
Multi-Objective SPIBB: Seldonian Offline Policy Improvement with Safety Constraints in Finite MDPs

Harsh Satija

McGill University, Mila
harsh.satija@mail.mcgill.ca

Philip S. Thomas

University of Massachusetts
pthomas@cs.umass.edu

Joelle Pineau

McGill University, Mila, Facebook AI Research
jpineau@cs.mcgill.ca

Romain Laroche

Microsoft Research
romain.laroche@microsoft.com

Abstract

We study the problem of Safe Policy Improvement (SPI) under constraints in the offline Reinforcement Learning (RL) setting. We consider the scenario where: (i) we have a dataset collected under a known baseline policy, (ii) multiple reward signals are received from the environment inducing as many objectives to optimize. We present an SPI formulation for this RL setting that takes into account the preferences of the algorithm’s user for handling the trade-offs for different reward signals while ensuring that the new policy performs at least as well as the baseline policy along each individual objective. We build on traditional SPI algorithms and propose a novel method based on Safe Policy Iteration with Baseline Bootstrapping (SPIBB, [Laroche et al., 2019](#)) that provides high probability guarantees on the performance of the agent in the true environment. We show the effectiveness of our method on a synthetic grid-world safety task as well as in a real-world critical care context to learn a policy for the administration of IV fluids and vasopressors to treat sepsis.

1 Introduction

Reinforcement Learning (RL) as a paradigm for sequential decision-making ([Sutton, 1988](#)) has shown tremendous success in a variety of simulated domains ([Mnih et al., 2015](#); [Silver et al., 2017](#); [OpenAI, 2018](#)). However, there are still quite a few challenges between the traditional RL research and real-world tasks. Most of these challenges stem from assumptions that are rarely satisfied in practice ([Dulac-Arnold et al., 2019](#)), or the inability of the algorithm’s user to specify the desired behavior of the agent without being a domain expert ([Thomas et al., 2019](#)). We focus on the real-world application point of view and posit the following requirements:

- **Multiple reward functions:** Traditional RL methods assume a single scalar reward is present in the environment. However, most real-world tasks, have multiple (possibly conflicting) objectives or constraints that need to be taken into consideration together, such as the signals related to the safety (physical well-being of the agent or the environment), budget utilization (energy or maintenance costs), etc.
- **Stakeholder control of the trade-off:** The ML practitioners should have the ability to control the different trade-offs the agent is making and choose the one they consider best for the task at hand.
- **Offline setting:** In many real-world domains (e.g., healthcare, finance or autonomous vehicles), there is an abundance of data, collected under a sub-optimal policy, but training the agent directly via interactions with the environment is expensive and risky. We assume that we only have access to a dataset of past trajectories that can be used for training ([Lange et al., 2012](#)).

- **Preventing unintended behavior:** We want the agent to be robust to both extrapolation errors from offline RL and misaligned objectives that are poor proxy of the user’s intentions and algorithm’s actual performance (Ng et al., 1999; Amodei et al., 2016). We consider the case where the user can specify undesirable behavior in the context of the performance observed in the batch.
- **Practical guarantees:** We want guarantees about the undesirable behavior that might be caused by the agent in the real-world. We care about the results that can be obtained using the finite amount of samples we have in the batch, and aim to provide some measure of confidence in deploying the agents in the environment.

To achieve this set of properties, we adopt the Seldonian framework (Thomas et al., 2019), which is a general algorithm design framework that allows high-confidence guarantees for constraint satisfaction in a multi-objective setting. Based on the above specifications, we seek to answer the question: *if we are given a batch of data collected under some (suboptimal) behavioral policy and some user preference, can we build a policy improvement algorithm that returns a policy with practical high-confidence guarantees on the performance of the policy w.r.t. the behavioral policy?*

We acknowledge that there are other important challenges in RL, such as partial observability, safe exploration, non-stationary environments and function approximation in high-dimensional spaces, that also stand in the way of making RL a more applicable paradigm. These challenges are beyond the scope of this work, which should rather be thought of as taking a step towards this broader goal.

In Section 2, we present our contribution positioned with respect to other related work. In Section 3, we formalize the setting and then extend traditional SPI algorithms to this setting. We then show it is possible to extend the previous work on Safe Policy Iteration (SPI), particularly Safe Policy Iteration with Baseline Bootstrapping (SPIBB, Laroche et al., 2019), for the design of agents that satisfy the above requirements. We show that the resulting algorithm is theoretically-grounded and provides practical high-probability guarantees. We extensively test our approach on a synthetic safety-gridworld task in Section 4 and show that the proposed algorithm achieves better data efficiency than the existing approaches. Finally, we show its benefits on a critical-care task in Section 5. The accompanying codebase is available at <https://github.com/hercky/mo-spibb-codebase>.

2 Related work

Multi-Objective RL (MORL): Traditional multi-objective approaches (Mannor and Shimkin, 2004; Roijers et al., 2013; Liu et al., 2014) focus on finding the Pareto-frontier of optimal reward functions that gives all the possible trade-offs between different objectives. The user can then select a policy from the solution set according to their arbitrary preferences. In practice, an alternate trial and error based approach of scalarization is used to transform the multiple reward functions into a scalar reward based on preferences across objectives (usually, by taking a linear combination). Most traditional MORL approaches have focused on the online, interactive settings where the agent has access to the environment. While some recent approaches are based on off-policy learning methods (Lizotte et al., 2012; Van Moffaert and Nowé, 2014; Yang et al., 2019; Abdolmaleki et al., 2020), they lack guarantees. In contrast, our work focuses exclusively on learning in the offline setting and gives high-probability guarantees on the performance in the environment.

Constrained-RL: RL under constraints frameworks, such as Constrained MDPs (CMDPs, Altman, 1999), present an alternative way to define preferences in the form of constraints over policy’s returns. Here, the user assigns a single reward function to be the primary objective (to maximize) and hard constraints are specified for the others. The major limitation of this setting is that it assumes the thresholds for the constraints are known a priori. Le et al. (2019) study offline policy learning under constraints and provide performance guarantees w.r.t. the optimal policy, but their work relies on the concentrability assumption (Munos, 2003).

Concentrability is a strong assumption that upper bounds the ratio between the future state-action distributions of any non-stationary policy and the baseline policy under which the dataset was generated by some constant. From a practical perspective, it is unclear how to get a tractable estimate of this constant, as the space of future state-action distributions of non-stationary policies is vast. Thus, this constant can be arbitrarily huge, potentially even infinite when the baseline policy fails to cover the support of the space of all non-stationary policies (such as in the low-data regime), leading to the performance bounds given by these methods to blow up (and even be unbounded). Additionally, the guarantees in Le et al. (2019) are only valid with respect to the performance of the optimal policy.

In this work, we instead focus on the performance guarantees based on returns observed in the dataset, as it does not require making any of the above assumptions.

