

Fixed Points of Set-Based Bellman Operator

Sarah H.Q. Li * Assalé Adjé ** Pierre-Loïc Garoche ***
Behçet Açıkmeşe *

* *William E. Boeing Department of Aeronautics and Astronautics, University of Washington, Seattle, USA. (e-mail: sarahli@uw.edu, behcet@uw.edu).*

** *LAMPS, Université de Perpignan Via Domitia, Perpignan, France (e-mail: assale.adje@univ-perp.fr)*

*** *ONERA – The French Aerospace Lab, Univ. of Toulouse, France, (e-mail: pierre-loic.garoche@onera.fr)*

Abstract: Motivated by uncertain parameters encountered in Markov decision processes (MDPs), we study the effect of parameter uncertainty on Bellman operator-based methods. Specifically, we consider a family of MDPs where the cost parameters are from a given compact set. We then define a Bellman operator acting on an input set of value functions to produce a new set of value functions as the output under all possible variations in the cost parameters. Finally we prove the existence of a *fixed point* of this set-based Bellman operator by showing that it is a contractive operator on a complete metric space.

Keywords: Markov decision process, stochastic control, game theory

1. INTRODUCTION

Markov decision process (MDP) is a widely used mathematical framework for control design in stochastic environments, eg. density control of a swarm of agents (Açıkmeşe and Bayard, 2012; Demir et al., 2015). It is also a fundamental framework for reinforcement learning, robotic motion planning and stochastic games (Filar and Vrieze, 2012; Li et al., 2019). An MDP can be solved for different objectives including minimum average cost, minimum discounted cost, reachability, among others (Puterman, 2014). Given an objective, solving an MDP is equivalent to computing the optimal policy and the optimal *value function* of a decision maker over the state space. Among different algorithms for computing the optimal policy, most are based on the Bellman equation that characterizes the value function as its fixed point.

In applications, it is common to encounter MDPs with uncertainties. When modeling an environment as a stochastic process, sampling techniques are often used to determine process parameters such as process costs or probabilities; such models are inherently uncertain. In stochastic games, the cost and probability parameters change with respect to another decision maker's strategy. While existing works focus on certain perturbations in MDPs (Bielecki and Filar, 1991; Altman and Gaitsgory, 1993; Abbad and Filar, 1992), these results do not generalize to the analysis of overall behaviour of the MDP under all possible cost parameters in a compact set.

Additionally, how uncertainty in MDP cost parameters affect the outcome of value iteration type methods is not well studied. Dynamic programming on bounded MDPs is studied in Givan et al. (2000) for specifically interval sets, however convergence over general compact sets is not considered. While computation of the fixed points of Bellman operator is the topic of numerous studies (Delage and Mannor, 2010), most focus on the convergence analysis of value iteration and its stopping criteria (Ashok et al., 2017; Eisentraut et al., 2019). However, they do not consider the relationship between bounds on the optimal

value function and the uncertainty in cost. Similarly motivated, Haddad and Monmege (2018) analyzes entry-wise uncertain transition kernels by using graph-based MDP transformations. While we also derive bounds of an MDP due to uncertain parameters, we differ in our approach: our set-based framework allows for direct extraction of the value iteration trajectories with respect to the set of cost parameters. This differentiates our work from Haddad and Monmege (2018) due to their graphical abstraction of MDP, which allows for derivation of bounds but not extraction of value function trajectories.

Contributions: We characterize the solutions of a family of MDPs at once, represented as *sets* of MDPs. More specifically, we: (i) develop a characterization of MDPs with uncertain cost parameters; (ii) propose a set-based Bellman operator over non-empty compact sets; (iii) establish the contractivity of this set-based Bellman operator with the existence of a unique compact fixed point set.

2. REVIEW OF MDPS AND BELLMAN OPERATOR

Notation: Sets of N elements are given by $[N] = \{0, \dots, N - 1\}$. We denote the set of matrices of i rows and j columns with real (non-negative) valued entries as $\mathbb{R}^{i \times j}(\mathbb{R}_+^{i \times j})$. Elements of sets and matrices are denoted by capital letters, X , while sets are denoted by cursive letters, \mathcal{X} . The ones column vector is denoted by $\mathbf{1}_N = [1, \dots, 1]^T \in \mathbb{R}^{N \times 1}$.

