
Clinician-in-the-Loop Decision Making: Reinforcement Learning with Near-Optimal Set-Valued Policies

Shengpu Tang¹ Aditya Modi¹ Michael W. Sjoding^{2,3} Jenna Wiens¹

Abstract

Standard reinforcement learning (RL) aims to find an optimal policy that identifies the best action for each state. However, in healthcare settings, many actions may be near-equivalent with respect to the reward (e.g., survival). We consider an alternative objective – learning set-valued policies to capture near-equivalent actions that lead to similar cumulative rewards. We propose a model-free algorithm based on temporal difference learning and a near-greedy heuristic for action selection. We analyze the theoretical properties of the proposed algorithm, providing optimality guarantees and demonstrate our approach on simulated environments and a real clinical task. Empirically, the proposed algorithm exhibits good convergence properties and discovers meaningful near-equivalent actions. Our work provides theoretical, as well as practical, foundations for clinician/human-in-the-loop decision making, in which humans (e.g., clinicians, patients) can incorporate additional knowledge (e.g., side effects, patient preference) when selecting among near-equivalent actions.

1. Introduction

In the standard RL setup, one aims to find an optimal policy, which identifies the action for each state that maximizes some discounted expected cumulative reward. However, in healthcare, the reward can be difficult to define; e.g., one might want to optimize for long-term quality of life vs. short-term stabilization of the symptoms (without treating the underlying disease process). Past work has augmented

reward signals via reward shaping (Lizotte et al., 2012; Raghu et al., 2017a; Nemati et al., 2016). Still, designing a single reward function that captures the goals and objectives across different individuals remains challenging. As a result, when applying RL in healthcare, survival is often used as the reward, since it represents a clear goal and is straightforward to measure (Raghu et al., 2017b; Komorowski et al., 2018; Li et al., 2018).

While using survival as the reward signal can simplify the RL setup, we hypothesize that it induces many near-equivalent actions that could otherwise differ. For example, different doses of a drug might perform similarly in terms of keeping the patient alive, yet doses that are too large might cause severe side effects. In other cases, patients may opt for less invasive treatments, if they are likely to yield similar outcomes to more invasive ones. In such cases, learning a single best action and ignoring near-equivalent actions may be undesirable, as important considerations not captured through the reward signal can impact decision making.

Thus, we consider the task of learning a “set-valued policy” (SVP), which returns a set of near-equivalent actions rather than a single optimal action. This allows for interaction between the clinician (or patient) and the decision support system (i.e., clinician/patient-in-the-loop). Such a setup provides clinicians/patients with an opportunity to incorporate additional information (e.g., patient preferences, adverse drug reactions, costs/availability of equipment) when choosing among near-equivalent actions.

We study a particular formalization of SVPs introduced by Fard & Pineau (2011) where they consider the problem of computing SVPs (planning). In contrast to their exhaustive search based approach, which requires a model of the underlying environment, we propose a *model-free* learning algorithm. The sequential nature of decisions makes learning such policies non-trivial, as near-equivalent actions should be not only similar in the short term (immediate reward), but also similar for all possible future trajectories (expected cumulative rewards). Our contributions are:

- we propose a new algorithm based on temporal difference methods and a near-greedy heuristic for learning near-optimal set-valued policies,

¹Department of Electrical Engineering & Computer Science, University of Michigan, Ann Arbor, US ²Department of Internal Medicine, Michigan Medicine, University of Michigan, Ann Arbor, US ³Institute for Healthcare Policy & Innovation, University of Michigan, Ann Arbor, US. Correspondence to: Shengpu Tang <tangsp@umich.edu>, Jenna Wiens <wiensj@umich.edu>.

- we investigate its convergence behavior and optimality using a worst-case analysis, providing theoretical guarantees in a directed acyclic graph (DAG) setting,
- we demonstrate the algorithm’s behavior in both DAG and non-DAG synthetic environments, and
- on a clinical task, we demonstrate how the algorithm can help discover clinically meaningful near-equivalencies among treatment actions.

Our work provides both theoretical and practical foundations for learning near-optimal SVPs, and represents an important step towards clinician/human-in-the-loop decision support systems. Beyond applications in healthcare, this framework could also be applied to other domains involving human-machine-cooperative decision making, including intelligent tutoring systems and self-driving cars. The code to reproduce our experiments is available online¹.

2. Background

We consider finite Markov decision processes (MDPs) defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$, where \mathcal{S} is the set of states (finite or infinite), \mathcal{A} is a finite set of actions, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ defines the transition model with $p(s'|s, a)$ specifying the probability of moving from state s to s' given action a , and $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ defines the reward function (can be stochastic in general) with $r(s, a)$ denoting the expected instantaneous reward obtained from taking action a in state s , and discount factor $\gamma \in [0, 1]$.

2.1. Set-Valued Policy

In this work, we focus on set-valued policies (SVP), first formalized in Fard & Pineau (2011)².

Definition 1. An SVP, π , is a function that maps each state to a non-empty subset of actions, $\pi : \mathcal{S} \rightarrow 2^{\mathcal{A}} \setminus \{\emptyset\}$.

SVPs have been explored as a way to encode alternative choices (Fard & Pineau, 2011; Lizotte et al., 2012) and to encapsulate an (approximately) exponential number of deterministic policies (Lizotte & Laber, 2016). In particular, Fard & Pineau (2011) proposed a mixed-integer programming (MIP) formulation to find SVP solutions for a finite-horizon planning problem. However, their approach involves an exhaustive search with exponential complexity over the state and action spaces. In contrast, we propose a scalable model-free learning algorithm that does not require knowledge of the MDP. Moreover, our proposed approach allows the extent of near-optimality to be set *a priori*, providing increased flexibility and optimality guarantees.