Reward design: Reward-design (Sorg et al., 2010) and reward-modelling approaches (Christiano et al., 2017; Littman et al., 2017; Leike et al., 2018) focus on designing suitable reward functions that are consistent with the user’s intentions. These approaches rely heavily on the human or simulator feedback, and thus do not carry over easily to the offline setting.

Seldonian-RL (and Safe Policy Improvement): The Seldonian framework (Thomas et al., 2019) is a general algorithm design framework that allows the user to design ML algorithms that can avoid undesirable behavior with high-probability guarantees. In the context of RL, the Seldonian framework allows to design policy optimization problems with multiple constraints, where the solution policies satisfy the constraints with high-probability. In the offline-RL setting, SPI refers to the objective of guaranteeing a performance improvement over the baseline with high-probability guarantees (Thomas et al., 2015a; Petrik et al., 2016; Laroche et al., 2019). Therefore, SPI algorithms are a specific setting that falls in the general category of Seldonian-RL algorithms.

We focus on two categories of SPI algorithms that provide practical error bounds on safety: SPIBB (Laroche et al., 2019) that provides Bayesian bounds and HCPI (Thomas et al., 2015a,b) that provides frequentist bounds. SPIBB methods constrain the change in the policy according to the local model uncertainty. SPIBB has been formulated in the context of a single reward function, and as such does not handle multiple rewards and by extension also lacks the ability for the user to specify preferences. Our primary focus is to provide a construction for extending the SPIBB methodology to the multi-objective setting that handles user preferences and provides high-probability guarantees.

Instead of relying on model uncertainty, HCPI methods utilize the high-confidence lower bounds on the Importance Sampling (IS) estimates of a target policy’s performance to ensure safety guarantees. HCPI has been applied to solve Seldonian optimization problems for constrained-RL setting using an enumerable policy class. Thomas et al. (2019) suggested using HCPI for the MORL setting, and we build on that idea. Particularly, we show how HCPI can be implemented with stochastic policies in the context of our setting with user preferences and baseline constraints.

3 Methodology

3.1 Setting

We consider the setting where the agent’s interactions with the environment can be modelled as a Markov Decision Process (MDP, Bellman, 1957). Let \mathcal{X} and \mathcal{A} respectively be the (finite) state and action spaces. Let $p^* : \mathcal{X} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{X})$ denote the true (unknown) transition probability function, where $\mathcal{P}(\mathcal{X})$ denotes the set of probability distributions on \mathcal{X} . Without loss of generality, we assume that the process deterministically begins in the state x_0 . We define $[N]$ to be the set $\{0, 1, \dots, N - 1\}$ for any positive integer N . Let there be d different reward signals and $\mathbf{r}^* = \{r_k\}_{k \in [d]} : \mathcal{X} \times \mathcal{A} \rightarrow [-r_\tau, r_\tau]^d$ be the true (unknown) stochastic multi-reward signal.¹ Finally, $\gamma = \{\gamma_k\}_{k \in [d]} \in [0, 1)^d$ is the multi-discount-factor.

The MDP, m^* , can now be defined with the tuple $(\mathcal{X}, \mathcal{A}, p^*, \mathbf{r}^*, \gamma, x_0)$. A policy $\pi : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{A})$ maps a state to a distribution over actions. We denote by Π the set of stochastic policies. We consider the infinite horizon discounted return setting. For any $k \in [d]$, the k^{th} reward value function $v_{m,k}^\pi(x) : \mathcal{X} \rightarrow \mathbb{R}$ denotes the expected discounted sum of rewards when following policy π in an MDP m starting from state x . Analogously, we define the state-action value functions for performing action a in state x in MDP m under π for rewards as $q_{m,k}^\pi(x, a)$. Let $\text{Adv}_{m,k}^\pi(x, a) = q_{m,k}^\pi(x, a) - v_{m,k}^\pi(x)$ denote the corresponding advantage function. The expected return of policy π w.r.t. the k^{th} reward in the true MDP m^* is denoted by $\mathcal{J}_{m^*,k}^\pi = v_{m^*,k}^\pi(x_0) = \mathbb{E}_{\pi, m^*}[\sum_{t=0}^\infty \gamma_k^t R_{k,t} \mid X_0 = x_0]$, where action $A_t \sim \pi(\cdot \mid X_t)$, immediate reward $R_{k,t} \sim r_k^*(\cdot \mid X_t, A_t)$, and state $X_{t+1} \sim p^*(\cdot \mid X_t, A_t)$.

We consider the offline setting, where instead of having access to the environment we have a pre-collected dataset of trajectories denoted by $\mathcal{D} = \{\tau_i\}_{i \in [|\mathcal{D}|]}$, where $|\mathcal{D}|$ denotes the number of trajectories in the dataset. A trajectory τ of length T is an ordered set of transition tuples of the form $\tau = \{x_i, a_i, x'_i, \mathbf{r}_i\}_{i \in [T]}$, where x'_i denotes the state at the next time-step. We denote the Maximum

¹Costs, which are meant to be minimized, can be expressed as negative rewards.

Likelihood Estimation (MLE) of the MDP with $\hat{m} = (\mathcal{X}, \mathcal{A}, \hat{p}, \hat{r}, \gamma, x_0)$, where \hat{p} and \hat{r} denote the transition and reward models estimated from the dataset’s statistics.

Assumption 3.1 (Baseline policy). We assume that we have access to the policy that generated the dataset. We call such policy the baseline policy and denote it by π_b .²

3.2 Problem formulation

We consider safe policy improvement with respect to the baseline according to the d dimensions of the multi-objective setting. Therefore, under a Bayesian approach, we search for target policies such that they perform better (up to a precision error ζ) than the baseline along every objective function with high probability $1 - \delta$, where ζ and δ are hyper-parameters controlled by the user, denoting the risk that the practitioner is willing to take. We denote by Π_A the set of admissible policies that satisfy:

$$\mathbb{P} \left(\forall k \in [d], \mathcal{J}_{m^*,k}^\pi - \mathcal{J}_{m^*,k}^{\pi_b} > -\zeta \mid \mathcal{D} \right) > 1 - \delta. \quad (1)$$

In the multi-objective case, there does not exist a single optimal value, but a Pareto frontier of optimal values. One way to evaluate the MORL problems is via the *multiple-policy* approaches (Vamplew et al., 2011; Roijers et al., 2013) that compute the policies that approximate the true optimal Pareto-frontier. However, note that optimality and safety are contradicting objectives. It is not clear how (and if) one can make claims about optimality in the offline setting without bringing in additional unrealistic assumptions (Section 2, MORL). Instead, we take an alternate approach inspired by another category of MORL methods called *single-policy* (Roijers et al., 2013; Van Moffaert and Nowé, 2014) where the trade-offs between different objectives are explicitly controlled by the user via providing a scalarization or preferences over objectives. We assume the user preference $\lambda = \{\lambda_k\}_{k \in [d]}$ is given as an input to our algorithms, and is used for scalarization of the objectives, where $\lambda_k \in \mathbb{R}^+$. Our objective therefore becomes

$$\arg \max_{\pi \in \Pi_A} \sum_{k \in [d]} \lambda_k \mathcal{J}_{m^*,k}^\pi. \quad (2)$$

