



Towards Robust Off-Policy Evaluation via Human Inputs

Harvineet Singh*
New York University
New York City, New York, USA

Shalmali Joshi
Finale Doshi-Velez
Himabindu Lakkaraju
Harvard University
Cambridge, Massachusetts, USA

ABSTRACT

Off-policy Evaluation (OPE) methods are crucial tools for evaluating policies in high-stakes domains such as healthcare, where direct deployment is often infeasible, unethical, or expensive. When deployment environments are expected to undergo changes (that is, dataset shifts), it is important for OPE methods to perform robust evaluation of the policies amidst such changes. Existing approaches consider robustness against a large class of shifts that can arbitrarily change any observable property of the environment. This often results in highly pessimistic estimates of the utilities, thereby invalidating policies that might have been useful in deployment. In this work, we address the aforementioned problem by investigating how domain knowledge can help provide more realistic estimates of the utilities of policies. We leverage human inputs on which aspects of the environments may plausibly change, and adapt the OPE methods to only consider shifts on these aspects. Specifically, we propose a novel framework, Robust OPE (ROPE), which considers shifts on a subset of covariates in the data based on user inputs, and estimates worst-case utility under these shifts. We then develop computationally efficient algorithms for OPE that are robust to the aforementioned shifts for contextual bandits and Markov decision processes. We also theoretically analyze the sample complexity of these algorithms. Extensive experimentation with synthetic and real world datasets from the healthcare domain demonstrates that our approach not only captures realistic dataset shifts accurately, but also results in less pessimistic policy evaluations.

CCS CONCEPTS

• **Computing methodologies** → *Sequential decision making*; Adversarial learning; *Batch learning*; Markov decision processes; **Causal reasoning and diagnostics**; *Model verification and validation*.

KEYWORDS

dataset shift, policy evaluation, robust learning, adversarial machine learning, human-in-the-loop

*Work done as a summer fellow at the Center for Research on Computation and Society, Harvard University. Correspondence at hs3673@nyu.edu

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](https://permissions.acm.org).
AIES'22, August 1–3, 2022, Oxford, United Kingdom

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-9247-1/22/08...\$15.00
<https://doi.org/10.1145/3514094.3534198>

ACM Reference Format:

Harvineet Singh, Shalmali Joshi, Finale Doshi-Velez, and Himabindu Lakkaraju. 2022. Towards Robust Off-Policy Evaluation via Human Inputs. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (AIES'22)*, August 1–3, 2022, Oxford, United Kingdom. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3514094.3534198>

1 INTRODUCTION

Off-policy evaluation (OPE) refers to the task of estimating the expected utility of a decision-making policy without having to deploy the policy [30]. Such an ability is critical for vetting policies in high-stakes decision problems such as in healthcare [16], where deploying a policy directly is often risky or unethical. Therefore, we must rely on existing data collected from alternate policies deployed possibly in a different environment. Accurate evaluation of a policy is important so that stakeholders can identify beneficial policies and discard the harmful ones.

In real world applications, OPE is a challenging task since the deployment environments often undergo changes (i.e., dataset shifts). It is critical for OPE methods to evaluate policies in a way that is robust to these changes. Prior work has proposed different solutions to address this problem. While some approaches address the scenario where potential shifts in the data are fully known in advance [26, 63, 67], others focus on the case where there is little to no knowledge of potential shifts in advance [33, 57] since this is more common in real world applications. Prior works that focus on robust OPE under unseen dataset shifts predominantly model the shifts by considering bounded perturbations to the joint distribution of the data [17, 54, 68], inspired by adversarial machine learning literature. However, prior research has also demonstrated that such shifts can be overly conservative, and often result in pessimistic estimates of policy utilities [11, 46]. Moreover, accounting for many irrelevant shifts may result in poor performance on shifts of interest. For instance, consider adversarial training methods that perturb training data in an ℓ_p -norm ball and minimize the worst-case error across the perturbations [15, 34, 55]. Existing work shows that robustness to shifts modeled as ℓ_p perturbations does not transfer well to real world shifts [61] or even to other ℓ_p norms [36]. Such methods can degrade performance on train distribution [49] and lead to degenerate solutions [20].

While prior research has often considered bounded perturbations to the joint distribution of the data, this is quite uncommon in practice and does not represent realistic dataset shifts [57, 61]. To illustrate, we plotted the percentage of features that shift in three different real world datasets comprising of census, loan, and medical data (See Figure 1). For each of these datasets, we often observe that only a subset of the covariates (and not the joint distribution of the data) undergo a change in case of dataset shifts. For example, only 4

(14%) out of the 28 total features changed between pre 2006 and post 2006 loan datasets (Figure 1 – middle). In order to model realistic dataset shifts, it therefore becomes important to exploit domain knowledge and inputs from human experts which can guide us to the plausible subset of features that are likely to shift. While such inputs have been incorporated in learning or evaluating classifiers that are robust to realistic dataset shifts [33, 57], there is little to no work in the OPE literature that leverages domain knowledge and/or human inputs when modeling shifts in the data.

In this work, we address the aforementioned challenges by investigating how domain knowledge can help with providing more realistic estimates of the utilities of policies. To this end, we leverage human inputs on which aspects of the environments may plausibly change and adapt the OPE methods to only consider shifts on these aspects. The hope is that this enables a domain expert to constrain the shifts to only the most relevant or plausible ones. Then, we leverage the framework of *distributionally robust optimization* [DRO, 10] for carrying out robust policy evaluation for contextual bandits and Markov decision processes. More specifically, we make the following key contributions:

- We propose a novel framework, Robust OPE (ROPE), which considers shifts on a subset of covariates in the data based on user inputs and estimates worst-case utility under these shifts.
- We develop computationally efficient algorithms for robust OPE via human inputs in case of contextual bandits and Markov decision processes.
- We theoretically analyze the sample complexity of the proposed algorithms.
- We carry out extensive experiments with synthetic and real world datasets from the healthcare domain. Our results demonstrate that ROPE can successfully tackle the over-conservatism of existing robust policy evaluation methods.

Our work paves the way for modeling realistic dataset shifts in the context of off-policy evaluation in reinforcement learning.

2 RELATED WORK

As robustness has been extensively studied in a variety of learning settings, we review only closely related works on approaches for handling unseen shifts, i.e., adversarial robustness, DRO and causal robustness.

Adversarial Robustness. Adversarial shifts modeled as ℓ_p -norm perturbations have been widely considered to learn models robust to adversarial attacks [15, 34]. However, such methods provide limited gains in robustness to real world shifts [61]. Recent work compensates for the non-realistic shifts by combining multiple ℓ_p norm balls [36], considering shifts in a perceptual distance [28], or augmenting with additional datasets from adjacent domains [37, 61]. Similar methods have been extended to RL under perturbations to transition dynamics [47, 55], or horizon length and initial state distribution [48].

Distributionally Robust Optimization. DRO generalizes adversarial shifts to perturbations in the distributions rather than data points [10]. The primary mechanism of DRO is to specify *uncertainty sets* that encode the uncertainty about potential test distributions. These sets can be defined over joint distribution of the

data [2–4, 10], marginals [11, 33], or conditionals [57], and are well explored for supervised learning. Interestingly, adversarial training can be understood as solving DRO with a Wasserstein metric-based set [55, 56]. Applications of DRO have been explored in contextual bandits for policy learning [13, 38, 54] and evaluation [22, 63]. In robust MDPs, sets based on KL-divergence, L_1 , L_2 and L_∞ norms have been studied [21, 40]. Some approaches iteratively refine the sets with newly observed data, but the sets are still constructed using L_1 norm balls [46]. Thus, DRO methods (including adversarial training) lack ways to add domain knowledge and constrain the uncertainty sets, excluding some recent work [57].

Causal Robustness. Causal methods provide robustness to arbitrarily strong shifts by leveraging properties of the data generating process. For instance, using only features that cause the outcome leads to a robust model against arbitrary shifts in the features (under some structural assumptions), as shown in recent works [35, 45, 51, 52, 58]. A relaxation to bounded shifts has been proposed [6, 42, 52] for supervised learning, but in the special cases of additive shifts or linear Gaussian causal models. In contrast, we do not make parametric assumptions on the shifts. Under bounded shifts in conditional distributions, which is a broader class than what we consider, Subbaswamy et al. [57] propose methods for evaluating the performance of a given classification model.

Importantly, there are gaps in the literature in applying these ideas beyond supervised learning. In RL, Hatt et al. [17], Si et al. [54] perform OPE restricted to contextual bandits and Zhou et al. [68] generalizes the DRO approach to MDPs. However, all of these works only consider shifts in the joint distribution and do not investigate the use of domain knowledge to restrict the uncertainty sets. As RL starts to be deployed in critical applications such as for mechanical ventilation in ICUs [44], faithful evaluation of RL policies under plausible data shifts is an important need. Thus, extending ROPE to RL and demonstrating its utility in realistic settings is a useful contribution.

3 PRELIMINARIES

We first introduce the robust evaluation framework based on distributionally robust optimization, then give necessary background on decision-making problems modelled as Contextual Bandits (CB) and Markov Decision Processes (MDPs).

Notation We denote the random variable for all observable properties of a domain by V . The outcome variable is denoted by Y . This can represent labels in supervised learning, or rewards and states in RL. Features are denoted by X . For a subset of features $Z \subseteq X$, the remaining features are denoted by $X \setminus Z$. Train and test distributions over V are denoted by P and Q respectively (using the same notation for their densities). An uncertainty set w.r.t. a distribution P is denoted by \mathcal{U}_P .

