

# Approximate Information States for Worst-Case Control and Learning in Uncertain Systems

Aditya Dave, *Member, IEEE*, Nishanth Venkatesh, *Student Member, IEEE*, and Andreas A. Malikopoulos, *Senior Member, IEEE*

**Abstract**—In this paper, we investigate discrete-time decision-making problems in uncertain systems with partially observed states. We consider a non-stochastic model, where uncontrolled disturbances acting on the system take values in bounded sets with unknown distributions. We present a general framework for decision-making in such problems by using the notion of the information state and approximate information state, and introduce conditions to identify an uncertain variable that can be used to compute an optimal strategy through a dynamic program (DP). Next, we relax these conditions and define *approximate information states* that can be learned from output data without knowledge of system dynamics. We use approximate information states to formulate a DP that yields a strategy with a bounded performance loss. Finally, we illustrate the application of our results in control and reinforcement learning using numerical examples.

**Index Terms**—Uncertain systems, worst-case control, approximate dynamic programming, robust reinforcement learning

## I. INTRODUCTION

Decision-making under incomplete information is a fundamental problem in modern engineering applications involving cyber-physical systems [1], e.g., connected and automated vehicles [2] and social media platforms [3]. In such applications, an agent must sequentially select control inputs to a dynamic system using only partial observations while accounting for uncontrolled disturbances that interfere with the system's evolution. The most common modeling paradigm for such decision-making problems is the *stochastic approach*, where all disturbances to the system are considered as random variables with known distributions, and the agent selects a decision-making strategy to minimize the *expected incurred cost* [4]. Stochastic models have been utilized for problems in both control theory [5]–[11] and reinforcement learning [12]–[16]. A strategy derived using the stochastic approach performs optimally on average across numerous operations of the system. However, this performance degrades with any mismatch between the distributions assumed in modeling and the frequency of realizations encountered in implementation [17]. Furthermore, many safety-critical applications require

This research was supported by NSF under Grants CNS-2149520 and CMMI-2219761.

The authors are with the School of Civil and Environmental Engineering, Cornell University, Ithaca, NY (email: a.dave@cornell.edu; ns924@cornell.edu; amaliko@cornell.edu).

guarantees on the agent's performance during each operation [18], making expected cost an inadequate measure.

The *non-stochastic approach* is an alternate modeling paradigm for safety-critical systems, where all disturbances are considered to belong to known sets with unknown distributions. The agent aims to select a decision-making strategy that minimizes the *worst-case incurred cost* across a time horizon [19]. Because this approach focuses on robustness against worst-case disturbances, the resulting strategy yields more conservative decisions than the stochastic approach and yields guarantees for each system operation. Thus, this approach has been applied to adversarial settings, e.g., cyber-security [20] or cyber-physical systems [21], and failure-critical settings, e.g., water reservoirs [22] or power systems [23].

In this paper, we propose a framework for non-stochastic decision-making using only partial observations in a dynamic system. When the system's dynamics are known to the agent, this problem falls under the purview of control theory [24]. However, many applications involve decision-making with incomplete knowledge of the dynamics as, e.g., mixed traffic driving [25] and human-robot coordination [26], [27], or decision-making without a reliable state-space model, e.g., healthcare [28]. These restrictions yield a reinforcement learning problem [29], [30]. To account for both cases, we formulate our problem using only output variables without assuming a known state-space model. In our exposition, we present rigorous definitions for the notions of *information states* and *approximate information states*. Using these notions, a surrogate state-space model can be constructed from output variables. This surrogate model can be used to formulate a control problem with full-state observation, whose solution yields a strategy of the original problem. In reinforcement learning problems, the surrogate model can be learned from output data. Given a surrogate model, the agent can derive a strategy using standard techniques [31].

### A. Related Work

1) *Control theory*: There have been numerous research efforts in control theory to study dynamic decision-making problems given the system dynamics. For both stochastic and non-stochastic models, an agent can derive an optimal decision-making strategy offline using a dynamic programming (DP) decomposition [32]. For systems with perfectly observed states, the agent's optimal action at each instance of time is simply a function of the state. Using this property in a

DP facilitates the efficient computation of an optimal control strategy [33], [34]. In contrast, for partially observed systems, any optimal action is a function of the agent's entire memory of past observations and actions, which grows in size with time [35]. Subsequently, the domain of the optimal strategy grows with time, and the DP for the problem requires a large number of computations [36]. This concern is alleviated using an *information state* instead of the memory [37].

The most commonly used information state in stochastic control is the *belief state*, i.e., a distribution on the state space conditioned on the agent's memory [38], [39]. A general notion of information states for stochastic control was recently defined in [40]. For non-stochastic control problems, the DP decomposition has been simplified using two well-known information states: (1) the *conditional range*, which is the set of feasible states at any time consistent with the agent's memory [41] and can be used in both terminal cost [42]–[45] and instantaneous cost problems [46]–[48]; and (2) the *maximum cost-to-come*, which is the maximum accrued cost functional [49] used in additive cost problems [50]–[52].

A general notion of an information state for non-stochastic terminal cost problems was presented in [53]. Information states have also been derived for risk-sensitive formulations [54] and distributionally robust formulations [55]. A related notion of symbolic abstractions has been developed in non-stochastic settings by drawing upon system approximation using bisimulation [56]. Typically, symbolic abstractions discretize the state and action spaces of continuous systems [57]. Such abstractions and their worst-case performance were studied for perfectly observed control problems in [58], [59].

The advantage of using information states (or other state abstractions) is that they result in a smaller approximate space. Thus, they generally yield a more efficient DP decomposition than the entire memory. However, in problems with large state spaces, utilizing information states may still not be practical [60], [61], creating a need for principled approximations.

2) *Reinforcement learning*: The literature on reinforcement learning is concerned with decision-making when the agent does not have prior knowledge of the system's dynamics [62]. For systems with perfectly observed states, these problems have been addressed using a variety of approaches [63]. In the stochastic formulation, both model-based [64], [65] and model-free approaches [66] have been utilized. In the non-stochastic formulation, the worst-case reinforcement learning problem was formulated and analyzed in [67]. Worst-case Q-learning was proposed for reinforcement learning problems in [68]–[70] and extended to problems with output feedback and known dynamics in [71]. Actor-critic methods [72] and model-based off-policy learning approaches [73] have also been developed for robust control. Alternate approaches using online adaptive algorithms were proposed in [74], [75]. However, in general, reinforcement learning with partial observations remains challenging, especially for non-stochastic formulations.

For stochastic formulations, the notion of *approximate information states* was presented in [76] to address the challenges of control and learning with partial observations. Approximate information states can improve the computational tractability of control problems at the cost of a bounded loss in

performance. In reinforcement learning, they can be learned from output data to facilitate standard learning algorithms. Their performance has been empirically validated in robotics [77] and healthcare [78]. No general theory of approximate information states exists yet for non-stochastic formulations.

### B. Contributions and Organization

In this paper, we develop a non-stochastic theory of approximate information states for both instantaneous and terminal cost problems to facilitate control and reinforcement learning using partial observations. Our contributions are: (1) the introduction of a general notion of *information states* (Definition 1) and an optimal DP decomposition for worst-case control (Theorem 1); (2) the introduction of *approximate information states* (Definition 2) that can either be constructed or learned from output data (Subsection IV-D); (3) the formulation of an approximate DP (Theorem 3) to compute a control strategy with a bounded optimality gap (Theorem 4); and (4) the exposition of examples (Subsection IV-C) with corresponding theoretical guarantees (Theorems 5 - 6). We also illustrate our results using numerical examples (Subsection V).

Note that while our theory shares conceptual similarities to that for stochastic problems in [76], our focus on non-stochastic problems necessitates the use of a distinct formulation with uncertain variables [79]. Our results bound the worst-case approximation loss rather than the expected loss. Our mathematical approach in the non-stochastic framework differs from [76] and results in novel theoretical results, including the Lipschitz continuity of the approximate value functions in Theorem 2, novel properties of the Hausdorff metric in Appendix B to prove Theorem 4, and novel worst-case performance bounds on approximate DPs using state-quantization in Subsection IV-C. We reported preliminary results for terminal-cost control problems in [53]. This paper extends the preliminary work as follows: (1) we consider worst-case instantaneous cost problems which subsume terminal cost problems; (2) we allow all variables to take values in continuous spaces; (3) we derive explicit bounds using state-discretization in both perfectly and partially observed systems; and (4) we illustrate the application of our results to a reinforcement learning problem.

The remainder of the paper proceeds as follows. In Section II, we present our problem formulation. In Section III, we define the notion of information states and prove the optimality of the corresponding DP decomposition. In Section IV, we present the notion of approximate information states, a resulting approximate DP, and theoretical bounds on the approximation loss. In Section V, we present a numerical example to illustrate the application of our results. In Section VI, we draw concluding remarks and discuss future work.

## II. MODELING FRAMEWORK

### A. Preliminaries

1) **Uncertain Variables**: In this paper, we utilize the mathematical framework for *uncertain variables* from [79]. An uncertain variable is a non-stochastic analogue of a random variable with set-valued uncertainty. For a sample space  $\Omega$  and a set  $\mathcal{X}$ , an uncertain variable is a mapping  $X : \Omega \rightarrow \mathcal{X}$ .

For any  $\omega \in \Omega$ , it has the realization  $X(\omega) = x \in \mathcal{X}$ . The *marginal range* of  $X$  is the set  $[[X]] := \{X(\omega) \mid \omega \in \Omega\}$ . For two uncertain variables  $X \in \mathcal{X}$  and  $Y \in \mathcal{Y}$ , their *joint range* is  $[[X, Y]] := \{(X(\omega), Y(\omega)) \mid \omega \in \Omega\}$ . For a given realization  $y$  of  $Y$ , the *conditional range* of  $X$  is  $[[X|y]] := \{X(\omega) \mid Y(\omega) = y, \omega \in \Omega\}$  and, generally,  $[[X|Y]] := \bigcup_{y \in [[Y]]} [[X|y]]$ .

**2) Hausdorff Distance:** Consider that the  $\mathcal{X}, \mathcal{Y}$  are nonempty subsets of a metric space  $(\mathcal{S}, \eta)$ , where  $\eta(\cdot, \cdot)$  is the metric. Then, the *Hausdorff distance* between  $\mathcal{X}$  and  $\mathcal{Y}$  is

$$\mathcal{H}(\mathcal{X}, \mathcal{Y}) := \max \left\{ \sup_{x \in \mathcal{X}} \inf_{y \in \mathcal{Y}} \eta(x, y), \sup_{y \in \mathcal{Y}} \inf_{x \in \mathcal{X}} \eta(x, y) \right\}. \quad (1)$$

When the two sets  $\mathcal{X}, \mathcal{Y}$  are bounded, the Hausdorff distance in (1) constitutes a pseudo-metric, i.e.,  $\mathcal{H}(\mathcal{X}, \mathcal{Y}) = 0$  if and only if  $\text{closure}(\mathcal{X}) = \text{closure}(\mathcal{Y})$  [80, Appendix]. When both  $\mathcal{X}, \mathcal{Y}$  are compact, the Hausdorff distance is a metric, i.e.,  $\mathcal{H}(\mathcal{X}, \mathcal{Y}) = 0$  if and only if  $\mathcal{X} = \mathcal{Y}$  [81, Chapter 1.12].

**3) L-invertible Functions:** Consider a function  $f : \mathcal{X} \rightarrow \mathcal{Y}$ . For any  $y \in \mathcal{Y}$ , the pre-image of the function is  $f^{-1}(y) = \{x \in \mathcal{X} \mid f(x) = y\}$ . Then, the function  $f$  is called *L-invertible* if there exists a constant  $L_{f^{-1}} \in \mathbb{R}_{\geq 0}$  such that

$$\mathcal{H}(f^{-1}(y^1), f^{-1}(y^2)) \leq L_{f^{-1}} \cdot \eta(y^1, y^2), \quad \forall y^1, y^2 \in \mathcal{Y}. \quad (2)$$

For uncertain variables  $X \in \mathcal{X}$  and  $Y \in \mathcal{Y}$  such that  $Y = f(X)$ , the pre-image of  $f$  given a realization  $y \in [[Y]]$  is  $[[X|y]]$ , i.e.,  $f^{-1}(y) = [[X|y]]$ . Thus, if  $f$  is *L-invertible*, we equivalently state that for all  $y^1, y^2 \in [[Y]]$  and  $L_{X|Y} = L_{f^{-1}}$ :

$$\mathcal{H}([[X|y^1]], [[X|y^2]]) \leq L_{X|Y} \cdot \eta(y^1, y^2). \quad (3)$$

**4) Notation:** Uncertain variables will be denoted by uppercase letters and their realizations by lower-case letters. Variables and functions with a bar on them, such as  $\bar{V}(\cdot)$ , pertain to information states, and those with a hat on them, such as  $\hat{V}(\cdot)$ , pertain to approximate information states.