¹<https://gitlab.eecs.umich.edu/MLD3/RL-Set-Valued-Policy>

²In Fard & Pineau (2011), this was referred to as “non-deterministic policy”; however, since the policy is indeed a deterministic mapping, we prefer the term “set-valued policy” to avoid confusion.

2.2. Value Functions for SVPs

In contrast to the standard RL setting, here, a learning agent can suggest a set of actions. Thus, our notion of “value” must account for all possible policies consistent with the proposed set of actions, regardless of which proposed action is selected. To this end, we consider a worst-case analysis (Fard & Pineau, 2011), where the value of a state s is taken as the worst case over all the actions in the set $\pi(s)$.

Definition 2. The worst-case value functions of an SVP π are defined as

$$V^\pi(s) = \min_{a \in \pi(s)} \{Q^\pi(s, a)\},$$

$$Q^\pi(s, a) = r(s, a) + \gamma \mathbb{E}_{s'|s, a} [V^\pi(s')].$$

2.3. Near-Optimal SVPs

We quantify the “goodness” of an SVP according to how far it is from the optimal value function, V^* . This gives rise to the definition of near-optimal SVPs. In some of the MDP literature, near-optimality has been formalized with an additive constraint: $\forall s, V^\pi(s) \geq V^*(s) - \epsilon$ (Kearns & Singh, 2002), which specifies a constant margin of sub-optimality across all states. This could lead to conservative action choices in some states, as ϵ is fixed but the magnitude of $V^\pi(s)$ and $V^*(s)$ could vary across different states. We argue that in healthcare settings this is not a suitable formalization, since at any point of decision making (in any particular state), the SVP should be near-optimal with respect to what we could achieve in that state; accepting the same value margin in a “healthy” state (larger value) vs. a “sick” state (smaller value) may not be desirable, because a margin that leads to acceptable outcomes in a “healthy” state can have a larger relative impact (perhaps devastating) in a “sick” state. We consider the additive near-optimality formalization as a baseline, providing the derivations of this setting in Appendix A. However, given the limitations, we focus on a multiplicative constraint for near-optimality, which accounts for differences in values at different states.

Definition 3. An SVP π is ζ -optimal, with $\zeta \in [0, 1]$, if

$$V^\pi(s) \geq (1 - \zeta)V^*(s), \quad \forall s \in \mathcal{S}.$$

Here, ζ is a hyperparameter that defines the sub-optimality margin, quantifying the trade-off between action variety and optimality.

Remark. Note that this definition requires $V^*(s) \geq 0, \forall s \in \mathcal{S}$. A sufficient (though not necessary) condition to ensure this is to enforce $r(s, a) \geq 0$. In experiments with clinical data (Section 4), we discuss practical considerations to deal with problem domains having negative rewards.

As a naïve solution, one might construct π as:

$$\pi(s) = \{a : Q^*(s, a) \geq (1 - \zeta)V^*(s)\}.$$

However, this construction fails to satisfy near-optimality. By using $Q^*(s, a) = r(s, a) + \gamma \mathbb{E}_{s'|s, a} V^*(s')$, the optimal Q-function, it assumes all future returns $V^*(s')$ are obtained following π^* as opposed to π . During the roll-out of a policy, this fails to account for the fact that the future must be consistent with the policy. Alternatively, one might consider an exponentially large action space, $\tilde{\mathcal{A}} = 2^{\mathcal{A}} \setminus \{\emptyset\}$ and apply standard value-based methods to learn Q-values defined over $\mathcal{S} \times \tilde{\mathcal{A}}$, but analysis shows that this naïve approach defaults to the greedy optimal policy (Appendix B). In the sections that follow, we propose a new approach to learn SVPs that does not violate the ζ -optimal constraint in Definition 3 and focuses on value functions in the original action space \mathcal{A} .

3. Methods

We present an algorithm that jointly learns SVPs and their value functions. First, we provide two heuristics to construct near-optimal SVPs given the value functions. Using these heuristics as action selection strategies, we describe a variant of temporal difference (TD) learning and provide a theoretical analysis of its convergence and optimality.

3.1. Heuristics for Constructing Near-Optimal SVPs

To guarantee near-optimality, we start with a *conservative* approach based on a loose lower bound of future returns. Then, we improve on this approach and propose a *near-greedy* heuristic that leverages the learned policy.

Conservative. Assuming the future follows a ζ -optimal policy, one could construct π as:

$$\pi(s) = \{a : \check{Q}_\zeta^*(s, a) \geq (1 - \zeta)V^*(s)\}, \quad (1)$$

where $\check{Q}_\zeta^*(s, a) = r(s, a) + \gamma(1 - \zeta)\mathbb{E}_{s'|s, a}[V^*(s')] \leq Q^*(s, a)$ is the action-value function using a loose lower bound for near-optimal future returns.

The conservative heuristic is consistent with the definition by Fard & Pineau (2011); however, one key difference is that we provide an explicit way to construct a conservative SVP given an oracle for the optimal value function V^* . Though this will not violate the near-optimality bound, it may limit action diversity.

To encourage action diversity while satisfying the near-optimality criteria, we can formulate a fixed-point equation for computing a near-optimal SVP. Recall that, in a standard RL setup, the optimal policy π^* is the unique fixed-point solution to the following equation: $\pi^*(s) = \arg \max_a Q^{\pi^*}(s, a) \forall s \in \mathcal{S}$ which applies a greedy action selection over optimal Q-values. For a near-optimal SVP, we seek the fixed-point solution to a similar equation with a near-greedy action selection.

Near-Greedy. Consider the fixed-point solution to the following equation:

$$\pi(s) = \{a : Q^\pi(s, a) \geq (1 - \zeta)V^*(s)\}, \quad (2)$$

where $Q^\pi(s, a)$ is the action-value function for policy π as computed via Definition 2.