The above formulation gives freedom to the user in terms of what particular quantity they want to optimize via λ , but still ensures that the solution policy performs as well as the baseline policy across all d objectives. Note that our explicit goal is to maximize the objective specified by the user. However, the user might make mistakes in specifying this objective (Section 2, Reward design), and the above formulation offers guarantees that prevent deteriorating the performance of the policy across any of the d objectives. This allows the user to experiment with different reward design strategies in safety-critical settings without worrying about the risks of ill-defined scalarizations. A naïve approach would be applying the user scalarization to also define the safety constraints. However, this construction fails to prevent undesirable behavior for the individual objectives (shown in Appendix A).

3.3 Multi-Objective SPIBB (MO-SPIBB)

Robust MDPs (Iyengar, 2005; Nilim and El Ghaoui, 2005) can be regarded as an approximation of the Bayesian formulation by partitioning the MDP space \mathcal{M} into two subsets: the subset of plausible MDPs Ξ and the subset of implausible MDPs. The plausible set is classically constructed from concentration bounds over the reward and transition function:

$$\Xi = \left\{ m, \text{ s.t. } \forall x, a, \begin{cases} \|p(\cdot|x, a) - \hat{p}(\cdot|x, a)\|_1 \leq e(x, a), \\ \|\mathbf{r}(x, a) - \hat{\mathbf{r}}(x, a)\|_\infty \leq e(x, a)r_\top \end{cases} \right\},$$

where e is an upper bound on the state-action error function of the model that are classically obtained with concentration bounds, such that the true environment $m^* \in \Xi$ with high probability $1 - \delta$. In the single objective framework, Laroche et al. (2019) empirically show that optimising the worst-case performance policy in Ξ provides policies that are too conservative. Petrik et al. (2016) prove that it is NP-hard to find the policy π that maximises the worst-case policy improvement over Ξ .

Instead, the SPIBB methodology (Laroche et al., 2019) consists in searching for a policy that maximizes the safe policy improvement in the MLE MDP, under some policy constraints: SPIBB

²Simão et al. (2020) proved that SPIBB/Soft-SPIBB bounds may be obtained with an estimate of π_b .

and Soft-SPIBB (Nadjahi et al., 2019) policy search constraints both revolve around the idea that we must only consider policies for which the policy improvement may be accurately estimated. Using π_b as reference, SPIBB allows policy changes only in state-action pairs for which more than n_\wedge samples have been collected. Soft-SPIBB extends this by applying soft constraints that allow slight changes in the policy for the uncertain state-action pairs, which are controlled by an error bound related to model uncertainty. As such, on low-confidence transitions, this class of methods provides a mechanism that prevents the agent from deviating too much from π_b . In this work, we build on Soft-SPIBB because it has yielded better empirical results. Formally, its constraint on the policy class is defined by:

$$\Pi_s = \left\{ \pi, \text{ s.t. } \forall x, \sum_a e(x, a) |\pi(a|x) - \pi_b(a|x)| \leq \epsilon \right\},$$

where ϵ is a hyper-parameter that controls the deviation from the baseline policy.

We define $q_{m,\lambda}^\pi(x, a) = \sum_{k \in [d]} \lambda_k q_{m,k}^\pi(x, a)$ to be the state-action value function associated with the linearized λ parameters. The same notation extension is used for $v_{m,\lambda}^\pi(x, a)$ and $\mathcal{J}_{m,\lambda}^\pi$. The application of Soft-SPIBB to multi-objective safe policy improvement is therefore direct:

$$\arg \max_{\pi \in \Pi_A \cap \Pi_s} \mathcal{J}_{m,\lambda}^\pi, \quad (3)$$

which is always realizable since $\pi_b \in \Pi_A \cap \Pi_s$.

We show that the construction of the plausible set required for the application of SPIBB is technically sound by deriving the concentration bounds for the multi-objective case. In Appendix B.1, we show with Hoeffding’s inequality that e grows as the square root of the logarithm of d (the number of reward functions), *i.e.* almost imperceptibly. From there, all the SPIBB theoretical results from Laroche et al. (2019); Nadjahi et al. (2019); Simão et al. (2020) may be generalized at a negligible SPI guarantee cost to the multi-objective setting, by applying their theorems separately to every objective function.

Now, the problem in Equation (3) can be transformed into a policy improvement procedure that solves for every state $x \in \mathcal{X}$ the following optimization problem³:

$$\begin{aligned} \pi_{\text{S-OPT}} &= \arg \max_{\pi \in \Pi} \langle \pi(\cdot|x), q_{m,\lambda}^\pi(x, \cdot) \rangle && (\text{S-OPT}) \\ \text{s.t. } &\sum_{a \in \mathcal{A}} e(x, a) |\pi(a|x) - \pi_b(a|x)| \leq \epsilon, && (\pi \in \Pi_s) \\ &\forall k \in [d], \sum_{a \in \mathcal{A}} \pi(a|x) \text{Adv}_{m,k}^{\pi_b}(x, a) \geq 0. && (\pi \in \Pi_A) \end{aligned}$$

The above procedure requires us to make additional algorithmic modifications that are not present in the original SPIBB algorithms. In particular, we need to explicitly incorporate advantage constraints for safety-guarantees for the individual objectives (proof given in Appendix B.2). The classic single-objective SPIBB algorithms do not need to check the advantage conditions because it is automatically guaranteed by the $\arg \max$ and the fact that $\pi_b \in \Pi_s$.

Using the construction above, we directly get the following result on the performance guarantees for each objective function that satisfies the desired property in Equation (1):

Proposition 3.1. The policy π returned from solving the S-OPT satisfies the following property in every state $x \in \mathcal{X}$ with probability at least $(1 - \delta)$:

$$\forall k \in [d], v_{m^*,k}^\pi(x) - v_{m^*,k}^{\pi_b}(x) \geq -\frac{\epsilon v_{\max}}{1 - \gamma}, \quad (4)$$

where $v_{\max} \leq \frac{r_{\max}}{1 - \gamma}$ is the maximum of the value function.