### B. Problem Formulation

We consider an agent that seeks to control an uncertain system over  $T \in \mathbb{N}$  discrete time steps. At each time  $t = 0, \dots, T$ , the agent receives an observation from the system, denoted by the uncertain variable  $Y_t \in \mathcal{Y}_t$ , and generates a control action denoted by the uncertain variable  $U_t \in \mathcal{U}_t$ . After generating the action at each  $t$ , the agent incurs a cost denoted by the uncertain variable  $C_t \in \mathcal{C}_t \subset \mathbb{R}_{\geq 0}$ . To account for a possibly unknown state-space model, we describe the dynamics using an *input-output model*, as follows. At each  $t$ , the system receives two inputs: the control action  $U_t$ , and an uncontrolled disturbance  $W_t \in \mathcal{W}_t$ . The uncontrolled disturbances  $\{W_t \mid t = 0, \dots, T\}$  constitute a sequence of independent uncertain variables that enter the system through a channel separate from the control action. After receiving the inputs at each  $t$ , the system generates two outputs:

$$Y_{t+1} = h_{t+1}(W_{0:t}, U_{0:t}), \quad (4)$$

$$C_t = d_t(W_{0:t}, U_{0:t}), \quad (5)$$

using some observation function  $h_{t+1}$  and cost function  $d_t$ . The initial observation is generated as  $Y_0 = h_0(W_0)$ .

The agent has a perfect recall of the history of observations and control actions. The memory of the agent at each  $t$  is

denoted by the uncertain variable  $M_t := (Y_{0:t}, U_{0:t-1})$ , which takes values in the set  $\mathcal{M}_t := \prod_{\ell=0}^t \mathcal{Y}_\ell \times \prod_{\ell=0}^{t-1} \mathcal{U}_\ell$ . The agent uses the memory  $M_t$  and a control law  $g_t : \mathcal{M}_t \rightarrow \mathcal{U}_t$  at each  $t$  to generate the action  $U_t = g_t(M_t)$ . We denote the control strategy by  $\mathbf{g} := (g_0, \dots, g_T)$  and the set of all feasible control strategies by  $\mathcal{G}$ . The performance of a strategy  $\mathbf{g} \in \mathcal{G}$  is measured by the *worst-case or maximum instantaneous cost*

$$\mathcal{J}(\mathbf{g}) := \max_{t=0, \dots, T} \sup_{w_{0:t} \in [[W_{0:t}]]} C_t. \quad (6)$$

**Problem 1.** The agent's robust control problem is to derive a control strategy  $\mathbf{g}^* \in \mathcal{G}$ , such that  $\mathcal{J}(\mathbf{g}^*) \leq \mathcal{J}(\mathbf{g})$  for all  $\mathbf{g} \in \mathcal{G}$ , given the marginal ranges  $\{[[U_t]], [[W_t]], [[C_t]], [[Y_t]] \mid t = 0, \dots, T\}$  and the functions  $\{h_t, d_t \mid t = 0, \dots, T\}$ .

We seek to tractably compute an optimal strategy if one exists. In our framework, we impose the following assumptions:

**Assumption 1.** We consider that the sets  $\{\mathcal{U}_t, \mathcal{W}_t, \mathcal{Y}_t \mid t = 0, \dots, T\}$  and  $\{\mathcal{C}_t \mid t = 0, \dots, T\}$  are all bounded subsets of a metric space  $(\mathcal{S}, \eta)$  and  $\mathbb{R}_{\geq 0}$ , respectively.

Assumption 1 allows for both continuous and finite valued feasible sets while ensuring that the problem is well-posed.

**Assumption 2.** The observation functions  $\{h_t \mid t = 0, \dots, T\}$  of the system are both Lipschitz and *L-invertible*, whereas the cost functions  $\{d_t \mid t = 0, \dots, T\}$  are Lipschitz continuous.

Assumption 2 is satisfied by a large class of observation functions, including (1) all functions with compact domains and finite co-domains and (2) bi-Lipschitz functions with compact domains and compact co-domains (see Appendix A). We require both assumptions our main results.

**Remark 1.** In our exposition, we also consider a special case of (6), called the *maximum terminal cost criterion*, given by

$$\mathcal{J}^{\text{tm}}(\mathbf{g}) := \sup_{w_{0:T} \in [[W_{0:T}]]} C_T. \quad (7)$$

In addition to results for Problem 1, we often present specializations to terminal cost problems that are easier to interpret.

**Remark 2.** Problem 1 is a robust control problem with known dynamics. Our results are derived primarily within this setting. However, our main results in Section IV can also be used in learning problems where  $h_t$  and  $d_t$  are unknown for all  $t$ , as explained in Subsections IV-D and V-B.

### III. DYNAMIC PROGRAMS AND INFORMATION STATES

In this section, we first present a memory-based DP to compute the optimal solution to Problem 1. This will serve as a reference for analyzing subsequent DPs. Then, we highlight the DP's computational challenges and introduce information states in Subsections III-A and III-B to alleviate them. Subsection III-C presents examples of information states.

To arrive at the memory-based DP, we construct a "new" perfectly observed system whose state at each  $t$  is the memory  $M_t$ , evolving as  $M_{t+1} = (M_t, U_t, Y_{t+1})$ . For realizations  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$ , the maximum cost at time  $t$  is

$$\sup_{w_{0:t} \in [[W_{0:t}]]} C_t = \sup_{c_t \in [[C_t]]^{\mathbf{g}}} c_t = \sup_{m_t, u_t \in [[M_t, U_t]]^{\mathbf{g}}} \sup_{c_t \in [[C_t|m_t, u_t]]^{\mathbf{g}}} c_t,$$

for all  $t = 0, \dots, T$ , where  $[[C_t]]^g$ ,  $[[M_t, U_t]]^g$  and  $[[C_t|m_t, u_t]]^g$  are the respective marginal ranges and the conditional range induced by strategy  $g$ . Recall that  $m_t = (y_{0:t}, u_{0:t-1})$  and thus, we can expand the conditional range as

$$\begin{aligned} [[C_t|m_t, u_t]]^g &= \{c_t \in \mathcal{C}_t \mid \exists w_{0:t} \in [[W_{0:t}]] \text{ such that} \\ &c_t = d_t(w_{0:t}, u_{0:t}), y_0 = h_0(w_0), y_\ell = h_\ell(w_{0:\ell}, u_{0:\ell-1}), \\ &\forall \ell = 1, \dots, t\} = [[C_t|m_t, u_t]], \end{aligned} \quad (8)$$

where, the set of feasible disturbances  $[[W_{0:t}]]$ , cost function  $d_t(\cdot)$ , and observation functions  $h_\ell(\cdot)$  for all  $\ell = 0, \dots, t$  are independent of the choice of strategy  $g$ . Thus,  $[[C_t|m_t, u_t]]^g$  is independent of the choice of strategy  $g$  and we can drop  $g$ . Next, we define  $e_t(m_t, u_t) := \sup_{c_t \in [[C_t|m_t, u_t]]} c_t$ , independent of  $g$ , and state that

$$\begin{aligned} \sup_{m_t, u_t \in [[M_t, U_t]]^g} \sup_{c_t \in [[C_t|m_t, u_t]]^g} c_t &= \sup_{m_t, u_t \in [[M_t, U_t]]^g} e_t(m_t, u_t) \\ &= \sup_{w_{0:t} \in [[W_{0:t}]]} e_t(M_t, U_t). \end{aligned} \quad (9)$$

Since  $e_t(M_t, U_t)$  is a function of the new state  $M_t$  and control action  $U_t$ , it serves as an incurred cost at each  $t = 0, \dots, T$  in our new perfectly observed system [47]. The new instantaneous performance criterion is  $\mathcal{E}(g) := \max_{t=0, \dots, T} \sup_{w_{0:t} \in [[W_{0:t}]]} e_t(M_t, U_t)$  and from (9),  $\mathcal{E}(g) = \mathcal{J}(g)$  for any  $g$ . Subsequently, any optimal strategy in the new system is optimal for Problem 1. If such an optimal strategy exists, we can compute it using a standard DP for perfectly observed systems, as follows. For all  $t$ , for each  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$  we recursively define the *value functions*

$$Q_t(m_t, u_t) := \max \left\{ \sup_{c_t \in [[C_t|m_t, u_t]]} c_t, \sup_{m_{t+1} \in [[M_{t+1}|m_t, u_t]]} V_{t+1}(m_{t+1}) \right\}, \quad (10)$$

$$V_t(m_t) := \inf_{u_t \in [[U_t]]} Q_t(m_t, u_t), \quad (11)$$

where  $V_{T+1}(m_{T+1}) := 0$ , identically. We define the extra value function  $V_{T+1}$  to ensure that the right-hand side (RHS) of (10) is well defined at time  $T$ . Then, we can show using standard arguments [42], [47] that the optimal value of Problem 1 is  $\inf_{g \in \mathcal{G}} \mathcal{J}(g) = \sup_{m_0 \in [[M_0]]} V_0(m_0)$ . Furthermore, if there exists an action  $u_t^* \in [[U_t]]$  which achieves the infimum in the RHS of (11), then  $g_t^*(m_t) := \arg \min_{u_t \in [[U_t]]} Q_t(m_t, u_t)$  gives an optimal control law at each  $t = 0, \dots, T$  and the control strategy  $g^* = g_{0:T}^*$  is optimal for Problem 1.

**Remark 3.** The DP (10) - (11) can be specialized to the terminal cost criterion (7) by defining for all  $t = 0, \dots, T-1$ ,

$$Q_t^{\text{tm}}(m_t, u_t) := \sup_{m_{t+1} \in [[M_{t+1}|m_t, u_t]]} V_{t+1}^{\text{tm}}(m_{t+1}), \quad (12)$$

$$V_t^{\text{tm}}(m_t) := \inf_{u_t \in [[U_t]]} Q_t^{\text{tm}}(m_t, u_t), \quad (13)$$

where  $Q_T^{\text{tm}}(m_T, u_T) := \sup_{c_T \in [[C_T|m_T, u_T]]} c_T$  and  $V_T^{\text{tm}}(m_T) := \inf_{u_T \in [[U_T]]} Q_T^{\text{tm}}(m_T, u_T)$ . We will use this terminal cost DP to simplify the exposition in Section IV.

**Remark 4.** A valid argument attaining the minimum in the RHS of (11) at each  $t = 0, \dots, T$  is both a necessary and

sufficient condition to ensure the existence of an optimal control strategy in Problem 1 [42], [82]. This can be satisfied if the marginal ranges of all uncertain variables are *compact* rather than just bounded. Recall from Assumption 2 that the observation and cost functions at each  $t$  are Lipschitz. Then, using (12) - (13), we can show that the value functions are continuous and the conditional ranges are compact for all  $t$ , implying that the minimum is achieved in (11). However, for generality, we continue using  $\sup$  and  $\inf$  in our exposition.

**Remark 5.** In the RHS of (11) at each  $t$ , we are required to solve an optimization for each  $m_t \in [[M_t]]$ . This is computationally challenging for longer horizons as the size of the set  $[[M_t]]$  increases with time  $t$ . This concern necessitates an alternate DP decomposition that can derive an optimal control strategy while potentially achieving more favorable computational properties. We present such a DP decomposition in Subsection III-A by identifying an uncertain variable, called an *information state*, which can be used to generate an optimal control action at each time step instead of the memory.

### A. Information States

In this subsection, we define the notion of information states, use them in a DP decomposition, and prove it is optimal.

**Definition 1.** An *information state* at each  $t = 0, \dots, T$  is an uncertain variable  $\Pi_t = \sigma_t(M_t)$  taking values in a bounded set  $\mathcal{P}_t$ , generated by a function  $\sigma_t : \mathcal{M}_t \rightarrow \mathcal{P}_t$ , and for all  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$ , satisfying the following properties:

1) *Sufficient to evaluate cost*: An information state can replace the memory to compute the worst-case cost, i.e.,

$$\sup_{c_t \in [[C_t|m_t, u_t]]} c_t = \sup_{c_t \in [[C_t|\sigma_t(m_t), u_t]]} c_t. \quad (14)$$

2) *Sufficient to predict its own evolution*: An information state can replace the memory to compute its own conditional range at the next time step, independent of the strategy  $g$ , i.e.,

$$[[\Pi_{t+1}|m_t, u_t]] = [[\Pi_{t+1}|\sigma_t(m_t), u_t]]. \quad (15)$$

We can use the information states from Definition 1 directly in a DP, as follows. For all  $t = 0, \dots, T$ , for all  $\pi_t \in [[\Pi_t]]$  and  $u_t \in [[U_t]]$ , we recursively define the value functions

$$\bar{Q}_t(\pi_t, u_t) := \max \left\{ \sup_{c_t \in [[C_t|\pi_t, u_t]]} c_t, \sup_{\pi_{t+1} \in [[\Pi_{t+1}|\pi_t, u_t]]} \bar{V}_{t+1}(\pi_{t+1}) \right\}, \quad (16)$$

$$\bar{V}_t(\pi_t) := \inf_{u_t \in [[U_t]]} \bar{Q}_t(\pi_t, u_t), \quad (17)$$

where  $\bar{V}_{T+1}(\pi_{T+1}) := 0$  identically. If the minimum in the RHS of (17) exists, this DP yields a control law as  $\bar{g}_t^*(\pi_t) := \arg \min_{u_t \in [[U_t]]} \bar{Q}_t(\pi_t, u_t)$  at each  $t$ . Next, we prove that this DP computes the same value as the optimal DP (10) - (11).