Depending on the dynamics of the MDP and the true optimal value function V^* , it is possible that no solution exists for Eqn. (2) (see Appendix C for an example). Thus, for a general MDP, directly applying this heuristic might not generate the desired near-optimal SVP. In Section 3.3, we discuss the sufficient conditions for the existence and uniqueness of SVPs constructed according to Eqn. (2).

For the same optimality threshold ζ , compared to the conservative SVP, the near-greedy SVP is more likely to contain more actions, due to its consideration of the policy-dependent worst-case future $V^\pi(s')$, rather than a loose lower bound $(1 - \zeta)V^*(s')$, when deciding whether an action should be included.

Algorithm 1 TD learning for near-greedy ζ -optimal SVP

```

1: Input: step size  $\alpha_t \in (0, 1]$ ,
           optimal value function  $V^*$  where  $V^*(s') \geq 0 \forall s \in \mathcal{S}$ ,
           sub-optimality margin  $\zeta \in [0, 1]$ .
2: Initialize  $Q(s, a) = 0$  for all  $s \in \mathcal{S}, a \in \mathcal{A}$ 
3: for each episode do
4:   Initialize  $s$ 
5:   for each step do
6:     Choose action  $a$  using an exploratory policy (e.g.,  $\epsilon$ -greedy)
7:     Take action  $a$ , observe  $r, s'$ 
8:      $\pi(s') = \{a' : Q(s', a') \geq (1 - \zeta)V^*(s')\}$ 
9:     if  $\pi(s') \neq \emptyset$ 
10:       $Q(s, a) \leftarrow Q(s, a) + \alpha_t[r + \gamma \min_{a' \in \pi(s')} Q(s', a') - Q(s, a)]$ 
11:    else
12:       $Q(s, a) \leftarrow Q(s, a) + \alpha_t[r + \gamma \max_{a' \in \mathcal{A}} Q(s', a') - Q(s, a)]$ 
13:     $s \leftarrow s'$ 
14:   end for
15: end for

```

3.2. Learning Near-Optimal SVPs

As stated, a fixed-point solution to Eqn. (2) might not exist for a general MDP. However, the near-greedy construction can be used to devise a model-free learning algorithm in the TD-learning framework (Algorithm 1). Specifically, to compute the TD update target, we temporarily construct a candidate SVP $\pi(s')$ for the next state s' using the current estimates of the Q-values (line 8). With $V^*(s') \geq 0$, when we have a non-empty set of near-optimal actions $\pi(s')$, we compute the update target by using the worst near-greedy action (line 10). Otherwise, when $\pi(s')$ is empty, we use the standard greedy target (line 12). Finally, the algorithm outputs estimates of Q^π , as well as the SVP π constructed according to Eqn. (2). Note that the algorithm requires the optimal value function V^* as an input. In practice, one can

run a separate Q-learning procedure to learn a good estimate of Q^* and thus of V^* , or learn estimates of Q^* and Q^π simultaneously.

3.3. Theoretical Analysis

Here, we discuss the sufficient conditions for which the near-greedy SVP exists and is unique and for which the near-greedy TD algorithm converges under the same conditions for Q-learning. For completeness, in [Appendix D](#), we show that the conservative heuristic leads to a stochastic approximation algorithm that converges to a unique SVP solution for any MDP with non-negative rewards.

Theorem 1. *The near-greedy ζ -optimal SVP exists and is unique, if the MDP is a directed acyclic graph (DAG) with non-negative rewards.*

Proof. We show this by explicit construction of π . Note that the states in a DAG form a topological sort tree, where s precedes s' if and only if there is a transition from s to s' .

Base case. Consider every terminal state $s_\infty \in \mathcal{S}_\infty$, for which there is no immediate reward and no future time step starting from this state, i.e., $V^\pi(s_\infty) = 0$. We set $\pi(s_\infty) = \mathcal{A}$ trivially.

Inductive step. Given a state s , consider its successor states s' . Assuming we have every $\pi(s')$ and their associated value functions satisfying $V^\pi(s') \geq (1 - \zeta)V^*(s')$, we calculate $Q^\pi(s, a)$ for all $a \in \mathcal{A}$:

$$\begin{aligned} Q^\pi(s, a) &= r(s, a) + \gamma \mathbb{E}_{s'|s, a} V^\pi(s') \\ &\geq r(s, a) + \gamma \mathbb{E}_{s'|s, a} (1 - \zeta)V^*(s') \\ &= \zeta r(s, a) + (1 - \zeta)[r(s, a) + \gamma \mathbb{E}_{s'|s, a} V^*(s')] \\ &= \zeta r(s, a) + (1 - \zeta)Q^*(s, a) \\ &\geq (1 - \zeta)Q^*(s, a). \end{aligned}$$

Using $Q^\pi(s, a)$, we construct $\pi(s)$ according to [Eqn. \(2\)](#). Importantly, $\pi(s)$ is non-empty because the optimal action a^* is always included. Given $\pi(s)$, $V^\pi(s) = \min_{a \in \pi(s)} Q^\pi(s, a) \geq (1 - \zeta)V^*(s)$ satisfies the near-optimality constraint by construction.

Enumerating the states in reverse topological order starting from the terminal states, we follow this process until we have $\pi(s)$ for all $s \in \mathcal{S}$. Each step of the construction process is unique, hence overall the policy is unique. \square

Since the proof relies on the topological ordering of states in a DAG MDP, it holds when we are only given V^* , as we can obtain V^π from the base case for terminal states and from the inductive step for non-terminal states. The above theorem provides a *sufficient* condition for a near-greedy SVP to exist: the environment is a DAG with non-negative rewards. An MDP with cycles and/or negative rewards may still have a near-greedy ζ -optimal SVP, depending on the sub-optimality margin ζ and the MDP parameters.

Similarly, we can show the following convergence result for [Algorithm 1](#) (proof provided in [Appendix E](#)).