2.1 MDP

We consider a *discounted infinite-horizon MDP* defined by $([S], [A], P, C, \gamma)$ for a decision maker, where

- (1) $[S]$ denotes the finite set of states.
- (2) $[A]$ denotes the finite set of actions. Without loss of generality, assume that every action is admissible from each state $s \in [S]$.
- (3) $P \in \mathbb{R}^{S \times SA}$ denotes the transition kernel. Each component $P_{s',sa}$ is the probability of arriving in state s' by tak-

ing state-action (s, a) . Matrix P is column stochastic and element-wise non-negative — i.e. $\sum_{(s,a) \in [S] \times [A]} P_{s',sa} = 1$, $P_{s',sa} \geq 0$, $\forall s', s \in [S], a \in [A]$.

- (4) $C \in \mathbb{R}^{S \times A}$ denotes the cost of each pair (s, a) .
(5) $\gamma \in (0, 1)$ denotes the discount factor.

At each time step t , the decision maker chooses an action a based on its current state s . The state-action pair (s, a) induces a probability distribution $P_{(\cdot),sa} \in \mathbb{R}^S$, where $P_{s',sa}$ is the probability that the decision maker arrives at s' at time step $t + 1$. The state-action (s, a) also induces a cost C_{sa} that must be paid by the decision maker.

At each time step, the decision maker chooses a *policy* that dictates the action chosen at each state s . We denote policy as a function $\pi : S \times A \rightarrow \mathbb{R}_+$, where $\pi(s, a)$ denotes the probability that action a is chosen at state s . We denote the set of all feasible policies of an MDP by Π . In our context, it suffices to consider only *deterministic, stationary* policies i.e. $\pi(s, a)$ is a time invariant function that returns 1 for exactly one action, and 0 for all other possible actions.

We denote the policy matrix induced by a policy as $M_\pi \in \mathbb{R}^{S \times SA}$, where

$$(M_\pi)_{s',sa} = \begin{cases} \pi(s, a) & s' = s \\ 0 & s' \neq s \end{cases}$$

Every stationary policy induces a stationary *Markov chain* (El Chamie et al., 2018), given by PM_π^T . Each stationary policy also induces a stationary cost given by

$$C(\pi) = \sum_{i \in [S]} e_i e_i^T M_\pi (\mathbf{1}_S \otimes I_A) C^T e_i, \quad C(\pi) \in \mathbb{R}^S, \quad (1)$$

where $e_i \in \mathbb{R}^S$ is the unit vector pointing in the i^{th} coordinate.

For an MDP $([S], [A], P, C, \gamma)$, we are interested in minimizing the *discounted infinite horizon expected cost*, defined with respect to a policy π as

$$V_{s_0}^* = \min_{\pi \in \Pi} \mathbb{E}_{s_0}^\pi \left\{ \sum_{t=0}^{\infty} \gamma^t C_{s^t a^t} \right\}, \quad \forall s_0 \in [S] \quad (2)$$

where $\gamma \in (0, 1)$ is the discount factor of future cost, s^t and a^t are the state and action taken at time step t , and s^0 is the state that the decision maker starts from at $t = 0$.

The minimum expected cost V_s^* is called the *optimal value function*. The policy π^* that achieves the optimal value function is called an *optimal policy*. In general, V_s^* is unique while π^* is not. It is well known that the set of optimal policies always includes at least one *deterministic* stationary policy (Puterman, 2014, Thm 6.2.11) — i.e. for each s , $\pi(s, a)$ returns 1 for exactly one action, and 0 for all other possible actions.

2.2 Bellman Operator

Determining the optimal value function of a given MDP is equivalent to finding the fixed point of the associated Bellman operator, for which a myriad of techniques exists (Puterman, 2014). We introduce the Bellman operator here as well as relating its fixed point to the corresponding MDP problem.