**Theorem 1.** Let  $\Pi_t = \sigma_t(M_t)$  be an information state at any  $t$ . Then, for all  $t$ , and for all  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$ ,

$$Q_t(m_t, u_t) = \bar{Q}_t(\sigma_t(m_t), u_t) \text{ and } V_t(m_t) = \bar{V}_t(\sigma_t(m_t)). \quad (18)$$

*Proof.* Let  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$  be given realizations of  $M_t$  and  $U_t$ , respectively, for all  $t = 0, \dots, T$ . We

prove the result by mathematical induction starting at the last time step. At time  $T + 1$ , (18) holds trivially because  $V_{T+1}(m_{T+1}) = \bar{V}_{T+1}(\sigma_{T+1}(m_{T+1})) = 0$ . This forms the basis of our induction. Next, for any  $t = 0, \dots, T$ , we consider the induction hypothesis that  $V_{t+1}(m_{t+1}) = \bar{V}_{t+1}(\sigma_{t+1}(m_{t+1}))$ . Given the hypothesis, we first prove that  $Q_t(m_t, u_t) = \bar{Q}_t(\sigma_t(m_t), u_t)$  by comparing the RHS of (10) to the RHS of (16) term by term. The first terms are equal by direct application of (14) from Definition 1. Next, we use the induction hypothesis for the second term in the RHS of (10), to state that  $\sup_{m_{t+1} \in [[M_{t+1}|m_t, u_t]]} V_{t+1}(m_{t+1}) = \sup_{m_{t+1} \in [[M_{t+1}|m_t, u_t]]} \bar{V}_{t+1}(\sigma_{t+1}(m_{t+1})) = \sup_{\sigma_{t+1}(m_{t+1}) \in [[\Pi_{t+1}|\sigma_t(m_t), u_t]]} \bar{V}_{t+1}(\sigma_{t+1}(m_{t+1}))$ , where, in the second equality, we use the fact that  $[[\Pi_{t+1}|m_t, u_t]] = \{\sigma_{t+1}(m_{t+1}) \in \mathcal{P}_{t+1} | m_{t+1} \in [[M_{t+1}|m_t, u_t]]\}$  and (15) from Definition 1. This establishes that the second term in the RHS of (10) equals the second term in the RHS of (16) and given the induction hypothesis for time  $t + 1$ , we have  $Q_t(m_t, u_t) = \bar{Q}_t(\sigma_t(m_t), u_t)$ . Next, we minimize both sides of the equality with respect to  $u_t \in [[U_t]]$ , and use the definitions in (11) and (17) to write that  $V_t(m_t) = \inf_{u_t \in \mathcal{U}_t} Q_t(m_t, u_t) = \inf_{u_t \in \mathcal{U}_t} \bar{Q}_t(\sigma_t(m_t), u_t) = V_t(\sigma_t(m_t))$ , which proves the hypothesis at time  $t$ . The result follows for all  $t$  using mathematical induction.  $\square$

Theorem 1 implies that (16) - (17) is an optimal DP decomposition for Problem 1, i.e., if an optimal strategy exists for this DP, it is an optimal solution to Problem 1 as follows. Consider a control strategy  $\bar{g}^* = \bar{g}_{0:T}^*$  computed using (16) - (17). A corresponding memory-based strategy  $\mathbf{g} = g_{0:T}$  contains  $g_t(m_t) := \bar{g}_t^*(\sigma_t(m_t))$  for all  $m_t \in [[M_t]]$  and  $t = 0, \dots, T$ . Then, from Theorem 1,  $\mathbf{g}$  achieves the infimum value at each  $t$  to form an optimal solution to Problem 1.

**Remark 6.** In practice, using an information state to construct the DP decomposition is useful computationally only if, for most time steps in  $t = 0, \dots, T$ , either the value functions in (16) - (17) have useful properties like concavity, or the set  $\mathcal{P}_t$  is *smaller* than  $\mathcal{M}_t$  for some measure of size, such as the cardinality, diameter, and dimension. We present some examples of information states for different systems in Subsection III-C.

### B. Alternate Characterization of Information States

When exploring whether an uncertain variable is a valid candidate for an information state, it may be difficult to verify (15) in Definition 1. In this subsection, we present two *stronger* conditions to replace (15). Specifically, at each  $t = 0, \dots, T$ , to establish that  $\Pi_t = \sigma_t(M_t)$  is a valid information state, it is sufficient to satisfy the following conditions instead of (15):

1) *State-like evolution:* There exists a function  $\bar{f}_t : \mathcal{P}_t \times \mathcal{U}_t \times \mathcal{Y}_{t+1} \rightarrow \mathcal{P}_{t+1}$ , independent of the strategy  $\mathbf{g}$ , such that

$$\Pi_{t+1} = \bar{f}_t(\Pi_t, U_t, Y_{t+1}). \quad (19)$$

2) *Predicts observations:* For all  $m_t \in \mathcal{M}_t$  and  $u_t \in \mathcal{U}_t$ ,

$$[[Y_{t+1}|m_t, u_t]] = [[Y_{t+1}|\sigma_t(m_t), u_t]], \quad (20)$$

where, both conditional ranges in (20) can be evaluated independently of the choice of strategy  $\mathbf{g}$ .

Next, we prove that these two conditions, in addition to (14), are sufficient to identify an information state.

**Lemma 1.** For all  $t = 0, \dots, T$ , if an uncertain variable  $\Pi_t = \sigma_t(M_t)$  satisfies (19) - (20), it also satisfies (15).

*Proof.* For all  $t = 0, \dots, T$  and  $m_t \in \mathcal{M}_t$ , suppose that  $\pi_t = \sigma_t(m_t)$  satisfy (19) - (20). Then, we substitute (19) into the left hand side (LHS) of (15) to state that

$$\begin{aligned} [[\Pi_{t+1}|m_t, u_t]] &= [[\bar{f}_t(\sigma_t(m_t), u_t, Y_{t+1}) | m_t, u_t]] \\ &= \{\bar{f}_t(\sigma_t(m_t), u_t, y_{t+1}) | y_{t+1} \in [[Y_{t+1}|m_t, u_t]]\}, \end{aligned} \quad (21)$$

where, in the second equality, we write the conditional range as a set. Next, using (20) on the range of observations in the conditioning of (21), we can state that  $\{\bar{f}_t(\sigma_t(m_t), u_t, y_{t+1}) | y_{t+1} \in [[Y_{t+1}|m_t, u_t]]\} = \{\bar{f}_t(\sigma_t(m_t), u_t, y_{t+1}) | y_{t+1} \in [[Y_{t+1}|\sigma_t(m_t), u_t]]\} = [[\bar{f}_t(\sigma_t(m_t), u_t, Y_{t+1}) | \sigma_t(m_t), u_t]] = [[\Pi_{t+1}|\sigma_t(m_t), u_t]]$ , which is equal to the RHS of (15).  $\square$

### C. Examples of Information States

In this subsection, we present examples of information states for systems with given state-space models.

1) *Systems with perfectly observed states:* At each  $t = 0, \dots, T$ , let the system's state be denoted by the uncertain variable  $X_t \in \mathcal{X}_t$ , where  $\mathcal{X}_t$  is a known state space. The agent's observation is  $Y_t = X_t$  and the agent incurs a cost  $C_t = d_t(X_t, U_t)$  when they implement an action  $U_t \in \mathcal{U}_t$ . Starting at  $X_0 \in \mathcal{X}_0$ , the state evolves as  $X_{t+1} = f_t(X_t, U_t, W_t)$  for all  $t$ . Each uncertain variable in  $\{X_0, W_t | t = 0, \dots, T\}$  is independent of all other uncertain variables in that set. Then, an information state at each  $t$  is  $\Pi_t = X_t$ , i.e., the state itself [42]. It takes values in the set  $\mathcal{X}_t$  and satisfies (14) - (15) for all  $t$ . Note that it is always computationally advantageous to construct a DP decomposition using the state instead of the entire memory.

2) *Systems with delayed observations:* Consider the dynamics in Case 1. However, at each  $t = 0, \dots, T$ , the agent observes the state with a delay of  $n \in \mathbb{N}$  time steps, i.e.,  $Y_t = X_{t-n}$ . Then, an information state at each  $t$  is  $\Pi_t = (X_{t-n}, U_{t-n:t-1})$  that takes values in the set  $\mathcal{X}_t \times \prod_{\ell=n}^{t-1} \mathcal{U}_\ell$ . It can be verified that this information state satisfies (14) - (15) for all  $t$  and that it is computationally advantageous to utilize it instead of the complete memory of the agent.

3) *Systems with partially observed states:* Consider the dynamics in Case 1. However, at each  $t = 0, \dots, T$ , the agent's observation is  $Y_t = h_t(X_t, N_t)$ , where  $N_t \in \mathcal{N}_t$  is a noise in observation and each uncertain variable in  $\{X_0, W_t, N_t | t = 0, \dots, T\}$  is independent of all other uncertain variables in that set. Then, an information state at each  $t = 0, \dots, T$  is the conditional range  $\Pi_t = [[X_t|M_t]]$  [44], [47]. This is a set-valued uncertain variable whose realization at time  $t$  is a subset of the state space containing values of  $x_t$  consistent with a given realization of the memory  $m_t$ . Explicitly, for a given realization of the memory  $m_t \in \mathcal{M}_t$  at time  $t$ , the conditional range takes the realization  $P_t := \{x_t \in \mathcal{X}_t | \exists x_0 \in \mathcal{X}_0, w_{0:t-1} \in \prod_{\ell=0}^{t-1} \mathcal{W}_\ell, n_{0:t} \in \prod_{\ell=0}^t \mathcal{N}_\ell \text{ such that } y_t = h_t(x_t, n_t), x_{\ell+1} = f_\ell(x_\ell, u_\ell, w_\ell), y_\ell = h_\ell(x_\ell, n_\ell) \text{ for all } \ell = 1, \dots, t-1\}$ .

$0, \dots, t-1\}$ . To establish that the conditional range is a valid information state, it is easier to verify the alternate conditions (19) and (20) instead of the property (15) in Definition 1. Generally, it is computationally advantageous to construct a DP decomposition using the conditional range instead of the memory for systems with longer time horizons.

**Remark 7.** For systems with known dynamics, any uncertain variable can be proposed as an information state using the conditions in Definition 1 to simplify the DP decomposition. However, systems with large state spaces may require further improvement in tractability, even at the cost of optimality. Moreover, in certain applications, we need to learn a representation of the information state using incomplete knowledge of the dynamics, unable to satisfy the conditions exactly. In Section IV, we introduce approximate information states that can address the above concerns.

#### IV. APPROXIMATE INFORMATION STATES

In this section, we define approximate information states by relaxing the conditions in Definition 1, and utilize them to develop an approximate DP decomposition which computes a sub-optimal control strategy for Problem 1. In Subsection IV-A, we prove Lipschitz continuity of approximate value functions and guarantees on the worst-case performance of resulting approximate control strategies. Next, we present an alternate characterization of approximate information states in Subsection IV-B. Finally, in Subsection IV-C, we present examples of these performance bounds when approximate information states are constructed using state quantization.

**Definition 2.** An *approximate information state* for Problem 1 at each  $t = 0, \dots, T$  is an uncertain variable  $\hat{\Pi}_t = \hat{\sigma}_t(M_t)$  taking values in a bounded set  $\hat{\mathcal{P}}_t$  and generated by an  $L$ -invertible function  $\hat{\sigma}_t : \mathcal{M}_t \rightarrow \hat{\mathcal{P}}_t$ . Furthermore, for all  $t = 0, \dots, T$ , there exist parameters  $\epsilon_t, \delta_t, \lambda_t \in \mathbb{R}_{\geq 0}$  such that for all  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$ , it satisfies the properties:

1) *Sufficient to approximately evaluate worst-case cost:*

$$\left| \sup_{c_t \in [[C_t | m_t, u_t]]} c_t - \sup_{c_t \in [[C_t | \hat{\sigma}_t(m_t), u_t]]} c_t \right| \leq \epsilon_t. \quad (22)$$

2) *Sufficient to approximate its own evolution:* We define the sets  $\mathcal{K}_{t+1} := [[\hat{\Pi}_{t+1} | m_t, u_t]]$  and  $\hat{\mathcal{K}}_{t+1} := [[\hat{\Pi}_{t+1} | \hat{\sigma}_t(m_t), u_t]]$ , each independent of strategy  $\mathbf{g}$ . Then

$$\mathcal{H}(\mathcal{K}_{t+1}, \hat{\mathcal{K}}_{t+1}) \leq \delta_t, \quad (23)$$

where recall that  $\mathcal{H}$  is the Hausdorff distance in (1).