3.1 Robust Evaluation using DRO

Say, we want to evaluate a decision-making model, parameterized by θ , on a test distribution Q . For a given reward function $r(\theta, V)$, this means, we want to find the value of the expected reward $\mathbb{E}_{V \sim Q}[r(\theta, V)]$. However, we do not know the test distribution Q a priori. Robust evaluation methods (e.g. [33, 57]) address this




Figure 1: Prevalence of shifts on subset of features. Plots show the percentage of features that shift across domains within three different datasets (US Census, Loan, Medical). More details are given in Appendix E.1. The three datasets have 10, 2, and 11 domains, respectively. The x and y-axes show the different domains (10 US state codes in US Census, 2 time periods in Loan, and 11 hospitals in Medical data). Each cell denotes the percentage of the total features (shown in each subtitle) that shift for the corresponding pair of domains. We detect and count feature shifts using conditional independence tests [27], where we check if $P_{\text{train}}(X_i|X \setminus X_i) = P_{\text{test}}(X_i|X \setminus X_i)$. Firstly, plots show that distribution shifts are quite prevalent. Secondly, we observe that only a few domain pairs in US Census undergo shift in all features (value of 100% in the leftmost plot). In all other pairs, only a fraction of the features shift (similar to Loan and Medical data). This shows that shifts on feature subsets are prevalent in the real world datasets and methods that assume shifts in *all* features are aiming for robustness to unrealistic shifts.

challenge by performing a worst-case evaluation of the model. Instead of finding the expected reward under Q , the *robust* value of a model is defined as its worst-case reward across a set of distributions \mathcal{U}_P .

$$\mathcal{R}(\theta, \mathcal{U}_P) := \inf_{\tilde{P} \in \mathcal{U}_P} \mathbb{E}_{V \sim \tilde{P}} [r(\theta, V)] \quad (1)$$

where \mathcal{U}_P is referred to as the uncertainty set. The key property of the solution to (1) is that if the \mathcal{U}_P is suitably chosen such that it contains the test distribution Q , then $\mathcal{R}(\theta, \mathcal{U}_P) \leq \mathbb{E}_{V \sim Q} [r(\theta, V)]$. Thus, we get the guarantee that the reward of the model on Q is at least as good as what our robust evaluation finds. Naturally, a major part of this framework is the choice of uncertainty set \mathcal{U}_P . If we precisely choose $\mathcal{U}_P = \{Q\}$, then we recover the test set reward. But as we keep on increasing the uncertainty set, in the hope of including Q , we gradually get worse values due to the worst-case nature of the optimization (1). Next we describe a broad class of uncertainty sets popular in past work.

Divergence-based Uncertainty Sets. The uncertainty set \mathcal{U}_P is all distributions lying in a δ -ball around the train distribution, defined using a divergence metric \mathcal{D} :

$$[\text{Joint DRO set}] \quad \mathcal{U}_P^{\text{div}} := \{Q \ll P \text{ s.t. } \mathcal{D}(Q \parallel P) \leq \delta\} \quad (2)$$

where $Q \ll P$ denotes absolute continuity i.e. $P(V) = 0$ implies $Q(V) = 0$. An example is the set with:

$$\mathcal{D}_{\text{CVaR}}(Q \parallel P) = \log \sup_{V \in \text{dom}(V)} \frac{Q(V)}{P(V)},$$

where CVaR stands for Conditional Value at Risk [50].

Types of shifts The above definition of $\mathcal{D}_{\text{CVaR}}$ assumes that P and Q may differ in distribution over all variables in V , namely

shifts in the *joint distribution*. There are many ways to limit the distributions considered by $\mathcal{U}_P^{\text{div}}$. *Covariate shift* (e.g. Duchi et al. [11]) assumes that only the covariate distribution $P(X)$ may change at test time to $Q(X)$, while the conditional distribution $P(Y|X)$ remains the same. A slight generalization is *subcovariate shift*, considered by us, that posits that $P(Z)$ may change for a subset of covariates Z keeping the conditional distribution constant. This reduces the size of the set given by $\mathcal{D}_{\text{CVaR}}$ since many of the Q are removed from the divergence set (2) due to the constraint of matching the corresponding distributions in P . We leverage human input to specify (sub)covariate shifts and propose OPE methods in reinforcement learning. Next, we provide preliminaries for OPE in two widely-studied settings in RL – contextual bandits and MDPs.

3.2 Contextual Bandits

Here we have access to n tuples $\{(Z_i, T_i, Y_i)\}_i$ collected with a known stochastic policy that applies treatment T_i in context Z_i and observes the corresponding outcome Y_i . Further, it is assumed that the tuples are i.i.d. and the joint distribution P factorizes as $\prod_i P(Z_i)P(T_i|Z_i)P(Y_i|Z_i, T_i)$.

OPE in CBs Given data sampled from P , the goal in OPE is to evaluate the expected outcome $\mathbb{E}_Q[Y]$ under a distribution Q induced by following a new policy in the *same* environment [12, 62]. Importantly, the difference between the two distributions is assumed to be only due to different policies i.e. $P(T|Z) \neq Q(T|Z)$ and the rest of the environment-related factors are the same. That is, only shifts on T are considered as the domain expert intends to implement a different policy, $Q(T|Z)$ instead of $P(T|Z)$.

3.3 Markov Decision Process

Markov Decision Processes are characterized by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$, where \mathcal{S} is the state-space, \mathcal{A} is the action-space, $\mathcal{P} \triangleq P(\cdot|s, a)$ characterizes the dynamics, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward model and $\gamma \in [0, 1)$ is the discount factor. We restrict ourselves to finite state, finite action infinite-horizon MDPs in this work. The the reward model and transition model are assumed to be fixed. Value of a policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ is defined as the expected cumulative reward received starting from s_0 ,

$$V^\pi(s_0) = \mathbb{E}_{\pi, P} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 \right].$$

OPE in MDPs In off-policy evaluation, the goal is to estimate the value of a policy π given that we have collected a batch of trajectories obtained using an alternative policy μ . We motivate the need for robustness in the OPE task in Section 4.2.

4 OUR FRAMEWORK

To evaluate policies, we rely on user input to characterize which covariates indeed shift across domains. We then incorporate this domain knowledge in the form of specifying an uncertainty set which characterizes potential shifts relative to the nominal training distribution. Our goal is then to provide realistic off-policy estimates of policy performance under the potential shifts.

Let $Z \subseteq X$ be the subset of variables that shift across domains. In this case, we assume that the uncertainty set $\mathcal{U}_P^{\text{sub}}$ contains all distributions resulting from shifts in subcovariates Z which are bounded in some metric \mathcal{D} . Formally,

[Subcovariate DRO set]

$$\mathcal{U}_P^{\text{sub}} := \left\{ \nu(Z) \ll P(Z) \text{ s.t. } \mathcal{D}(\nu(Z) \| P(Z)) \leq \delta \right\} \quad (3)$$

A key question is the choice of divergence metric \mathcal{D} that can allow us to incorporate expert knowledge. While it is easy for a domain expert to provide precise information on which covariates may shift across domains, the magnitude of shift may not be precisely clear. We capture this complexity using the $\mathcal{D}_{\text{CVaR}}$ -based uncertainty set. Existing literature [11, 33, 57] has leveraged $\mathcal{D}_{\text{CVaR}}$ for specifying magnitude of shifts and compute worse-case loss in case of supervised learning. Specifically, Duchi et al. [11] builds the uncertainty set in such a way that the training population contains at least δ proportion of the test population.

$$\mathcal{U}_P^{\text{sub}} := \{Q(Z)P(V \setminus Z|Z) \text{ s.t. } P(Z) = \alpha'Q(Z) + (1 - \alpha')Q'(Z), \alpha \leq \alpha' \leq 1\} \quad (4)$$

where $Q'(Z)$ is any distribution and $\alpha \in (0, 1]$ determines the minimum size of the subpopulation shared between train and test. Thus, $\mathcal{U}_P^{\text{sub}}$ constrains the ratio $\frac{Q(Z)}{P(Z)} \leq \frac{1}{\alpha}$ for all values of Z . Equivalently, the set can be expressed using $\mathcal{D}_{\text{CVaR}}(Q \| P) = \log \sup_{Z \in \text{dom}(Z)} \frac{Q(Z)}{P(Z)} \leq \delta$ for $\delta = \frac{1}{\alpha}$. Using the subcovariate uncertainty set $\mathcal{U}_P^{\text{sub}}$, the worst-case value $\mathcal{R}(\theta, \mathcal{U}_P^{\text{sub}})$ in Eq. (1) can be found using techniques from convex duality Duchi et al. [11]. Details of the resulting estimator of $\mathcal{R}(\theta, \mathcal{U}_P^{\text{sub}})$ are presented in Appendix A. For the remaining discussion, we will assume that




Figure 2: Graphical model for shifts in contextual bandits. Expert input (shaded node Z) denotes that $P(Z)$ changes at test time and $P(Y|X, T)$ remains the same. For example, treatments T are prescribed based on age Z and have different outcomes Y by age. In practice, only the age distribution may change across environments e.g. across hospitals (as opposed to the joint distribution over age, treatment, and outcome). This motivates a less conservative approach to robustness by focusing on marginal shifts in age.

we can solve the worst-case optimization problem resulting from subcovariate shifts and focus on how to leverage this optimizer for OPE.

We now present our main technical results, particularly the framework ROPE for off-policy evaluation (OPE) under realistic shifts where human input is incorporated in terms of knowledge of (sub)covariate shifts we anticipate in practice.

4.1 Robust OPE in Contextual Bandits

As suggested in Section 3.2, the bandit setup typically assumes that the joint distribution P changes only due to the change in policy. OPE tasks in CBs focus on evaluating the value of the policy under this assumption. Departing from this assumption, we evaluate the utility of a policy under a *new* environment in which the domain expert anticipates shifts in Z . We characterize these by unknown and bounded shifts in Z . The following assumption is the CB analogue of the covariate shift assumption in supervised learning.

ASSUMPTION 1 (EXPERT INPUT FOR CBs). *Suppose the human expert specifies the subset of features $Z \subseteq X$, such that across environments, only $P(Z)$ changes while $P(Y|X, T)$ remains the same.*

Following this input, the joint distribution at test environment factorizes as $Q(X, T, Y) = Q(Z)Q(T, X \setminus Z|Z)P(Y|X, T)$. For simplicity, we will group the variables $T, X \setminus Z$ into T . The rest of the discussion remains the same and an extra factor of $P(X \setminus Z|Z)$ will be suppressed.

Thus, we define the uncertainty set containing distributions Q that shift in the marginal distribution of Z as,

$$\mathcal{U}_P^{\text{CB}} = \{Q \ll P \text{ s.t. } Q = \nu(Z)Q(T|Z)P(Y|Z, T), \mathcal{D}(\nu(Z) \| P(Z)) \leq \delta\} \quad (5)$$

The robust OPE problem aims to find the *worst-case* average outcome under $\mathcal{U}_P^{\text{CB}}$ instead of the average,

$$\mathcal{R}(\mathcal{U}_P^{\text{CB}}) = \inf_{Q \in \mathcal{U}_P^{\text{CB}}} \mathbb{E}_{(Z, T, Y) \sim Q} [Y], \quad (6)$$

Algorithm 1 Robust OPE in CBs

Input: Data $\{Z_i, T_i, Y_i\}_i$, Target policy $Q(T|Z)$, behavior policy $P(T|Z)$, hyperparameters $\delta, L, 1r$.