Definition 1. (Standard Bellman Operator). For a discounted infinite horizon MDP $([S], [A], P, C, \gamma)$, its associated Bellman operator $f_C : \mathbb{R}^S \rightarrow \mathbb{R}^S$ is given component-wise by

$$(f_C(V))_s := \min_{a \in [A]} C_{sa} + \gamma \sum_{s' \in [S]} P_{s',sa} V_{s'}, \quad \forall s \in [S].$$

The fixed point of the Bellman operator is a value function $V \in \mathbb{R}^S$ that is invariant under the operator.

Definition 2. (Fixed Point). Let $F : \mathcal{X} \rightarrow \mathcal{X}$ be an operator on the metric space \mathcal{X} , $V^* \in \mathcal{X}$ is a fixed point of F if it satisfies

$$V^* = F(V^*). \quad (3)$$

In order to show that the Bellman operator has a unique fixed point, we consider the following operator property.

Definition 3. (Contraction Operator). Let (\mathcal{X}, d) be a complete metric space. An operator $F : \mathcal{X} \rightarrow \mathcal{X}$ is a contraction operator if it satisfies

$$d(F(V), F(V')) < d(V, V'), \quad \forall V, V' \in \mathcal{X}.$$

The Bellman operator is known to be a contraction operator on the complete metric space $(\mathbb{R}^S, \|\cdot\|_\infty)$. From the Banach fixed point theorem (Puterman, 2014), it has a unique fixed point. Because the optimal value function V^* is given by the unique fixed point of the associated Bellman operator, we use the terms optimal value function and fixed point of f_C interchangeably.

In addition to obtaining V^* , MDPs are also solved to determine the *optimal policy*, π^* . We note that because every feasible policy π induces a Markov chain, π also induces a unique stationary value function V which satisfies

$$V = C(\pi) + \gamma M_\pi P^T V. \quad (4)$$

Given a feasible policy π , we can equivalently solve for the stationary value function V as $V = (I - \gamma M_\pi P^T)^{-1} C(\pi)$. From this perspective, the optimal value function V^* is the minimum vector among the finite set of stationary value functions generated by the set of policies Π .

From the optimal value function V^* , we can also derive a deterministic optimal policy $\pi^*(s, a)$ from the Bellman operator for $(s, a) \in [S] \times [A]$ as

$$\pi^*(s, a) = \begin{cases} 1 & a = \underset{\bar{a} \in [A]}{\operatorname{argmin}} C_{s\bar{a}} + \gamma \sum_{s' \in [S]} P_{s',s\bar{a}} V_{s'}^* \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

where we note that when multiple actions achieve the minimum of $C_{s\bar{a}} + \gamma \sum_{s' \in [S]} P_{s',s\bar{a}} V_{s'}^*$ for state s , the operator argmin returns only one of the actions that achieves the minimum; the choice of which action does not matter. While the optimal policy does not need to be deterministic and stationary, the optimal policy π^* derived from (5) will always be deterministic.

2.3 Termination Criteria for Value Iteration

Among different algorithms to determine the fixed point of the Bellman operator, *value iteration* (VI) is a commonly used and simple technique in which the Bellman operator is iteratively applied until the optimal value is reached — i.e. starting from any value function $V^0 \in \mathbb{R}^S$ and $k = 1, \dots$, we apply

$$V_s^{k+1} = \min_{a \in [A]} C_{sa} + \gamma \sum_{s' \in [S]} P_{s',sa} V_{s'}^k, \quad \forall s \in [S]. \quad (6)$$

The iteration scheme given by (6) converges to the fixed point of the corresponding discounted infinite horizon MDP. The *stopping criteria* of VI can be considered the over-approximation of the optimal value function.

Lemma 4. (Puterman, 2014, Thm. 6.3.1) For any initial value function $V^0 \in \mathbb{R}^S$, let $\{V^k\}_{k \in \mathbb{N}}$ be the value function trajectory from (6). Whenever there exists $\epsilon > 0$, such that $\|V^{k+1} - V^k\|_\infty < \epsilon \frac{(1-\gamma)}{2\gamma}$, then V^{k+1} is within $\epsilon/2$ of the fixed point V^* , i.e. $\|V^{k+1} - V^*\|_\infty < \frac{\epsilon}{2}$.