Theorem 2. *The near-greedy TD algorithm converges to the unique solution if the MDP is a DAG with non-negative rewards, under the same convergence conditions for regular TD learning: rewards have bounded variance, each (s, a) is updated infinitely many times, $\sum_t \alpha_t = \infty$, $\sum_t \alpha_t^2 < \infty$.*

Remark. The near-greedy heuristic from [Section 3.1](#) can be used as the policy improvement step (replacing the greedy action selection) in any value-based generalized policy iteration algorithm. For instance, when we are given the MDP model, one can derive a version of near-greedy value iteration with similar theoretical guarantees on DAG environments. In that case, it is most efficient to learn the value functions in reverse topological order of the states (as per the proof). However, in the more general case (where the underlying MDP is unknown, or when we have a continuous state space), we require a different approach. In particular, here we present an algorithm based on TD learning. Though the theoretical analyses only hold for DAG environments, in our experiments we will demonstrate the algorithm’s behavior in the more general setting of non-DAG environments.

4. Experiments

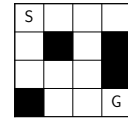
We consider a set of experiments to i) demonstrate the ability of the proposed algorithm to learn SVPs for different values of ζ and ii) characterize the algorithm’s empirical convergence behavior. Throughout, we compare to alternative approaches to learning SVPs (described in [Section 4.2](#)).



(a) Chain-5



(b) CyclicChain-5



(c) FrozenLake-4x4

Figure 1. Synthetic environments include (a) a simple DAG, and (b)(c) two non-DAG environments.

4.1. Environments

We consider both synthetic and real environments.

Chain- k ([Figure 1a](#)). First, as a sanity check, we consider a DAG with k sequentially connected states, with s_1 the starting state and s_k the terminal state. There are four actions at each state s_i that transition to s_{i+1} , and intermediate rewards $r_1, r_2, r_3, r_4 \in \{0, 0.01, 0.02, 0.03, 0.04, 0.05\}$ that are predetermined, but randomly assigned. Transitions reaching the terminal state result in an additional reward of +1. In this simple setting, we test the ability of the proposed algorithm to identify near-equivalent actions.

Next, we investigate the empirical performance of the algorithm in two non-DAG settings: *CyclicChain-k* and *FrozenLake*. These represent more complex environments, testing the generalizability of the proposed algorithm.

CyclicChain-k (Figure 1b): An extension of *RandomWalk* (Sutton & Barto, 2018), similar to *Chain-k*, except that of the four actions from s_i , two lead to s_{i-1} and two lead to s_{i+1} . We set $\gamma < 1$ to encourage the agent to reach the terminal state quickly and avoid cycling.

FrozenLake (Figure 1c). This is a standard discrete space path-finding problem from OpenAI Gym (Brockman et al., 2016). The agent controls the movement of a character in a grid world. Some tiles of the grid lead the agent to fall into the water. The agent is rewarded +1 for finding a path to a goal tile. We used the standard 4×4 and 8×8 maps. Multiple paths exist from the starting tile to the goal tile; these paths take the same number of steps and are thus equivalent. To introduce near-equivalent actions, we added a small reward to all actions from every non-terminal state so that the four actions vary slightly in value. For the 4×4 map, we added $r_i \in \{0.01, 0.02, 0.03, 0.04\}$, while for the 8×8 map, we added $r_i \in \{0.001, 0.002, 0.003, 0.004\}$.

Finally, we explore a challenging clinical task using observational patient data. In contrast to the environments above, in this setting, i) we do not have access to the underlying MDP, ii) transitions are stochastic, and iii) the reward signal is sparse.

MIMIC-sepsis. This is a previously studied RL task in the healthcare domain, in which the goal is to learn optimal treatment strategies for patients with sepsis in the ICU. Sepsis is defined as severe infection leading to life-threatening organ dysfunction and is one of the leading causes of mortality in hospitals (Gotts & Matthay, 2016; Liu et al., 2014). While a lot of work has focused on sepsis prediction (Henry et al., 2015; Reyna et al., 2019; Bedoya et al., 2020), the management of intravenous (IV) fluids and vasopressors in sepsis treatment still represents a key clinical challenge (Byrne & Van Haren, 2017). We based our analysis, in part, on the setup described in Komorowski et al. (2018), and used the same data and preprocessing steps, outlined below. Patient data are 48-dimensional time series (Appendix H) collected at 4h intervals, consisting of measurements from 24h preceding until 48h following the time of sepsis onset. Similar to in Komorowski et al. (2018), we consider 750 discrete health states obtained from clustering the training set using k -means. Additionally, 2 terminal states are added to represent death and discharge. Actions pertain to treatment decisions in each 4h interval, representing total volume of IV fluids and amount of vasopressors administered. Though we consider the same number of discrete actions (25), the corresponding IV fluid doses differ substantially from those considered by Komorowski et al. (2018). Specifically, we

updated the five levels of IV fluids to use the following bins [0, <500mL, 500mL–1L, 1–2L, >2L] to represent more clinically relevant fluid boluses. We made this modification based on feedback from a critical care physician. Furthermore, the actions available at each state, $A(s)$, are restricted to those observed ≥ 5 times (in training data; the most frequent action is used if no action occurs ≥ 5 times). Rewards are sparse and only assigned at the end of each trajectory: +100 for survival (and discharge), -100 for in-hospital death; all intermediate rewards are 0. In applying Algorithm 1, when $V^*(s) < 0$ (due to negative rewards), we fall back to the greedy update target in line 12. γ is set to 0.99 to place nearly as much importance on late deaths as early deaths. Applying the specified inclusion and exclusion criteria (Komorowski et al., 2018) to the MIMIC-III database (Johnson et al., 2016), we identified a cohort of 20,940 patients with sepsis (Table 1). The cohort was split into 70% training, 10% validation and 20% test.

