Regret Minimization for Causal Inference on Large Treatment Space

Akira Tanimoto NEC Corporation, Kyoto University, RIKEN

Takashi Takenouchi Future University Hakodate, RIKEN

Abstract

Predicting which action (treatment) will lead to a better outcome is a central task in decision support systems. To build a prediction model in real situations, learning from observational data with a sampling bias is a critical issue due to the lack of randomized controlled trial (RCT) data. To handle such biased observational data, recent efforts in causal inference and counterfactual machine learning have focused on debiased estimation of the potential outcomes on a binary action space and the difference between them, namely, the conditional average treatment effect. When it comes to a large action space (e.g., selecting an appropriate combination of medicines for a patient), however, the regression accuracy of the potential outcomes is no longer sufficient in practical terms to achieve a good decision-making performance. This is because a high mean accuracy on the large action space does not guarantee the nonexistence of a single potential outcome misestimation that misleads the whole decision. Our proposed loss minimizes the classification error of whether or not the action is relatively good for the individual target among all feasible actions, which further improves the decision-making performance, as we demonstrate. We also propose a network architecture and a regularizer that extracts a debiased representation not only from the feature but also from the biased action for better generalization on large action spaces. Extensive experiments on synthetic and semi**Tomoya Sakai** NEC Corporation, RIKEN

> Hisashi Kashima Kyoto University

synthetic datasets demonstrate the superiority of our method for large combinatorial action spaces.

1 INTRODUCTION

Predicting individualized causal effects is an important issue in many domains for decision-making. For example, a doctor considers which medication would be the most effective for a patient, a teacher considers which problems are most effective for helping student learn, and a retail store manager considers which assortment of items would improve the overall store sales. To support such decision-making, we advocate providing a prediction of which actions will lead to better outcomes.

Recent efforts in causal inference and counterfactual machine learning have focused on making predictions of the potential outcomes that correspond to each action for each individual target on the basis of observational data. Observational data consists of features of targets, past actions actually taken, and their outcomes. We have no direct access to the past decisionmakers' policies, i.e., the mechanism of how to choose an action under a given target feature. Unlike in normal prediction problems, pursuing high-accuracy predictions only with respect to the historical data carries the risk of incorrect estimates due to the sampling bias in the past policies. This bias may cause spurious correlation (Simon, 1954; Pearl, 2009), which might mislead the decision-making. For those cases where real-world experiments such as randomized controlled trials (RCTs) or multi-armed bandit are infeasible or too expensive, causal inference methods provide debiased estimation of potential outcomes from observational data.

While most of the existing approaches assume limited action spaces such as a binary one, as in conditional average treatment effect (CATE) estimation, there are

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

many real-world situations where the number of options is large. For example, doctors need to consider which combination of medicines will best suit a patient.

For such cases, it is difficult to apply existing methods (as in (Shalit et al., 2017; Yoon et al., 2018; Schwab et al., 2018; Lopez et al., 2020)) for two reasons. One is the issue of sample-efficiency for large action spaces. Since the sample sizes for each action would be limited, building models for each action (or using a multi-head neural network), which existing methods adopt, is not sample-efficient. The other reason is the gap between the decision-making performance and the regression accuracy of the potential outcome. Even if we manage to achieve the same level of regression accuracy as when the action space is limited, the same decisionmaking performance is no longer guaranteed in a large action space, as we demonstrate in Section 4. This is because, in a nutshell, the overestimated potential outcome of only a single action may mislead the decision, even though it has only a small impact on the mean regression accuracy over all actions.

To achieve informative causal inference for decisionmaking in a large action space, we propose solutions for the above two issues. For the sample-efficiency, we propose extracting representations not only from features but also from actions. We extend two existing representation-based causal effect inference methods, respectively, to balance the representation distribution to be similar to that in the randomized trials.

For the gap between the decision performance and the regression accuracy, we prove that we can further improve the decision performance by minimizing the classification error of whether or not each action is relatively good for each target, in addition to the regression error (MSE). Unlike the recommendation problems in which ranking losses can be used, we cannot directly observe whether the action is relatively good or not since only one action and its outcome is observed for each target. We therefore propose a proxy loss that compares the observed outcome to the estimated conditional average performance of the past decision-makers, which is estimated by regular supervised learning.

In summary, our proposed method minimizes both the classification error and the MSE by using debiased representations of both the features and the actions. We demonstrate the effectiveness of our method through extensive experiments with synthetic and semi-synthetic datasets.

x	а		Ya						y		
		a_0	0				1				
		<i>a</i> 1	0 1		0		1				
		<i>a</i> ₂	0	1	0	1	0	1	0	1	
x_1	(0, 0, 1)		-	1	_	_	_	_	_	_	1
<i>x</i> ₂	(0, 1, 0)		-	_	3	_	_	-	-	-	3
<i>x</i> ₃	(0, 0, 0)		4	-	-	-	-	_	-	-	4
<i>x</i> ₄	(1,0,1)		-	_	_	_	_	6	-	-	6

Figure 1: An example data table for our causal inference on a combinatorial action space. Dashes indicate missing entries. Only factual outcomes are observed (when a = a', $y_{a'}$ is observed) and the counterfactual records $\{y_a\}_{a\neq a'}$ are missing.

2 PROBLEM SETTING

In this section, we formulate our problem and define a decision-focused performance metric. Our aim is to build a predictive model to inform decision-making. Given a feature vector $x \in \mathcal{X} \subset \mathbb{R}^d$, the learned predictive model f is expected to correctly predict which action $a \in \mathcal{A}(x)$ leads to a better outcome $y \in \mathcal{Y} \subset \mathbb{R}$, where $\mathcal{A}(x)$ is a feasible subset of a finite action space \mathcal{A} given x. We hereafter assume that the feasible action space does not depend on the feature, i.e., $\mathcal{A}(x) = \mathcal{A}$, for simplicity. A typical case of large action spaces is when an action consists of multiple causes, i.e., $\mathcal{A} = \{0,1\}^m$ (combinatorial action space).