Lemma 4 connects the sequence $\{V^k\}_{k \in \mathbb{N}}$'s relative convergence to its absolute convergence towards V^* by showing that the former implies the latter. In general, the stopping criteria differ for different MDP objectives (see Haddad and Monmege (2018) for recent results on stopping criteria for reachability).

3. SET-BASED BELLMAN OPERATOR

The standard Bellman operator with respect to a fixed cost parameter C is well studied. Motivated by a *family* of MDPs corresponding to a compact set of cost parameters $\mathcal{C} \subseteq \mathbb{R}^{S \times A}$ with all other data parameters remaining identical, we *lift* the Bellman operator to operate on sets rather than individual vectors in \mathbb{R}^S . For the set-based operator, we analyze its set-based domain and prove that it is a contraction operator. We also prove the existence of a unique fixed point *set* \mathcal{V}^* for a set-based Bellman operator and relate its properties to the fixed point of the standard Bellman operator.

3.1 Set-based operator properties

We define a new metric space $(H(\mathbb{R}^S), d_H)$ based on the Banach space $(\mathbb{R}^S, \|\cdot\|_\infty)$ to serve as our set-based operator domain (Rudin et al., 1964), where $H(\mathbb{R}^S)$ is the collection of non-empty compact subsets of \mathbb{R}^S equipped with *partial order*: for $\mathcal{V}, \mathcal{V}' \in H(\mathbb{R}^S)$, $\mathcal{V} \preceq \mathcal{V}'$ if $\mathcal{V} \subseteq \mathcal{V}'$ — i.e. if \mathcal{V} is a subset of \mathcal{V}' . The metric d_H is the *Hausdorff distance* (Henrikson, 1999) defined as

$$d_H(\mathcal{V}, \mathcal{V}') = \max\left\{ \sup_{V \in \mathcal{V}} \inf_{V' \in \mathcal{V}'} \|V - V'\|_\infty, \sup_{V' \in \mathcal{V}'} \inf_{V \in \mathcal{V}} \|V - V'\|_\infty \right\}. \quad (7)$$

Since $(\mathbb{R}^S, \|\cdot\|_\infty)$ is a complete metric space, $H(\mathbb{R}^S)$ is a complete metric space with respect to d_H .

Lemma 5. (Henrikson, 1999, Thm 3.3) If \mathcal{X} is a complete metric space, then its induced Hausdorff metric space $(H(\mathcal{X}), d_H)$ is a complete metric space.

On the metric space $H(\mathbb{R}^S)$, we define a *set-based Bellman operator*.

Definition 6. (Set-based Bellman Operator). For a family of MDP problems, $([S], [A], P, \mathcal{C}, \gamma)$, where $\mathcal{C} \subseteq \mathbb{R}^{S \times A}$ is a compact set, its associated set-based Bellman operator is given by

$$F_{\mathcal{C}}(\mathcal{V}) = \text{cl} \bigcup_{(C, V) \in \mathcal{C} \times \mathcal{V}} f_C(V), \quad \forall \mathcal{V} \in H(\mathbb{R}^S)$$

where cl is the closure operator.

As we take the union of uncountably many bounded sets, the resulting set may not be bounded, and therefore it is not immediately obvious that $F_{\mathcal{C}}(\mathcal{V})$ maps into the metric space $H(\mathbb{R}^S)$. We show this is true in Proposition 7.

Proposition 7. If \mathcal{C} is compact, then $F_{\mathcal{C}}(\mathcal{V}) \in H(\mathbb{R}^S)$, $\forall \mathcal{V} \in H(\mathbb{R}^S)$.

Proof. For a non-empty set \mathcal{A} of some finite dimensional real vector space, let us define its diameter to be denoted as $\text{diam}(\mathcal{A}) = \sup_{x, y \in \mathcal{A}} \|x - y\|_\infty$. The diameter of any compact set in a metric space is bounded.