Compute importance weights $\{W_i = \frac{Q(T_i|Z_i)}{P(T_i|Z_i)}\}_i$.

Create dataset $\{V_i = (Z_i, W_i \times Y_i)\}_i$.

Estimate $\hat{R}(\mathcal{U}_P^{\text{CB}})$ (Eq. (6)) with the worst-case risk estimator in Eq. (13) for the dataset.

return $\hat{R}(\mathcal{U}_P^{\text{CB}})$

where $Q(T|Z)$ is the policy to be evaluated and is considered to be known and fixed. Consider each distribution in the set $\mathcal{U}_P^{\text{CB}}$, $Q(\cdot) = \nu(Z)Q(T|Z)P(Y|Z, T)$, which differs from the train distribution in the factors for Z and $T|Z$.

$$\mathcal{R}(\mathcal{U}_P^{\text{CB}}) = \inf_{Q \in \mathcal{U}_P^{\text{CB}}} \mathbb{E}_{Z \sim Q(Z)} \mathbb{E}_{P(Y|T, Z)Q(T|Z)} [Y] \quad (7)$$

$$\stackrel{(7a)}{=} \inf_{Q \in \mathcal{U}_P^{\text{CB}}} \mathbb{E}_{Z \sim Q(Z)} \mathbb{E}_{P(Y|T, Z)P(T|Z)} \left[\frac{Q(T|Z)}{P(T|Z)} Y \right] \quad (8)$$

$$\stackrel{(7b)}{=} \inf_{\eta \in \mathbb{R}} \frac{1}{\delta} \mathbb{E}_{Z \sim P(Z)} [(\mathbb{E}[W \times Y|Z] - \eta)_+] + \eta$$

To solve (6) for this Q , we first use importance sampling to account for the change in $T|Z$ due to the known policy $Q(T|Z)$, step (7a). As a result, the set $\mathcal{U}_P^{\text{CB}}$ now consists of shifts on Z alone. Thus, the robust OPE problem reduces to solving $\sup_{Q \in \mathcal{U}_P^{\text{CB}}} \mathbb{E}_{V \sim Q} [W \times Y]$ where W are the importance sampling weights:

$$W(T, Z) = Q(T|Z) / P(T|Z)$$

Using convex duality arguments Shapiro et al. [53] for our choice of uncertainty set $\mathcal{D}_{\text{CVaR}}$, we can obtain (7b). This motivates the full optimization procedure summarized in Algorithm 1. We first compute importance sampling weights W_i , and create a re-weighted dataset $\{V_i = (Z_i, W_i \times Y_i)\}_i$. The risk defined in Eq. (6) is approximated by an estimate of its upper bound given in Eq. (13) when Z are continuous valued (see Appendix A for the detailed derivation).

4.2 Robust OPE in MDPs

Off-policy evaluation is critical in sequential decision-making settings, often encountered in human-centered domains such as health. Often such environments are best modeled as a Markov Decision Process (MDP) as introduced in Section 3. In this case, we have to consider shifts in the transition dynamics across environments which can invalidate OPE methods for MDPs as they often assume stationary dynamics.

Our goal is to evaluate *robust* value for a given policy π to be deployed in a new environment with unknown transition dynamics. Hence, the uncertainty set for each state-action pair is defined over $P(s'|s, a)$, denoted by $\mathcal{U}(s, a)$.¹ Specifically, we want to estimate the robust value for π starting from s_0 as $V^\pi(s_0) = \inf_{P \in \mathcal{U}} \mathbb{E}_{\pi, P} [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0]$. Iyengar [21] proves that $V^\pi(\cdot)$

¹We drop the dependence on the train environment’s transition probabilities for conciseness and use P to denote target probabilities instead of Q to avoid confusion with the Q -value function in RL.




Figure 3: Graphical model representing a Markov Decision Process. Expert input (shaded nodes in red) denotes the variables whose distribution changes across environments. Here, only $P(s_{t+1}^1 | s_t, a_t)$ changes. Rewards r_t are not shown for simplicity. Each r_t has directed edges from s_t and a_t .

is the solution to the following fixed-point equation (namely, robust Bellman equation) if we assume that uncertainty sets for each state-action pair are constructed independently (known as *SA-rectangularity*),

$$V^\pi(s) = r(s, \pi(s)) + \inf_{P \in \mathcal{U}(s, \pi(s))} \gamma \mathbb{E}_{s' \sim P(s' | s, \pi(s))} [V^\pi(s')] \quad (10)$$

Intuitively, SA-rectangularity implies that the uncertainty sets are constructed across time steps independently. This property yields a tractable method to compute the value function estimates using dynamic programming [21]. We state the assumption formally in Appendix B for completeness.

Eq. (10) can be solved iteratively by dynamic programming [59]. Given the value function at any iteration, we additionally have to solve the minimization problem over $\mathcal{U}(s, a)$. Thus, the robust OPE problem in MDPs reduces to solving multiple DRO subproblems with a chosen uncertainty set.

Applying ROPE to OPE in MDPs. Past work has only considered uncertainty sets for the joint distribution [21, 46, 60, 68]. In contrast, we consider sets based on the shifts on parts of the state space of the MDP i.e. leveraging human input. Figure 3 pictorially represents the probabilistic model for the MDP denoting the shifting state features. The following assumption is the RL analogue of the subcovariate shifts in supervised learning.

For any time step $t > 0$ in the MDP, consider a partitioning of the state feature vector s_t into two feature sets, (s_t^1, s_t^2) . The factors of the joint distribution in any environment are given by:

$$P(s_{t+1}, r_{t+1} | s_t, a_t) = P(s_{t+1}^1 | s_t, a_t) P(s_{t+1}^2 | s_t, a_t, s_{t+1}^1) P(r_{t+1} | s_t, a_t, s_{t+1})$$

ASSUMPTION 2 (EXPERT INPUT FOR MDPs). *Suppose the human expert specifies the following,*

- (a) $P(s_{t+1}^1 | s_t, a_t)$ can shift across environments independently of state-action pairs at any other time step, while
- (b) $P(s_{t+1}^2 | s_t, a_t, s_{t+1}^1)$ and $P(r_{t+1} | s_t, a_t, s_{t+1})$ are the same as that in the train environment.

Remark. Assumption 2 implies that the MDP satisfies SA-rectangularity. For any state-action pair (s_t, a_t) at time step t , the shifts that result

Algorithm 2 Robust OPE in MDPs

Input: Trajectories $\{(s, a, s', r)\}$ sampled using policy μ and transition model P_0 , Target policy π , Discount factor γ , Robustness level δ .

Learn transition models with sample averages across observed trajectories $\widehat{P}_0(s^1|s, a) = \frac{\text{Count}\{(s, a, (s^1, *, *))\}}{\text{Count}\{(s, a, (*, *, *))\}}$ and $\widehat{P}_0(s^2|s, a, s^1) = \frac{\text{Count}\{(s, a, (s^1, s^2, *))\}}{\text{Count}\{(s, a, (s^1, *, *))\}}$, where the wildcard $(s, a, *, *)$ denotes transitions from trajectories that match the state-action pair s, a .

Learn reward model with sample averages across observed trajectories $\widehat{r}(s, a) = \frac{\sum_{(s, a, r) \in \{(s, a, *, *)\}} r}{\text{Count}\{(s, a, *, *)\}}$.

Initialize $V^\pi(s) = 0$, for all $s \in \mathcal{S}$.

repeat

for $s \in \mathcal{S}$ **do**

Update $V^\pi(s)$ using Eq. (20) in Appendix D with \widehat{P}_0 .

end for

until V^π converges

return V^π

in the uncertainty set $\mathcal{U}^{\text{MDP}}(s_t, a_t)$ are independent of the state-action pairs at other time steps. This implies that $\mathcal{U}^{\text{MDP}}(s_t, a_t)$ is determined independently of the uncertainty sets at other time steps. Thus, the collection \mathcal{U}^{MDP} is all possible combinations of sets $\mathcal{U}^{\text{MDP}}(s_t, a_t)$. A scenario where SA-rectangularity does not hold is when uncertainty sets are constructed adaptively based on previous state-action pairs, which we do not address.

Thus, the uncertainty set needs to be defined only for $P(s_{t+1}^1|s_t, a_t)$, denoted by $\mathcal{U}^{\text{MDP}}(s, a)$:

$$\mathcal{U}^{\text{MDP}}(s, a) := \left\{ P(\cdot|s, a) \ll P_0(\cdot|s, a) \quad \text{s.t.} \right. \\ \left. \forall s', P(s'|s, a) = v(s^1|s, a)P_0(s^2|s, a, s^1), \right. \\ \left. \mathcal{D}(v(s^1|s, a) \parallel P_0(s^1|s, a)) \leq \delta \right\}$$

where $P_0(\cdot)$ denotes the distribution of the train environment. With \mathcal{U}^{MDP} , the DRO subproblem in (10) reduces to,

$$\inf_{P \in \mathcal{U}^{\text{MDP}}(s, \pi(s))} \mathbb{E}_P[V^\pi(s')] \\ = \inf_{P \in \mathcal{U}^{\text{MDP}}(s, \pi(s))} \mathbb{E}_{P(s^1|s, \pi(s))} \left[\mathbb{E}_{P_0(s^2|s, \pi(s), s^1)} [V^\pi(s')] \right]$$

We estimate the inner expectation with Monte-Carlo averaging on batch data, followed by solving the DRO problem as done in bandits (7). Here, we use the maximum likelihood estimate of the transition model P_0 as we do not have access to the true model. Steps for estimation are outlined in Algorithm 2.

Given enough samples of next state for each state-action pair, the robust value can still be estimated accurately. Suppose $V_{\mathcal{U}_{P_0}}^\pi$ be the robust value corresponding to the true transition model P_0 which we want to show is close to the robust value corresponding to the *estimated* transition model \widehat{P}_0 denoted by $V_{\mathcal{U}_{\widehat{P}_0}}^\pi$. To show this, we

will restrict to uncertainty sets in Eq. (3) defined by KL-divergence $\mathcal{D}(Q \parallel P) = \sum_z Q(z) \log\left(\frac{Q(z)}{P(z)}\right)$ being smaller than δ .