We assume there exists a joint distribution $p(x, a, y_1, \ldots, y_{|\mathcal{A}|}) = p(x)\mu(a|x)p(y_1, \ldots, y_{|\mathcal{A}|}|x)$, where $\mu(a|x)$ is the unknown decision-making policy of past decision-makers, called propensity, and $y_1, \ldots, y_{|\mathcal{A}|}$ are the potential outcomes corresponding to each action. The observed (factual) outcome yis the one corresponding to the observed action a, i.e., a training instance is the triplet (x_n, a_n, y_{a_n}) , where n denotes the instance index, and the other (counterfactual) potential outcomes are regarded as missing, as shown in Fig. 1.

We make the following assumptions on the distributions of the observational data.

- $(y_1, \ldots, y_{|\mathcal{A}|}) \perp a | x \text{ (unconfoundedness)}$
- $\forall a \in \mathcal{A} \text{ and } \forall x, 0 < \mu(a|x) < 1 \text{ (overlap)}$

These are commonly required to identify causal effects (Imbens and Wooldridge, 2009; Pearl, 2009).



Figure 2: Example scatter plots of true vs. predicted potential outcomes for a target (fixed x) for different models. Each point corresponds to an action. ER^{u}_{μ} corresponds to the rate of instances in the shaded areas. Assuming that the predicted best action $\hat{a}^* := \arg \max_{a} f(x, a)$ is adopted, minimizing the difference between its outcome $y_{\hat{a}^*}$ and the true optimal outcome y_{a^*} (regret) is our aim (see the definition in Section 4).

3 REGRET MINIMIZATION NETWORK: DEBIASED POTENTIAL OUTCOME REGRESSION AND CLASSIFICATION

For this problem of estimating the action evaluation model, we propose our regret minimization network (RMNet), which consists of two parts: 1) a decisionfocused risk to reduce the gap between the decisionmaking performance and the regression accuracy, and 2) representation balancing methods for debiased and sample-efficient learning.

3.1 Decision-Focused Risk

Most of the existing causal effect inference methods aim at minimizing the MSE of the treatment effect (a.k.a. the precision in estimation of heterogeneous effect (PEHE) (Hill, 2011)) in the binary treatment setting. In multiple treatment settings, a typical performance measure is the MSE averaged uniformly over all the actions (Schwab et al., 2018; Yoon et al., 2018):

$$\mathrm{MSE}^{u}(f) = \mathbb{E}_{x} \left[\frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} \mathbb{E}_{y_{a}|x} [(y_{a} - f(x, a))^{2}] \right].$$
(1)

We refer to MSE^u as MSE, or specifically the uniform MSE, in this paper.

On the other hand, there is a gap between the decision performance and the regression accuracy (MSE^{u}). Specifically, we do not necessarily have to accurately estimate the outcomes of candidate actions, but it is enough to identify better actions among others to achieve a higher decision-making performance. This is analogous to the personalized ranking approach in recommender systems (Rendle et al., 2009), in which pairwise comparison of the item preference for each target user is considered.

The pairwise ranking approach (Joachims, 2002; Burges et al., 2005) measures the consistency between the actual and predicted orders by means of the error rate of pairwise comparison as

$$\operatorname{ER}_{\operatorname{rank}}(f) = \mathop{\mathbb{E}}_{i,j} \left[I(y_i \ge y_j \oplus f(x_i) \ge f(x_j)) \right],$$

where \oplus denotes the logical XOR. However, we cannot apply a regular pairwise loss, since we typically only have the outcome for one action observed among the feasible actions. Instead, we propose minimizing the following comparison loss to the average performance of the past decision-makers as the personalized baseline for the target (x):

$$\operatorname{ER}^{u}_{\mu}(f) = \operatorname{\mathbb{E}}_{x} \left[\frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} I(y_{a} \ge \bar{y} \oplus f(x, a) \ge \bar{y}) \right], \quad (2)$$

where $\bar{y} = \mathbb{E}_{a \sim \mu(a|x)} [Y_a|x]$ is the average performance of the past decision-makers under x. As shown in Fig. 2, minimizing ER^u_{μ} leads to better models in terms of decision performance. The MSE is the same in Fig. 2(a) and Fig. 2(b), and thus MSE cannot be used to determine which of these prediction models is better. Minimizing ER^u_{μ} enables us to correctly choose the model in Fig. 2(a) with a high decision performance (small regret).

Replacing the expected value \bar{y} with its estimation g(x) and the 0-1 loss with cross entropy, we get the following risk:

$$\widetilde{\mathrm{ER}}_{g}^{u}(f) = \\ \mathbb{E}\left[-\frac{1}{|\mathcal{A}|}\sum_{a\in\mathcal{A}} \{s\log v + (1-s)\log(1-v)\}\right], \quad (3)$$

where $s := I(y - g(x) \ge 0)$ and $v := \sigma(f(x, a) - g(x))$, and $g(x) \simeq \mathbb{E}_{a \sim \mu(a|x)}[Y_a|x]$ is the estimated average performance of the past decision-makers. We first fit gwith the standard supervised learning procedure from $\{(x_n, y_n)\}$ and then plug it into (3).