The proof is presented in Appendix B.3. The solution of S-OPT is computed by solving the Linear Program using standard solvers, such as cvxpy (Diamond and Boyd, 2016). There is an increase in the computational cost proportional to the number of reward functions. Compared to Soft-SPIBB,

³In practice, we also need to check $\forall x$ that $\pi(\cdot|x)$ is a valid probability distribution: positive and sums to 1.

the value and advantage functions estimation cost increases by a factor of d : respectively $\mathcal{O}(d|\mathcal{X}|^3)$ and $\mathcal{O}(d|\mathcal{A}||\mathcal{X}|^2)$. There is $\mathcal{O}(|\mathcal{D}|)$ cost for estimating the error bounds, and we also require solving a Linear Program for each state that approximately amounts to an additional $\mathcal{O}(|\mathcal{X}||\mathcal{A}|^2(|\mathcal{A}| + d))$ steps to the total computational cost (Boyd et al., 2004).

Remark (Extension to Constrained-RL). The above methodology can also be extended to the Constrained-RL setting for offline policy improvement in general CMDPs. Recall that SPIBB algorithms offer guarantees in the form of: $v_t - v_b \geq \hat{v}_t - \hat{v}_b - \xi$, where v_t and v_b are respectively the true values of the target and baseline policies, \hat{v}_t and \hat{v}_b are their estimates in the MLE MDP, and ξ is an error term due to parametric uncertainty. As a consequence, any constraint c such that $c \leq v_b + \hat{v}_t - \hat{v}_b - \xi$ may be guaranteed ($v_b - \hat{v}_b$ may easily be bounded with Hoeffding’s inequality), and when c is larger, we can return no solution found as with other Seldonian algorithms.

3.4 Multi-Objective HCPI (MO-HCPI)

We briefly recall how the HCPI methodology (Thomas et al., 2015a,b) can be applied directly for solving the objective in Equation (2). For a target policy, π_t , we use $\text{IS}_k(\mathcal{D}, \pi_t, \pi_b)$ to denote the estimated returns for the k th reward component (r_k) using any IS based off-policy estimator (Precup, 2000). A high-confidence lower bound on $\mathcal{J}_{m^*,k}^{\pi_t}$ can be defined as:

$$\Pr\left(\mathcal{J}_{m^*,k}^{\pi_t} \geq \text{IS}_k(\mathcal{D}, \pi_t, \pi_b) - \text{CI}_k(\mathcal{D}, \delta/d)\right) \geq 1 - \delta/d, \quad (5)$$

where $\text{CI}_k(\mathcal{D}, \delta) \geq 0$ denotes the terms associated with the choice of concentration inequality employed (and typically $\lim_{|\mathcal{D}| \rightarrow \infty} \text{CI}_k(\mathcal{D}, \delta) = 0$).

The dataset \mathcal{D} is first split into train (\mathcal{D}_{tr}) and test (\mathcal{D}_s) sets by the user. Let IS_λ denote the IS estimator associated with the user-specified reward scalarization λ . Given the user specified parameters: $\lambda, \delta, \text{CI}, \text{IS}, \mathcal{D}_{tr}, \mathcal{D}_s$ and π_b , the policy improvement problem in Equation (2) is transformed to the following optimization problem:

$$\begin{aligned} \pi_{\text{H-OPT}} &= \arg \max_{\pi \in \Pi} \text{IS}_\lambda(\mathcal{D}_{tr}, \pi, \pi_b) && \text{(H-OPT)} \\ \text{s.t. } &\forall k \in [d], \text{IS}_k(\mathcal{D}_s, \pi, \pi_b) - \text{CI}_k(\mathcal{D}_s, \delta/d) \geq \mu_k, \end{aligned}$$

where μ_k denote the empirical returns for r_k under π_b . The policy π returned by H-OPT will only violate the safety guarantees with probability at most δ . Proof of this claim and additional details are provided in Appendix C.

Although we only focus on finite MDPs in this work, the HCPI based approach relies on IS estimates and therefore it can also be used for infinite MDPs or POMDPs. Unfortunately, the IS estimates are typically known to suffer from high variance (Guo et al., 2017). Furthermore, the optimization problem in H-OPT is more challenging, and we need to resort to regularization based heuristics.

4 Synthetic Experiments

The main benefits of working in a synthetic domain are: (i) we can evaluate the performance on the true MDP instead of relying on off-policy evaluation (OPE) methods, (ii) we have control over the quality of the dataset. We test both MO-SPIBB (S-OPT) and MO-HCPI (H-OPT) on a variety of parameters: the amount of data, quality of baseline and different user reward scalarizations.

Env details: We take a standard CMDP benchmark (Leike et al., 2017; Chow et al., 2018) which consists of a 10×10 grid. From any state, the agent can move to the adjoining cells in the 4 directions using the 4 actions. The transitions are stochastic, with some probability α (generated randomly for each state-action for every environment instance) the agent is successfully able to reach the next state, and with $(1 - \alpha)$ the agent stays in the current state. The agent starts at the bottom-right corner, and the goal is to reach the opposite corner (top-left). The pits are spawned randomly with some uniform probability ($\eta_{\text{pit}} = 0.3$) for each cell. The reward vector consists of two rewards signals. A primary reward r_0 that is related the goal and is +1000.0 on reaching the goal and -1.0 at every other time-step. The secondary reward r_1 is related to pits, for which the agent gets -1.0 for any action taken in the pit. The constraint threshold for this CMDP is -2.0 and $\gamma = 0.99$. Maximum length of an episode is 200 steps. Therefore, the task objective is to reach the goal in the least number of steps, such that the agent does not spend more than 2 time-steps in the pit cells.

Dataset collection procedure: For every random CMDP generated, we first find the optimal policy π^* by using the procedure described in Appendix D.1. The baseline policy is generated using a convex combination of the optimal policy and a uniform random policy (π_{rand}), i.e., $\pi_b = \rho\pi^* + (1-\rho)\pi_{rand}$, where ρ controls how close π_b 's performance is to π^* . Different datasets with varying sizes and ρ are then collected under π_b and given as input to the methods.

Baselines: We compare against the following baselines:

- **Linearized:** This baseline transforms the rewards into a single scalar using λ and then applies the traditional policy improvement methods on the linearized objective, i.e., $\arg \max_{\pi \in \Pi} \mathcal{J}_{\hat{m}, \lambda}^{\pi}$.
- **Adv-Linearized:** This method has the same objective as the Linearized baseline, with the additional constraints based on advantage estimators built from \hat{m} , i.e. $\forall x \in \mathcal{X}$:

$$\begin{aligned} & \arg \max_{\pi \in \Pi} \langle \pi(\cdot|x), q_{\hat{m}, \lambda}^{\pi}(x, \cdot) \rangle & (6) \\ \text{s.t. } & \forall k \in [d], \sum_{a \in \mathcal{A}} \pi(a|x) \text{Adv}_{\hat{m}, k}^{\pi_b}(x, a) \geq 0. \end{aligned}$$

Evaluation: Using m^* , we can directly calculate the returns for any solution policy. Only tracking the scalarized objective can be misleading, so we track the following metrics:

- **Improvement over π_b :** This denotes the difference between the scalarized return of the solution policy and the baseline policy, i.e., $\mathcal{J}_{m^*, \lambda}^{\pi} - \mathcal{J}_{m^*, \lambda}^{\pi_b}$. Mean improvement over π_b captures on average improvement over π_b in terms of the scalarized objective.
- **Failure-rate:** The failure rate over n runs captures the number of times, on average, the solution policy ends up violating the safety constraints in Equation (1), and thus performs worse than the baseline. In the context of this task, safety constraints are violated if either the agent takes longer to reach the goal, or it steps into more number of pits compared to π_b .