THEOREM 1 (ROBUST OPE ESTIMATION ERROR). *Given at least n samples from $P_0(\cdot|s, a)$ for all s, a , assuming that the rewards are bounded $r \in [0, r_{\max}]$, and that \mathcal{U}^{MDP} are defined by KL-divergence, then with probability at least $1 - \alpha$,*

$$\|V_{\mathcal{U}_{P_0}}^\pi - V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty \leq O\left(\frac{\gamma r_{\max} |\mathcal{S}|}{(1-\gamma)^2} \sqrt{\frac{1}{n} \log\left(\frac{4|\mathcal{S}| \times |\mathcal{A}| \times |\mathcal{S}|}{\alpha}\right)}\right)$$

Thus, the error in evaluating the robust value with the estimated transition model as compared to the true model converges at the rate $\tilde{O}\left(\frac{\gamma|\mathcal{S}|}{\sqrt{n}(1-\gamma)^2}\right)$, ignoring logarithmic factors. The dependence on $1/\sqrt{n}$ matches that for the non-robust case but is sub-optimal in $|\mathcal{S}|$ [31]. The proof is included in Appendix C in which we show an analogue of the *simulation lemma* [25] for robust MDPs. In terms of time complexity, each iteration of dynamic programming with DRO can be computed in $O(|\mathcal{A}||\mathcal{S}|^2 \log|\mathcal{S}|)$ which is only a $\log|\mathcal{S}|$ factor more than the non-robust solution.

5 EXPERIMENTAL EVALUATION

In this section, we study the robustness-utility trade-offs achieved by the proposed shifts and compare these with the existing approaches². More specifically, we study whether optimizing for marginal shifts with ROPE improves model performance under those shifts in comparison with optimizing for broader classes of shifts using existing methods. We start with subcovariate shifts in bandit settings as described in Sec. 5.1. We specifically consider the partial feedback problem in the contextual bandits setup since we only get to observe the feedback (outcomes) for the actions taken in the collected data. Thus, we tackle robustness under partial feedback. We show that our method provides more faithful estimates for efficacy of a drug dosing policy under patient population shifts. In Sec. 5.2, we consider sequential decision problems under shifts in environment dynamics for MDPs. Within a simulated Sepsis environment, we show that the proposed sets provide significantly less conservative value estimates for a treatment policy.

5.1 Robust OPE in CBs

We present results on a synthetic and a real world dataset.

Synthetic. We generate data with two features $Z := (Z_1, Z_2)$, binary treatment T , and continuous outcome Y . Additional details on how data is simulated are deferred to Appendix E.3. Here we assume changes the marginal distribution of Z_1 for evaluation. We simulate $n=2000$ samples in the train environment following a logistic policy and use it to estimate the robust value of a different policy in shifted environments.

Baselines include: Standard returns the average value assuming no shift. Inverse Probability Weighting (IPW) only corrects for shift in policy using importance sampling. JointDRO accounts for shifts in all variables V , as done in past work [54].

In Figure 4a, we plot the MSE between the estimated and the true policy value, evaluated using 20000 samples from the test environment. We observe that when the test environment is close

²Code for replicating the experiments is at <https://github.com/AI4LIFE-GROUP/rise-against-distribution-shift>




Figure 4: (a) CB, Synthetic. MSE in value estimates for test sets (y-axis) with varying levels of shift in Z_1 (x-axis). ROPE performs well for moderate shifts. **(b) CB, Warfarin.** MSE in value estimates for test sets with shift in the race distribution. ROPE achieves the right level of conservatism to match the value at test. Curves for Standard and JointDRO are not visible as they have high error (around 0.1 and 0.7, respectively) and lie outside the plotted y-axis range. Robustness level in Eq. (5) is set to $\delta = 0.8$. Error bars are computed over 5 random initializations.

to the train one, not accounting for the shift (Standard) performs well. But, as the shift increases, our approach (ROPE) does better. With large shifts, larger uncertainty sets are required. For large shifts, JointDRO does better than the other methods. In summary, ROPE performs well when the shift is significant but not too large. This highlights the importance of choosing the desired robustness level appropriately which is a challenging open problem for DRO methods.

Warfarin Dosing Policy. Warfarin is an oral anticoagulant drug. Optimal dosage to assign to a patient while initiating Warfarin therapy has been a subject of multiple clinical trials [18]. Using the public PharmGKB dataset [7] of 5528 patients, Bastani and Bayati [1] learn contextual bandit policies that adapt doses based on patient covariates like demographics and clinical information.³ Reward, either 0 or 1, is defined as a policy’s accuracy in making the correct dosing decision. However, the value of the policies is suspect when applied to patient populations different from the development cohort. Thus, we estimate robust value of a policy under shifts in race distribution. Note that the ground truth optimal dose for each patient is available in the data which enables evaluating different policies. We learn a dosing policy with linear regression on held-out data and estimate it on test data with shifted race distributions. Specifically, we subsample (without replacement) fewer patients with a recorded race into our analysis set. The policy has lower performance on patients with Unknown race. Thus, the value of the policy decreases with increasing shift as the relative proportion of this group increases. To estimate this value correctly, the robust method must consider the right level of conservativeness. Figure 4b shows that MSE between the estimated and true average reward is lower for ROPE than for JointDRO and IPW, as it constructs the sets for marginal shifts alone.

5.2 Robust OPE in MDPs

We present results on two simulated RL domains.

³A preprocessed version of the Warfarin dataset was downloaded from <https://github.com/khashayarkhv/contextual-bandits/blob/master/datasets/warfarin.csv>.




Figure 5: Illustration of the Cliffwalking domain. Plot shows (part of) the value function estimated using robust Bellman equation with ROPE. Start position is (5,0), goal is (5,5), and cliff corresponds to the row (5,0) to (5,5). Agent slips downward by 1 cell with probability 0.1 when taking actions in any of the columns except first and last columns. Results in Table 1a report the value at the start position (5,0), which is -1182.45 here, averaged over 10 random initializations of the domain.

Cliffwalking Domain. We consider a 6×6 gridworld which an agent navigates from a start to a goal position avoiding a cliff [59, Ex. 6.6]. Please refer to Appendix E.4 for more details. An illustration of the gridworld is provided in Figure 5. With a constant shift probability, the agent slips towards the cliff instead of taking the prescribed action. The slip probability varies across environments, changing the transition dynamics and necessitating robustness in policy evaluation. We evaluate the value estimate for an agent following uniform random policy using dynamic programming with the standard Bellman equation (Standard) or the robust one in (10)

δ	Standard	JointDRO		ROPE	
	-	0.4	0.8	0.4	0.8
mean	-1136.43	-1448.16	-1221.78	-1416.39	-1190.57
std.	6.22	6.32	5.28	6.70	6.33
median	-1136.64	-1449.91	-1222.76	-1417.12	-1190.92

(a) Cliffwalking domain

δ	Standard	JointDRO	ROPE
	-	0.8	0.8
mean	-0.037	-0.939	-0.662
std.	0.008	0.006	0.148
median	-0.039	-0.941	-0.705

(b) Sepsis domain

Table 1: Robust OPE in MDP. Estimated value (10 random runs) with standard (that is, non-robust) and robust dynamic programming. ROPE provides less conservative value than JointDRO meaning it has a smaller decrease in value from Standard.

(JointDRO, ROPE). To simulate subcovariate shifts, we duplicate the state features s^1, s^2 such that s^1 follows the agent’s actions while s^2 is random noise. Since agent’s actions affect only s^1 , ROPE correctly constructs sets on $P(s^1|s, a)$. In contrast, JointDRO ignores this structure and constructs sets on both the features, $P(s^1, s^2|s, a)$. This is the same setting as considered in past work [68] which used the KL divergence to define uncertainty sets. Table 1a reports the value estimates for the start position. We observe that for a high level of desired robustness $\delta = 0.4$, JointDRO decreases the value by 27.4% (from -1136 to -1448) while ROPE only decreases it by 24.6% (to -1416). This validates that both DRO methods have the expected behavior in the MDP.

Sepsis Treatment Evaluation. Sepsis simulator [41] is a domain with more involved transition dynamics and has been used to test treatment policies [14, 26, 39, 41]. It has a total of 1440 states which includes 4 vital signs (blood pressure, glucose levels, heart rate, and oxygen concentration) and diabetic status. Actions correspond to 3 treatment combinations (antibiotics, mechanical vents, and vasopressors). Terminal states, discharge from ICU or death, have rewards $+1$ or -1 , respectively. Glucose levels fluctuate more for diabetics than non-diabetics. Our goal is to evaluate a policy learned using policy iteration, namely RL policy, on a dataset with 20% diabetics. We consider a setting where the percentage of diabetics and the fluctuation in their glucose levels varies in the test environment. To interrogate the RL policy for possible deployment, we find its robust value accounting for these shifts. JointDRO constructs sets based on the full 1440 states, while ROPE considers uncertainty only in glucose level dynamics for diabetics and non-diabetics. Thus, ROPE represents the actual shifts more faithfully compared to JointDRO. We check how conservative the OPE estimates from assumptions of joint shifts are relative to leveraging domain knowledge to restrict to subcovariate shifts. Table 1b reports the value estimates for the RL policy obtained with standard and robust dynamic programming. We observe that JointDRO reports a decrease in value by 21 times as compared to the value at train environment and ROPE only reports a decrease in value by 14 times. Thus, experiments demonstrate the benefit of curating the uncertainty sets using domain knowledge to balance utility and robustness.

Optimization and system details. For the CB experiments, we perform gradient descent with Adagrad [9] implementation in PyTorch [43] to solve Eq. 13. In Synthetic CB dataset, we use learning rate 0.5 for ROPE and 0.1 for JointDRO, Lipschitz constant 1, and number of steps 100. In Warfarin dataset, we use learning rate 0.1

for both ROPE and JointDRO, Lipschitz constant 0.01, and number of steps 200. While solving the dynamic program iteratively in the MDP experiments, the convergence condition for V^π is taken as maximum change in $V^\pi(\cdot) < 10^{-2}$ after one repetition across all states. The minimization over η while solving Eq. (12) is performed using Brent’s method, implemented in the package `scipy` [65]. For the MDP experiments, the only hyperparameter is δ which is set to 0.8 throughout or to 0.4 in Table 1a. Experiments were run on a compute cluster using a single node with a 2.50 GHz Intel processor, 2 GPUs with 6 GB memory each, and 256 GB system memory. None of the datasets include personally identifiable information.