Not only the classification error but also the regression error (MSE) matters to the decision-making performance. This is because even with high classification accuracy, decisions might be misleading if only one misclassified action a is predicted as the best $(\hat{y}_a$ is the highest among others $\{\hat{y}_{a'}\}_{a'}$ but is actually quite bad $(y_a$ is quite low), as in Fig. 2(c).

Therefore, we propose minimizing a combination of both of the regression and classification risks, i.e., the geometric mean of them.

$$L^{u}(f;g) = \sqrt{\widetilde{\mathrm{ER}}_{g}^{u}(f) \cdot \mathrm{MSE}^{u}(f)}.$$
 (4)

The reason we chose the geometric mean will be explained theoretically in Section 4. Intuitively, it is sufficient to make one of these losses small, e.g., if the classification loss is zero, good decisions can be made even if the MSE is large. As shown in Fig. 2(a), a model that achieves $\text{ER}^{u}_{\mu} = 0$ (thus the geometric mean is also zero) can at least outperform the past decision-makers on average $(y_{\hat{a}^*} \geq \bar{y})$ no matter how large the MSE^u is.

3.2 Debiased and Sample-Efficient Learning

While accessible observational data taken from p(x, a)is biased by the propensity $\mu(a|x)$, our target expected risk $L^u(f;g)$ is averaged over all actions uniformly, i.e., $p^u(x,a) = p(x)p^u(a)$, where $p^u(a) = \text{Unif}(\mathcal{A})$ is the discrete uniform distribution. In this section, therefore, we construct two debiasing methods for the sampling bias that performs domain adaptation from p(x, a) to $p^u(x, a)$ as extensions of two existing approaches. Also, we propose network architectures that extract representations from both the feature and the action for better generalization in a large action space.

There are two major approaches for debiased learning in individual-level causal inference. One is a density estimation-based method called inverse probability weighting using propensity score (IPW) (Austin, 2011), in which each instance is weighted by the inverse propensity $1/\mu(a_n|x_n)$. Since the expected risk matches that of the RCT, a good performance can be expected asymptotically under accurate estimation of μ or when it is recorded as in logged bandit problems. However, in observational studies where the propensity has to be estimated and plugged in, its efficacy would easily decrease (Kang et al., 2007). It becomes further difficult when it comes to a large treatment space. Zou et al. (2020) proposed assuming an intrinsic low-dimensional structure for combinatorial treatment assignments (bundle treatments) $a \in \{0, 1\}^p$ and estimating weights on the latent space. While in this study we examine a general case of large treatment spaces without additional assumptions, it may be necessary to introduce such assumptions to consider such a huge treatment space of combinatorial interventions.

The other approach is representation balancing (Shalit et al., 2017; Johansson et al., 2016; Lopez et al., 2020), in which a representation extractor of the feature ϕ_x is trained to eliminate the effect of confounding as well as to preserve the relation to the outcome. Shalit et al. (2017); Johansson et al. (2016) proposed regularizing the conditional probabilities of representations $\{p(\phi_x|a)\}_a$ to be similar to each other by means of the integral probability metric (IPM) regularizer (Müller, 1997; Sriperumbudur et al., 2012) (as in Fig. 3(a)) for limited action spaces such as the binary space $\mathcal{A} = \{0, 1\}$. Lopez et al. (2020) proposed regularizing the representation ϕ_x to be independent from the action a by means of the Hilbert-Schmidt Independence Criterion (HSIC) (Gretton et al., 2005, 2008) for realvalued action space $\mathcal{A} \subset \mathbb{R}$. We extend this approach to large treatment spaces.

To deal with a large treatment space, we propose performing representation extraction from the treatment a as well as the feature x. RMNet-IPM (Fig. 3(b)) extracts the joint representation $\phi_{x,a}$ from x and a, which is regularized to be distributionally similar to that of the RCTs $p^u(\phi_{x,a})$. That is, IPM measures the discrepancy between the distributions

$$p(\phi_{x,a}) := \int \sum_{a'} p(\phi_{x,a}|x',a')\mu(a'|x')p(x')dx',$$
$$p^{u}(\phi_{x,a}) := \int \sum_{a'} p(\phi_{x,a}|x',a')p^{u}(a')p(x')dx',$$

where $p(\phi_{x,a}|x',a') = \delta(\phi_{x,a} - \phi(x',a'))$. IPM is defined for a pair of distributions (p_1, p_2) over S and a function family G as follows.

$$\operatorname{IPM}_{G}(p_{1}, p_{2}) = \sup_{g \in G} \left| \int_{\mathcal{S}} g(s) \left(p_{1}(s) - p_{2}(s) \right) \mathrm{d}s \right|$$

We adopt the set of 1-Lipschitz functions as G (as in (Shalit et al., 2017)), after which IPM is equivalent to the Wasserstein distance. Specifically, we use an entropy relaxation of the exact Wasserstein distance, called Sinkhorn distance (Cuturi, 2013), to ensure the compatibility with the gradient-based optimization. This discrepancy upper-bounds the gap between our target risk (4), which is averaged over the uniform distribution with respect to action $p^u(x, a) = p(x)p^u(a)$, and the one of observational distribution p(x, a). The-



Figure 3: Network structures of counterfactual regression for CATE (Shalit et al., 2017; Schwab et al., 2018) and our proposed methods. A broken line indicates no backpropagation.

oretical analysis for this point can be found in Appendix B.

Note that minimizing the discrepancy between $p(\phi_{x,a})$ and $p^u(\phi_{x,a})$ and preserving the causal relation are not necessarily incompatible. In this sense, our approach, which directly regularizes the representation distribution $p(\phi_{x,a})$ to be similar to that taken from RCTs $p^u(\phi_{x,a})$, provides a weaker and sufficient condition for this domain adaptation problem. We discuss this point in Appendix C.