We test on different combinations of user preference (λ) and baseline's quality (ρ) on 100 randomly generated CMDPs, where $\lambda_i \in \{0, 1\}$, $\rho \in \{0.1, 0.4, 0.7, 0.9\}$ and $|D| \in \{10, 50, 500, 2000\}$. We evaluate under two settings: (i) we use a fixed set of parameters across different (λ, ρ) combinations, where we run **S-OPT** with $\epsilon \in \{0.01, 0.1, 1.0\}$ and **H-OPT** with Doubly Robust IS estimator (Jiang and Li, 2015) and Student's t-test concentration inequality; (ii) we treat them as hyper-parameters that can be optimized for a particular (λ, ρ) combination. The best hyper-parameters are tuned in a single environment instance and then they are used to benchmark the results on 100 random CMDPs.

Results: The mean results with fixed parameters and $\delta = 0.1$ can be found in Figure 1a. The high failure rate of Linearized baseline, regardless of the size of the dataset, is expected as it optimizes the scalarized reward directly and is agnostic of the individual rewards. Adv-Linearized performs better, but in the low data-regime, we see a high failure rate that eventually decreases as the size of dataset increases. This is expected because with more data, more reliable advantage functions estimates are calculated that are representative of the underlying CMDP. Compared to the baselines, both **S-OPT** and **H-OPT** maintain a failure rate below the required confidence parameter δ , regardless of the amount of data. Also, as the size of dataset increases, we see an increase in improvement over π_b , that makes sense as the methods only deviate from baseline when they are sure of the performance guarantees. We expect **S-OPT** to violate the constraints with increasing value of ϵ , as it relaxes the constraint on the policy-class (Section 3.3) and leads to a looser guarantee on performance. This again is reflected in our experiments where **S-OPT** with $\epsilon = 1.0$ has a higher failure-rate than $\epsilon = 0.1$. We observed similar trends for different δ values. A more detailed plot corresponding to different λ and ρ combinations as well as results for a riskier value of $\delta = 0.9$ are given in Appendix D.2.

The results with optimized hyper-parameters can be found in Figure 1b. We notice that when the ϵ parameter is tuned properly, **S-OPT** has better performance in terms of improvement over π_b for the same amount of samples when compared to **H-OPT**, while still ensuring the failure rate is less than δ . These observations are consistent with the results in the single-objective setting in the original SPIBB works (Laroche et al., 2019; Nadjahi et al., 2019). The general trends and observations from the fixed-parameter case are also valid here. Additional details, including results for λ, ρ combinations, hyper-parameters considered and qualitative analysis can be found in Appendix D.3.

We also compare our methods against Le et al. (2019) in Appendix D.4. We show the advantage of our approach over Le et al. (2019), particularly in the low-data regime, where our methods can

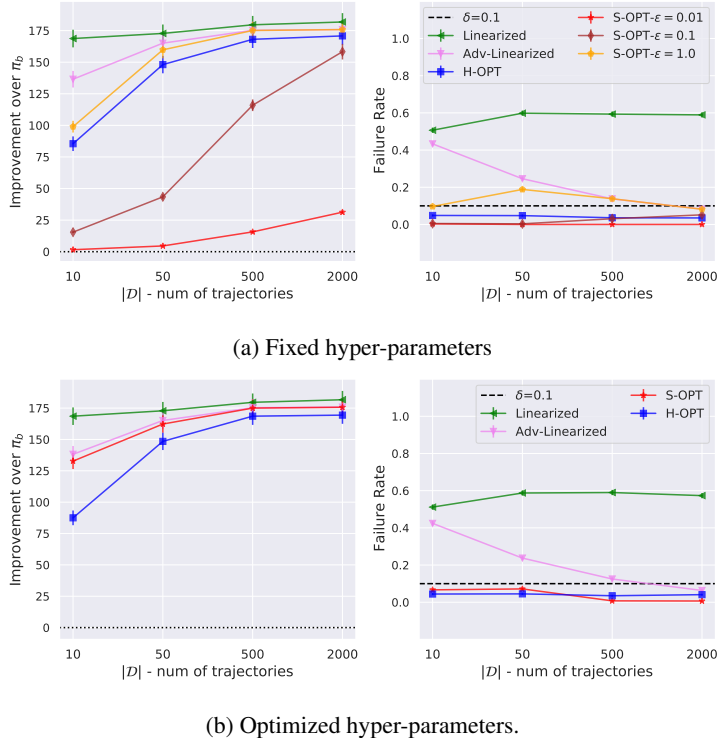


Figure 1: Results on 100 random CMDPs for different λ and ρ combinations with $\delta = 0.1$. The different agents are represented by different markers and colored lines. Each point on the plot denotes the mean (with standard error bars) for 12 different λ, ρ combinations for the 100 randomly generated CMDPs (1200 datapoints). The x-axis denotes the amount of data the agents were trained on. The y-axis for left subplot in each sub-figure represents the improvement over baseline and the right subplot denotes the failure rate. The dotted black line in the right subplots represents the high-confidence parameter $\delta = 0.1$. Figure 1a denotes when the hyper-parameters are fixed $\epsilon = \{0.01, 0.1, 1.0\}$ and IS = Doubly Robust (DR) estimator with student’s t-test concentration inequality. Figure 1b is the version with tuned hyper-parameters for each combination.

improve over the baseline policy while ensuring a low failure rate. This makes sense as the method in [Le et al. \(2019\)](#) relies on the concentrability coefficient which can be arbitrarily high in the low data setting, and therefore their performance guarantees do not hold anymore. We also provide experiments on the scalability of methods with the number of objectives d in [Appendix D.5](#).

5 Real-world experiment

In order to validate the applicability of our methods on a real-world task, we consider recent works on sepsis management via RL, where we only have access to a pre-collected patient dataset and goal is to recommend treatment strategies for patients with sepsis in the ICU ([Komorowski et al., 2018](#); [Tang et al., 2020](#)). Sepsis is defined as a life-threatening organ dysfunction caused by a dysregulated host response to an infection ([Singer et al., 2016](#)). The main treatment method of sepsis involves the repeated administration of intravenous (IV) fluids and vasopressors, but how to manage their appropriate doses at the patient level is still a key clinical challenge ([Rhodes et al., 2017](#)).