6 CONCLUSIONS AND FUTURE WORK

In this work we are focused on leveraging human expertise to provide value estimates of ML policies under realistic distribution shifts. Here, we propose to represent domain knowledge via uncertainty sets over sub-covariate shifts. We argue that this enables representing more realistic shifts and lead to less conservative solutions. We then provide novel estimators for robust OPE in contextual bandits and MDPs leveraging the distributionally robust optimization framework. Future directions include expanding human input to address shifts in conditional distributions for OPE (e.g. see [57] for supervised learning). Finally, applying the robust OPE method to continuous state-action spaces with function approximators (e.g. [60]) is an interesting direction of future work. We hope that the perspective of leveraging human input to define uncertainty sets for robustness in off-policy evaluation opens up more possibilities to tackle over-conservatism of robust learning as well as the challenging problem of model selection in off-policy evaluation.

Broader implications. Although it is challenging to foresee the impact of using robust methods in real-world applications, one negative is that from over-reliance on results without adequate scrutiny of the assumptions like causal knowledge and uncertainty about future deployment settings. Adversaries can manipulate the deployment environments such that the estimated robust values of policies have high errors. The design of uncertainty sets that account for such adversarial changes should be informed and validated by domain experts.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their feedback and all the funding agencies listed below for supporting this work. This work is supported in part by the Center for Research on Computation and Society (CRCS) at Harvard University,

NSF awards #IIS-2008461, #IIS-2040989, and #1922658, and research awards from Amazon, Harvard Data Science Institute, Bayer, and Google. The authors would also like to thank Rediet Abebe, Sara Kingsley, Elita Lobo, Aviva Prins, and Milind Tambe for early feedback on this work. Lastly, HL would like to thank Sujatha and Mohan Lakkaraju for their continued support and encouragement. The views expressed here are those of the authors and do not reflect the official policy or position of the funding agencies.

REFERENCES

- [1] Hamsa Bastani and Mohsen Bayati. 2015. Online decision-making with high-dimensional covariates. *Forthcoming in Operations Research* (2015).
- [2] Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenare, Bertrand Melenberg, and Gijss Rennen. 2013. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science* 59, 2 (2013), 341–357.
- [3] Dimitris Bertsimas, Vishal Gupta, and Nathan Kallus. 2018. Data-driven robust optimization. *Mathematical Programming* 167, 2 (2018), 235–292.
- [4] Jose Blanchet and Karthyek Murthy. 2019. Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research* 44, 2 (2019), 565–600.
- [5] Clément L. Canonne. 2022. A short note on an inequality between KL and TV. <https://doi.org/10.48550/ARXIV.2202.07198>
- [6] Rune Christiansen, Niklas Pfister, Martin Emil Jakobsen, Nicola Gnecco, and Jonas Peters. 2020. A causal framework for distribution generalization. *arXiv e-prints* (2020), arXiv:2006.
- [7] International Warfarin Pharmacogenetics Consortium. 2009. Estimation of the warfarin dose with clinical and pharmacogenetic data. *New England Journal of Medicine* 360, 8 (2009), 753–764.
- [8] Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. 2021. Retiring Adult: New Datasets for Fair Machine Learning. In *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (Eds.). https://openreview.net/forum?id=bYi_2708mKK
- [9] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research* 12, 61 (2011), 2121–2159. <http://jmlr.org/papers/v12/duchi11a.html>
- [10] John Duchi and Hongseok Namkoong. 2018. Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750* (2018).
- [11] John C Duchi, Tatsunori Hashimoto, and Hongseok Namkoong. 2019. Distributionally robust losses against mixture covariate shifts. *Under review* (2019).
- [12] Miroslav Dudík, Dumitru Erhan, John Langford, Lihong Li, et al. 2014. Doubly robust policy evaluation and optimization. *Statist. Sci.* 29, 4 (2014), 485–511.
- [13] Louis Fauray, Ugo Tanielian, Elvis Dohmatob, Elena Smirnova, and Flavian Vasile. 2020. Distributionally robust counterfactual risk minimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 3850–3857.
- [14] Joseph Futoma, Michael C Hughes, and Finale Doshi-Velez. 2020. Popcorn: Partially observed prediction constrained reinforcement learning. *arXiv preprint arXiv:2001.04032* (2020).
- [15] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [16] Omer Gottesman, Fredrik Johansson, Joshua Meier, Jack Dent, Donghun Lee, Srivatsan Srinivasan, Linying Zhang, Yi Ding, David Wihl, Xuefeng Peng, Jiayu Yao, Isaac Lage, Christopher Mosch, Li-wei H. Lehman, Matthieu Kumorowski, Matthieu Kumorowski, Aldo Faisal, Leo Anthony Celi, David Sontag, and Finale Doshi-Velez. 2018. Evaluating Reinforcement Learning Algorithms in Observational Health Settings. <https://doi.org/10.48550/ARXIV.1805.12298>
- [17] Tobias Hatt, Daniel Tschernutter, and Stefan Feuerriegel. 2021. Generalizing Off-Policy Learning under Sample Selection Bias. arXiv:2112.01387 [stat.ML]
- [18] Carl Heneghan, Sally Tyndel, Clare Bankhead, Yi Wan, David Keeling, Rafael Perera, and Alison Ward. 2010. Optimal loading dose for the initiation of warfarin: a systematic review. *BMC cardiovascular disorders* 10, 1 (2010), 1–12.
- [19] Wassily Hoeffding. 1994. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*. Springer, 409–426.
- [20] Weihua Hu, Gang Niu, Issei Sato, and Masashi Sugiyama. 2018. Does distributionally robust supervised learning give robust classifiers?. In *International Conference on Machine Learning*. 2029–2037.
- [21] Garud N Iyengar. 2005. Robust dynamic programming. *Mathematics of Operations Research* 30, 2 (2005), 257–280.
- [22] Sookyo Jeong and Hongseok Namkoong. 2020. Robust causal inference under covariate shift via worst-case subpopulation treatment effects. In *Conference on Learning Theory*. PMLR, 2079–2084.
- [23] Nan Jiang. 2020. Notes on Tabular Methods. <https://nanjiang.cs.illinois.edu/files/cs598/note3.pdf>. Accessed: 2021–06-04.
- [24] Alistair E. W. Johnson, Tom J. Pollard, and Tristan Naumann. 2018. Generalizability of predictive models for intensive care unit patients. arXiv:1812.02275 [cs.LG]
- [25] Michael Kearns and Satinder Singh. 2002. Near-optimal reinforcement learning in polynomial time. *Machine learning* 49, 2 (2002), 209–232.
- [26] Taylor W Killian, Marzyeh Ghassemi, and Shalmali Joshi. 2020. Counterfactually Guided Policy Transfer in Clinical Settings. *arXiv preprint arXiv:2006.11654* (2020).
- [27] Sean Kulinski, Saurabh Bagchi, and David I. Inouye. 2020. Feature Shift Detection: Localizing Which Features Have Shifted via Conditional Distribution Tests. In *Neural Information Processing Systems (NeurIPS)*.
- [28] Cassidy Laidlaw, Sahil Singla, and Soheil Feizi. 2021. Perceptual Adversarial Robustness: Defense Against Unseen Threat Models. In *ICLR*.
- [29] David A Levin and Yuval Peres. 2017. *Markov chains and mixing times*. Vol. 107. American Mathematical Soc.
- [30] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. 2020. Offline Reinforcement Learning: Tutorial, Review, and Perspectives on Open Problems. arXiv:2005.01643 [cs.LG]
- [31] Gen Li, Yuting Wei, Yuejie Chi, Yuantao Gu, and Yuxin Chen. 2020. Breaking the Sample Size Barrier in Model-Based Reinforcement Learning with a Generative Model. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 12861–12872. <https://proceedings.neurips.cc/paper/2020/file/96e64f3a1aa2fd00c7faacfb8ac9-Paper.pdf>
- [32] Min Li, Amy Mickel, and Stanley Taylor. 2018. “Should This Loan be Approved or Denied?": A Large Dataset with Class Assignment Guidelines. *Journal of Statistics Education* 26, 1 (2018), 55–66. <https://doi.org/10.1080/10691898.2018.1434342>
- [33] Mike Li, Hongseok Namkoong, and Shangzhou Xia. 2021. Evaluating model performance under worst-case subpopulations. In *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (Eds.). <https://openreview.net/forum?id=nehzxAdyJxJF>
- [34] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*.
- [35] Sara Magliacane, Thijs van Ommen, Tom Classens, Stephan Bongers, Philip Versteeg, and Joris M Mooij. 2018. Domain adaptation by using causal inference to predict invariant conditional distributions. In *Advances in Neural Information Processing Systems*. 10846–10856.
- [36] Pratyush Maini, Eric Wong, and Zico Kolter. 2020. Adversarial Robustness Against the Union of Multiple Perturbation Models. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 6640–6650. <http://proceedings.mlr.press/v119/maini20a.html>
- [37] John Miller, Karl Krauth, Benjamin Recht, and Ludwig Schmidt. 2020. The Effect of Natural Distribution Shift on Question Answering Models. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 6905–6916. <http://proceedings.mlr.press/v119/miller20a.html>
- [38] Weibin Mo, Zhengling Qi, and Yufeng Liu. 2020. Learning optimal distributionally robust individualized treatment rules. *J. Amer. Statist. Assoc.* (2020), 1–16.
- [39] Hongseok Namkoong, Ramtin Keramati, Steve Yadlowsky, and Emma Brunskill. 2020. Off-policy Policy Evaluation For Sequential Decisions Under Unobserved Confounding. In *Advances in Neural Information Processing Systems*.
- [40] Arnab Nilim and Laurent El Ghaoui. 2005. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research* 53, 5 (2005), 780–798.
- [41] Michael Oberst and David Sontag. 2019. Counterfactual Off-Policy Evaluation with Gumbel-Max Structural Causal Models. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 4881–4890. <http://proceedings.mlr.press/v97/oberst19a.html>
- [42] Michael Oberst, Nikolaj Thams, Jonas Peters, and David Sontag. 2021. Regularizing towards Causal Invariance: Linear Models with Proxies. *arXiv preprint arXiv:2103.02477* (2021).
- [43] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035. <https://proceedings.neurips.cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf>
- [44] Arne Peine, Ahmed Hallawa, Johannes Bickenbach, Guido Dartmann, Lejla Begic Fazlic, Anke Schmeink, Gerd Ascheid, Christoph Thiemermann, Andreas Schuppert, Ryan Kindle, et al. 2021. Development and validation of a reinforcement learning algorithm to dynamically optimize mechanical ventilation in critical care. *NPJ digital medicine* 4, 1 (2021), 1–12.
- [45] Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. 2016. Causal inference by using invariant prediction: identification and confidence intervals. *Journal*