RMNet-HSIC (Fig. 3(c)) extracts each representation ϕ_x and ϕ_a from x and a separately, and they are regularized to be independent from each other by minimizing $\text{HSIC}(p(\phi_x), p(\phi_a))$. HSIC can be defined as a special case of the (squared) maximum mean discrepancy (MMD), which is an instance of the IPM with the class of norm-1 reproducing kernel Hilbert space (RKHS) functions, as follows:

$$\operatorname{HSIC}(p(\phi_x), p(\phi_a)) = \operatorname{MMD}^2(p(\phi_x, \phi_a), p(\phi_x)p(\phi_a))$$

Algorithm 1 Regret minimization network

Input: Observational data $D = \{(x_n, a_n, y_n)\}_n$, a hyperparameter α

- **Output:** Trained network parameter W
- 1: Train g by an arbitrary supervised learning method with $D' = \{(x_n, y_n)\}_n$, e.g.: $g = \arg \min \sum (y_n - g'(x_n))^2$.
- 2: if Method is RMNet-HSIC then
- 3: Set weight $\beta_n = 1/\hat{p}(a_n)$ for each instance, where $\hat{p}(a_n)$ is the count $|\{n \in D \mid a = a_n\}|$.
- 4: **else**
- 5: Set $\beta_n = 1$ for all n.
- 6: **end if**
- 7: while Convergence criteria is not met do
- 8: Sample mini-batch $\{n_1, \ldots, n_b\} \subset \{1, \ldots, N\}.$
- 9: Calculate the gradient of the supervised loss L in (5):

 $g_1 = \nabla_W \frac{1}{b} \sum L(f(x_{n_i}, a_{n_i}; W), y_{n_i}; g(x_{n_i}), \beta_n).$

- 10: Calculate the gradient of the representation balancing regularizer:
 - $g_2 = \nabla_W D_{\text{bal}}(\{\phi(x_{n_i}, a_{n_i}; W)\}).$ Obtain step size η with an optimizer (e.g., Adam
- 11: Obtain step size η with an optimizer (e.g., Adam (Kingma and Ba, 2015)).
- 12: $W \leftarrow [W \eta(g_1 + \alpha g_2)].$
- 13: Check convergence criterion.
- 14: end while
- 15: return W

This means the joint distribution is being separated, i.e., $p(\phi_x, \phi_a) = p(\phi_x)p(\phi_a)$, but it does not mean the consistency with the RCTs $p^u(\phi_{x,a}) = p(\phi_x)p^u(\phi_a)$. To compensate $p(\phi_a)$, we weight the loss according to the estimated marginal probability of the actions $\beta = 1/\hat{p}(a)$.

The resulting objective function is

$$\min_{f} \frac{1}{N} \sum_{n} L(f(x_n, a_n), y_n; g(x_n), \beta_n)
+ \alpha \cdot D_{\text{bal}} \left(\{ \phi(x_n, a_n) \}_n \right) + \Re(f),$$
(5)

where L is the empirical instance-wise version of (4), D_{bal} is the balancing regularizer (IPM or HSIC), and \mathfrak{R} is a regularizer. The resulting learning flow is shown in Algorithm 1.

4 RELATION BETWEEN PREDICTION ACCURACY AND DECISION-MAKING PERFORMANCE

In this section, we analyze our decision-focused performance metric. This analysis demonstrates the difficulty of maximizing the decision performance only by minimizing the regression error when the action space is large. At the same time, however, it is shown that we can further minimize the upper-bound of the regret by minimizing a classification error, which justifies our proposed loss (4) in Section 3.1.

Here we define the decision performance of a model f as the simple average of the potential outcomes for the top-k predicted actions by f. We call that performance metric the mean cumulative gain (mCG), and also define its difference from the oracle's performance (regret).

$$\mathrm{mCG}_k(f) := \frac{1}{k} \mathop{\mathbb{E}}_{x} \left[\sum_{a:\mathrm{rank}(f(x,a)) \le k} y_a \right], \tag{6}$$

$$\operatorname{Regret}_{k}(f) := \frac{1}{k} \mathop{\mathbb{E}}_{x} \left[\sum_{a:\operatorname{rank}(y_{a}) \leq k} y_{a} \right] - \operatorname{mCG}_{k}(f), \quad (7)$$

where rank(·) is the rank among all the feasible actions, e.g., rank $(f(x,a); \{f(x,a')\}_{a'}) := |\{a' \mid f(x,a') \ge f(x,a), a' \in \mathcal{A}\}|$. Here, $(1 - \text{mCG}_{k=1}(f))$ is known as the policy risk (Shalit et al., 2017). Since the first term in (7) is constant with respect to f, the mCG and the regret are two sides of the same coin as the performance metrics of a model.

The relation between the regret and the regression and classification accuracies is the following (full proof and analysis on the tightness can be found in Appendix A). **Proposition 4.1.** The regret in (7) will be bounded

$$\operatorname{Regret}_{k}(f) \leq \frac{|\mathcal{A}|}{k} \sqrt{\operatorname{ER}_{k}^{u}(f) \cdot \operatorname{MSE}^{u}(f)}, \quad (8)$$

where $\operatorname{ER}_{k}^{u}(f)$ is the top-k classification error rate, *i.e.*,

$$\begin{split} & \mathrm{ER}_k^u(f) := \\ & \mathbb{E}_x \left[\frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} I\left((\mathrm{rank}(y_a) \leq k) \oplus (\mathrm{rank}(f(x,a)) \leq k) \right) \right]. \end{split}$$