Table 1. Cohort statistics of MIMIC-sepsis patients.

	N	% Female	Mean Age	Mean Hours in ICU
Survivors	18,057	44.3%	64	56.6
Non-survivors	2,883	42.7%	69	60.9

In an effort to explore the stability of training with function approximation, instead of the tabular lookup algorithm, we implemented a linear approximator for the Q-function with a one-hot state feature encoding (based on the clustering results), where we aimed to minimize the mean squared TD error (Sutton & Barto, 2018). This setup allows the implementation to be readily extended to any linear (or possibly non-linear) approximation of the Q-function.

4.2. Baselines

We compare our proposed SVP learning algorithm based on a near-greedy heuristic to one based on a conservative heuristic. In addition, we compare to three other baselines:

- *Q*-based* (Section 2.3). This approach assumes that the future follows an optimal policy, which can result in $\pi(s)$ including arbitrarily bad actions, especially in complex environments with long horizons.
- *Q-based*. In Algorithm 1, replace the policy construction step (line 8) of update rule with:

$$\pi(s) = \{a : Q(s, a) \geq (1 - \zeta)V(s)\},$$

where $V(s) = \max_{a \in \mathcal{A}} Q(s, a)$ replaces $V^*(s)$. This is similar to the near-greedy algorithm, except that it uses V , the worst-case value of the SVP as the baseline for near-optimality, instead of V^* in the definition of ζ -optimal, so the optimality constraint may be violated.

- *Additive.* After learning Q^* , we construct π following the additive constraint definition in [Appendix A](#):

$$\pi(s) = \{a : Q^\pi(s, a) \geq V^*(s) - \epsilon\},$$

where $\epsilon = \zeta(1 - \gamma)\|V^*\|_\infty$. Using this selection criteria, we can guarantee that $\forall s : V^\pi(s) \geq V^*(s) - \zeta\|V^*\|_\infty$.

In addition to comparing to these model-free alternatives, we compare to the MIP approach proposed by [Fard & Pineau \(2011\)](#) in [Appendix F](#) (their method requires knowledge of the MDP model). Another alternative involves first learning a (parameterized) probabilistic policy $\pi_\theta(a|s)$ and then converting it to an SVP by thresholding based on the action probabilities, $\pi(s) = \{a : \pi_\theta(a|s) \geq \tau\}$. Relevant methods to learn probabilistic policies include: diversity-inducing policy gradient ([Masood & Doshi-Velez, 2019](#)) and maximum entropy approaches (e.g., soft actor-critic ([Haarnoja et al., 2018](#))). However, such approaches rely on probabilistic assumptions when computing Q-values and require one to choose a probability threshold τ , which does not align with our definition of near-optimality.

4.3. Evaluation

We perform qualitative evaluations by inspecting the learned sets of near-equivalent actions. In particular, we visualize the different routes induced by learned SVPs in the **Frozen-Lake** environment, as well as near-equivalent treatment actions in **MIMIC-sepsis**.

To evaluate the quality of learned policies, we use standard Q-learning to establish the optimal Q-function Q^* as the baseline for deciding near-optimality. For synthetic environments, we characterize SVPs in terms of:

- Average policy size: the expected number of actions to which the SVP maps a state, $\frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} |\pi(s)|$. A larger average policy size means more actions are considered near-equivalent within the sub-optimality margin ζ .
- Worst-case near-optimality: we consider the largest deviation from V^* among all states, $1 - \min_{s \in \mathcal{S}} \frac{V^\pi(s)}{V^*(s)}$. This represents to what extent, in the worst-case, optimality is sacrificed in exchange for more choice. The value of each state is found by running a modified version of policy evaluation for the returned SVP (see [Appendix G](#) for pseudo-code and convergence result).

Evaluating the learned policies on the real data task **MIMIC-sepsis** presents significant challenges due to the present limitations of off-policy evaluation methods ([Imbens & Rubin, 2015](#); [Thomas & Brunskill, 2016](#); [Gottesman et al., 2018](#)). However, our main focus is on testing whether or not we can learn reasonable near-equivalent actions, rather than on learning the best policy for sepsis treatment. Still,

for completeness, we evaluated the quality of the learned policies using multiple off-policy evaluation estimators and considered a non-deterministic variation of the learned policies ([Gottesman et al., 2018](#)). For both the estimated optimal policy and the learned SVPs, we follow [Komorowski et al. \(2018\)](#) and evaluate softened policies: 99% probability is distributed among actions in the recommended set; the remaining 1% is distributed to non-suggested actions. This allows us to use nearly all sample trajectories in the test set, maintaining a large effective sample size. We applied two off-policy evaluation methods, doubly-robust estimator (DR) ([Jiang & Li, 2016](#)) and weighted doubly robust estimator (WDR) ([Thomas & Brunskill, 2016](#)), to evaluate the softened policies derived from the learned policies, computing empirical error bars based on 1,000 bootstraps of the test set.

5. Results

Across environments, our proposed approach is able to learn SVPs and discover near-equivalent actions. Empirically, we observe good convergence with near-optimal behavior under reasonable settings of γ and ζ even within non-DAG environments. The near-greedy algorithm outperforms the baselines, while the learned SVPs induce diverse behavior and meaningful alternative treatment recommendations.

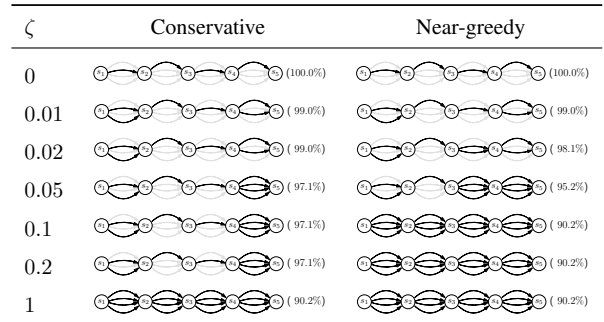


Figure 2. SVPs learned by the near-greedy and conservative algorithms on **Chain-5** at different ζ s. Parenthesized percentages denote the worst-case near-optimality.

