluation.

Criteo data: This data set consists of banner advertisement log data collected by Criteo [18] available at <http://www.cs.cornell.edu/~adith/Criteo/>. The task is to select a product to be displayed in a given banner so that the click rate will be maximized. We only use records for banners with only one advertisement slot. Each display banner has 10 features, and each product has 35 features. We take the 12th feature of a product as a treatment variable merely because it is a well-balanced binary variable. The outcome is whether the displayed advertisement was clicked. We treat the data set as the population although it is biased from the actual population since non-clicked impressions were randomly sub-sampled down to 10% to reduce the data set size. We made two subsets with different treatment policies by appropriately sub-sampling according to the predefined treatment policies (see Appendix L in the supplementary material). We set $p_k(t | \mathbf{x})$ as $p_1(t | \mathbf{x}) = 1/(1 + \exp(-t\mathbf{1}^\top \mathbf{x}))$ and $p_2(t | \mathbf{x}) = 1/(1 + \exp(t\mathbf{1}^\top \mathbf{x}))$, where $\mathbf{1} := (1, \dots, 1)^\top$.

7.2 Experimental Settings

We conduct experiments under the following settings.

Methods compared: We compare the proposed method with baselines that separately estimate the four conditional expectations in Eq. (6). In the case of binary outcomes, we use the logistic-regression-based (denoted by FourLogistic) and a neural-network-based method trained with the soft-max cross-entropy loss (denoted by FourNNC). In the case of real-valued outcomes, the ridge-regression-based (denoted by FourRidge) and a neural-network-based method trained with the squared loss (denoted by FourNNR). The neural networks are fully connected ones with two hidden layers each with 10 hidden units. For the proposed method, we use the linear-in-parameter models with Gaussian basis functions centered at randomly sub-sampled training data points (see Appendix K for more details).

Performance evaluation: We evaluate trained uplift models by the area under the uplift curve (AUUC) estimated on test samples with joint labels as well as *uplift curves* [26]. The uplift curve of an estimated individual uplift is the trajectory of the average uplift when individuals are gradually moved from the control group to the treated group in the descending order according to the ranking given by the estimated individual uplift. These quantities can be estimated when data are randomized experiment ones. The Criteo data are not randomized experiment data unlike other data sets, but there are accurately logged propensity scores available. In this case, uplift curves and the AUUCs can be estimated using the inverse propensity scoring [3, 20]. We conduct 50 trials of each experiment with different random seeds.

7.3 Results

The results on the synthetic data are summarized in Figure 1. From the plots, we can see that all methods perform relatively well in terms of AUUCs when the policies are distant from each other (i.e., b is larger). However, the performance of the baseline methods immediately declines as the treatment policies get closer to each other (i.e., b is smaller).³ In contrast, the proposed method maintains its performance relatively longer until b reaches the point around 2. Note that the two policies would be identical when $b = 0$, which makes it impossible to identify the individual uplift from their samples by any method since the system in Eq. (5) degenerates. Figure 2 highlights their performance in terms of the squared error. For FourNNC, test points with small policy difference $|p_1(t = 1 | \mathbf{x}) - p_2(t = 1 | \mathbf{x})|$ (colored darker) tend to have very large estimation errors. On the other hand, the proposed method has relatively small errors even for such points. Figure 3 shows results on real data sets. The proposed method and the baseline method with logistic regressors both performed better than the baseline method with neural nets on the Email data set (Figure 3a).

³The instability of performance of FourLogistic can be explained as follows. FourLogistic uses linear models, whose expressive power is limited. The resulting estimator has small variance with potentially large bias. Since different b induces different $u(\mathbf{x})$, the bias depends on b . For this reason, the method works well for some b but poorly for other b .




Figure 1: Results on the synthetic data. The plot shows the average AUUCs obtained by the proposed method and the baseline methods for different b . $p_1(t | \mathbf{x})$ and $p_2(t | \mathbf{x})$ are closer to each other when b is smaller.




Figure 2: The plots show the squared errors of the estimated individual uplifts on the synthetic data with $b = 1$. Each point is darker-colored when $|p_1(t = 1 | \mathbf{x}) - p_2(t = 1 | \mathbf{x})|$ is smaller, and lighter-colored otherwise.