Proof Sketch. Let $s(x,a) := I(\operatorname{rank}(y_a) \leq k) - I(\operatorname{rank}(f(x,a)) \leq k)$ denote the classification error. Then, we have

$$k \cdot \operatorname{Regret}_{k}(f) = |\mathcal{A}| \underset{x,a \sim p^{u}(x,a)}{\mathbb{E}} [s(x,a)y_{a}] \leq |\mathcal{A}| \underset{x,a \sim p^{u}(x,a)}{\mathbb{E}} [s(x,a)(y_{a} - f(x,a))]$$
(9)

$$\leq |\mathcal{A}| \sqrt{\underset{x,a \sim p^{u}(x,a)}{\mathbb{E}}} \left[s(x,a)^{2} \right] \underset{x,a \sim p^{u}(x,a)}{\mathbb{E}} \left[(y_{a} - f(x,a))^{2} \right]$$
(10)

$$= |\mathcal{A}| \sqrt{\mathrm{ER}_k^u(f) \cdot \mathrm{MSE}^u(f)}.$$

with uniform MSE in (1) as

Equation (9) is from the definition of s(x, a) and (10) is from the Cauchy-Schwarz inequality. By dividing both sides by k, we get the proposition.

Since $\operatorname{ER}_k^u(f) \leq 1$ for any f, we see that only minimizing the uniform MSE as in existing causal inference methods leads to minimizing the regret. However, if $|\mathcal{A}|/k$ is large, the bound would be loose, and only unrealistically small MSE^u provide a meaningful guarantee for the regret.

At the same time, we see that the bound can be further improved by minimizing the uniform top-kclassification error rate $ER_{k}^{u}(f)$ simultaneously, which leads to our proposed method. Let k' be the past decision-makers' average performance, i.e., $y_{a_{\mu'+1}^*} \leq$ $\mathbb{E}_{a \sim \mu(a|x)}[y_a|x] \leq y_{a_{k'}^*}$. Then, the proposed method can be interpreted as minimizing the upper-bound of $\operatorname{Regret}_{k'}$. While training a model for a particular k is an interesting direction, the proposed method is not so sensitive to the difference between the decision-making performance of the data k' and the actual k to be evaluated, as we will see in Section 5. Another interesting direction is optimizing k or the decision-making policy. The mCG_k can be interpreted as the expected performance (reward) of the following plug-in policy that takes an action uniformly at random from the predicted top-k actions.

$$\pi_k^f(a|x) := \begin{cases} 1/k & \text{if } \operatorname{rank}(f(x,a); \{f(x,a')\}_{a'}) \leq k \\ 0 & \text{otherwise,} \end{cases}$$

Therefore, choosing k means choosing a policy. If we choose k greater than 1, the oracle's performance (the first term in (7)) would be smaller, but the upper bound of the regret (8) would be larger. Thus there may exist an optimal k > 1 that maximizes the overall performance of the decision-making.

5 EXPERIMENTS

We investigated the effectiveness of our method through synthetic and semi-synthetic experiments. We newly designed both datasets for the problem setting with a large action space.

5.1 Experimental Setup

Compared Methods We compared our proposed method (RMNet) with ridge linear regression (OLS), random forests (Breiman, 2001) (RF), k-nearest neighbor (kNN), Bayesian additive regression trees (BART) (Hill, 2011), naive deep neural network (S-DNN), naive DNN with multi-head architecture for each action (M-DNN) (a.k.a. TARNET (Shalit et al., 2017)), RankNet (Burges et al., 2005), and a straightforward extension of the existing action-wise representationbalancing method (counterfactual regression network (CFRNet)) (Shalit et al., 2017). We also made an ablation study to clarify the contributions of each component. The strength of representation-balancing regularizer α in CFRNet and the proposed method was selected from [0.1, 0.3, 1.0, 3.0, 10.0]. Other specifications of the DNN parameters can be found in Appendix D.

Evaluation We used the normalized mean gain (NMG) as the main metric, defined as follows.

NMG :=
$$\sum_{x} y_{\hat{a}^*}(x) / \sum_{x} y_{a^*}(x),$$

where \hat{a}^* and a^* are the predicted and true best actions for each x, respectively. The NMG is proportional to the mean CG (k = 1) (6). We can see NMG ≤ 1 . Since we have standardized the outcome, the chance rate is NMG = 0. In addition to NMG, we have also evaluated with respect to MSE^{*u*} and ER^{*u*}_{*k*=1}. The validation and the model selection were based on the NMG. For those cases where the complete validation dataset to compute NMG is not accessible, an alternative validation strategy needs to be considered, e.g., imputing missing values by 1-NN or BART (as in (Hassanpour and Greiner, 2019)) or constructing a special method (such as the counterfactual cross-validation in (Saito and Yasui, 2020)).

Infrastructure All the experiments were run on a machine with 28 CPUs (Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz), 250GB memory, and 8 GPUs.