3) *Lipschitz-like evolution:* For all  $\hat{\pi}_t^1, \hat{\pi}_t^2 \in [[\hat{\Pi}_t]]$ ,

$$\mathcal{H}([[\hat{\Pi}_{t+1} | \hat{\pi}_t^1, u_t]], [[\hat{\Pi}_{t+1} | \hat{\pi}_t^2, u_t]]) \leq \lambda_t \cdot \eta(\hat{\pi}_t^1, \hat{\pi}_t^2), \quad (24)$$

where  $\eta$  is an appropriate metric on  $\hat{\mathcal{P}}_t$ .

Using the approximate information state in Definition 2, we can construct a DP as follows. For all  $t$ , for all  $\hat{\pi}_t \in [[\hat{\Pi}_t]]$  and  $u_t \in [[U_t]]$ , we recursively define the value functions

$$\hat{Q}_t(\hat{\pi}_t, u_t) := \max \left\{ \begin{array}{l} \sup_{c_t \in [[C_t | \hat{\pi}_t, u_t]]} c_t, \\ \sup_{\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1} | \hat{\pi}_t, u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1}) \end{array} \right\}, \quad (25)$$

$$\hat{V}_t(\hat{\pi}_t) := \inf_{u_t \in [[U_t]]} \hat{Q}_t(\hat{\pi}_t, u_t), \quad (26)$$

where  $\hat{V}_{T+1}(\hat{\pi}_{T+1}) := 0$  identically. If there exists a minimizing argument in the RHS of (26) at each  $t = 0, \dots, T$ , then  $\hat{g}_t^*(\hat{\pi}_t) := \arg \min_{u_t \in \mathcal{U}_t} \hat{Q}_t(\hat{\pi}_t, u_t)$  constitutes an approximate control law at time  $t$ . Furthermore, we call  $\hat{g}^* = \hat{g}_{0:T}^*$  an approximately optimal strategy for Problem 1. In Subsection IV-A, we derive approximation bounds for the same.

**Remark 8.** As we showed in Section III, we can specialize this DP for terminal cost problems, with the value functions for all  $t = 0, \dots, T-1$  given by

$$\hat{Q}_t^{\text{tm}}(\hat{\pi}_t, u_t) := \sup_{\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1} | \hat{\pi}_t, u_t]]} \hat{V}_{t+1}^{\text{tm}}(\hat{\pi}_{t+1}), \quad (27)$$

$$\hat{V}_t^{\text{tm}}(\hat{\pi}_t) := \inf_{u_t \in \mathcal{U}_t} \hat{Q}_t^{\text{tm}}(\hat{\pi}_t, u_t), \quad (28)$$

and  $\hat{Q}_T^{\text{tm}}(\hat{\pi}_T, u_T) := \sup_{c_T \in [[C_T | \hat{\pi}_T, u_T]]} c_T$  and  $\hat{V}_T^{\text{tm}}(\hat{\pi}_T) := \inf_{u_T \in \mathcal{U}_T} \hat{Q}_T^{\text{tm}}(\hat{\pi}_T, u_T)$  at time  $T$ .

**Remark 9.** The conditions in Definition 2 can be investigated using only output variables. Thus, using them to construct a training loss, an approximate information state model can be learned from output data without knowledge of dynamics. This is illustrated in Subsections IV-D and V-B.

**Remark 10.** The notion of an approximate information state is related to approximate bisimulations [80]. Recall that in Section III, we constructed a memory-based system that yielded an optimal solution to Problem 1. Then, with the worst-case cost as each system's output, (14) - (15) in Definition 2 define an approximate bisimulation between an approximate system with state  $\hat{\Pi}_t$  and the memory-based system with state  $M_t$  at each  $t$ . Together, (14) - (15) can be used to bound the branching distance between the two systems across  $t = 0, \dots, T$  [80, Subsection VI-B]. Instead, our subsequent results focus on the worst-case performance of approximate control strategies to facilitate approximately robust control.

#### A. Properties of Approximate Information States

In this subsection, we present several properties of the approximate DP (25) - (26) using preliminary results detailed in Appendix B. To begin, we prove in Theorem 2 that each approximate value function is Lipschitz continuous. This property subsequently allows us to establish error bounds.

**Theorem 2.** In the approximate DP (25) - (26), the value functions  $\hat{Q}_t(\hat{\pi}_t, u_t)$  and  $\hat{V}_t(\hat{\pi}_t)$  are Lipschitz continuous with respect to  $\hat{\pi}_t \in [[\hat{\Pi}_t]]$  for all  $u_t \in [[U_t]]$  and  $t = 0, \dots, T$ .

*Proof.* We prove the Lipschitz continuity of the value functions by constructing a valid candidate for the Lipschitz constant  $L_{\hat{V}_t}$  at each  $t = 0, \dots, T$ , using mathematical induction. At time  $T+1$ , recall that  $\hat{V}_{T+1}(\hat{\pi}_{T+1}) = 0$  identically and thus,  $\hat{V}_{T+1}(\hat{\pi}_{T+1})$  is trivially Lipschitz continuous with a constant  $L_{\hat{V}_{T+1}} = 0$ . This forms the basis of our induction. Then, at each  $t = 0, \dots, T$ , we consider the induction hypothesis that  $\hat{Q}_{t+1}(\hat{\pi}_{t+1}, u_{t+1})$  and  $\hat{V}_{t+1}(\hat{\pi}_{t+1})$  are Lipschitz continuous with respect to  $\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1}]]$  for all  $u_{t+1} \in [[U_{t+1}]]$ , and denote the constant by  $L_{\hat{V}_{t+1}} \in \mathbb{R}_{\geq 0}$ .

At time  $t$ , we first prove the result for the value function  $\hat{Q}_t(\hat{\pi}_t, u_t)$ . Let  $\hat{\pi}_t^1, \hat{\pi}_t^2 \in [[\hat{\Pi}_t]]$  be two possible realizations of  $\hat{\Pi}_t$ . Then, using the definition (25) of  $\hat{Q}_t(\hat{\pi}_t, u_t)$  and (61) from Lemma 8, we state that

$$|\hat{Q}_t(\hat{\pi}_t^1, u_t) - \hat{Q}_t(\hat{\pi}_t^2, u_t)| \leq \max \left\{ \left| \sup_{c_t^1 \in [[C_t | \hat{\pi}_t^1, u_t]]} c_t^1 \right. \right. \\ \left. \left. - \sup_{c_t^2 \in [[C_t | \hat{\pi}_t^2, u_t]]} c_t^2 \right|, \left| \sup_{\hat{\pi}_{t+1}^1 \in [[\hat{\Pi}_{t+1} | \hat{\pi}_t^1, u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1}^1) \right. \right. \\ \left. \left. - \sup_{\hat{\pi}_{t+1}^2 \in [[\hat{\Pi}_{t+1} | \hat{\pi}_t^2, u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1}^2) \right| \right\}. \quad (29)$$

We consider the RHS of (29) term by term. In the first term, we note that for all  $\hat{\pi}_t \in [[\hat{\Pi}_t]]$ ,

$$\sup_{c_t \in [[C_t | \hat{\pi}_t, u_t]]} c_t = \sup_{m_t \in [[M_t | \hat{\pi}_t]]} \left( \sup_{c_t \in [[C_t | m_t, u_t]]} c_t \right). \quad (30)$$

In the RHS of (30), recall from Assumption 2 that the uncertain variable  $C_t$  is a Lipschitz function of  $(W_{0:t}, U_{0:t})$ , and  $(M_t, U_t)$  is an  $L$ -invertible function of  $(W_{0:t}, U_{0:t})$ . Thus, using (57) from Lemma 6, there exists a constant  $L_{C|M,U}$  such that  $\mathcal{H}([[C_t | m_t^1, u_t]], [[C_t | m_t^2, u_t]]) \leq L_{M|C,U} \cdot \eta(m_t^1, m_t^2)$  for all  $m^1, m^2 \in [[M_t]]$ . Furthermore, we use (63) to state that

$$\left| \sup_{c_t^1 \in [[C_t | m_t^1, u_t]]} c_t^1 - \sup_{c_t^2 \in [[C_t | m_t^2, u_t]]} c_t^2 \right| \\ \leq L_{M|C,U} \cdot L_{c_t} \cdot \eta(m_t^1, m_t^2). \quad (31)$$

Then, consider a function  $e_t : \mathcal{M}_t \times \mathcal{U}_t \rightarrow \mathbb{R}_{\geq 0}$  defined as  $e_t(m_t, u_t) := \sup_{c_t \in [[C_t | m_t, u_t]]} c_t$ . As a direct consequence of (31),  $e_t$  is Lipschitz continuous with respect to  $m_t$  with a constant  $L_{e_t} := L_{M|C,U} \cdot L_{c_t}$ . Using (30) and the definition of  $e_t$  in the first term in the RHS of (29),

$$\left| \sup_{c_t^1 \in [[C_t | \hat{\pi}_t^1, u_t]]} c_t^1 - \sup_{c_t^2 \in [[C_t | \hat{\pi}_t^2, u_t]]} c_t^2 \right| \\ = \left| \sup_{m_t^1 \in [[M_t | \hat{\pi}_t^1]]} e_t(m_t^1, u_t) - \sup_{m_t^2 \in [[M_t | \hat{\pi}_t^2]]} e_t(m_t^2, u_t) \right|. \quad (32)$$

In (32), recall that the uncertain variable  $\hat{\Pi}_t$  is an  $L$ -invertible function of  $M_t$  and thus, the conditional range  $[[M_t | \hat{\pi}_t]]$  satisfies (3). Then, we use (63) once more to state that

$$\left| \sup_{c_t^1 \in [[C_t | \hat{\pi}_t^1, u_t]]} c_t^1 - \sup_{c_t^2 \in [[C_t | \hat{\pi}_t^2, u_t]]} c_t^2 \right| \leq L_{\hat{\sigma}_t^{-1}} \cdot L_{e_t} \cdot \eta(\hat{\pi}_t^1, \hat{\pi}_t^2). \quad (33)$$

In the second term in the RHS of (29), we use the induction hypothesis and (62) from Lemma 9 to write that

$$\left| \sup_{\hat{\pi}_{t+1}^1 \in [[\hat{\Pi}_{t+1} | \hat{\pi}_t^1, u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1}^1) - \sup_{\hat{\pi}_{t+1}^2 \in [[\hat{\Pi}_{t+1} | \hat{\pi}_t^2, u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1}^2) \right| \\ \leq L_{\hat{V}_{t+1}} \cdot \mathcal{H}([[\hat{\Pi}_{t+1} | \hat{\pi}_t^1, u_t]], [[\hat{\Pi}_{t+1} | \hat{\pi}_t^2, u_t]]) \\ \leq L_{\hat{V}_{t+1}} \cdot \lambda_t \cdot \eta(\hat{\pi}_t^1, \hat{\pi}_t^2), \quad (34)$$

where, in the second inequality, we use the third property (24) of approximate information states in Definition 2. Then, the proof for  $\hat{Q}_t(\hat{\pi}_t, u_t)$  is complete by substituting (33) and (34) into the RHS of (29) and defining

$$L_{\hat{Q}_t} := \max \{L_{\hat{\sigma}_t^{-1}} \cdot L_{e_t}, L_{\hat{V}_{t+1}} \cdot \lambda_t\}. \quad (35)$$

To prove the result for  $\hat{V}_t(\hat{\pi}_t)$ , we use (60) from Lemma 7 to state that  $|\hat{V}_t(\hat{\pi}_t^1) - \hat{V}_t(\hat{\pi}_t^2)| = |\inf_{u_t \in [[U_t]]} \hat{Q}_t(\hat{\pi}_t^1, u_t) - \inf_{u_t \in [[U_t]]} \hat{Q}_t(\hat{\pi}_t^2, u_t)| \leq \sup_{u_t \in [[U_t]]} |\hat{Q}_t(\hat{\pi}_t^1, u_t) - \hat{Q}_t(\hat{\pi}_t^2, u_t)| \leq L_{\hat{Q}_t} \cdot \eta(\hat{\pi}_t^1, \hat{\pi}_t^2)$ , which proves the induction hypothesis at time  $t$ . Thus, the result holds using mathematical induction.  $\square$

**Remark 11.** From (35), we can construct  $L_{\hat{Q}_t}$  and  $L_{\hat{V}_t}$  at any  $t$  as functions of the Lipschitz constants of the underlying dynamics, the  $L$ -invertible constants  $(L_{\hat{\sigma}_t^{-1}}, \dots, L_{\hat{\sigma}_t^{-1}})$ , and  $(\lambda_t, \dots, \lambda_T)$  from (24). The Lipschitz constant  $L_{\hat{V}_t}$  is not an explicit function of the bounds  $\epsilon_t$  in (22) and  $\delta_t$  in (23). Furthermore, a smaller value of  $L_{\hat{V}_t}$  can be found in special cases, such as problems where  $\hat{\mathcal{P}}_t$  and  $\mathcal{U}_t$  are finite for all  $t$ .