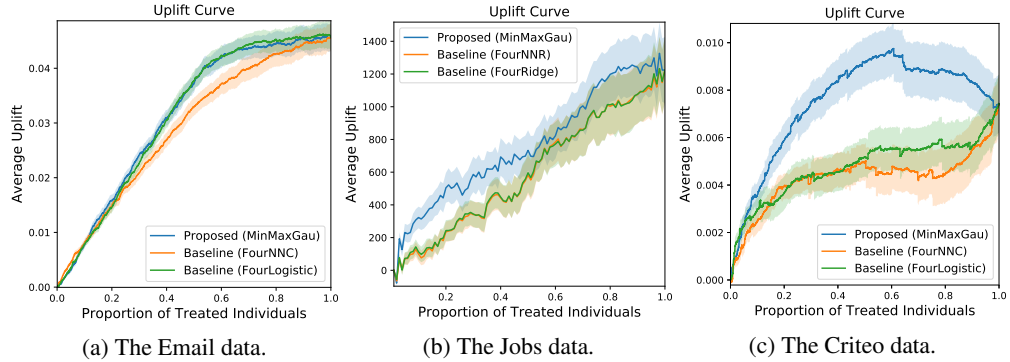


Figure 3: Average uplifts as well as their standard errors on real-world data sets.

On the Jobs data set, the proposed method again performed better than the baseline methods with neural networks. For the Criteo data set, the proposed method outperformed the baseline methods (Figure 3c). Overall, we confirmed the superiority of the proposed both on synthetic and real data sets.

8 Conclusion

We proposed a theoretically guaranteed and practically useful method for uplift modeling or individual treatment effect estimation under the presence of systematic missing labels. The proposed method showed promising results in our experiments on synthetic and real data sets. The proposed framework is model-independent: any models can be used to approximate the individual uplift including ones tailored for specific problems and complex models such as neural networks. On the other hand, model selection may be a challenging problem due to the min-max structure. Addressing this issue would be important research directions for further expanding the applicability and improving the performance of the proposed method.

Acknowledgments

We are grateful to Marthinus Christoffel du Plessis and Takeshi Teshima for their inspiring suggestions and for the meaningful discussions. We would like to thank the anonymous reviewers for their helpful comments. IY was supported by JSPS KAKENHI 16J07970. JA and FY would like to thank Adway for its support. MS was supported by the International Research Center for Neurointelligence (WPI-IRCN) at The University of Tokyo Institutes for Advanced Study.

References

- [1] E. Abrahams and M. Silver. The Case for Personalized Medicine. *Journal of Diabetes Science and Technology*, 3(4):680–684, July 2009.
- [2] S. Athey, J. Tibshirani, and S. Wager. Generalized Random Forests. *arXiv:1610.01271 [econ, stat]*, October 2016. arXiv: 1610.01271.
- [3] P. C. Austin. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate behavioral research*, 46(3):399–424, 2011.
- [4] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., 2006.
- [5] P. Gutierrez and J. Y. Gérardy. Causal inference and uplift modelling: A review of the literature. In *Proceedings of The 3rd International Conference on Predictive Applications and APIs*, volume 67 of *Proceedings of Machine Learning Research*, pages 1–13. PMLR, 11–12 Oct 2017.
- [6] J. Hartford, G. Lewis, K. Leyton-Brown, and M. Taddy. Deep IV: A Flexible Approach for Counterfactual Prediction. In *International Conference on Machine Learning*, pages 1414–1423, July 2017.
- [7] J. L. Hill. Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, January 2011.
- [8] K. Hillstrom. The minethatdata e-mail analytics and data mining challenge, 2008. <https://blog.minethatdata.com/2008/03/minethatdata-e-mail-analytics-and-data.html>.
- [9] K. Imai and M. Ratkovic. Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics*, 7(1):443–470, March 2013.
- [10] G. W. Imbens. Instrumental Variables: An Econometrician’s Perspective. *Statistical Science*, 29(3): 323–358, August 2014.
- [11] M. Jaskowski and S. Jaroszewicz. Uplift modeling for clinical trial data. In *ICML Workshop on Clinical Data Analysis*, 2012.
- [12] F. D. Johansson, U. Shalit, and D. Sontag. Learning Representations for Counterfactual Inference. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 3020–3029. JMLR, 2016.
- [13] S. H. Katsanis, G. Javitt, and K. Hudson. A Case Study of Personalized Medicine. *Science*, 320(5872): 53–54, April 2008.
- [14] R. Kohavi, R. Longbotham, D. Sommerfield, and R. M. Henne. Controlled experiments on the web: survey and practical guide. *Data Mining and Knowledge Discovery*, 18(1):140–181, February 2009.
- [15] V. Koltchinskii. Rademacher penalties and structural risk minimization. *IEEE Transactions on Information Theory*, 47(5):1902–1914, July 2001.
- [16] S. R. Künzel, J. S. Sekhon, P. J. Bickel, and B. Yu. Meta-learners for Estimating Heterogeneous Treatment Effects using Machine Learning. *arXiv:1706.03461 [math, stat]*, June 2017. arXiv: 1706.03461.
- [17] R. J. LaLonde. Evaluating the econometric evaluations of training programs with experimental data. *The American economic review*, pages 604–620, 1986.
- [18] D. Lefortier, A. Swaminathan, X. Gu, T. Joachims, and M. de Rijke. Large-scale validation of counterfactual learning methods: A test-bed. In *NIPS Workshop on "Inference and Learning of Hypothetical and Counterfactual Interventions in Complex Systems"*, 2016.
- [19] G. Lewis and V. Syrgkanis. Adversarial Generalized Method of Moments. *arXiv:1803.07164 [cs, econ, math, stat]*, March 2018. arXiv: 1803.07164.