5.2 Synthetic Experiment

Dataset We prepared four biased datasets with sampling bias in total to examine the robustness of the proposed and baseline methods. For a detailed description of the generation process, see Appendix D. The feature space and the action space are fixed to \mathbb{R}^5 and $\{0,1\}^5$, respectively. The true causal models are set as follows. Three settings (called Quadratic) have a relation $y_a(x) = a_{\Upsilon}^2 - 2x_{\Upsilon} + \varepsilon$, where $a_{\Upsilon} = w_a^{\top} a$ and $x_{\Upsilon} = w_x^{\top} x$ are the one-dimensional representations of a and x, respectively, and where $w_a, w_x \sim N(0, 1/5)^5$. The last setting (called Bilinear) has a bilinear relation $y = x^{\top}Wa + \varepsilon$, where $W \sim N(0, 1/25)^{5 \times 5}$. For training, only one action and the corresponding outcome for each x are sampled as $p(a|x) \propto \exp(10|x_{\Sigma} - a_{\Sigma}|)$, where x_{Σ} and a_{Σ} are additional representations of x and a. The three settings for the quadratic patterns correspond to the relation between \cdot_{Σ} and \cdot_{Υ} as illustrated in Fig. 4(a)–(c), i.e., $x_{\Sigma} = x_{\Upsilon}$ (=: x_{Δ}) in Setups A and C, and $a_{\Sigma} = a_{\Upsilon}$ (=: a_{Δ}) in Setups B and C. These relations of variables were designed



Figure 4: Data generation models for synthetic experiment. Shaded variables denote the accessible variables in training. Non-shaded variables are latent onedimensional representations of x and a.

Table 1: Synthetic results on NMG (larger is better and the maximum is one) and its standard error in ten data generations. Best and second-best methods are in bold.

Method	Quadratic-A	Quadratic-B	Quadratic-C	Bilinear
OLS	0.35 ± 0.13	0.74 ± 0.10	0.73 ± 0.12	0.02 ± 0.02
RF	0.71 ± 0.08	0.24 ± 0.02	0.91 ± 0.04	0.67 ± 0.03
kNN	0.58 ± 0.05	0.33 ± 0.04	0.53 ± 0.07	0.59 ± 0.03
BART	0.53 ± 0.12	$0.91\ \pm\ 0.05$	$0.99\ \pm\ 0.00$	0.14 ± 0.07
M-DNN	0.46 ± 0.09	0.42 ± 0.12	0.57 ± 0.12	-0.01 ± 0.04
S-DNN	0.63 ± 0.08	0.43 ± 0.07	0.60 ± 0.08	0.58 ± 0.09
CFRNet	0.46 ± 0.08	0.43 ± 0.12	0.63 ± 0.13	-0.01 ± 0.04
RankNet	0.62 ± 0.09	$0.70~\pm~0.05$	0.68 ± 0.08	$0.74\ \pm\ 0.04$
RMNet-IPM	0.86 ± 0.04	0.84 ± 0.03	0.82 ± 0.05	0.77 ± 0.04
RMNet-HSIC	$0.90\ \pm\ 0.02$	$0.88\ \pm\ 0.05$	$0.86\ \pm\ 0.07$	$0.14~\pm~0.03$

to reproduce spurious correlations, which mislead the decision-making as follows. In Setup A, a_{Σ} would have dependence on y through its dependence on x_{Δ} despite a_{Σ} itself having no causal relation to y. In the same manner, in Setup B, x_{Σ} would have dependence on y through a_{Δ} , and the causal effect of a_{Δ} may appear discounted. Setup C has both effects. The sample sizes for x were 1,000 for training, 100 for validation, and 200 for testing.

Results The results listed in Table 1 show that our proposed method achieved the best or comparable performance under all settings, while the other methods varied in performance across settings. We analyze the reason of the poor performance of RMNet-HSIC in Bilinear in the ablation study in Section 5.4.

5.3 Semi-Synthetic Experiment

Dataset (GPU Kernel Performance) For the semi-synthetic experiment, we used the SGEMM GPU kernel performance dataset (Nugteren and Codreanu, 2015; Ballester-Ripoll et al., 2019), which has 14 feature attributes of GPU kernel parameters and four target attributes of elapsed times in milliseconds for four independent runs of each combination of parameters.

	Normalized mean gain				MSE^{u}				$\mathrm{ER}_{k=1}^{u}$			
$ \mathcal{A} $	8	16	32	64	8	16	32	64	8	16	32	64
Method												
OLS	-0.04 ± 0.15	-0.08 ± 0.20	-0.10 ± 0.13	-0.01 ± 0.10	1.12	1.89	1.70	5.86	0.221	0.116	0.061	0.031
\mathbf{RF}	0.24 ± 0.08	0.33 ± 0.07	0.33 ± 0.05	$0.38~\pm~0.05$	1.03	0.87	0.93	1.07	0.214	0.114	0.059	0.030
kNN	0.35 ± 0.04	0.39 ± 0.04	0.33 ± 0.04	0.39 ± 0.02	0.59	0.64	0.64	0.63	0.211	0.113	0.059	0.030
BART	-0.05 ± 0.13	0.13 ± 0.13	0.13 ± 0.10	0.04 ± 0.09	1.06	1.05	1.15	1.63	0.222	0.116	0.060	0.031
M-DNN	0.40 ± 0.05	0.48 ± 0.06	0.30 ± 0.07	$0.37~\pm 0.05$	0.78	0.83	0.82	0.84	0.211	0.113	0.059	0.030
S-DNN	0.28 ± 0.09	0.25 ± 0.10	0.32 ± 0.07	0.45 ± 0.05	0.75	0.64	0.74	0.74	0.212	0.114	0.059	0.029
CFRNet	0.50 ± 0.06	0.39 ± 0.14	0.39 ± 0.10	$0.35~\pm0.05$	0.78	0.80	0.87	0.86	0.210	0.113	0.058	0.030
RankNet	$0.35~\pm~0.07$	0.29 ± 0.09	0.38 ± 0.06	$0.45~\pm~0.05$	6.08	10.13	8.47	2.42	0.210	0.113	0.058	0.029
RMNet-IPM	$\textbf{0.68}~\pm~\textbf{0.01}$	$0.61~\pm~0.05$	$0.61\ \pm\ 0.04$	$0.51~\pm~0.06$	0.76	0.81	0.85	0.75	0.204	0.109	0.055	0.029
RMNet-HSIC	$0.59\ \pm\ 0.04$	0.57 ± 0.06	$0.55~\pm~0.06$	$0.69~\pm~0.06$	0.48	0.66	0.61	0.39	0.207	0.109	0.056	0.028

Table 2: Semi-synthetic results on NMG with the standard error in ten different samplings of the training data. The MSE^u and $ER_{k=1}^u$ are also shown. Best and second-best methods are in bold.