5.1. Applied to Synthetic Data

In the simple DAG environment, **Chain-5** with $\gamma = 0.9$, we recover near-equivalent actions. When $\zeta = 0.01$ is small, $\pi(s_1)$ and $\pi(s_4)$ for both the near-greedy and conservative heuristics include two actions ([Figure 2](#)). As ζ increases, so does the number of actions, resulting in diversity and choice among near-equivalent actions. As expected, for the same ζ , the conservative approach contains fewer actions compared to the near-greedy SVP. Moreover, for a wide range of $\zeta \geq 0.05$, the conservative approach underestimates the near-optimality, failing to produce a diversity of choice.

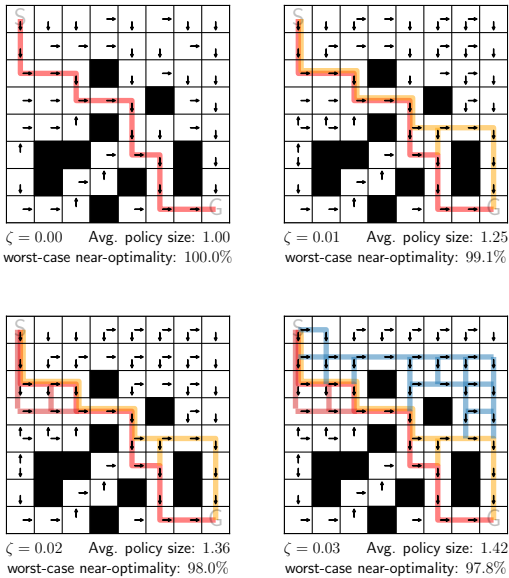


Figure 3. On the FrozenLake-8x8 environment, at different levels of near-optimality, ζ , the near-greedy algorithm learns different but near-equivalent routes to the goal, G. The learned SVPs conform to the near-optimality margin in the worst-case for the range of ζ values shown.

On FrozenLake-8x8 with $\gamma = 0.9$, our proposed near-greedy algorithm discovers different routes to the goal tile (Figure 3). Due to subtle differences in instantaneous rewards (see Section 4.1), these routes are not exactly the same, but are near-equivalent.

How does the algorithm perform empirically on non-DAGs? Here, we investigate the empirical performance of the near-greedy algorithm in non-DAG environments, namely, CyclicChain-5 and FrozenLake-4x4. We substituted the near-greedy heuristic as the policy improvement step in the value iteration algorithm and monitor whether the SVP (derived from the learned value estimate V) stabilizes towards the end of training. For small γ and small ζ , the near-greedy algorithm demonstrates good convergence (Figure 4). Under settings with large γ or large ζ (close to 1), the algorithm displays instability. For large γ , the effective horizon is longer and the worst-case value function must account for a longer future. When ζ is large, we permit more sub-optimal actions, and the local effect of adding an action to $\pi(s)$ could be so large that the action is no longer near-optimal, resulting in oscillating behavior and non-convergence. Notably, for the region of parameter values where the algorithm converges, the learned SVPs are non-trivial solutions (they include more than one action for certain states), as indicated by lighter colors in Figure 4. Empirically, for many operating regimes with reasonable settings of γ and ζ (e.g., small ζ for near-optimality), the near-greedy algorithm exhibits good convergence behavior.




Figure 4. Empirical behavior of the near-greedy value iteration algorithm, visualized as a heatmap of policy size at the end of training. White \times 's indicate non-convergence (i.e., the SVP derived from the learned V does not stabilize towards the end of training). In regions of the parameter space where the algorithm converges, we are able to discover non-trivial solutions (learned policies include more than one action for some states).

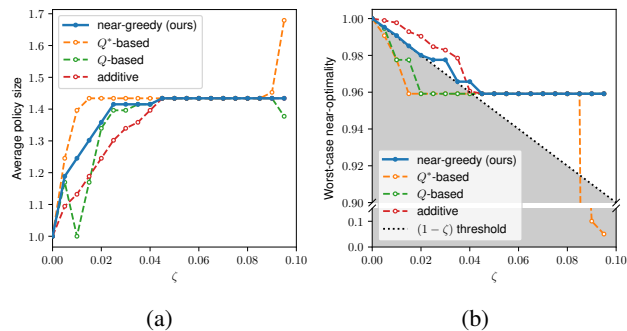


Figure 5. Comparison of average policy size and worst-case near-optimality of SVPs for different values of ζ on FrozenLake-8x8.

How does the proposed algorithm compare to the alternatives? Using the FrozenLake-8x8 environment with $\gamma = 0.9$, we compare the proposed approach to the three baselines in terms of average policy size (Figure 5a) and worst-case near-optimality (Figure 5b). We focused on small values of ζ , given the goal is to learn near-optimal behavior. For this particular environment, the conservative algorithm completely fails to discover near-equivalent actions. As ζ increases, the near-greedy finds solutions of comparable policy size to the baselines. However, holistically, near-greedy performs best because it respects the predefined optimality threshold (lies to the upper right of the shaded region) while maximizing the policy size. For $0 < \zeta < 0.04$, both Q^* -based and Q -based baselines violate the predefined near-optimality threshold while only yielding marginally larger policy sizes. For $\zeta \geq 0.9$, the Q^* -based approach finds an arbitrarily bad solution with a worst-case near-optimality of $< 10\%$. This simple environment illustrates the importance of considering a worst-case future during learning and accounting for the inter-dependence of including/excluding actions at different states. On this environment, we used

$\|V^*\|_\infty$ to define the margin ϵ for *additive*, and the solution includes fewer actions and does not make good use of the allowable sub-optimality compared to *near-greedy* for $\zeta \leq 0.04$ (for other values of ζ it leads to the same behavior as *near-greedy*).