The problem is safety-critical as our methods need to be cautious about using the data that was possibly collected under unobservable confounders and that can lead to biased model estimates. For instance, a study by [Ji et al. \(2020\)](#) of the model used in [Komorowski et al. \(2018\)](#) found that the learned model suggests clinically implausible behavior in the form of unexpectedly aggressive treatments. We show that our methodology can be applied here to prevent such behavior that results from small sample sizes. We propose to do so by incorporating safety constraints to prevent recommending the treatment decisions that were never or rarely performed in the dataset.

Data and MDP Construction: We use the publicly available ICU dataset MIMIC-III (Johnson et al., 2016), with the setup described by Komorowski et al. (2018); Tang et al. (2020) and build on top of their data pre-processing and MDP construction methodology.⁴ This leaves us with a cohort of 20,954 unique patients. The state-space consisting of 48 clinical variables summarizing features like demographics, physiological condition, laboratory values, etc., is discretized using a k-means based clustering algorithm to map the states to 750 clusters. The actions include administration of IV fluids and vasopressors, which are categorized into 5 dosage bins each, leading to a total of $|\mathcal{A}| = 25$. The γ is set to 0.99. The reward is based on patient mortality. The agent gets a reward, r_0 , of ± 100 at the end of the episode based on the survival of the patient. More details can be found in Appendix E.1.

In the original work, the rare state-actions taken by the clinicians (state-action pairs observed infrequently in the training set) are removed from the dataset. Instead of removing them, we define an additional reward, r_1 , based on the rarity of the state-action pair. We define rare state-action pairs to be those that are taken less than 10 times throughout training dataset, and the agent gets a reward of -10 for every such rare state-action taken, i.e., $r_1(x, a) = -10.0$ if $\text{Count}(x, a) < 10$. The final task objective then becomes to suggest treatments that handles the trade-off between prioritizing improving the survival vs prioritizing commonly used treatment decisions.

Evaluation: We compare our approach with the same baselines from Section 4 on different λ combinations. We run our methods for 10 runs with different random seeds, where for each run the cohort dataset was split into train/valid/test sets in the ratios of 0.7/0.1/0.2. We evaluate the performance of the solution policies returned by different methods on the test sets using two different OPE methods, Doubly Robust (DR) (Jiang et al., 2015) and Weighted Doubly Robust (WDR) (Thomas and Brunskill, 2016). We acknowledge that these methods are a proxy of the actual performance of the deployed policies. Hence, these results should not be misinterpreted as us claiming that the policies returned by our methods are now ready to be used in the ICU.

Table 1: Performance of various methods using DR and WDR estimators with mean and standard deviation on 10 random splits of the cohort dataset. The red cells denote the corresponding safety constraint violation, i.e, either $\mathcal{J}_0^\pi < \mathcal{J}_0^{\pi^b}$ or $-\mathcal{J}_1^\pi > -\mathcal{J}_1^{\pi^b}$.

User preference (λ)	Policy	Survival return (\mathcal{J}_0)		Rare-treatment return ($-\mathcal{J}_1$)	
		DR	WDR	DR	WDR
	Clinician’s (π_b)	64.78 \pm 0.90	64.78 \pm 0.90	13.58 \pm 0.19	13.58 \pm 0.19
$[\lambda_0 = 1, \lambda_1 = 0]$	Linearized	97.68 \pm 0.22	97.58 \pm 0.20	27.64 \pm 1.11	27.84 \pm 1.09
	Adv-Linearized	91.62 \pm 0.46	92.68 \pm 0.23	15.18 \pm 0.59	13.56 \pm 0.42
	S-OPT	66.11 \pm 0.87	66.05 \pm 0.86	13.42 \pm 0.20	13.46 \pm 0.20
	H-OPT	65.95 \pm 0.00	65.95 \pm 0.00	13.37 \pm 0.00	13.37 \pm 0.00
$[\lambda_0 = 1, \lambda_1 = 1]$	Linearized	87.17 \pm 0.48	89.11 \pm 0.37	2.41 \pm 0.47	1.52 \pm 0.41
	Adv-Linearized	86.77 \pm 0.49	88.58 \pm 0.25	2.53 \pm 0.50	1.57 \pm 0.43
	S-OPT	86.77 \pm 0.49	88.58 \pm 0.25	2.53 \pm 0.50	1.57 \pm 0.43
	H-OPT	86.37 \pm 0.00	88.03 \pm 0.00	2.58 \pm 0.00	1.43 \pm 0.00
$[\lambda_0 = 0, \lambda_1 = 0]$	Linearized	-89.39 \pm 0.43	-90.90 \pm 0.29	22.99 \pm 0.40	22.81 \pm 0.30
	Adv-Linearized	60.27 \pm 0.49	61.44 \pm 0.85	18.40 \pm 0.27	15.36 \pm 0.58
	S-OPT	67.73 \pm 0.82	67.22 \pm 0.88	13.24 \pm 0.24	13.55 \pm 0.33
	H-OPT	65.95 \pm 0.00	65.95 \pm 0.00	13.37 \pm 0.00	13.37 \pm 0.00
$[\lambda_0 = 0, \lambda_1 = 1]$	Linearized	58.27 \pm 2.18	60.52 \pm 2.07	0.04 \pm 0.03	0.02 \pm 0.01
	Adv-Linearized	76.05 \pm 0.65	76.85 \pm 0.72	0.07 \pm 0.05	0.04 \pm 0.03
	S-OPT	76.07 \pm 0.65	76.87 \pm 0.73	0.07 \pm 0.05	0.04 \pm 0.03
	H-OPT	76.54 \pm 0.00	77.55 \pm 0.00	0.09 \pm 0.00	0.05 \pm 0.00

Results: We refer to the return associated with the mortality reward (r_0) as survival return (\mathcal{J}_0), and the negative return associated with rare state-action reward (r_1) as rare-treatment return ($-\mathcal{J}_1$). Higher survival return implies more successful discharges, and lower rare-treatment return implies more adherence to common practice treatment decisions. We present the results on survival and rare-treatment returns in Table 1. As expected, we observe both the Linearized and Adv-Linearized baselines violates constraints across different λ , whereas S-OPT and H-OPT are able to respect the

⁴A caveat here is regarding the underlying assumption that the MDP construction methodology by Komorowski et al. (2018); Tang et al. (2020) maintains the Markovian property in the discretized state-space.

safety constraints irrespective of the λ .⁵ The validation set was used to tune the hyper-parameters, and we report how the performance varies with different hyper-parameters in Appendix E.2.