We used the inverse of the mean elapsed times as the outcome, resulting in 241.6k instances in total. By treating some of the feature attributes as action dimensions, we obtained a *complete* dataset, which has all the entries (potential outcomes) in Fig. 1 observed. Then we composed our semi-synthetic dataset by biased subsampling of only one action a and the corresponding potential outcome y_a for each x. The details of this preprocess can be found in Appendix D.

The sampling policy in the training data was $p(a|x, y) \propto \exp(-10|y - [x^{\top}, a^{\top}]^{\top}w|)$, where w is sampled from $\mathcal{N}(0, 1)^{d+m}$. This policy reproduces a spurious correlation; that is, a random projection of the feature and the action $[x^{\top}, a^{\top}]^{\top}w$ is likely to have little causal relation with y but does have a strong correlation due to the sampling policy. This policy also depends on y, which violates the unconfoundedness assumption. However, the dataset we used has a low noise level, i.e., $y \simeq g(x, a)$ for some function g, and thus the violation is limited, i.e., $p(a|x, y) \simeq p(a|x, g(x))$.

We split the feature set $\{x_n\}_n$ into 80% for training, 5% for validation, and 15% for testing. Then, for the training set, only one action a and the corresponding outcome y was taken for each x. The resulting training sample size for each setting of m is listed in Table 5 in Appendix D. We repeated the training and evaluation process ten times for different splits and samplings of a.

Results The results listed in Table 2 show that our proposed methods outperformed the others in NMG in all cases. The decision performance (NMG) was more consistent with ER than MSE, indicating that ER as well as MSE needs to be considered. The performance of multi-head DNNs (M-DNN and CFRNet) decreased in larger action spaces, while single-head DNNs (S-DNN and the proposed methods) maintained their performance. This demonstrates the importance

of sample efficiency by extracting the representation of both the feature and the action.

5.4 Ablation Study

We examined the effect of each component of the proposed method, i.e., the balancing regularizer (D_{bal}) , each component of the risk (MSE and ER), and the representation extraction from the action (ϕ_a) and the reweighting with respect to the marginal distribution of the action (β) for RMNet-HSIC. Table 3 shows the results.

The effectiveness of D_{bal} was verified in the setting of $|\mathcal{A}| = 32$. Also, the effectiveness of ER was significant in the Bilinear setting. Extracting representation from the action (ϕ_a) was quite effective in Semi-synthetic settings. The reweighting (β) was also effective in the Semi-synthetic settings, while it decreased the performance in the Bilinear setting. A possible reason is the estimation variance induced by plugging the estimated marginal distribution of the action $\hat{p}(a)$ into weights as its inverse, which is the same issue as the inverse propensity score weighting approach.

6 SUMMARY

In this paper, we have investigated causal inference on a large action space with a focus on the decisionmaking performance. We analyzed the decisionmaking performance brought about by a model through a simple prediction-based decision-making policy. We showed that the bound with only the regression accuracy (MSE) gets looser as the action space gets large, which demonstrates the difficulty of utilizing causal inference in decision-making in a large action space. At the same time, however, our bound indicates that minimizing not only the regression loss but also the classification loss leads to a better performance. From this viewpoint, our proposed methods

Table 3: Ablation study of the proposed methods (indicated by (\dagger)) on semi-synthetic dataset. D_{bal} indicates the type of balancing regularizer. MSE and ER are the used loss. ϕ_a indicates whether or not the representation is also extracted from the action, i.e., if ϕ_a is not checked, identity function is used for ϕ_a (i.e., $\phi_a = a$). β indicates the reweighting with $1/\hat{p}(a)$, which is needed only in the HSIC-based methods (as explained in Section 3.2). Best and second-best methods are in bold.

					Normalized mean gain				
					Synthetic Semi-sy		Inthetic		
$D_{\rm bal}$	MSE	\mathbf{ER}	ϕ_a	β	Bilinear	$ \mathcal{A} = 32$	$ \mathcal{A} = 64$		
\dagger IPM	\checkmark	\checkmark	\checkmark		$\textbf{0.77}~\pm~\textbf{0.04}$	$0.61\ \pm\ 0.04$	$0.51~\pm~0.06$		
IPM		\checkmark	\checkmark		$0.73\ \pm\ 0.03$	$0.61~\pm~0.05$	0.58 ± 0.05		
IPM	\checkmark		\checkmark		0.55 ± 0.10	0.55 ± 0.05	0.49 ± 0.05		
None	\checkmark	\checkmark	\checkmark		0.72 ± 0.03	0.39 ± 0.07	0.49 ± 0.06		
†HSIC	\checkmark	\checkmark	\checkmark	\checkmark	0.14 ± 0.03	0.55 ± 0.06	$0.69\ \pm\ 0.06$		
HSIC		\checkmark	\checkmark	\checkmark	0.11 ± 0.02	$0.56~\pm~0.07$	$0.72\ \pm\ 0.02$		
HSIC	\checkmark		\checkmark	\checkmark	0.16 ± 0.05	0.59 ± 0.05	0.68 ± 0.06		
HSIC	\checkmark	\checkmark		\checkmark	0.04 ± 0.03	0.31 ± 0.08	0.23 ± 0.09		
HSIC	\checkmark	\checkmark	\checkmark		0.51 ± 0.07	0.38 ± 0.07	0.49 ± 0.06		
HSIC	\checkmark	\checkmark			0.63 ± 0.05	0.29 ± 0.07	0.22 ± 0.09		