- of the Royal Statistical Society: Series B (Statistical Methodology) 78, 5 (2016), 947–1012.
- [46] Marek Petrik and Reazul Hasan Russel. 2019. Beyond confidence regions: Tight bayesian ambiguity sets for robust mdps. In *Advances in Neural Information Processing Systems*. 7049–7058.
- [47] Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. 2017. Robust Adversarial Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, 2817–2826. <http://proceedings.mlr.press/v70/pinto17a.html>
- [48] Zhengling Qi and Peng Liao. 2020. Robust Batch Policy Learning in Markov Decision Processes. *arXiv preprint arXiv:2011.04185* (2020).
- [49] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John Duchi, and Percy Liang. 2020. Understanding and Mitigating the Tradeoff between Robustness and Accuracy. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 7909–7919. <https://proceedings.mlr.press/v119/raghunathan20a.html>
- [50] R Tyrrell Rockafellar, Stanislav Uryasev, et al. 2000. Optimization of conditional value-at-risk. *Journal of risk* 2 (2000), 21–42.
- [51] Mateo Rojas-Carulla, Bernhard Schölkopf, Richard Turner, and Jonas Peters. 2018. Invariant models for causal transfer learning. *The Journal of Machine Learning Research* 19, 1 (2018), 1309–1342.
- [52] Dominik Rothenhäusler, Nicolai Meinshausen, Peter Bühlmann, and Jonas Peters. 2018. Anchor regression: heterogeneous data meets causality. *arXiv preprint arXiv:1801.06229* (2018).
- [53] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. 2014. *Lectures on stochastic programming: modeling and theory*. SIAM.
- [54] Nian Si, Fan Zhang, Zhengyuan Zhou, and Jose Blanchet. 2020. Distributional Robust Batch Contextual Bandits. *arXiv preprint arXiv:2006.05630* (2020).
- [55] Aman Sinha, Hongseok Namkoong, and John Duchi. 2018. Certifying Some Distributional Robustness with Principled Adversarial Training. In *International Conference on Learning Representations*.
- [56] Matthew Staib and Stefanie Jegelka. 2017. Distributionally robust deep learning as a generalization of adversarial training. In *NIPS workshop on Machine Learning and Computer Security*.
- [57] Adarsh Subbaswamy, Roy Adams, and Suchi Saria. 2021. Evaluating Model Robustness and Stability to Dataset Shift. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 130)*, Arindam Banerjee and Kenji Fukumizu (Eds.). PMLR, 2611–2619. <http://proceedings.mlr.press/v130/subbaswamy21a.html>
- [58] Adarsh Subbaswamy, Peter Schulam, and Suchi Saria. 2019. Preventing failures due to dataset shift: Learning predictive models that transport. In *The 22nd International Conference on Artificial Intelligence and Statistics*. 3118–3127.
- [59] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [60] Aviv Tamar, Shie Mannor, and Huan Xu. 2014. Scaling up robust MDPs using function approximation. In *International Conference on Machine Learning*. 181–189.
- [61] Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. 2020. Measuring Robustness to Natural Distribution Shifts in Image Classification. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 18583–18599. <https://proceedings.neurips.cc/paper/2020/file/d8330f857a17c53d217014ee776bfd50-Paper.pdf>
- [62] Philip Thomas and Emma Brunskill. 2016. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*. 2139–2148.
- [63] Masatoshi Uehara, Masahiro Kato, and Shota Yasui. 2020. Off-Policy Evaluation and Learning for External Validity under a Covariate Shift. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 49–61. <https://proceedings.neurips.cc/paper/2020/file/0084ae4bc24c0795d1e6a4f58444d39b-Paper.pdf>
- [64] Tim van Erven and Peter Harremoës. 2014. Rényi Divergence and Kullback-Leibler Divergence. *IEEE Transactions on Information Theory* 60, 7 (2014), 3797–3820. <https://doi.org/10.1109/TIT.2014.2320500>
- [65] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- [66] Wolfram Wiesemann, Daniel Kuhn, and Berç Rustem. 2013. Robust Markov Decision Processes. *Mathematics of Operations Research* 38, 1 (2013), 153–183. <https://doi.org/10.1287/moor.1120.0566> arXiv:<https://doi.org/10.1287/moor.1120.0566>
- [67] Amy Zhang, Clare Lyle, Shagun Sodhani, Angelos Filos, Marta Kwiatkowska, Joelle Pineau, Yarin Gal, and Doina Precup. 2020. Invariant Causal Prediction for Block MDPs. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 11214–11224. <http://proceedings.mlr.press/v119/zhang20t.html>
- [68] Zhengqing Zhou, Zhengyuan Zhou, Qinxun Bai, Linhai Qiu, Jose Blanchet, and Peter Glynn. 2021. Finite-Sample Regret Bound for Distributionally Robust Offline Tabular Reinforcement Learning. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 130)*, Arindam Banerjee and Kenji Fukumizu (Eds.). PMLR, 3331–3339. <https://proceedings.mlr.press/v130/zhou21d.html>

A IMPLEMENTATION DETAILS OF SOLVING DRO GIVEN A FINITE SAMPLE

The DRO problem for $\mathcal{U}_P^{\text{div}}$ solves,

$$\arg \min_{\theta \in \Theta} \left\{ \mathcal{R}(\theta, \mathcal{U}_P^{\text{div}}) := \sup_{Q \in \mathcal{U}_P^{\text{div}}} \mathbb{E}_{V \sim Q} [\ell(\theta, V)] \right\} \quad (11)$$

Using convex duality arguments [53], one can write the worst-case risk alternatively as,

$$\mathcal{R}(\theta, \mathcal{U}_P^{\text{div}}) = \sup_{Q \in \mathcal{U}_P^{\text{div}}} \mathbb{E}_{Z \sim Q(Z)} \mathbb{E}_{P(V|Z)} [\ell(\theta, V)|Z] = \inf_{\eta \in \mathbb{R}} \frac{1}{\delta} \mathbb{E}_P [(\mathbb{E}_P [\ell(\theta, V)|Z] - \eta)_+] + \eta \quad (12)$$

where $(\cdot)_+ = \max(\cdot, 0)$. Note that this requires estimating $\mathbb{E}_{P(V|Z)}[\cdot]$ which may be hard if Z is continuous-valued or we do not have enough data for all possible values of Z . Assuming smoothness of this conditional loss, Duchi et al. [11, Lemma 4.2] gives an upper bound for the worst-case risk with the empirical version as,

$$\widehat{\mathcal{R}}(\theta, \mathcal{U}_P^{\text{div}}) = \inf_{\substack{\eta \geq 0, \\ B \in \mathbb{R}_+^{L \times n}}} \left\{ \frac{1}{\delta} \left(\frac{1}{n} \sum_{i=1}^n \left(\ell(\theta, V_i) - \frac{1}{n} \sum_{j=1}^n (B_{ij} - B_{ji}) - \eta \right)_+ \right)^2 \right\}^{\frac{1}{2}} + \frac{L}{\epsilon n^2} \sum_{i,j=1}^n \sum_{O \in \mathcal{Z}} \|O_i - O_j\| B_{ij} + \eta \quad (13)$$

for any $\epsilon > 0$, where η, B are dual variables. We use this estimator while solving the DRO problems in the contextual bandit experiments. The minimization is performed using gradient descent as the objective is convex in both η and B .

B ASSUMPTIONS FOR MDPS – SA-RECTANGULARITY

We assume that the uncertainty sets are SA-rectangular [21], which intuitively implies that the uncertainty sets are constructed across time steps independently. This property yields a tractable method to compute the value function estimates using dynamic programming [21].

DEFINITION 1 (SA-RECTANGULARITY). For any state-action pair (s, a) , consider the uncertainty set $\mathcal{U}^{\text{MDP}}(s, a)$ of plausible test-time transition models. Denote the collection of all such uncertainty sets by $\mathcal{U}^{\text{MDP}} \ni \mathcal{U}^{\text{MDP}}(s, a)$. Then, SA-rectangularity means that the collection \mathcal{U}^{MDP} is constructed by taking the Cartesian product of the individual sets $\mathcal{U}^{\text{MDP}}(s, a)$.

$$\mathcal{U}^{\text{MDP}} = \prod_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathcal{U}^{\text{MDP}}(s, a)$$

For example, take an MDP with one state and two actions, $\mathcal{S} = \{0\}$, $\mathcal{A} = \{0, 1\}$. If $\{\{P_{00}^1, P_{00}^2\}, \{P_{01}^1, P_{01}^2\}\}$ are the two uncertainty sets for the two state-action pairs $\{00, 01\}$ respectively, then the four possible uncertainty sets for the robust MDP are $\{\{P_{00}^1, P_{01}^1\}, \{P_{00}^1, P_{01}^2\}, \{P_{00}^2, P_{01}^1\}, \{P_{00}^2, P_{01}^2\}\}$.

We note that some policy evaluation methods for robust MDPS require weaker assumptions than SA-rectangularity [66].

C PROOF OF THEOREM 1 – ERROR BOUND FOR ROBUST OPE WITH ESTIMATED MODEL

We solve the robust OPE problem using dynamic programming with the *estimated* transition and reward models instead of assuming access to the true models. Hence, we will incur an error due to using finite samples to estimate the models and the value. Here, we will show that this error is of the order of, as in the non-robust OPE case, $\tilde{O}\left(\frac{\gamma|\mathcal{S}|}{\sqrt{n}(1-\gamma)^2}\right)$ (ignoring logarithmic factors).

Notation. The transition model for the training environment is written as P_0 . We will abbreviate the uncertainty set defined with respect to P_0 for each state-action pair, that is $\mathcal{U}_{P_0}(s, a)$, as \mathcal{U}_{P_0} . For finite state-action space, a policy's value V^π is a $|\mathcal{S}|$ dimensional vector. Thus, the L_∞ norm for this Banach space is $\|V^\pi\|_\infty = \max_s |V^\pi(s)|$. The robust Bellman operator is defined as $\mathcal{T}_{\mathcal{U}_{P_0}} V^\pi = r + \inf_{P \in \mathcal{U}_{P_0}} \gamma \langle P, V^\pi \rangle$. The reward and transition models are denoted by $M := (r_0, P_0)$, which are $|\mathcal{S}| \times |\mathcal{A}|$ and $|\mathcal{S}| \times |\mathcal{A}| \times |\mathcal{S}|$ tensors respectively. Denote the estimated models by $\widehat{M}_0 := (\widehat{r}_0, \widehat{P}_0)$. The robust value for policy π and state s is denoted by $V_{\mathcal{U}_{P_0}}^\pi(s)$. We exclude \widehat{r}_0 to avoid excessive notation.