We take any non-empty compact set $\mathcal{V} \in H(\mathbb{R}^S)$. As $F_{\mathcal{C}}(\mathcal{V}) \subseteq \mathbb{R}^S$, it suffices to prove that $F_{\mathcal{C}}(\mathcal{V})$ is closed and bounded. The closedness is guaranteed by the closure operator. A subset of a metric space is bounded iff its closure is bounded.

Hence, to prove the boundedness, it suffices to prove that $\text{diam}(\bigcup_{(C, V) \in \mathcal{C} \times \mathcal{V}} f_C(V)) < +\infty$. Consider any two cost-value function pairs, $(C, V), (C', V') \in \mathcal{C} \times \mathcal{V}$, they must satisfy

$$f_C(V) - f_{C'}(V') = (f_C(V) - f_{C'}(V)) + (f_{C'}(V) - f_{C'}(V')),$$

where the norm of the second term $\|f_{C'}(V) - f_{C'}(V')\|_\infty$ must be upper bounded by $\|V - V'\|_\infty$ due to contraction properties of $f_{C'}$. To bound the first term, we note that for any two vectors $a, b \in \mathbb{R}^S$, $\|a - b\|_\infty = \max\{\max(a - b), \max(b - a)\}$ and let π to be the optimal policy of $f_C(V)$,

$$\begin{aligned} & \max(f_{C'}(V) - f_C(V)) \\ & \leq \max(\nu'(\pi) + \gamma M_\pi P^T V - \nu(\pi) - \gamma M_\pi P^T V) \\ & \leq \max(\nu'(\pi) - \nu(\pi)) \\ & \leq \sum_{i \in [S]} \|e_i^T\|_\infty \|M_\pi\|_\infty \|\mathbf{1}_S \otimes I_A\|_\infty \|(C' - C)^T\|_\infty \|e_i\|_\infty^2. \end{aligned}$$

Since $\|\mathbf{1}_S \otimes I_A\|_\infty = \|e_i\|_\infty = \|e_i^T\|_\infty = \|M_\pi\|_\infty = 1$ for any $\pi \in \Pi$, $\max(f_{C'}(V) - f_C(V)) \leq S \text{diam}(\mathcal{C}^T)$. The result $\max(f_C(V) - f_{C'}(V)) \leq S \text{diam}(\mathcal{C}^T)$ can be similarly derived. Therefore, $\|f_C(V) - f_{C'}(V')\|_\infty < S \text{diam}(\mathcal{C}^T) + \text{diam}(\mathcal{V})$. Since it holds for all $(C, V), (C', V') \in \mathcal{C} \times \mathcal{V}$ then $\text{diam}(\bigcup_{(C, V) \in \mathcal{C} \times \mathcal{V}} f_C(V)) \leq S \text{diam}(\mathcal{C}^T) + \text{diam}(\mathcal{V}) < +\infty$ as \mathcal{C}^T and \mathcal{V} are bounded. \square

Proposition 7 shows that $F_{\mathcal{C}}$ is an operator from $H(\mathbb{R}^S)$ to $H(\mathbb{R}^S)$. Having established its range space, we can draw many parallels between $F_{\mathcal{C}}$ and f_C . Similar to the existence of a unique fixed point V^* for f_C , we consider whether a *fixed point set* of $F_{\mathcal{C}}$ which satisfies $F_{\mathcal{C}}(\mathcal{V}^*) = \mathcal{V}^*$ exists, and if it is unique. To take the comparison further, since V^* is the optimal value function for an MDP problem defined by $([S], [A], P, C, \gamma)$, how does \mathcal{V}^* relate to the *family* of optimal solutions that corresponds to the MDP family $([S], [A], P, \mathcal{C}, \gamma)$?

To prove the unique existence of \mathcal{V}^* , we utilize the Banach fixed point theorem (Puterman, 2014), which states that a unique fixed point must exist for all contraction operators on complete metric spaces. First, we show that $F_{\mathcal{C}}$ is a contraction as defined in Definition 3 on the complete metric space $(H(\mathbb{R}^S), d_H)$.