Next, we establish an upper bound on the approximation error when the value functions of the optimal DP (10) - (11) are estimated using the approximate DP (25) - (26) at each  $t$ .

**Theorem 3.** Let  $L_{\hat{V}_{t+1}}$  be the Lipschitz constant of  $\hat{V}_{t+1}$  for all  $t = 0, \dots, T$ . Then, for all  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$ ,

$$|Q_t(m_t, u_t) - \hat{Q}_t(\hat{\sigma}_t(m_t), u_t)| \leq \alpha_t, \quad (36)$$

$$|V_t(m_t) - \hat{V}_t(\hat{\sigma}_t(m_t))| \leq \alpha_t, \quad (37)$$

where  $\alpha_t = \max(\epsilon_t, \alpha_{t+1} + L_{\hat{V}_{t+1}} \cdot \delta_t)$  for all  $t$ , with  $\alpha_{T+1} = 0$ .

*Proof.* For all  $t = 0, \dots, T$ , let  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$  be realizations of  $M_t$  and  $U_t$ , respectively. We prove both results by mathematical induction, starting with time step  $T+1$ . At  $T+1$ , by definition,  $V_{T+1}(m_{T+1}, u_{T+1}) = V_{T+1}(\hat{\sigma}_{T+1}(m_{T+1})) = 0$ . This forms the basis of our mathematical induction. Then, at each  $t = 0, \dots, T$ , we consider the induction hypothesis  $|V_{t+1}(m_{t+1}) - \hat{V}_{t+1}(\hat{\sigma}_{t+1}(m_{t+1}))| \leq \alpha_{t+1}$ . At time  $t$ , we first prove (36). Using (61) from Lemma 8 in the LHS of (36) to state that

$$|Q_t(m_t, u_t) - \hat{Q}_t(\hat{\sigma}_t(m_t), u_t)| \leq \max \left\{ \left| \sup_{c_t \in [[C_t | m_t, u_t]]} c_t \right. \right. \\ \left. \left. - \sup_{c_t \in [[C_t | \hat{\sigma}_t(m_t), u_t]]} c_t \right|, \left| \sup_{m_{t+1} \in [[M_{t+1} | m_t, u_t]]} V_{t+1}(m_{t+1}) \right. \right. \\ \left. \left. - \sup_{\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1} | \hat{\sigma}_t(m_t), u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1}) \right| \right\}. \quad (38)$$

We consider the RHS of (38) term-by-term. By direct application of (22) in Definition 2, the first term in the RHS satisfies

$$\left| \sup_{c_t \in [[C_t | m_t, u_t]]} c_t - \sup_{c_t \in [[C_t | \hat{\sigma}_t(m_t), u_t]]} c_t \right| \leq \epsilon_t. \quad (39)$$

For the second term in the RHS of (38), we use the triangle inequality to write that

$$\left| \sup_{m_{t+1} \in [[M_{t+1} | m_t, u_t]]} V_{t+1}(m_{t+1}) - \sup_{\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1} | \hat{\sigma}_t(m_t), u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1}) \right| \\ \leq \left| \sup_{m_{t+1} \in [[M_{t+1} | m_t, u_t]]} V_{t+1}(m_{t+1}) \right. \\ \left. - \sup_{\hat{\sigma}_{t+1}(m_{t+1}) \in [[\hat{\Pi}_{t+1} | m_t, u_t]]} \hat{V}_{t+1}(\hat{\sigma}_{t+1}(m_{t+1})) \right| + \\ \left| \sup_{\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1} | m_t, u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1}) - \sup_{\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1} | \hat{\sigma}_t(m_t), u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1}) \right|. \quad (40)$$

For the first term in the RHS of (40), we first note that  $\sup_{\hat{\sigma}_{t+1}(m_{t+1}) \in [[\hat{\Pi}_{t+1}|m_t, u_t]]} \hat{V}_{t+1}(\hat{\sigma}_{t+1}(m_{t+1})) = \sup_{m_{t+1} \in [[M_{t+1}|m_t, u_t]]} \hat{V}_{t+1}(\hat{\sigma}_{t+1}(m_{t+1}))$  because  $[[\hat{\Pi}_{t+1}|m_t, u_t]] = \{\hat{\sigma}_{t+1}(m_{t+1}) \in \hat{\mathcal{P}}_t \mid m_{t+1} \in [[M_{t+1}|m_t, u_t]]\}$ . Then, we can state that  $|\sup_{m_{t+1} \in [[M_{t+1}|m_t, u_t]]} V_{t+1}(m_{t+1}) - \sup_{\hat{\sigma}_{t+1}(m_{t+1}) \in [[\hat{\Pi}_{t+1}|m_t, u_t]]} \hat{V}_{t+1}(\hat{\sigma}_{t+1}(m_{t+1}))| \leq \sup_{m_{t+1} \in [[M_{t+1}|m_t, u_t]]} |V_{t+1}(m_{t+1}) - \hat{V}_{t+1}(\hat{\sigma}_{t+1}(m_{t+1}))| \leq \alpha_{t+1}$ , where, in the first inequality, we use (59) from Lemma 7; and, in the second inequality, we use the induction hypothesis for time  $t+1$ . The second term in the RHS of (40) satisfies  $|\sup_{\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1}|m_t, u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1}) - \sup_{\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1}|\hat{\sigma}_t(m_t), u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1})| \leq L_{\hat{V}_{t+1}} \cdot \delta_t$  using (62) from Lemma 9 and (23) from Definition 2. Substituting the respective inequalities for each term in the RHS of (40) yields  $|\sup_{m_{t+1} \in [[M_{t+1}|m_t, u_t]]} V_{t+1}(m_{t+1}) - \sup_{\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1}|\hat{\sigma}_t(m_t), u_t]]} \hat{V}_{t+1}(\hat{\pi}_{t+1})| \leq \alpha_{t+1} + L_{\hat{V}_{t+1}} \cdot \delta_t$ . We complete the proof for (36) by substituting the inequalities in the RHS of (39) and (40) into the RHS of (38). Next, we prove (37) at time  $t$ . Using the definition of the value functions in the LHS of (37), we write that

$$|V_t(m_t) - \hat{V}_t(\hat{\sigma}_t(m_t))| = \left| \inf_{u_t \in [[U_t]]} Q_t(m_t, u_t) - \inf_{u_t \in [[U_t]]} \hat{Q}_t(\hat{\sigma}_t(m_t), u_t) \right| \leq \sup_{u_t \in [[U_t]]} |Q_t(m_t, u_t) - \hat{Q}_t(\hat{\sigma}_t(m_t), u_t)| \leq \max\{\epsilon_t, \alpha_{t+1} + L_{\hat{V}_{t+1}} \cdot \delta_t\}, \quad (41)$$

where in the first inequality, we use (60) from Lemma 7; and in the second inequality, we use (36). Thus, the results hold for all  $t = 0, \dots, T$  using mathematical induction.  $\square$

After bounding the approximation error for value functions, we also seek to bound the maximum performance loss in the implementation of an approximately optimal strategy. Consider an approximate strategy  $\hat{g}^* := \hat{g}_{0:T}^*$  computed using (25) - (26), where  $\hat{g}_t^*(\hat{\pi}_t) = \arg \min_{u_t \in [[U_t]]} \hat{Q}_t(\hat{\pi}_t, u_t)$  for all  $t$ . We can construct an approximate memory-based strategy  $g^{\text{ap}} = g_{0:T}^{\text{ap}}$  by selecting the control law  $g_t^{\text{ap}}(m_t) := \hat{g}_t^*(\hat{\sigma}_t(m_t))$  for all  $t$ . Note that  $g^{\text{ap}}$  is equivalent to  $\hat{g}^*$  because they generate the same actions at each  $t$  and subsequently, yield the same performance. Thus, we evaluate the performance of  $g^{\text{ap}}$  to determine the quality of approximation. To this end, for all  $t = 0, \dots, T$ , for all  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$ , we define

$$\Theta_t(m_t, u_t) := \max \left\{ \sup_{c_t \in [[C_t|m_t, u_t]]} c_t, \sup_{m_{t+1} \in [[M_{t+1}|m_t, u_t]]} \Lambda_{t+1}(m_{t+1}) \right\}, \quad (42)$$

$$\Lambda_t(m_t) := \Theta_t(m_t, g_t^{\text{ap}}(m_t)), \quad (43)$$

where  $\Lambda_{T+1}(m_{T+1}) := 0$ , identically. Then, the performance of the memory-based approximate strategy  $g^{\text{ap}}$  is  $\Lambda_0(m_0)$ . In contrast, recall that the performance of an optimal strategy  $g^*$  is the optimal value  $V_0(m_0)$  computed using (10) - (11). Next, we bound the difference in performance between  $g^{\text{ap}}$  and  $g^*$ .

**Theorem 4.** Let  $L_{\hat{V}_{t+1}}$  be the Lipschitz constant of  $\hat{V}_{t+1}$  for all  $t = 0, \dots, T$ . Then, for all  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$ ,

$$|Q_t(m_t, u_t) - \Theta_t(m_t, u_t)| \leq 2\alpha_t, \quad (44)$$

$$|V_t(m_t) - \Lambda_t(m_t)| \leq 2\alpha_t. \quad (45)$$

where  $\alpha_t = \max(\epsilon_t, \alpha_{t+1} + L_{\hat{V}_{t+1}} \cdot \delta_t)$  for all  $t$ , with  $\alpha_{T+1} = 0$ .

*Proof.* We begin by recursively defining the value functions that compute the performance of the strategy  $\hat{g}$ . For all  $t = 0, \dots, T$  and for each  $\hat{\pi}_t \in [[\hat{\Pi}_t]]$  and  $u_t \in [[U_t]]$ , let

$$\hat{\Theta}_t(\hat{\pi}_t, u_t) := \max \left\{ \sup_{c_t \in [[C_t|\hat{\pi}_t, u_t]]} c_t, \sup_{\hat{\pi}_{t+1} \in [[\hat{\Pi}_{t+1}|\hat{\pi}_t, u_t]]} \hat{\Lambda}_{t+1}(\hat{\pi}_{t+1}) \right\}, \quad (46)$$

$$\hat{\Lambda}_t(\hat{\pi}_t) := \hat{\Theta}_t(\hat{\pi}_t, \hat{g}_t(\hat{\pi}_t)), \quad (47)$$

where  $\hat{\Lambda}_{T+1}(\hat{\pi}_{T+1}) := 0$ , identically. Note that

$$\hat{\Theta}_t(\hat{\pi}_t, u_t) = \hat{Q}_t(\hat{\pi}_t, u_t) \quad \text{and} \quad \hat{\Lambda}_t(\hat{\pi}_t) = \hat{V}_t(\hat{\pi}_t), \quad (48)$$

for all  $t = 0, \dots, T$ , since  $\hat{g}_t(\hat{\pi}_t) = \arg \min_{u_t \in U_t} \hat{Q}_t(\hat{\pi}_t, u_t)$ .