When this value is computed instead with the estimated models \widehat{M}_0 , as done in our experiments, we will denote the value by $V_{\widehat{\mathcal{U}}_{\widehat{P}_0}}^\pi(s)$.

Thus, our goal is to bound the error between $V_{\mathcal{U}_{P_0}}^\pi(s)$ and $V_{\widehat{\mathcal{U}}_{\widehat{P}_0}}^\pi(s)$. For the non-robust MDP case, this result is referred to as *simulation lemma* [25]. It bounds the error in value estimates incurred from using the estimated models instead of the true ones. An alternative statement of the lemma and an accessible proof is provided in [23, Lemma 1]. In Lemma 1, we will first bound the difference between $V_{\mathcal{U}_{P_0}}^\pi(s)$ and $V_{\widehat{\mathcal{U}}_{\widehat{P}_0}}^\pi(s)$ assuming that the estimation error for the reward and transition models is bounded. Then, we will use standard concentration inequalities to prove the error bounds. Before proving the result, we make some observations required in the proof.

Remark 1 (Upper bound of value). Assume that the rewards in the MDP are bounded, $r \in [0, r_{\max}]$. Then, the robust value is upper bounded by $r_{\max}/(1-\gamma)$. This follows from the definition of the robust value as the infimum,

$$V_{\mathcal{U}_{P_0}}^\pi(s) = \inf_{P \in \mathcal{U}_{P_0}} \mathbb{E}_{\pi, P} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 \right] \leq \mathbb{E}_{\pi, P_0} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 \right] \leq \sum_{t=0}^{\infty} \gamma^t r_{\max} = \frac{r_{\max}}{1-\gamma}$$

Remark 2 (Contraction property). From Iyengar [21, Theorem 5], we know that the robust Bellman operator $\mathcal{T}_{\mathcal{U}_{P_0}}$ is a *contraction mapping* in L_∞ norm. That is, for any two value functions U and V ,

$$\|\mathcal{T}_{\mathcal{U}_{P_0}} V - \mathcal{T}_{\mathcal{U}_{P_0}} U\|_\infty \leq \gamma \|V - U\|_\infty \quad (14)$$

Now, we prove a lemma from which the error bound follows directly.

LEMMA 1 (SIMULATION LEMMA FOR ROBUST MDPs). *Given $\max_{s,a} |r(s,a) - \widehat{r}(s,a)| \leq \epsilon_r$ and $\max_{s,a} \|P_0(\cdot|s,a) - \widehat{P}_0(\cdot|s,a)\|_1 \leq \epsilon_P$, then for any policy π ,*

$$\|V_{\mathcal{U}_{P_0}}^\pi - V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty \leq \frac{\epsilon_r}{1-\gamma} + \frac{\gamma \epsilon_P r_{\max}}{(1-\gamma)^2}$$

PROOF. The proof follows the exposition of the simulation lemma by Jiang [23] for non-robust MDPs. The steps involve using the robust Bellman operator, then noting that it is a contraction, and finally, using the fact that robust value which is defined as infimum over the uncertainty set is lower than the value under the training transition model.

Recall that the robust value with respect to an uncertainty set defined using P_0 satisfies the following robust Bellman equation,

$$V_{\mathcal{U}_{P_0}}^\pi(s) = \sum_a \pi(a|s) \left(r(s,a) + \inf_{P \in \mathcal{U}_{P_0}} \gamma \langle P, V_{\mathcal{U}_{P_0}}^\pi \rangle \right)$$

Here, the policy π can be either probabilistic or deterministic, in which case $\pi(a|s) = 1$ for one of the actions a .

For all $s \in \mathcal{S}$, consider the difference between robust values defined with respect to P_0 and \widehat{P}_0 ,

$$\begin{aligned} \left| V_{\mathcal{U}_{P_0}}^\pi(s) - V_{\mathcal{U}_{\widehat{P}_0}}^\pi(s) \right| &= \left| \sum_a \pi(a|s) (r(s,a) + \inf_{P \in \mathcal{U}_{P_0}} \gamma \langle P, V_{\mathcal{U}_{P_0}}^\pi \rangle) - \sum_a \pi(a|s) (\widehat{r}(s,a) + \inf_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle) \right| \\ &\leq \left| \sum_a \pi(a|s) (r(s,a) - \widehat{r}(s,a)) \right| + \left| \sum_a \pi(a|s) \left(\inf_{P \in \mathcal{U}_{P_0}} \gamma \langle P, V_{\mathcal{U}_{P_0}}^\pi \rangle - \inf_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle \right) \right| \\ &\leq \epsilon_r + \underbrace{\sum_a \pi(a|s) \left| \inf_{P \in \mathcal{U}_{P_0}} \gamma \langle P, V_{\mathcal{U}_{P_0}}^\pi \rangle - \inf_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle \right|}_{\text{Part (I)}} \end{aligned} \quad (15)$$

where we used the bound on estimated rewards from the assumption. We now bound Part (I) for each pair (s, a) .

$$\begin{aligned} \text{Part (I)} &\leq \left| \left(\inf_{P \in \mathcal{U}_{P_0}} \gamma \langle P, V_{\mathcal{U}_{P_0}}^\pi \rangle - \inf_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle \right) + \left(\inf_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle - \inf_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle \right) \right| \\ &= \left| (\mathcal{T}_{\mathcal{U}_{P_0}} V_{\mathcal{U}_{P_0}}^\pi - \mathcal{T}_{\mathcal{U}_{\widehat{P}_0}} V_{\mathcal{U}_{\widehat{P}_0}}^\pi) + \left(\inf_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle - \inf_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle \right) \right| \\ &\leq \gamma \|V_{\mathcal{U}_{P_0}}^\pi - V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty + \left| \left(\inf_{P \in \mathcal{U}_{P_0}} \gamma \langle P, V_{\mathcal{U}_{P_0}}^\pi \rangle - \inf_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle \right) \right| \end{aligned} \quad (16)$$

$$\leq \gamma \|V_{\mathcal{U}_{P_0}}^\pi - V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty + \underbrace{\left| \left(\gamma \langle P_0, V_{\mathcal{U}_{P_0}}^\pi \rangle - \inf_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle \right) \right|}_{\text{Part (II)}} \quad (17)$$

$$\text{Part (II)} \quad (18)$$

where we used the contraction property of $\mathcal{T}_{\mathcal{U}_{P_0}}$ from Remark 2 in (16) and replaced \inf with a member from the set in (17).

Using the Hölder's inequality and the upper bound from Remark 1, we can write Part (II) as,

$$\begin{aligned} \text{Part (II)} &= \left| \sup_{P \in \mathcal{U}_{\widehat{P}_0}} \gamma \langle P_0 - P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle \right| \leq \gamma \sup_{P \in \mathcal{U}_{\widehat{P}_0}} |\langle P_0 - P, V_{\mathcal{U}_{\widehat{P}_0}}^\pi \rangle| \\ &\leq \gamma \sup_{P \in \mathcal{U}_{\widehat{P}_0}} \|P_0 - P\|_1 \|V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty \\ &= \gamma \sup_{P \in \mathcal{U}_{\widehat{P}_0}} \|P_0 - P\|_1 \frac{r_{\max}}{1-\gamma} \leq \gamma \sup_{P \in \mathcal{U}_{\widehat{P}_0}} (\|P_0 - \widehat{P}_0\|_1 + \|\widehat{P}_0 - P\|_1) \frac{r_{\max}}{1-\gamma} \end{aligned}$$

We re-write the ℓ_1 norm in terms of total variation distance (TV) [29, Prop 4.2] and use Pinsker's inequality [5, Lemma 2],

$$\text{Part (II)} \leq \gamma \epsilon_P \frac{r_{\max}}{1-\gamma} + \gamma \sup_{P \in \mathcal{U}_{\widehat{P}_0}} 2 \text{TV}(\widehat{P}_0, P) \frac{r_{\max}}{1-\gamma} \leq \gamma \epsilon_P \frac{r_{\max}}{1-\gamma} + \gamma \sup_{P \in \mathcal{U}_{\widehat{P}_0}} 2 \sqrt{\frac{1}{2} \text{KL}(\widehat{P}_0, P)} \frac{r_{\max}}{1-\gamma}$$

By the definition of the uncertainty set $\mathcal{U}_{\widehat{P}_0}$, all elements P satisfy $\text{KL}(\widehat{P}_0, P) \leq \delta$.

$$\text{Part (II)} \leq \gamma \epsilon_P \frac{r_{\max}}{1-\gamma} + \gamma \sqrt{2\delta} \frac{r_{\max}}{1-\gamma} = \gamma(\epsilon_P + \sqrt{2\delta}) \frac{r_{\max}}{1-\gamma} =: \gamma(\epsilon'_P) \frac{r_{\max}}{1-\gamma}$$

where we define $\epsilon'_P := \epsilon_P + \sqrt{2\delta}$. Now, we plug this back in (17) and use $\sum_a \pi(a|s) = 1$ in (15) to get,

$$\left| V_{\mathcal{U}_{P_0}}^\pi(s) - V_{\mathcal{U}_{\widehat{P}_0}}^\pi(s) \right| \leq \epsilon_r + \gamma \|V_{\mathcal{U}_{P_0}}^\pi - V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty + \gamma \epsilon'_P \frac{r_{\max}}{1-\gamma}$$

Since the above holds for any $s \in \mathcal{S}$, taking maximum over \mathcal{S} , we have that

$$\|V_{\mathcal{U}_{P_0}}^\pi - V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty \leq \epsilon_r + \gamma \|V_{\mathcal{U}_{P_0}}^\pi - V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty + \gamma \epsilon'_P \frac{r_{\max}}{1-\gamma}$$

$$\|V_{\mathcal{U}_{P_0}}^\pi - V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty \leq \frac{\epsilon_r}{1-\gamma} + \frac{\gamma \epsilon'_P r_{\max}}{(1-\gamma)^2}$$

■

Restating the result that we want to prove,

THEOREM 1 (ESTIMATION ERROR FOR ROBUST OPE). *Given at least n samples from $P_0(\cdot|s, a)$ for all s, a , assuming that the rewards are bounded $r \in [0, r_{\max}]$, and that \mathcal{U}^{MDP} are defined by KL-divergence, then with probability at least $1 - \alpha$,*

$$\|V_{\mathcal{U}_{P_0}}^\pi - V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty \leq O\left(\frac{\gamma r_{\max} |\mathcal{S}|}{(1-\gamma)^2} \sqrt{\frac{1}{n} \log\left(\frac{4|\mathcal{S} \times \mathcal{A} \times \mathcal{S}|}{\alpha}\right)}\right)$$

PROOF. We use Lemma 1 along with Hoeffding's inequality [19] to bound the error in terms of number of samples and state-action pairs.