minimize both the MSE and the classification loss of whether or not the outcome is better than the average performance of the past decision-makers. Specifically, we adopt the cross-entropy with a teacher label indicating whether an observed outcome is better than the estimated average decision performance of the past decision-makers under a given feature. For the sample efficiency in a large treatment space, we proposed extracting representations from both the feature and the action. To generalize in the distribution of RCTs, we proposed two balancing regularizers that encourage the representation distribution to be similar to that of RCTs as extensions of existing approaches. Experiments on synthetic and semi-synthetic datasets. which were designed to have misleading spurious correlations, demonstrated the superior performance of the proposed methods with respect to the decision performance.

Acknowledgements

TT was partially supported by JSPS KAKENHI Grant Numbers 20K03753 and 19H04071. HK was supported by the JSPS KAKENHI Grant Number 20H04244.

References

- P. C. Austin. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate behavioral re*search, 46(3):399–424, 2011.
- R. Ballester-Ripoll, E. G. Paredes, and R. Pajarola. Sobol tensor trains for global sensitivity analysis. *Reliability Engineering & System Safety*, 183:311– 322, 2019.
- L. Breiman. Random forests. *Machine learning*, 45(1): 5–32, 2001.

- C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender. Learning to rank using gradient descent. In *Proceedings of the* 22nd international conference on Machine learning, pages 89–96, 2005.
- M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In Advances in neural information processing systems, pages 2292–2300, 2013.
- A. Gretton, O. Bousquet, A. Smola, and B. Schölkopf. Measuring statistical dependence with hilbertschmidt norms. In *International conference on algorithmic learning theory*, pages 63–77. Springer, 2005.
- A. Gretton, K. Fukumizu, C. H. Teo, L. Song, B. Schölkopf, and A. J. Smola. A kernel statistical test of independence. In *Advances in neural information processing systems*, pages 585–592, 2008.
- N. Hassanpour and R. Greiner. Counterfactual regression with importance sampling weights. In *IJCAI*, pages 5880–5887, 2019.
- J. L. Hill. Bayesian nonparametric modeling for causal inference. Journal of Computational and Graphical Statistics, 20(1):217–240, 2011.
- G. W. Imbens and J. M. Wooldridge. Recent developments in the econometrics of program evaluation. *Journal of economic literature*, 47(1):5–86, 2009.
- T. Joachims. Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM* SIGKDD international conference on Knowledge discovery and data mining, pages 133–142, 2002.
- F. Johansson, U. Shalit, and D. Sontag. Learning representations for counterfactual inference. In *International conference on machine learning*, pages 3020– 3029, 2016.

- F. D. Johansson, D. Sontag, and R. Ranganath. Support and invertibility in domain-invariant representations. In K. Chaudhuri and M. Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 527–536. PMLR, 16–18 Apr 2019.
- J. D. Kang, J. L. Schafer, et al. Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical science*, 22(4):523–539, 2007.
- D. P. Kingma and J. L. Ba. Adam: A method for stochastic optimization. In *ICLR 2015 : International Conference on Learning Representations* 2015, 2015.
- R. Lopez, C. Li, X. Yan, J. Xiong, M. Jordan, Y. Qi, and L. Song. Cost-effective incentive allocation via structured counterfactual inference. In AAAI 2020 : The Thirty-Fourth AAAI Conference on Artificial Intelligence, volume 34, pages 4997–5004, 2020.
- A. Müller. Integral probability metrics and their generating classes of functions. Advances in Applied Probability, pages 429–443, 1997.
- C. Nugteren and V. Codreanu. Cltune: A generic auto-tuner for opencl kernels. In Embedded Multicore/Many-core Systems-on-Chip (MCSoC), 2015 IEEE 9th International Symposium on, pages 195–202. IEEE, 2015.
- J. Pearl. Causality. Cambridge university press, 2009.
- S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of* the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, pages 452–461, 2009.
- Y. Saito and S. Yasui. Counterfactual cross-validation: Stable model selection procedure for causal inference models. In *ICML 2020: 37th International Conference on Machine Learning*, 2020.
- P. Schwab, L. Linhardt, and W. Karlen. Perfect match: A simple method for learning representations for counterfactual inference with neural networks. arXiv preprint arXiv:1810.00656, 2018.
- U. Shalit, F. D. Johansson, and D. Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3076–3085. JMLR. org, 2017.
- H. A. Simon. Spurious correlation: A causal interpretation. Journal of the American statistical Association, 49(267):467–479, 1954.
- B. K. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf, G. R. Lanckriet, et al. On the empiri-

cal estimation of integral probability metrics. *Electronic Journal of Statistics*, 6:1550–1599, 2012.

- J. Yoon, J. Jordon, and M. van der Schaar. GAN-ITE: Estimation of individualized treatment effects using generative adversarial nets. In *International Conference on Learning Representations*, 2018.
- Y. Zhang, A. Bellot, and M. van der Schaar. Learning overlapping representations for the estimation of individualized treatment effects. In *AISTATS*, pages 1005–1014, 2020.
- H. Zhao, R. T. D. Combes, K. Zhang, and G. Gordon. On learning invariant representations for domain adaptation. volume 97 of *Proceedings of Machine Learning Research*, pages 7523–7532, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- H. Zou, P. Cui, B. Li, Z. Shen, J. Ma, H. Yang, and Y. He. Counterfactual prediction for bundle treatment. Advances in Neural Information Processing Systems, 33, 2020.