We first prove (44) for all  $t = 0, \dots, T$ . At time  $t$ , using the triangle inequality and (48) in the LHS of (44):

$$|Q_t(m_t, u_t) - \Theta_t(m_t, u_t)| \leq |Q_t(m_t, u_t) - \hat{Q}_t(\hat{\sigma}_t(m_t), u_t)| + |\hat{Q}_t(\hat{\sigma}_t(m_t), u_t) - \hat{\Theta}_t(\hat{\sigma}_t(m_t), u_t)| + |\hat{\Theta}_t(\hat{\sigma}_t(m_t), u_t) - \Theta_t(m_t, u_t)| \leq \alpha_t + |\hat{\Theta}_t(\hat{\sigma}_t(m_t), u_t) - \Theta_t(m_t, u_t)|, \quad (49)$$

where, in the second inequality, we use (36) from Theorem 3. Then, to prove (44), it suffices to show that

$$|\hat{\Theta}_t(\hat{\sigma}_t(m_t), u_t) - \Theta_t(m_t, u_t)| \leq \alpha_t. \quad (50)$$

We use mathematical induction starting at time  $T+1$  to prove (50) in addition to  $|\hat{\Lambda}_t(\hat{\sigma}_t(m_t)) - \Lambda_t(m_t)| \leq \alpha_t$  for all  $t = 0, \dots, T$ . At time  $T+1$ , using the definitions it holds that  $\hat{\Lambda}_{T+1}(\hat{\sigma}_{T+1}(m_{T+1})) = \Lambda_{T+1}(m_{T+1}) = 0$ . This forms the basis of our induction. Next, for all  $t = 0, \dots, T$ , we consider the induction hypothesis that  $|\hat{\Lambda}_{t+1}(\hat{\sigma}_{t+1}(m_{t+1})) - \Lambda_{t+1}(m_{t+1})| \leq \alpha_{t+1}$ . Given the hypothesis, (50) holds at time  $t$  using the same sequence of arguments as in the proof for Theorem 3. Next, using the definitions of the value functions from (43) and (47), we write that

$$|\hat{\Lambda}_t(\hat{\sigma}_t(m_t)) - \Lambda_t(m_t)| = |\hat{\Theta}_t(\hat{\sigma}_t(m_t), \hat{g}_t(\hat{\sigma}_t(m_t))) - \Theta_t(m_t, g_t(m_t))| = |\hat{\Theta}_t(\hat{\sigma}_t(m_t), \hat{u}_t) - \Theta_t(m_t, \hat{u}_t)| \leq \alpha_t, \quad (51)$$

where, in the second equality, we use the definition of the control law to write that  $g_t(m_t) = \hat{g}_t(\hat{\sigma}_t(m_t)) =: \hat{u}_t$ ; and in the inequality, we use (50). This proves the induction hypothesis for time  $t$  given the hypothesis for time  $t+1$ . Thus, using mathematical induction (50) holds for all  $t = 0, \dots, T$ . Subsequently, we complete the proof for (44) for all  $t = 0, \dots, T$  by substituting (50) into the RHS of (49). Furthermore, note that (45) follows directly from (44) using the same sequence of arguments used to prove (51).  $\square$

**Remark 12.** The definition of  $\alpha_t$  at any  $t$  characterizes a trade-off in approximate information states between cost and observation approximation. In this definition, the observation prediction error  $\delta_t$  is scaled by  $L_{\hat{V}_t}$  and accumulated with  $\alpha_{t+1}$ . Thus, to minimize  $\alpha_t$ , we typically prioritize a reduction in  $\delta_t$  over a reduction in the cost approximation error  $\epsilon_t$ .

**Remark 13.** We can specialize the results of both Theorem 3 and Theorem 4 to terminal cost problems, where the optimal DP is given by (12) - (13) and the approximate DP is given by (27) - (28). Then, the approximation bounds hold with  $\alpha_t := \alpha_{t+1} + L_{\hat{V}_{t+1}^{\text{tm}}} \cdot \delta_t$  for all  $t = 0, \dots, T-1$  and  $\alpha_T := \epsilon_T$ .

### B. Alternate Characterization

In this subsection, we provide stronger but simpler conditions to identify an approximate information state as alternatives to (23) and (24). They prescribe that an approximate information state  $\hat{\Pi}_t = \hat{\sigma}_t(M_t)$  must satisfy for all  $t$ :

1) *State-like evolution:* There exists a Lipschitz continuous function  $\hat{f}_t : \hat{\mathcal{P}}_t \times \mathcal{U}_t \times \mathcal{Y}_{t+1} \rightarrow \mathcal{P}_{t+1}$ , such that

$$\hat{\Pi}_{t+1} = \hat{f}_t(\hat{\Pi}_t, U_t, Y_{t+1}). \quad (52)$$

2) *Sufficient to approximate observations:* For all  $m_t \in [[M_t]]$  and  $u_t \in [[U_t]]$ , we define the sets  $\mathcal{K}_{t+1}^{\text{ob}} := [[Y_{t+1} \mid m_t, u_t]]$  and  $\hat{\mathcal{K}}_{t+1}^{\text{ob}} := [[Y_{t+1} \mid \hat{\sigma}_t(m_t), u_t]]$ , each independent of the strategy  $\mathbf{g}$ . Then, for some  $\delta_t^{\text{ob}} \in \mathbb{R}_{\geq 0}$ :

$$\mathcal{H}(\mathcal{K}_{t+1}^{\text{ob}}, \hat{\mathcal{K}}_{t+1}^{\text{ob}}) \leq \delta_t^{\text{ob}}. \quad (53)$$

3) *Lipschitz-like observation prediction:* There exists a constant  $\lambda_t^{\text{ob}} \in \mathbb{R}_{\geq 0}$  such that for all  $\hat{\pi}_t^1, \hat{\pi}_t^2 \in [[\hat{\Pi}_t]]$ ,

$$\mathcal{H}([[Y_{t+1} \mid \hat{\pi}_t^1, u_t]], [[Y_{t+1} \mid \hat{\pi}_t^2, u_t]]) \leq \lambda_t^{\text{ob}} \cdot \eta(\hat{\pi}_t^1, \hat{\pi}_t^2), \quad (54)$$

where  $\eta$  is an appropriate metric on  $\hat{\mathcal{P}}_t$ .

In addition to (22) in Definition 2, the conditions (52) - (54) are sufficient to characterize an approximate information state.

**Lemma 2.** For all  $t = 0, \dots, T$ , if an uncertain variable  $\hat{\Pi}_t = \hat{\sigma}_t(M_t)$  satisfies (52) - (53), it also satisfies (23).

*Proof.* Let  $m_t \in [[M_t]]$  be a given realization of  $M_t$  and let  $\hat{\pi}_t = \hat{\sigma}_t(m_t)$  satisfy (52) - (53), for all  $t = 0, \dots, T$ . Then, using (52), we can write the LHS in (23) as  $\mathcal{H}(\mathcal{K}_{t+1}, \hat{\mathcal{K}}_{t+1}) = \mathcal{H}([[\hat{f}_t(\hat{\sigma}_t(m_t), u_t, Y_{t+1}) \mid m_t, u_t]], [[\hat{f}_t(\hat{\sigma}_t(m_t), u_t, Y_{t+1}) \mid \hat{\sigma}_t(m_t), u_t]]) = \max \{ \sup_{y_{t+1} \in \mathcal{K}_{t+1}^{\text{ob}}} \inf_{\hat{y}_{t+1} \in \hat{\mathcal{K}}_{t+1}^{\text{ob}}} \eta(\hat{f}_t(\hat{\sigma}_t(m_t), u_t, y_{t+1}), \hat{f}_t(\hat{\sigma}_t(m_t), u_t, \hat{y}_{t+1})), \sup_{\hat{y}_{t+1} \in \hat{\mathcal{K}}_{t+1}^{\text{ob}}} \inf_{y_{t+1} \in \mathcal{K}_{t+1}^{\text{ob}}} \eta(\hat{f}_t(\hat{\sigma}_t(m_t), u_t, y_{t+1}), \hat{f}_t(\hat{\sigma}_t(m_t), u_t, \hat{y}_{t+1})) \}$ , where, in the second equality, we use the definition of the Hausdorff distance from (1). Note that  $\hat{f}_t$  is globally Lipschitz from the alternate characterization of the approximate information state. This implies that  $\eta(\hat{f}_t(\hat{\sigma}_t(m_t), u_t, y_{t+1}), \hat{f}_t(\hat{\sigma}_t(m_t), u_t, \hat{y}_{t+1})) \leq L_{f_t} \cdot \eta(y_{t+1}, \hat{y}_{t+1})$ , and thus  $\mathcal{H}(\mathcal{K}_{t+1}, \hat{\mathcal{K}}_{t+1}) \leq L_{f_t} \max \{ \sup_{y_{t+1} \in \mathcal{K}_{t+1}^{\text{ob}}} \inf_{\hat{y}_{t+1} \in \hat{\mathcal{K}}_{t+1}^{\text{ob}}} \eta(y_{t+1}, \hat{y}_{t+1}), \sup_{\hat{y}_{t+1} \in \hat{\mathcal{K}}_{t+1}^{\text{ob}}} \inf_{y_{t+1} \in \mathcal{K}_{t+1}^{\text{ob}}} \eta(y_{t+1}, \hat{y}_{t+1}) \} = L_{f_t} \cdot \mathcal{H}(\mathcal{K}_{t+1}^{\text{ob}}, \hat{\mathcal{K}}_{t+1}^{\text{ob}}) \leq L_{f_t} \cdot \delta_t^{\text{ob}}$ .  $\square$

**Lemma 3.** For all  $t = 0, \dots, T$ , if an uncertain variable  $\hat{\Pi}_t = \hat{\sigma}_t(M_t)$  satisfies (52) - (54), it also satisfies (24).

*Proof.* Let  $\hat{\pi}_t^1, \hat{\pi}_t^2 \in [[\hat{\Pi}_t]]$  be two possible realizations of an approximate information state  $\hat{\Pi}_t$ , which satisfies (52) - (54), for all  $t = 0, \dots, T$ . Then, using (52), we can write the LHS in (24) as  $\mathcal{H}([[\Pi_{t+1} \mid \hat{\pi}_t^1, u_t]], [[\Pi_{t+1} \mid \hat{\pi}_t^2, u_t]]) = \mathcal{H}([[\hat{f}_t(\hat{\pi}_t^1, u_t, Y_{t+1}) \mid \hat{\pi}_t^1, u_t]], [[\hat{f}_t(\hat{\pi}_t^2, u_t, Y_{t+1}) \mid \hat{\pi}_t^2, u_t]]) \leq L_{\hat{f}_t} \cdot (\eta(\hat{\pi}_t^1, \hat{\pi}_t^2) + \mathcal{H}([[Y_{t+1} \mid \hat{\pi}_t^1, u_t]], [[Y_{t+1} \mid \hat{\pi}_t^2, u_t]])) \leq L_{\hat{f}_t} \cdot (1 + \lambda_t^{\text{ob}}) \cdot \eta(\hat{\pi}_t^1, \hat{\pi}_t^2)$ , where, in the first inequality, we use the

Lipschitz continuity of the function  $\hat{f}_t$  along with the triangle inequality, and in the second inequality, we use (54). This completes the proof by defining  $\lambda_t := L_{\hat{f}_t} \cdot (1 + \lambda_t^{\text{ob}})$ .  $\square$

### C. Examples

We present two examples approximate information states constructed using state-quantization [83] in systems with a known state-space model. Consider a system as described in Subsection III-C with compact feasible sets  $\{\mathcal{X}_t, \mathcal{N}_t, \mathcal{W}_t \mid t = 0, \dots, T\}$  in a metric space  $(\mathcal{S}, \eta)$ . Recall that  $\mathcal{X}_t$  is the state space at any  $t$ . A *quantized state space* is defined as a finite set  $\hat{\mathcal{X}}_t \subset \mathcal{X}_t$ , such that for a given quantization parameter  $\gamma_t \in \mathbb{R}_{\geq 0}$ , it holds that  $\max_{x_t \in \mathcal{X}_t} \min_{\hat{x}_t \in \hat{\mathcal{X}}_t} \eta(x_t, \hat{x}_t) \leq \gamma_t$ . This implies that for every point  $x_t \in \mathcal{X}_t$ , there exists a point in the quantized space  $\hat{\mathcal{X}}_t$  that is at most  $\gamma_t$  distance away from  $x_t$ . Then, a *quantization function*  $\mu_t : \mathcal{X}_t \rightarrow \hat{\mathcal{X}}_t$  maps each point in the state space to the closest point in the quantized space for all  $t = 0, \dots, T$ , and is defined as  $\mu_t(x_t) := \arg \min_{\hat{x}_t \in \hat{\mathcal{X}}_t} \eta(x_t, \hat{x}_t)$  for every  $x_t \in \mathcal{X}_t$ .

1) *Perfectly Observed Systems:* Consider a system where  $Y_t = X_t$  for all  $t = 0, \dots, T$ . Recall from Subsection III-C that the  $\Pi_t = X_t \in \mathcal{X}_t$  for all  $t$ . Building upon this information state, an approximate information state for such a system can be constructed using the state quantization function, i.e.,  $\hat{\Pi}_t := \mu_t(X_t)$ . This approximate information state Definition 2 with  $\epsilon_t = 2L_{d_t} \cdot \gamma_t$  and  $\delta_t = 2\gamma_{t+1} + 2L_{f_t} \cdot \gamma_t$ , where  $\gamma_{T+1} = 0$ , and  $L_{d_t}$  and  $L_{f_t}$  are the Lipschitz constants for  $d_t$  and  $f_t$ , respectively (proof in Appendix C). Note that because  $\hat{\Pi}_t$  takes values in a finite set, it trivially satisfies (24) in Definition 2.