Say, we have n samples to estimate $\widehat{r}_0(s, a), \widehat{P}_0(s, a)$. Then, after using Hoeffding's inequality for each state-action pair and then with the union bound, we have that with probability at least $1 - \alpha$,

$$\begin{aligned} \max_{s,a} |r(s, a) - \widehat{r}(s, a)| &\leq r_{\max} \sqrt{\frac{1}{2n} \log\left(\frac{4|\mathcal{S} \times \mathcal{A}|}{\alpha}\right)}, & \max_{s,a,s'} |P_0(s'|s, a) - \widehat{P}_0(s'|s, a)| &\leq \sqrt{\frac{1}{2n} \log\left(\frac{4|\mathcal{S} \times \mathcal{A} \times \mathcal{S}|}{\alpha}\right)}, \\ & \implies \max_{s,a} \|P_0(\cdot|s, a) - \widehat{P}_0(\cdot|s, a)\|_1 &\leq |\mathcal{S}| \sqrt{\frac{1}{2n} \log\left(\frac{4|\mathcal{S} \times \mathcal{A} \times \mathcal{S}|}{\alpha}\right)}. \end{aligned}$$

Thus, $\|V_{\mathcal{U}_{P_0}}^\pi - V_{\mathcal{U}_{\widehat{P}_0}}^\pi\|_\infty = \tilde{O}\left(\frac{\gamma |\mathcal{S}|}{\sqrt{n(1-\gamma)^2}}\right)$. ■

The dependence on $1/\sqrt{n}$ matches that in the non-robust case. The dependence on $|\mathcal{S}|$ can be improved to $\sqrt{|\mathcal{S}|}$ using a better bound for $\max_{s,a,s'} |P_0(s'|s, a) - \widehat{P}_0(s'|s, a)|$, as done in [23, Section 2.2]. We hope that further work will tighten this bound using techniques from [31]. For Lemma 1, we restricted to KL-divergence in order to use the Pinsker's inequality to bound the ℓ_1 distance $\|P_0 - P\|_1$. We are not aware of a Pinsker-type inequality for our preferred divergence measure $\mathcal{D}_{\text{CVaR}}$ (while one exists for Rényi divergences of order $\in (0, 1]$, but not for order ∞ which is the same as $\mathcal{D}_{\text{CVaR}}$ [64]). While the proposed method for MDPs works for KL-divergence as well, we use $\mathcal{D}_{\text{CVaR}}$ in the experiments due to its intuitive interpretation in terms of worst-case subpopulations.

D ALGORITHM FOR ROBUST OPE FOR MDPs

Robust OPE with dynamic programming amounts to solving the following fixed point equation iteratively,

$$V^\pi(s) = r(s, \pi(s)) + \inf_{P \in \mathcal{U}^{\text{MDP}}(s, \pi(s))} \gamma \mathbb{E}_{s' \sim P(s'|s, \pi(s))} [V^\pi(s')] \quad (19)$$

At each iteration, we have to solve the DRO problem,

$$\inf_{P \in \mathcal{U}^{\text{MDP}}(s, \pi(s))} \mathbb{E}_P[V^\pi(s')] = \inf_{P \in \mathcal{U}^{\text{MDP}}(s, \pi(s))} \mathbb{E}_{P(s^1|s, \pi(s))} \left[\mathbb{E}_{P_0(s^2|s, \pi(s), s^1)} [V^\pi(s')] \right] =: \mathcal{R}_{\mathcal{U}^{\text{MDP}}(s, \pi(s))} (V^\pi(s'))$$

We estimate the inner expectation w.r.t. P_0 with Monte-Carlo averaging on the batch data available to us as this data contains samples from P_0 . We then compute the value function update using

$$V^\pi(s) \leftarrow \sum_a \pi(a|s) \left(r(s, a) + \gamma \mathcal{R}_{\mathcal{U}^{\text{MDP}}(s, a)} (V^\pi(s')) \right). \quad (20)$$

As earlier, we choose the divergence metric needed to define $\mathcal{U}^{\text{MDP}}(s, a)$ to be $\mathcal{D}_{\text{CVaR}}$, and accordingly, solve the DRO problem using the form in Eq. (12). This requires access to the true transition model, P_0 , and the reward model, r . Since we only have samples from P_0 , we use the maximum likelihood estimate \widehat{P}_0 from the batch data in place of P_0 . For finite state-action space, this amounts to counting transitions for each tuple (s, a, s') and averaging. Here, an important requirement for the batch data collected using policy μ is that it has sufficient

exploration to evaluate the given policy π . Formally, $\pi(a|s) > 0 \implies \mu(a|s) > 0$ for all s, a . In words, the support of the batch policy μ contains the support of evaluation policy π . Further, the transition models in $\mathcal{U}^{\text{MDP}}(s, a)$ that we will have to evaluate are also absolutely continuous with respect to the training transition model P_0 , as assumed in the set definition (3). Thus, given enough samples in the batch data collected using μ and P_0 , we can estimate the reward and transition models for all s, a, s' needed in the value function updates for π and \mathcal{U}^{MDP} in (20). Algorithm 2 gives the overall procedure in detail.

Extension to continuous state space. Although we specify the method for finite state MDPs, it can be extended to continuous or large state spaces by leveraging linear function approximations along with robust dynamic programming, as proposed in [60]. Validating such approximations for ROPE is a fruitful direction for more work.

E EXPERIMENTAL DETAILS

E.1 More detail on feature shift detection

Details of the three datasets used in Figure 1 and the process followed for feature extraction is given in Adult Income [8], SBA [32], and eICU [24]. We use the score-based feature shift detection method (named MB-SM) by [27] available at <https://github.com/inouye-lab/feature-shift>.

E.2 Choosing the uncertainty set size δ

As our problem setup demands robustness to *unknown* shifts at the test time, the choice of uncertainty set size (i.e. robustness level δ) and other hyperparameters necessarily involves making assumptions about future data. For the case of $\mathcal{D}_{\text{CVaR}}$ based sets, the set size corresponds to the minimum possible proportion of the worst-performing subpopulation. The modeler can exercise domain knowledge in choosing such a proportion, for instance, using the summary statistics of demographic data for the intended test populations.

For example, if the modeler worries about shifts in the age distribution. Say the train distribution has 50-50% young-old population, and at the worst the test distribution has 80% overlap, i.e. old population is at the least 40% (=0.8*50%) at test time, then $\delta = 0.8$ is reasonable to guarantee robustness. Summary statistics of demographics for the test environments (age, gender, comorbidity distribution at test hospitals) may help judge the size of the population in attributes of concern.

E.3 Data generating process for synthetic data in CBs

We generate data according to the causal graph in Figure 2 with two features $Z := (Z_1, Z_2)$ in the context, binary treatment T and continuous outcome Y . Structural equations are:

$$E \sim \text{Bernoulli}(-1, 1, \delta_0), Z_1 \sim E \cdot \mathcal{N}(10, 1), Z_2 \sim \mathcal{N}(5, 1) \quad Y_{T=t} \sim \mathcal{N}(Z^T w_t, 0.1^2), t \in \{0, 1\}, \text{ where } w_0 = [0.1, 0.1], w_1 = [0.1, 0.5]$$

$$T \sim \text{Bernoulli}(\sigma(Z^T \beta + \beta_0)), \beta = [0.1, 0.1]$$

where $\sigma(z) = 1/(1 + e^{-z})$. The bias term in the train policy is $\beta_0 = -1$, whereas the policy to be evaluated in test environments has $\beta_0 = -0.5$. In addition, the marginal distribution of Z_1 is changed, via change in E , by increasing δ_0 in the test environment. We simulate $n = 2000$ samples in the train environment and use it for all methods. The objective is to estimate the robust value of the new policy.

E.4 More details on Cliffwalking domain

We consider a 6×6 gridworld (Figure 5) with start, goal positions, and a cliff on one edge [59, Ex. 6.6]. The agent incurs a rewards -100 on falling off the cliff, -1 for each step, and 0 on reaching the goal. To add stochasticity in transitions, we make two changes to the domain. With a constant shift probability, the agent slips down one step towards the cliff instead of taking the prescribed action. This shift probability varies across environments, thus, changing the transition dynamics, necessitating robustness in policy evaluation. Second, we consider two state features in the observation space of the agent. The first feature is the agent’s position on the grid, and the second is discrete random noise sampled uniformly from the set $\{1, 2, \dots, 6 \times 6\}$. The second feature has an associated reward sampled from $\text{Uniform}(-1, 0)$ for the 6×6 values it takes. Total reward is the sum of rewards from the two features (one based on grid position and one sampled uniformly between $-1, 0$).

Since agent’s actions affect only the first feature, ROPE correctly constructs uncertainty sets based on transition dynamics in the first feature alone, $P(s^1|s, a)$. In contrast, JointDRO ignores this structure and constructs uncertainty sets using both the features, $P((s^1, s^2)|s, a)$. We evaluate the value function estimate for an agent following uniform random policy using dynamic programming with the standard Bellman equation (Standard) or the robust one in Eq. (10) (JointDRO, ROPE).

E.5 More details on Sepsis domain

The RL policy to be evaluated is obtained as follows. The optimal policy for the Sepsis simulator, namely the physician policy, is found with the procedure used in [41]. Physician policy is made stochastic by taking random action with probability 0.05 . With this policy, a dataset with 1000 trajectories each with maximum length of 20 is sampled. Diabetic population is fixed to 20% while sampling. The RL policy used in evaluation is obtained by running policy iteration on this dataset. Then, RL policy is used to sample another dataset with 10000 trajectories each with maximum length of 20, and 20% diabetics. This data is used for the final evaluation. As the policy being evaluated is the same as the one used to sample the available data, the importance weighting done in Eq. (20) is not needed.