- [20] L. Li, W. Chu, J. Langford, and X. Wang. An unbiased offline evaluation of contextual bandit algorithms with generalized linear models. In *Proceedings of the Workshop on On-line Trading of Exploration and Exploitation 2*, volume 26 of *Proceedings of Machine Learning Research*, pages 19–36. PMLR, 02 Jul 2012.
- [21] S. Liu, A. Takeda, T. Suzuki, and K. Fukumizu. Trimmed density ratio estimation. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4518–4528. Curran Associates, Inc., 2017.
- [22] M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of machine learning*. MIT press, 2012.
- [23] X. Nguyen, M. J. Wainwright, and M. I. Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, Nov 2010.
- [24] J. Pearl. *Causality*. Cambridge university press, 2009.
- [25] N. Radcliffe. Using control groups to target on predicted lift: Building and assessing uplift model. *Direct Marketing Analytics Journal*, pages 14–21, 2007.
- [26] N. Radcliffe and P. Surry. Differential Response Analysis: Modeling True Responses by Isolating the Effect of a Single Action. *Credit Scoring and Credit Control IV*, 1999.
- [27] N. J. Radcliffe. Hillstrom’s minethatdata email analytics challenge: An approach using uplift modelling. *Technical report, Stochastic Solutions Limited*, 2008.
- [28] N. J. Radcliffe and P. D. Surry. Real-world uplift modelling with significance-based uplift trees. *White Paper TR-2011-1, Stochastic Solutions*, 2011.
- [29] R. Renault. Chapter 4 - Advertising in Markets. In Simon P. Anderson, Joel Waldfogel, and David Strömberg, editors, *Handbook of Media Economics*, volume 1 of *Handbook of Media Economics*, pages 121–204. North-Holland, January 2015.
- [30] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [31] D. B. Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- [32] P. Rzepakowski and S. Jaroszewicz. Decision trees for uplift modeling with single and multiple treatments. *Knowledge and Information Systems*, 32(2):303–327, 2012.
- [33] U. Shalit, F. D. Johansson, and D. Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3076–3085. PMLR, 06–11 Aug 2017.
- [34] M. Sugiyama, I. Takeuchi, T. Suzuki, T. Kanamori, H. Hachiya, and D. Okanohara. Conditional density estimation via least-squares density ratio estimation. In *International Conference on Artificial Intelligence and Statistics*, pages 781–788, 2010.
- [35] M. Sugiyama, T. Suzuki, and T. Kanamori. *Density Ratio Estimation in Machine Learning*. Cambridge University Press, 2012.
- [36] S. Tufféry. *Data mining and statistics for decision making*, volume 2. Wiley Chichester, 2011.
- [37] S. Wager and S. Athey. Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *arXiv:1510.04342 [math, stat]*, October 2015.
- [38] M. Yamada, T. Suzuki, T. Kanamori, H. Hachiya, and M. Sugiyama. Relative density-ratio estimation for robust distribution comparison. *Neural Computation*, 25(5):1324–1370, 2013.