Qualitative Analysis: We conclude with a qualitative analysis of the policies returned from our setting and the traditional RL approach of maximizing just the survival return. Ji et al. (2020) found that the RL-policies for sepsis-management task usually end up recommending aggressive treatments, particularly high vasopressor doses for states where the common practice (according to most frequent action chosen by the clinician for that state) is to give no vasopressors at all. The common practice involves giving zero vasopressors for 722 of the 750 states. However, the policy returned by the traditional single-objective RL baseline recommends vasopressors in 562 (77.84%) of those 722 states, with 295 of those recommendations being large doses (upper 50th percentile of nonzero amounts or $> 0.2 \mu\text{g}/\text{kg}/\text{min}$). We compare these statistics for two of the policies returned by MO-SPIBB that deviate the most from π_b . The policy returned by S-OPT ($\lambda = [1, 1]$) recommends vasopressors in only 93 of those states (12.88 %), with 47 of those recommendations belonging to high dosages. The other policy, S-OPT ($\lambda = [0, 1]$), recommends vasopressors in 134 (18.56 %) of those states and 70 of those recommendations fall in large dosages. Therefore, the policies returned by our approach, even when they deviate from the baseline, are less aggressive in recommending rare treatments. In Appendix E.3, we present an additional qualitative analysis that demonstrates our methods recommend lesser rare-action treatments than the traditional single-objective RL approach.

An argument can be made against the case when all rare state-action pairs are removed from the training data itself. This will ensure that any learned policy will have near 0 rare-treatment return. However, it is not always clear how to define the cut-off criteria for rare-actions, and it might be possible that some of these rare state-action pairs are actually crucial for finding a better policy. For instance, we did an experiment where we assigned state-actions pairs with frequency < 100 to be rare state-action pairs and filtered those from the training set. The clinician’s performance on the test set using a DR estimator for survival return is 65.95. In this case, the traditional single-objective RL baseline gives the survival return of 11.26, which shows that removing such transitions from the dataset actually hampers the solution quality. Our approach of assigning a separate reward for rare state-action pairs is able to find a solution with a survival return of 86.75 even in this scenario.

6 Conclusion

We present a new Seldonian RL algorithm that takes the user preference based scalarization into account while ensuring the solution policy performs reliably in context to the baseline policy across all objectives. On both synthetic and real-world tasks, we show that the proposed approach can improve the policy while ensuring the safety constraints are respected.

Our setting can accommodate any general form of scalarizations (e.g. non-linear or convex) as well as objectives (such as fairness), making it applicable to a wide variety of real-world tasks. The only assumption we made is regarding the dataset being collected under a single known baseline policy. An exciting line of future work can be to relax this assumption and consider the scenario where the dataset comes from a variety of unknown policies with different qualities. We did not make any claims about the optimality of the solutions as often optimality and safety are contradicting objectives. It is not clear how (and if) one can make claims about optimality in the offline setting without bringing in additional unrealistic assumptions (Section 2). The extension to infinite MDPs and the function-approximation setting is also left for future work. It is important to note that when it comes to practical application, it is not unusual for continuous domains to be discretized to enable better interpretability, especially when interactions with humans are necessary. If the Markovian property is valid in the discretized space, SPIBB-based guarantees will also hold true.

7 Acknowledgements

The authors would like to thank NSERC (Natural Sciences and Engineering Research Council), IVADO (Institut de valorisation des données) and CIFAR (Canadian Institute for Advanced Research) for funding to McGill in support of this research. Philip S. Thomas was funded in part by NSF award

⁵In Table 1, $\lambda = [1, 1]$ represents a rare case of reward scalarization that allows all the methods to find a good solution policy that satisfies the constraints. In general, it is difficult to find such scalarization parameters as seen in synthetic experiments (Appendix D.2).

#2018372. The computational component of this research was enabled in part by support provided by Calcul Québec (www.calculquebec.ca), Compute Canada (www.computeCanada.ca) and Mila’s IDT team.

We would also like to thank Emmanuel Bengio and Koustuv Sinha for many helpful discussions about the work, and the anonymous reviewers for providing constructive feedback.

References

- Abdolmaleki, A., Huang, S., Hasenclever, L., Neunert, M., Song, F., Zambelli, M., Martins, M., Heess, N., Hadsell, R., and Riedmiller, M. (2020). A distributional view on multi-objective policy optimization. In *International Conference on Machine Learning*, pages 11–22. PMLR.
- Altman, E. (1999). *Constrained Markov decision processes*, volume 7. CRC Press.
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., and Mané, D. (2016). Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.
- Bellman, R. (1957). A Markovian decision process. *Journal of mathematics and mechanics*, pages 679–684.
- Boyd, S., Boyd, S. P., and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.
- Chow, Y., Nachum, O., Duenez-Guzman, E., and Ghavamzadeh, M. (2018). A Lyapunov-based approach to safe reinforcement learning. *arXiv preprint arXiv:1805.07708*.
- Christiano, P., Leike, J., Brown, T. B., Martic, M., Legg, S., and Amodei, D. (2017). Deep reinforcement learning from human preferences. *arXiv preprint arXiv:1706.03741*.
- Diamond, S. and Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5.
- Dulac-Arnold, G., Mankowitz, D., and Hester, T. (2019). Challenges of real-world reinforcement learning. *arXiv preprint arXiv:1904.12901*.
- Guo, Z. D., Thomas, P. S., and Brunskill, E. (2017). Using options and covariance testing for long horizon off-policy policy evaluation. *arXiv preprint arXiv:1703.03453*.
- Hansen, N. (2006). The cma evolution strategy: A comparing review. In *Towards a New Evolutionary Computation*, pages 75–102. Springer.
- Iyengar, G. N. (2005). Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280.
- Ji, C. X., Oberst, M., Kanjilal, S., and Sontag, D. (2020). Trajectory inspection: A method for iterative clinician-driven design of reinforcement learning studies. *arXiv preprint arXiv:2010.04279*.
- Jiang, N., Kulesza, A., Singh, S., and Lewis, R. (2015). The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1181–1189.
- Jiang, N. and Li, L. (2015). Doubly robust off-policy value evaluation for reinforcement learning. *arXiv preprint arXiv:1511.03722*.
- Johnson, A. E., Pollard, T. J., Shen, L., Li-Wei, H. L., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., and Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3(1):1–9.
- Kakade, S. and Langford, J. (2002). Approximately optimal approximate reinforcement learning. In *In Proc. 19th International Conference on Machine Learning*. Citeseer.
- Kivinen, J. and Warmuth, M. K. (1997). Exponentiated gradient versus gradient descent for linear predictors. *information and computation*, 132(1):1–63.

- Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C., and Faisal, A. A. (2018). The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, 24(11):1716–1720.
- Lange, S., Gabel, T., and Riedmiller, M. (2012). Batch reinforcement learning. In *Reinforcement Learning*, pages 45–73. Springer.
- Laroche, R., Trichelair, P., and Tachet des Combes, R. (2019). Safe policy improvement with baseline bootstrapping. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*.
- Le, H., Voloshin, C., and Yue, Y. (2019). Batch policy learning under constraints. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 3703–3712, Long Beach, California, USA. PMLR.
- Leike, J., Krueger, D., Everitt, T., Martic, M., Maini, V., and Legg, S. (2018). Scalable agent alignment via reward modeling: A research direction. *arXiv preprint arXiv:1811.07871*.
- Leike, J., Martic, M., Krakovna, V., Ortega, P. A., Everitt, T., Lefrancq, A., Orseau, L., and Legg, S. (2017). AI safety gridworlds. *arXiv preprint arXiv:1711.09883*.
- Littman, M. L., Topcu, U., Fu, J., Isbell, C., Wen, M., and MacGlashan, J. (2017). Environment-independent task specifications via gtl. *arXiv preprint arXiv:1704.04341*.
- Liu, C., Xu, X., and Hu, D. (2014). Multiobjective reinforcement learning: A comprehensive overview. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 45(3):385–398.
- Lizotte, D. J., Bowling, M., and Murphy, S. A. (2012). Linear fitted-q iteration with multiple reward functions. *The Journal of Machine Learning Research*, 13(1):3253–3295.
- Mannor, S. and Shimkin, N. (2004). A geometric approach to multi-criterion reinforcement learning. *The Journal of Machine Learning Research*, 5:325–360.
- Maurer, A. and Pontil, M. (2009). Empirical Bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.
- Munos, R. (2003). Error bounds for approximate policy iteration. In *ICML*, volume 3, pages 560–567.
- Nadjahi, K., Laroche, R., and Combes, R. T. d. (2019). Safe policy improvement with soft baseline bootstrapping. *arXiv preprint arXiv:1907.05079*.
- Ng, A. Y., Harada, D., and Russell, S. (1999). Policy invariance under reward transformations: Theory and application to reward shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning*, volume 99, pages 278–287.
- Nilim, A. and El Ghaoui, L. (2005). Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798.
- OpenAI (2018). Openai five. <https://blog.openai.com/openai-five/>.
- Petrik, M., Chow, Y., and Ghavamzadeh, M. (2016). Safe policy improvement by minimizing robust baseline regret. *arXiv preprint arXiv:1607.03842*.
- Precup, D. (2000). Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, page 80.
- Rhodes, A., Evans, L. E., Alhazzani, W., Levy, M. M., Antonelli, M., Ferrer, R., Kumar, A., Sevransky, J. E., Sprung, C. L., Nunnally, M. E., et al. (2017). Surviving sepsis campaign: International guidelines for management of sepsis and septic shock: 2016. *Intensive Care Medicine*, 43(3):304–377.

- Roijers, D. M., Vamplew, P., Whiteson, S., and Dazeley, R. (2013). A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359.
- Simão, T. D., Laroche, R., and Tachet des Combes, R. (2020). Safe policy improvement with an estimated baseline policy. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1269–1277.
- Singer, M., Deutschman, C. S., Seymour, C. W., Shankar-Hari, M., Annane, D., Bauer, M., Bellomo, R., Bernard, G. R., Chiche, J.-D., Coopersmith, C. M., et al. (2016). The third international consensus definitions for sepsis and septic shock (Sepsis-3). *Jama*, 315(8):801–810.
- Sorg, J., Singh, S. P., and Lewis, R. L. (2010). Internal rewards mitigate agent boundedness. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 1007–1014.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44.
- Tang, S., Modi, A., Sjoding, M., and Wiens, J. (2020). Clinician-in-the-loop decision making: Reinforcement learning with near-optimal set-valued policies. In *International Conference on Machine Learning*, pages 9387–9396. PMLR.
- Thomas, P. and Brunskill, E. (2016). Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148.
- Thomas, P., Theocharous, G., and Ghavamzadeh, M. (2015a). High confidence policy improvement. In *International Conference on Machine Learning*, pages 2380–2388.
- Thomas, P. S., Castro da Silva, B., Barto, A. G., Giguere, S., Brun, Y., and Brunskill, E. (2019). Preventing undesirable behavior of intelligent machines. *Science*, 366(6468):999–1004.
- Thomas, P. S., Theocharous, G., and Ghavamzadeh, M. (2015b). High-confidence off-policy evaluation. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- Vamplew, P., Dazeley, R., Berry, A., Issabekov, R., and Dekker, E. (2011). Empirical evaluation methods for multiobjective reinforcement learning algorithms. *Machine learning*, 84(1):51–80.
- Van Moffaert, K. and Nowé, A. (2014). Multi-objective reinforcement learning using sets of pareto dominating policies. *The Journal of Machine Learning Research*, 15(1):3483–3512.
- Walpole, R. E., Myers, R. H., Myers, S. L., and Ye, K. (1993). *Probability and statistics for engineers and scientists*, volume 5. Macmillan New York.
- Yang, R., Sun, X., and Narasimhan, K. (2019). A generalized algorithm for multi-objective reinforcement learning and policy adaptation. *arXiv preprint arXiv:1908.08342*.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes] See Sections 3.3 and 3.4 for the methodology and theoretical claims, and Sections 4 and 5 for the empirical results.
 - (b) Did you describe the limitations of your work? [Yes] See Section 1 for the limitations and scope of this work.

- (c) Did you discuss any potential negative societal impacts of your work? [Yes] As we mentioned in Sections 1 and 3.2, our goal is to maximize the objective specified by the user while ensuring that the solution policy avoids causing harmful effects after deployment in the true environment in comparison to the existing baseline policy. We aim to bridge the gap between traditional RL methods and high-stake real-world applications, but, as with any general technology, we acknowledge that some RL applications can have the potential of misuse, and our methods do not prevent that.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We discuss the societal impacts in the point above. We use pre-existing publicly available data and libraries and give more details about them in the Point 4 below.
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 3.1 for the assumption regarding access to baseline policy.
 - (b) Did you include complete proofs of all theoretical results? [Yes] The complete proofs are provided in the Appendix B and Appendix C.
 3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code required to produce the results is provided in the supplementary material.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] The training details are mentioned in Sections 4 and 5 in the main text, and Appendices D and E.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] The details about the error bars are provided in Sections 4 and 5.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Details about the compute and resources can be found in appendices D and E.
 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] See Sections 4 and 5 for the appropriate references.
 - (b) Did you mention the license of the assets? [Yes] We use the (MIMIC-III, Johnson et al., 2016) dataset and provide the appropriate reference. The explicit link to the license is here: <https://physionet.org/content/mimiciii/view-license/1.4/>.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We provide the accompanying code for running the experiments in the supplemental material.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A] We are using publicly available libraries and dataset.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] We are using the publicly available dataset (MIMIC-III, Johnson et al., 2016) that already deidentifies the data in accordance with Health Insurance Portability and Accountability Act (HIPAA) standards using structured data cleansing and date shifting. We refer to the original text for more details on the deidentification process.
 5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]