Proposition 8. For any $\mathcal{V} \in H(\mathbb{R}^S)$ and $\mathcal{C} \subseteq \mathbb{R}^{S \times A}$ closed and bounded, $F_{\mathcal{C}}$ is a contraction operator under the Hausdorff distance.

Proof. Consider $\mathcal{V}, \bar{\mathcal{V}} \in H(\mathbb{R}^S)$, to see that $F_{\mathcal{C}}$ is a contraction, we need to show

$$\sup_{V \in F_{\mathcal{C}}(\mathcal{V})} \inf_{\bar{V} \in F_{\mathcal{C}}(\bar{\mathcal{V}})} \|V - \bar{V}\|_\infty < d_H(\mathcal{V}, \bar{\mathcal{V}}) \quad (8)$$

$$\sup_{\bar{V} \in F_{\mathcal{C}}(\bar{\mathcal{V}})} \inf_{V \in F_{\mathcal{C}}(\mathcal{V})} \|V - \bar{V}\|_\infty < d_H(\mathcal{V}, \bar{\mathcal{V}}) \quad (9)$$

First we note that taking \sup (\inf) of a continuous function over a set \mathcal{A} is equivalent to taking the \sup (\inf) over the closure of \mathcal{A} . Let $G_{\mathcal{C}}(\mathcal{V}) = \bigcup_{(C, V) \in \mathcal{C} \times \mathcal{V}} f_C(V)$ and $\text{cl} G_{\mathcal{C}}(\mathcal{V}) = F_{\mathcal{C}}(\mathcal{V})$, then due to continuity of norms (Rudin et al., 1964, Thm 4.16),

$$\sup_{V \in F_{\mathcal{C}}(\mathcal{V})} \inf_{\bar{V} \in F_{\mathcal{C}}(\bar{\mathcal{V}})} \|V - \bar{V}\| = \sup_{V \in G_{\mathcal{C}}(\mathcal{V})} \inf_{\bar{V} \in G_{\mathcal{C}}(\bar{\mathcal{V}})} \|V - \bar{V}\|.$$

Therefore, it suffices to prove

$$\sup_{f_C(V) \in G_{\mathcal{C}}(\mathcal{V})} \inf_{f_{\bar{C}}(\bar{V}) \in G_{\mathcal{C}}(\bar{\mathcal{V}})} \|f_C(V) - f_{\bar{C}}(\bar{V})\|_\infty < d_H(\mathcal{V}, \bar{\mathcal{V}}),$$

$$\sup_{f_{\bar{C}}(\bar{V}) \in G_{\mathcal{C}}(\bar{\mathcal{V}})} \inf_{f_C(V) \in G_{\mathcal{C}}(\mathcal{V})} \|f_C(V) - f_{\bar{C}}(\bar{V})\|_\infty < d_H(\mathcal{V}, \bar{\mathcal{V}}).$$

For any $V \in \mathcal{V}$, $C \in \mathcal{C}$,

$$\inf_{(\bar{C}, \bar{V}) \in \mathcal{C} \times \bar{\mathcal{V}}} \|f_C(V) - f_{\bar{C}}(\bar{V})\|_\infty \quad (10a)$$

$$= \inf_{(\bar{C}, \bar{V}) \in \mathcal{C} \times \bar{\mathcal{V}}} \|C(\pi) + \gamma M_\pi P^T V - (\bar{C}(\bar{\pi}) + \gamma M_{\bar{\pi}} P^T \bar{V})\|_\infty, \quad (10b)$$

$$\leq \inf_{\bar{V} \in \bar{\mathcal{V}}} \|C(\bar{\pi}) + \gamma M_{\bar{\pi}} P^T V - (C(\bar{\pi}) + \gamma M_{\bar{\pi}} P^T \bar{V})\|_\infty, \quad (10c)$$