5.2. Applied to Real Clinical Data

For the MIMIC-sepsis task, we apply our proposed approach with a linear function approximator by first running Q-learning to estimate Q^* , and then running the near-greedy algorithm with various values of ζ . During training, each episode is generated by randomly sampling a patient trajectory from the training set (with replacement). Given the complexity of this environment, to improve convergence, we exponentially decay the step size α every 1,000 episodes. We train the RL agent for 1,000,000 episodes, after which TD errors stabilize and the estimated Q-values reach plateaus. We consider ζ values from 0.0 to 0.1 to obtain SVPs that are near-optimal.

For illustrative purposes, we focus on an intermediate value of $\zeta = 0.05$, where 50% of the states are mapped to more than one near-optimal action. On the test set, the estimated value of the softened policy derived from π is within approximately 5–8% of the estimated optimal value (Table 2), which closely matches with the optimality margin $\zeta = 0.05$, given the complexity of this task and the noise in the data.

To better understand the near-equivalence relationships among actions, for each ‘optimal’ action, we count the number of times every other action is suggested by the learned SVP. These counts are aggregated over all states in the test set and normalized by the maximum count (red numbers next to each cell), visualized as a heatmap in Figure 6. In interpreting these results, we worked closely with our co-author, MWS, a critical care physician who frequently treats patients with sepsis.

By far, the most common actions correspond to IV fluids alone (region A), i.e., no vasopressors are used. We observe that actions with similar amounts of IV fluids are often considered near-equivalent. Since the differences in fluid volumes across these actions (adjacent cells) are non-trivial, we hypothesize that grouping those actions together makes sense for some (but not all) patient states. Similarly, in region B where actions correspond to a low dose of IV fluids with various amounts of vasopressors, actions with similar vasopressor doses are often considered near-equivalent. We observe that the near-equivalent action sets are not always contiguous across IV/vasopressor doses. This is due, in part, to the fact that we restricted the action space $A(s)$ for each state to only those that were taken frequently. Interestingly, in region C, actions with very high vasopressor and IV fluid doses are considered near-equivalent to the null action (i.e., do nothing). Typically, only the sickest patients

Table 2. Value estimates of the learned SVPs on MIMIC-sepsis, with standard errors from 1,000 bootstraps of the test set. Effective sample size (measured as: usable trajectories / total number of trajectories, of the test set): 2,801/4,189.

ζ	observed returns of test set	
	DR	WDR
0, π^*	91.6 \pm .31	92.2 \pm .23
0.05	84.3 \pm .63	89.7 \pm .32




Figure 6. Near-equivalent relationships among actions in SVPs on the MIMIC-sepsis task with $\zeta = 0.05$. Both the outer grid and the inner grid contain the 25 actions, corresponding to 5 doses of vasopressor and 5 doses of IV fluids. Red numbers indicate the frequency that each action is recommended as optimal.

are prescribed the highest doses. The near-equivalence of this action with the null action may be due to the fact that these patients are so critically ill that doing ‘everything’ or ‘nothing’ leads to a similar outcome.

In Figure 6, we computed the ‘average’ near-equivalencies across all states. However, different sets of actions could be considered near-equivalent for different states. To better understand how action near-equivalencies might differ for different group of patients, we grouped the states into quartiles based on the V^* values and calculated the average policy size (Figure 7). Compared to ‘sicker’ states (with a lower value), the ‘healthier’ states (with a higher value) had on average more near-equivalencies. We hypothesize that when a patient is stable, the choice of action has less impact on the final outcome compared to less stable states.

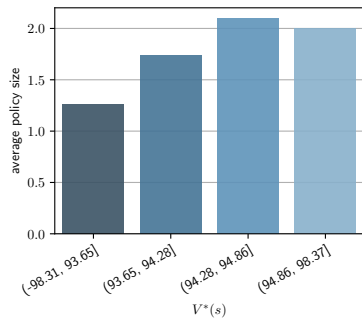


Figure 7. Average size of $\pi(s)$ at $\zeta = 0.05$ for states with different $V^*(s)$ values (grouped into quartiles) on MIMIC-sepsis.

6. Conclusion

In the context of learning SVPs for MDPs, we propose a model-free algorithm that is a variant of TD learning. On both synthetic and real RL tasks, our algorithm discovers meaningful action near-equivalencies, while maintaining overall near-optimality across states. Though the theoretical guarantees only hold for DAG settings, the near-greedy action selection heuristic can be easily extended to more complex settings involving non-DAGs and function approximation. In practice, to improve convergence and near-optimality guarantees, one could encode temporal information into the states for a discrete state space, converting a single ground state with different histories (e.g., visit number) into different states, effectively making the MDP a DAG. Despite current limitations, our proposed framework represents an important step toward clinician/human-in-the-loop decision making. Such a framework, in which both optimality guarantees and action choices are provided, allows clinicians (and patients) to incorporate additional information when making treatment decisions. Though motivated in a healthcare setting, our approach could apply more broadly to other application domains involving humans-in-the-loop, such as intelligent tutoring systems or self-driving cars.

Acknowledgments

This work was supported by the National Science Foundation (NSF award no. IIS-1553146) and the National Library of Medicine (NLM grant no. R01LM013325). The views and conclusions in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the National Science Foundation nor the National Library of Medicine. The authors would like to thank Satinder Singh, Brahmajee Nallamothu, Jessica Golbus, and members of the [MLD3 group](#) for helpful discussions regarding this work, as well as the reviewers for constructive feedback.