2) *Partially Observed Systems:* For a partially observed system, recall from Subsection III-C that an information state is given by the conditional range  $\Pi_t = [[X_t \mid m_t]] \in \mathcal{B}(\mathcal{X}_t)$ , where  $\mathcal{B}(\mathcal{X}_t)$  is the set of all compact subsets of  $\mathcal{X}_t$ . We can construct an approximate information state by mapping each element  $x_t \in \Pi_t$  to its closest element in the quantized space  $\hat{\mathcal{X}}_t$ . We denote this mapping by  $\nu_t : \mathcal{B}(\mathcal{X}_t) \rightarrow 2^{\hat{\mathcal{X}}_t}$ , where  $2^{\hat{\mathcal{X}}_t}$  is the power set of  $\hat{\mathcal{X}}_t$ , and note that  $\nu_t(\Pi_t) := \{\mu_t(x_t) \in \hat{\mathcal{X}}_t \mid x_t \in \Pi_t\}$  for any  $\Pi_t \in \mathcal{B}(\mathcal{X}_t)$ . Then,  $\hat{\Pi}_t = \nu_t(\Pi_t)$  is an approximate information state for the partially observed system for all  $t = 0, \dots, T$  with  $\epsilon_t = 2L_{d_t} \cdot \gamma_t$  and  $\delta_t = 2\gamma_{t+1} + 2L_{\bar{f}_t} \cdot L_{h_{t+1}} \cdot L_{f_t} \cdot \gamma_t$ , where  $\gamma_{T+1} = 0$ , and  $L_{\bar{f}_t}$ ,  $L_{h_{t+1}}$  and  $L_{f_t}$  are Lipschitz constants of  $\bar{f}_t$ ,  $h_{t+1}$ , and  $f_t$ , respectively (proof in Appendix D).

### D. Learning an approximate information state

In this subsection, we present an approach to learning an approximate information state from data. We assume access to multiple trajectories  $(Y_{t+1}, C_t, U_t : t = 0, \dots, T)$  generated using an exploratory control strategy. We denote the collection of such trajectories by the dataset  $\mathcal{D}$ . We begin by creating data-driven conditional ranges  $\mathcal{K}_{t+1}^{\text{ob}} = [[Y_{t+1} \mid M_t, U_t]]^{\mathcal{D}}$  for all  $t$  by combining the observations at time  $t+1$  for trajectories in  $\mathcal{D}$  that share common  $(M_t, U_t)$  at time  $t$ . Similarly, we estimate the maximum incurred cost from the data at each  $t$  as  $C_t^{\max} = \max_{c_t \in [[C_t \mid M_t, U_t]]^{\mathcal{D}}} c_t$ . Here note that the maximum exists because  $[[C_t \mid M_t, U_t]]^{\mathcal{D}}$  is a finite set.

Then, we select three function approximators for each  $t$ : **(1) the encoder**  $\psi_t : \hat{\mathcal{P}}_{t-1} \times \mathcal{U}_{t-1} \times \mathcal{Y}_t \rightarrow \hat{\mathcal{P}}_t$  that recursively

compresses the history into an approximate information state, **(2) the observation decoder**  $\phi_t : \hat{\mathcal{P}}_t \times \mathcal{U}_t \rightarrow \mathcal{B}(\mathcal{Y}_{t+1})$  that predicts next observations, where  $\mathcal{B}(\mathcal{Y}_{t+1})$  is a set of compact subsets of  $\mathcal{Y}$ , and **3) the cost decoder**  $\phi_t^c : \hat{\mathcal{P}}_t \times \mathcal{U}_t \rightarrow \mathcal{C}_t$  that predicts maximum cost. The encoder is typically constructed with a recurrent neural network (e.g., GRU) whose hidden state is treated as  $\hat{\Pi}_t = \psi_t(\hat{\Pi}_{t-1}, U_{t-1}, Y_t)$  at each  $t$ . The output of the encoder is augmented with the latest action  $U_t$  to predict the conditional range  $\hat{\mathcal{K}}_{t+1}^{\text{ob}} = [[Y_{t+1} \mid \hat{\Pi}_t, U_t]] = \phi_t(\hat{\Pi}_t, U_t)$ , that best approximates  $\mathcal{K}_{t+1}^{\text{ob}}$ , and the worst-case cost  $\hat{C}_t := \phi_t^c(\hat{\Pi}_t, U_t)$  that best approximates  $C_t^{\max}$ .

We also select two training losses  $L_t^c = \|\hat{C}_t - C_t^{\max}\|_2$  and  $L_t^{\text{ob}} = \hat{\mathcal{H}}(\hat{\mathcal{K}}_{t+1}^{\text{ob}}, \mathcal{K}_{t+1}^{\text{ob}})$ , where  $\hat{\mathcal{H}}$  refers to a differentiable surrogate for the Hausdorff metric. For  $\hat{\mathcal{H}}$ , we can use either the surrogate procedure in [84, Subsection II-B], or the average Hausdorff distance [85] between two sets  $\mathcal{X}, \mathcal{Y} \subset (\mathcal{S}, \eta)$ :

$$\hat{\mathcal{H}}(\mathcal{X}, \mathcal{Y}) := \left\{ \frac{\int x \min_{y \in \mathcal{Y}} \eta(x, y) dx}{\int dx} + \frac{\int y \min_{x \in \mathcal{X}} \eta(x, y) dy}{\int dy} \right\}.$$

Furthermore, we replace each min with a “softmin” to ensure differentiability of  $\hat{\mathcal{H}}$ . Then, the net training loss is defined at each  $t$  as  $L_t = \lambda \cdot L_t^c + L_t^{\text{ob}}$ , where  $\lambda \in \mathbb{R}_{\geq 0}$  is a weight.

To learn an approximate information state model, we train the complete network assembly over each trajectory with loss  $\sum_{t=0}^T L_t$  for one round of training. Further details on training are presented through a numerical example in Subsection V-B. Given such a model learned from data, we can use a standard DP decomposition to derive an approximate control strategy.

## V. NUMERICAL EXAMPLES

We present two numerical examples to illustrate our approach: **(1) The Wall Defense Problem**: a worst-case control problem with partial observations, used to illustrate state-quantization from Subsection IV-C, and **(2) The Pursuit Evasion Problem**: a worst-case reinforcement learning problem used to illustrate the learning approach in Subsection IV-D.

### A. The Wall Defense Problem

In the wall defense problem, we consider an agent who defends a wall in a  $5 \times 5$  grid world from an attacker over a time horizon  $T$ . The wall is located across the central row of the grid. We illustrate the wall defense problem for one initial condition in Fig. 1(a). Here, the black-colored cells constitute the wall, and the grey-hatched cells are adjacent to the wall. The solid blue triangle, solid red circle and red ring are the agent, attacker and observation, respectively, at  $t = 0$ . The pink cells are feasible positions of the attacker given the observation. The attacker moves within the bottom two rows of the grid and damages a wall cell when positioned in an adjacent cell. At each  $t = 0, \dots, T$ , we denote the position of the attacker by  $X_t^{\text{at}} \in \mathcal{X}^{\text{at}} = \{(-2, -1), \dots, (2, -1), (-2, -2), \dots, (2, -2)\}$ . In contrast, the agent moves within the top two rows of the grid and repairs a wall cell when positioned in an adjacent cell. At each  $t$ , we denote the position of the agent by  $X_t^{\text{ag}} \in \mathcal{X}^{\text{ag}} = \{(-2, 1), \dots, (2, 1), (-2, 2), \dots, (2, 2)\}$ . The state of the wall at each  $t$  is the accumulated damage denoted by  $D_t = (D_t^{-2}, \dots, D_t^2)$ , where  $D_t^i \in \mathcal{D}_t^i = \{0, 1, 2, 3\}$  for all  $i = -2, \dots, 2$  and  $\mathcal{D}_t = \times_{i=-2}^2 \mathcal{D}_t^i$ .

The attacker starts at the position  $X_0^{\text{at}} \in \mathcal{X}^{\text{at}}$ , which evolves for all  $t$  as  $X_{t+1}^{\text{at}} = \mathbb{I}(X_t^{\text{at}} + W_t \in \mathcal{X}^{\text{at}}) \cdot (X_t^{\text{at}} + W_t) + (1 - \mathbb{I}(X_t^{\text{at}} + W_t \in \mathcal{X}^{\text{at}})) \cdot X_t^{\text{at}}$ , where  $\mathbb{I}$  is the indicator function and  $W_t \in \mathcal{W}_t$  is an uncontrolled disturbance with  $\mathcal{W}_t = \{(-1, 0), (1, 0), (0, 0), (0, 1), (0, -1)\}$ . At each  $t$ , the agent observes their own position and the wall’s state. The agent also partially observes the attacker’s position as  $Y_t = \mathbb{I}(X_t^{\text{at}} + N_t \in \mathcal{X}^{\text{at}}) \cdot (X_t^{\text{at}} + N_t) + (1 - \mathbb{I}(X_t^{\text{at}} + N_t \in \mathcal{X}^{\text{at}})) \cdot X_t^{\text{at}}$ , where  $N_t \in \mathcal{N}_t = \{(0, 0), (0, 1)\}$  is the measurement noise. Given the history of observations, the agent selects an action  $U_t \in \mathcal{U}_t = \mathcal{W}_t$  at each  $t$ . Starting with  $X_0^{\text{ag}} \in \mathcal{X}^{\text{ag}}$ , the agent moves as  $X_{t+1}^{\text{ag}} = \mathbb{I}(X_t^{\text{ag}} + U_t \in \mathcal{X}^{\text{ag}}) \cdot (X_t^{\text{ag}} + U_t) + (1 - \mathbb{I}(X_t^{\text{ag}} + U_t \in \mathcal{X}^{\text{ag}})) \cdot X_t^{\text{ag}}$ . Starting with  $D_0 = (0, 0, 0, 0, 0)$ , the state of the wall evolves as  $D_{t+1}^i = \min \{3, \max \{0, D_t^i + \mathbb{I}(X_t^{\text{at}} = (i, -1)) - \mathbb{I}(X_t^{\text{ag}} = (i, 1))\}\}$  for all  $t$  and  $i = -2, \dots, 2$ . At each  $t$ , after selecting the action, the agent incurs a cost for the damage to the wall, i.e.,  $c_t(D_t) = \sum_{i=-2}^2 D_t^i$ . The agent’s aim is to minimize the maximum instantaneous damage to the wall, i.e.,  $\mathcal{J}(\mathbf{g}) = \max_{t=0, \dots, T} \max_{x_0, w_0, t, n_0, t} c_t(D_t)$ .

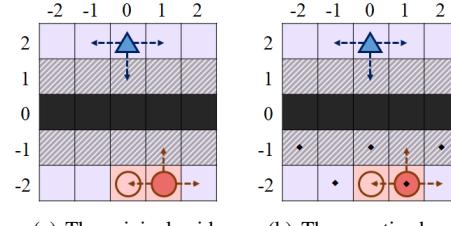


Fig. 1: The wall defense problem with the initial conditions  $x_0^{\text{ag}} = (0, 2)$  and  $y_0 = (0, -2)$ .

Since we are given a state-space model for this problem, this system is equivalent to a combination of Case 1 and Case 2 in Subsection III-C, i.e., the state of the agent and damage to the wall are perfectly observed, whereas the state of the attacker is partially observed. Thus, an information state at each  $t$  is given by a combination of the two information states  $\Pi_t = (X_t^{\text{ag}}, D_t, [[X_t^{\text{at}} | M_t]])$ . We construct an approximation of the conditional range  $[[X_t^{\text{at}} | M_t]]$  at time  $t$  using state quantization explained in Subsection IV-C, and define the approximate range  $\hat{A}_t = \{\mu_t(x_t) \in \hat{\mathcal{X}}^{\text{at}} | x_t \in [[X_t^{\text{at}} | M_t]]\}$ . The set of quantized cells  $\hat{\mathcal{X}}^{\text{at}}$ , with  $\gamma_t = 1$  for all  $t$ , is marked in Fig. 1(b) with dots. Recall that  $\mu_t(x_t) = \arg \min_{\hat{x}_t \in \hat{\mathcal{X}}} d(x_t, \hat{x}_t)$  and the approximate range at time  $t$  is  $\hat{A}_t = \{\mu_t(x_t) \in \hat{\mathcal{X}}^{\text{at}} | x_t \in [[X_t^{\text{at}} | M_t]]\}$ . We consider the approximate information state  $\hat{\Pi}_t = (X_t^{\text{ag}}, D_t, \hat{A}_t, Y_0)$  for all  $t$ . Here,  $\hat{\Pi}_t$  includes the initial observation  $Y_0$  along with the quantized information state  $\hat{A}_t$  because, in practice, retaining knowledge of  $Y_0$  helps improve the prediction of  $\hat{A}_{t+1}$ . Including this additional term does not violate Definition 2. For five initial conditions, we compute the best control strategy for  $T = 6$  using both the information state (IS) and the approximate information state (AIS). In Table I, we present the computational times (Run.) for both the DPs in seconds. Note that the approximate DP has a faster runtime in all cases. We also implement both strategies with random disturbances in the system with  $T = 6$ . In Table I, we also present the *actual* worst-case costs across  $5 \times 10^3$  implementations of both strategies and note that the AIS has a bounded deviation from the IS.