$$\leq \inf_{\bar{V} \in \bar{\mathcal{V}}} \|\gamma M_{\bar{\pi}} P^T (V - \bar{V})\|_\infty \leq \gamma \inf_{\bar{V} \in \bar{\mathcal{V}}} \|V - \bar{V}\|_\infty, \quad (10d)$$

where π corresponds to the optimal policy for the MDP $([S], [A], P, C, \gamma)$ and $\bar{\pi}$ corresponds to the optimal policy for the MDP $([S], [A], P, \bar{C}, \gamma)$ in (10b). In (10c) we replaced M_π by $M_{\bar{\pi}}$ by noting that π is optimal, therefore $\bar{\pi}$ must result in a larger value function (similar to the proof of Prop. 7). In (10d) we note that the infimum over set \mathcal{C} must be upper bounded by when $\bar{C} = C \in \mathcal{C}$, and used the fact that $\|M_{\bar{\pi}} P^T\|_\infty \leq 1$.

Taking the sup over $G_C(\mathcal{V})$ and $G_C(\bar{\mathcal{V}})$,

$$\sup_{V \in G_C(\mathcal{V})} \inf_{\bar{V} \in G_C(\bar{\mathcal{V}})} \|V - \bar{V}\|_\infty \leq \gamma \sup_{V \in \mathcal{V}} \inf_{\bar{V} \in \bar{\mathcal{V}}} \|V - \bar{V}\|_\infty,$$

$$\sup_{\bar{V} \in G_C(\bar{\mathcal{V}})} \inf_{V \in G_C(\mathcal{V})} \|V - \bar{V}\|_\infty \leq \gamma \sup_{V \in \mathcal{V}} \inf_{\bar{V} \in \bar{\mathcal{V}}} \|V - \bar{V}\|_\infty.$$

Therefore $d_H(F_C(\mathcal{V}), F_C(\bar{\mathcal{V}})) \leq \gamma d_H(\mathcal{V}, \bar{\mathcal{V}})$. Since $\gamma \in (0, 1)$, F_C is a contraction operator on $H(\mathbb{R}^S)$. \square

The contraction property of F_C implies that repeated application of the operator to any $\mathcal{V}^0 \in H(\mathbb{R}^S)$ will result in closer and closer sets in the Hausdorff sense of distance to a fixed point set. It is then natural to consider if there is a unique set which all $F_C^k(\mathcal{V}^0)$ converges to.

Theorem 9. There exists a unique fixed point \mathcal{V}^* to the set-based Bellman operator F_C as defined in Definition 6, such that $F_C(\mathcal{V}^*) = \mathcal{V}^*$, and \mathcal{V}^* is a closed and bounded subset of \mathbb{R}^S . Furthermore, for any iteration starting from arbitrary $\mathcal{V}^0 \in H(\mathbb{R}^S)$, $\mathcal{V}^{k+1} = F_C(\mathcal{V}^k)$, the sequence converges in the Hausdorff sense i.e. $\lim_{k \rightarrow \infty} d_H(F_C(\mathcal{V}^k), \mathcal{V}^*) = 0$.

Proof. As shown in Proposition 8, F_C is a contraction operator. From the Banach fixed point theorem (Puterman, 2014, Thm 6.2.3), there exists a unique fixed point \mathcal{V}^* , and any arbitrary $\mathcal{V}^0 \in H(\mathbb{R}^S)$ will generate a sequence $\{F_C(\mathcal{V}^k)\}_{k \in \mathbb{N}}$ that converges to the fixed point. \square

The fixed point V^* of Bellman operator f_C on metric space \mathbb{R}^S corresponds to the optimal value function of the MDP associated with cost parameter C . Because there is no direct association of an MDP problem to the set-based Bellman operator F_C , we cannot claim the same for \mathcal{V}^* . However, \mathcal{V}^* does have many interesting properties on $H(\mathbb{R}^S)$, in parallel to operator f_C on \mathbb{R}^S , especially in terms of the value iteration method (6). Suppose that instead of a fixed cost parameter, we have that at each iteration k , a C^k that is random chosen from a compact set of costs, $C^k \in \mathcal{C}$, then it is interesting to ask if \mathcal{V}^* contains all the limit points of $\lim_k f_{C^k}(V^k)$. Indeed, we can infer from Theorem 9 that the sequence $\{V_k\}$ converges to \mathcal{V}^* under the Hausdorff metric. Furthermore, even when V^k itself does not converge, it must converge to the set \mathcal{V}^* under the Hausdorff metric—i.e. $\lim_{k \rightarrow \infty} \inf_{V \in \mathcal{V}^*} \|V^k - V\|_\infty = 0$.