References

- Bedoya, A. D., Futoma, J., Clement, M. E., Corey, K., Brajer, N., Lin, A., Simons, M. G., Gao, M., Nichols, M., Balu, S., Heller, K., Sendak, M., and O’Brien, C. Machine learning for early detection of sepsis: an internal and temporal validation study. *JAMIA Open*, 3(2):252–260, Apr 2020. URL <https://doi.org/10.1093/jamiaopen/ooaa006>.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. OpenAI Gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Byrne, L. and Van Haren, F. Fluid resuscitation in human sepsis: time to rewrite history? *Annals of intensive care*, 7(1):4, 2017.
- Fard, M. M. and Pineau, J. Non-deterministic policies in Markovian decision processes. *Journal of Artificial Intelligence Research*, 40:1–24, 2011.
- Gottesman, O., Johansson, F., Meier, J., Dent, J., Lee, D., Srinivasan, S., Zhang, L., Ding, Y., Wihl, D., Peng, X., et al. Evaluating reinforcement learning algorithms in observational health settings. *arXiv preprint arXiv:1805.12298*, 2018.
- Gotts, J. E. and Matthay, M. A. Sepsis: pathophysiology and clinical management. *BMJ*, 353:i1585, 2016.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1861–1870, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL <http://proceedings.mlr.press/v80/haarnoja18b.html>.
- Henry, K. E., Hager, D. N., Pronovost, P. J., and Saria, S. A targeted real-time early warning score (TREWScore) for septic shock. *Science translational medicine*, 7(299):299ra122, 2015.
- Imbens, G. W. and Rubin, D. B. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- Jiang, N. and Li, L. Doubly robust off-policy value evaluation for reinforcement learning. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 652–661, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <http://proceedings.mlr.press/v48/jiang16.html>.

- 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. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3:160035, 2016.
- Kakade, S. and Langford, J. Approximately optimal approximate reinforcement learning. In *Proceedings of the Nineteenth International Conference on Machine Learning*, ICML '02, pp. 267–274, San Francisco, CA, USA, 2002. Morgan Kaufmann Publishers Inc. ISBN 1558608737.
- Kearns, M. and Singh, S. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2–3):209–232, 2002.
- Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C., and Faisal, A. A. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, 24(11):1716–1720, 2018.
- Li, L., Komorowski, M., and Faisal, A. A. The actor search tree critic (ASTC) for off-policy POMDP learning in medical decision making. *arXiv preprint arXiv:1805.11548*, 2018.
- Liu, V., Escobar, G. J., Greene, J. D., Soule, J., Whippy, A., Angus, D. C., and Iwashyna, T. J. Hospital deaths in patients with sepsis from 2 independent cohorts. *JAMA*, 312(1):90–92, 07 2014. URL <https://doi.org/10.1001/jama.2014.5804>.
- Lizotte, D. J. and Laber, E. B. Multi-objective Markov decision processes for data-driven decision support. *The Journal of Machine Learning Research*, 17(1):7378–7405, 2016.
- Lizotte, D. J., Bowling, M., and Murphy, S. A. Linear fitted-Q iteration with multiple reward functions. *Journal of Machine Learning Research*, 13(Nov):3253–3295, 2012.
- Masood, M. and Doshi-Velez, F. Diversity-inducing policy gradient: Using maximum mean discrepancy to find a set of diverse policies. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pp. 5923–5929. International Joint Conferences on Artificial Intelligence Organization, Jul 2019. URL <https://doi.org/10.24963/ijcai.2019/821>.
- Melo, F. S. Convergence of Q-learning: A simple proof. *Institute Of Systems and Robotics, Tech. Rep.*, pp. 1–4, 2001.
- Nemati, S., Ghassemi, M. M., and Clifford, G. D. Optimal medication dosing from suboptimal clinical examples: A deep reinforcement learning approach. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 2978–2981. IEEE, 2016.
- Raghu, A., Komorowski, M., Ahmed, I., Celi, L., Szolovits, P., and Ghassemi, M. Deep reinforcement learning for sepsis treatment. *Neural Information Processing Systems: workshop on Machine Learning for Health (ML4H)*, 2017a.
- Raghu, A., Komorowski, M., Celi, L. A., Szolovits, P., and Ghassemi, M. Continuous state-space models for optimal sepsis treatment: a deep reinforcement learning approach. In Doshi-Velez, F., Fackler, J., Kale, D., Ranganath, R., Wallace, B., and Wiens, J. (eds.), *Proceedings of the 2nd Machine Learning for Healthcare Conference*, volume 68 of *Proceedings of Machine Learning Research*, pp. 147–163, Boston, Massachusetts, 18–19 Aug 2017b. PMLR. URL <http://proceedings.mlr.press/v68/raghu17a.html>.
- Reyna, M. A., Josef, C., Seyedi, S., Jeter, R., Shashikumar, S. P., Westover, M. B., Sharma, A., Nemati, S., and Clifford, G. D. Early prediction of sepsis from clinical data: the physionet/computing in cardiology challenge 2019. In *2019 Computing in Cardiology (CinC)*, pp. Page–1. IEEE, 2019.
- Robbins, H. and Monro, S. A stochastic approximation method. *Ann. Math. Statist.*, 22(3):400–407, 09 1951. doi: 10.1214/aoms/1177729586. URL <https://doi.org/10.1214/aoms/1177729586>.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Thomas, P. and Brunskill, E. Data-efficient off-policy policy evaluation for reinforcement learning. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2139–2148, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <http://proceedings.mlr.press/v48/thomasa16.html>.
- Van Seijen, H., Van Hasselt, H., Whiteson, S., and Wiering, M. A theoretical and empirical analysis of Expected Sarsa. In *2009 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, pp. 177–184. IEEE, 2009.
- Watkins, C. J. and Dayan, P. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- Williams, D. *Probability with martingales*. Cambridge university press, 1991.