4. CONCLUSION

We summarize our results on set-based Bellman operator: for a compact cost function set \mathcal{C} , F_C converges to to a unique

compact set \mathcal{V}^* which contains all the fixed points of f_C for all fixed $C \in \mathcal{C}$. Furthermore, \mathcal{V}^* also contains the limit points of $f_{C^k}(V^k)$ for any $\{C^k\}_{k \in \mathbb{N}} \subseteq \mathcal{C}$, $V^0 \in \mathbb{R}^S$, given that $\lim_{k \rightarrow \infty} V^k$ converges. Even if the limit does not exist, V^k must asymptotically converge to \mathcal{V}^* in the Hausdorff sense. Future work includes extending the uncertainty analysis to consider uncertainty in the transition kernel to fully capture learning in a general stochastic game.

REFERENCES

- Abbad, M. and Filar, J.A. (1992). Perturbation and stability theory for markov control problems. *IEEE Trans. Autom. Control*.
- Açikmeşe, B. and Bayard, D.S. (2012). A markov chain approach to probabilistic swarm guidance. In *Amer. Control Conf.*, 6300–6307. IEEE.
- Altman, E. and Gaitsgory, V.A. (1993). Stability and singular perturbations in constrained markov decision problems. *IEEE Trans. Autom. Control*, 38(6), 971–975.
- Ashok, P., Chatterjee, K., Daca, P., Křetínský, J., and Meggen-dorfer, T. (2017). Value iteration for long-run average reward in markov decision processes. In *Int. Conf. Comput. Aided Verification*, 201–221. Springer.
- Bielecki, T.R. and Filar, J.A. (1991). Singularly perturbed markov control problem: Limiting average cost. *Ann. Op. Res.*, 28(1), 153–168.
- Delage, E. and Mannor, S. (2010). Percentile optimization for markov decision processes with parameter uncertainty. *Op. Res.*, 58(1), 203–213.
- Demir, N., Eren, U., and Açikmeşe, B. (2015). Decentralized probabilistic density control of autonomous swarms with safety. *Auton. Robots*, 39(4), 537–554.
- Eisentraut, J., Křetínský, J., and Rotar, A. (2019). Stopping criteria for value and strategy iteration on concurrent stochastic reachability games. *arXiv preprint arXiv:1909.08348*.
- El Chamie, M., Yu, Y., Açikmeşe, B., and Ono, M. (2018). Controlled markov processes with safety state constraints. *IEEE Trans. Autom. Control*, 64(3), 1003–1018.
- Filar, J. and Vrieze, K. (2012). *Competitive Markov decision processes*. Springer Science & Business Media.
- Givan, R., Leach, S., and Dean, T. (2000). Bounded-parameter markov decision processes. *Artif. Intell.*, 122(1-2), 71–109.
- Haddad, S. and Monmege, B. (2018). Interval iteration algorithm for mdps and imdps. *Theor. Comput. Sci.*, 735, 111–131.
- Henrikson, J. (1999). Completeness and total boundedness of the hausdorff metric. In *MIT Undergraduate J. Math.* Citeseer.
- Li, S.H.Q., Yu, Y., Calderone, D., Ratliff, L., and Açikmeşe, B. (2019). Tolling for constraint satisfaction in markov decision process congestion games. In *Amer. Control Conf.*, 1238–1243. IEEE.
- Puterman, M.L. (2014). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.
- Rudin, W. et al. (1964). *Principles of mathematical analysis*, volume 3. McGraw-hill New York.