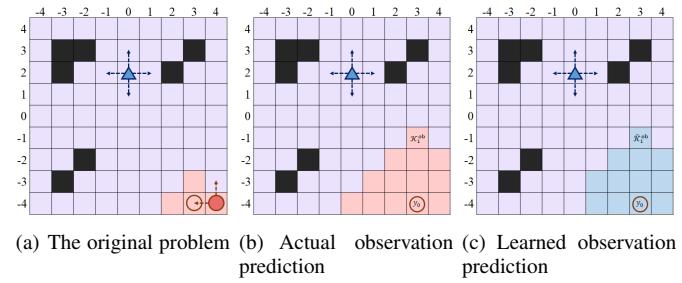
Initial Conditions		Strategy IS		Strategy AIS	
$x_0^{\text{ag}}$	$y_0$	Run (s)	Max. cost	Run (s)	Max. cost
(1, 2)	(1, -2)	354.2	2	221.2	2
(0, 1)	(0, -1)	1209.7	2	607.3	3
(-2, 2)	(-1, -1)	686.1	3	446.4	3
(2, 2)	(2, -2)	19.22	2	12.81	2
(-2, 2)	(2, -1)	582.3	3	362.1	3
(-1, 1)	(1, -1)	1551.5	2	1075.1	3

TABLE I: Results for  $5 \times 10^4$  simulations and  $T = 6$ .

### B. Pursuit Evasion Problem

In the pursuit-evasion problem, we consider an agent who chases a moving target in a  $9 \times 9$  grid world with static obstacles. The agent aims to get close to the target over a time horizon  $T$ . For each  $t = 0, \dots, T$ , we denote the position of the agent by  $X_t^{\text{ag}} \in \mathcal{X}$  and that of the target by  $X_t^{\text{ta}} \in \mathcal{X}$ , where  $\mathcal{X} = \{(-4, -4), \dots, (4, 4)\} \setminus \mathcal{O}$  is the set of feasible grid cells and  $\mathcal{O} \subset \mathcal{X}$  is the set of obstacles. The target starts at the position  $X_0^{\text{ta}} \in \mathcal{X}$ , which is updated as  $X_{t+1}^{\text{ta}} = \mathbb{I}(X_t^{\text{ta}} + W_t \in \mathcal{X}) \cdot (X_t^{\text{ta}} + W_t) + (1 - \mathbb{I}(X_t^{\text{ta}} + W_t \in \mathcal{X})) \cdot X_t^{\text{ta}}$ , where  $W_t \in \mathcal{W}_t = \{(-1, 0), (1, 0), (0, 0), (0, 1), (0, -1)\}$  is the disturbance. At each  $t$ , the agent perfectly observes their own position and nosily observes the target's position as  $Y_t = \mathbb{I}(X_t^{\text{ta}} + N_t \in \mathcal{X}) \cdot (X_t^{\text{ta}} + N_t) + (1 - \mathbb{I}(X_t^{\text{ta}} + N_t \in \mathcal{X})) \cdot X_t^{\text{ta}}$ , where  $N_t \in \mathcal{N}_t = \mathcal{W}_t$  is the measurement noise. Next, starting with  $X_0^{\text{ag}} \in \mathcal{X}$ , the agent selects an action  $U_t \in \mathcal{U}_t = \mathcal{W}_t$  to move as  $X_{t+1}^{\text{ag}} = \mathbb{I}(X_t^{\text{ag}} + U_t \in \mathcal{X}) \cdot (X_t^{\text{ag}} + U_t) + (1 - \mathbb{I}(X_t^{\text{ag}} + U_t \in \mathcal{X})) \cdot X_t^{\text{ag}}$ . At time  $T$ , the agent selects no action and observes the target's position  $X_T^{\text{ta}}$  and incurs a cost  $c_T(X_T^{\text{ta}}, X_T^{\text{ag}}) = \eta(X_T^{\text{ta}}, X_T^{\text{ag}}) \in \mathbb{R}_{\geq 0}$ , where  $\eta$  is the shortest distance between two cells while avoiding obstacles. The distance between two adjacent cells is 1 unit. The agent seeks to minimize the worst-case terminal cost *without* prior knowledge of either the observation function or the target's evolution dynamics. Note that this is a reinforcement learning generalization of Problem 1. We illustrate the grid and one initial setup in Fig. 2(a). Here, the black cells are obstacles. The solid blue triangle, solid red circle, and red ring are the agent, target, and observation, respectively, at  $t = 0$ . The pink cells are feasible positions of the target given the observation.

We approach this problem using the model-based reinforcement learning approach from Subsection IV-D, where we learn an approximate information state representation from  $3 \times 10^7$  trajectories and then compute a control strategy. We use the dataset  $\mathcal{D}$  to construct estimates of the conditional range  $\mathcal{K}_{t+1}^{\text{ob}} = [[Y_{t+1}|Y_{0:t}]]^{\mathcal{D}}$  for all  $t = 0, \dots, T-2$  and  $\mathcal{K}_T^{\text{ob}} = [[X_T^{\text{ta}}|Y_{0:T-1}]]^{\mathcal{D}}$ . Then, we set up a deep neural network with an encoder-decoder structure for each  $t = 0, \dots, T$ , as illustrated in Fig. 3. At each  $t$ , the encoder  $\psi_t$  comprises of 3 layer neural network with sizes  $(2, 14), (14, 12), (12+24, 24)$  and ReLU activation for the first two layers, where the inputs are a 2-d vector of coordinates for observation  $Y_t$  and a 24-d vector for the previous approximate information state  $\hat{\Pi}_{t-1}$ . The encoder compresses these inputs to a 24-d vector representing the approximate information state  $\hat{\Pi}_t$ . At each  $t$ , the decoder  $\phi_t$  is a 4-layer neural network of size  $(24, 48), (48, 56), (56, 64), (64, 74)$  with ReLU activation for the first three layers and sigmoid activation for the last layer. Its input is



(a) The original problem (b) Actual observation prediction (c) Learned observation prediction

Fig. 2: The pursuit evasion problem with the initial conditions  $x_0^{\text{ag}} = (0, 2)$  and  $y_0 = (3, -4)$ .

$\hat{\Pi}_t$ , and its output is a 74-d vector with each component taking values in  $[0, 1]$ . Each component of the 74-d output gives a set inclusion value for a specific feasible cell in the  $9 \times 9$  grid, excluding obstacles. The output is thus interpreted as the conditional range  $\hat{\mathcal{K}}_{t+1}^{\text{ob}} = [[Y_{t+1} | \hat{\Pi}_t]]$  for all  $t = 0, \dots, T-2$  and  $\hat{\mathcal{K}}_T^{\text{ob}} = [[X_T^{\text{ta}} | \hat{\Pi}_{T-1}]]$ . We consider a set-inclusion threshold of 0.5 for inclusion in  $\hat{\mathcal{K}}_{t+1}^{\text{ob}}$  at each  $t$ .

The learning objective of our neural network at each  $t$  is to minimize  $\hat{\mathcal{H}}(\mathcal{K}_{t+1}^{\text{ob}}, \hat{\mathcal{K}}_{t+1}^{\text{ob}})$ , where we use the surrogate function in [84, Subsection II-B]. Note that at the terminal time step, this objective also minimizes the difference in maximum costs. We train the network for 40 epochs using 90% of the available data with a learning rate of 0.0003 and test it against the other 10%. To illustrate the training results, consider an out-of-sample initial observation  $y_0 = (3, -4)$ . Then, the set  $\mathcal{K}_1^{\text{ob}}$  constructed using data is shown by pink cells in 2(b) and the set  $\hat{\mathcal{K}}_1^{\text{ob}}$  generated by the trained network is shown by blue cells in 2(c). Here, the trained network's output matches the conditional range constructed from data accurately except for one cell  $(0, -4)$ . We train a neural network for each  $t$  up to  $T = 8$  to learn a complete approximate information state representation for the problem. Then, at each  $t$ , the agent uses  $(X_t^{\text{ag}}, \hat{\Pi}_t)$  in the approximate DP (25) - (26).

We compare the performance of this approximate strategy with a baseline strategy that uses the observation  $Y_t$  at each  $t$  instead of  $\hat{\Pi}_t$ . Thus, for this baseline we train a network to match the prediction  $[[Y_{t+1} | Y_t]]$  to  $\mathcal{K}_{t+1}^{\text{ob}}$  for all  $t = 0, \dots, T-2$  and  $[[X_T^{\text{ta}} | Y_{T-1}]]$  to  $[[X_T^{\text{ta}} | Y_{0:T-1}]]$  at time  $T-1$ . The neural network structure is the same as before except for a lack of  $\hat{\Pi}_{t-1}$  in the encoder input at each  $t$ , and we use the same training hyper-parameters as before. Subsequently, the agent computes an approximately optimal strategy using the approximate DP with the state  $(X_t^{\text{ag}}, Y_t)$  at each  $t$ .

For six initial conditions, we present in Table II the worst case costs obtained when implementing both the approximately optimal strategy (with AIS) and the baseline strategy

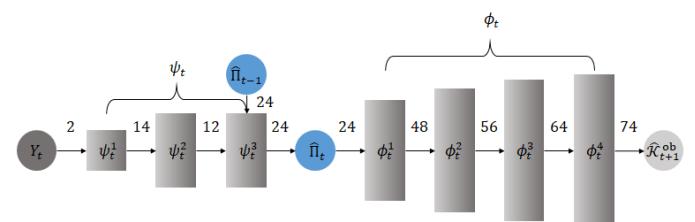


Fig. 3: The neural network architecture for approximate information states at any  $t = 0, \dots, T-1$ .

(without AIS) for  $T = 4$  and  $T = 8$ . Across  $10^3$  simulations with randomly generated uncertainties, we note that using the learned approximate information state consistently improves worst-case performance compared to the baseline. Furthermore, the approximate strategy's outperformance grows larger for  $T = 8$  as compared to  $T = 4$  in most cases. Typically, we expect this margin to grow for longer time horizons because there is more loss of information when using only the latest observation for decision-making. Thus, learning an approximate information state representation is a viable approach for robust model-based reinforcement learning.

Initial Conditions	$T = 4$		$T = 8$	
	With AIS	No AIS	With AIS	No AIS
$x_0^{\text{ag}}$	$y_0$			
(3,2)	(0,0)	2	2	2
(-4,4)	(4,-4)	12	12	8
(4,4)	(-4,-2)	11	12	8
(-2,-3)	(-4,3)	7	8	7
(-4,2)	(1,4)	7	7	7
(-3,1)	(4,-1)	8	9	4
				5

TABLE II: Worst-case costs for  $10^4$  simulations

## VI. CONCLUSION

In this paper, we proposed a rigorous framework for worst-case control and learning in partially observed systems using non-stochastic approximate information states. We first presented two sets of properties to characterize information states and used them to construct a DP that yields an optimal control strategy. Then, we proposed two sets of properties to characterize approximate information states that can be constructed from output variables with knowledge of the dynamics or learned from output data with incomplete knowledge of the dynamics. We proved that approximate information states can be used in a DP to compute an approximate control strategy with a bounded loss in performance. We also presented theoretical examples of this bound and numerical examples to illustrate the performance of our approach in both worst-case control and reinforcement learning.

Our ongoing work is to specialize the approach in this paper to robust control [86] and reinforcement learning [52] in additive cost problems with partial observations. While this paper presents a theoretical framework and preliminary numerical results for approximate information states, several key questions remain unanswered for future research. An explicit characterization of the trade-off between computational tractability and approximation quality is missing from the theory and would play a major role in practical application of approximate information states. It is also important to identify specific constructions of approximate information states beyond state quantization that improve performance in specific applications. Similarly, our reinforcement learning example is a toy problem with simple dynamics and a terminal cost criterion, presented with the goal of illustrating our theory. It is imperative that we study specific neural network architectures and learning algorithms to train efficient approximate models in realistic problems with large spaces and complex underlying interactions among various components.

## APPENDIX A – $L$ -INVERTIBLE FUNCTIONS

In this appendix, we present two classes of functions that are  $L$ -invertible: 1) all bi-Lipschitz functions that have a compact domain and a compact co-domain, and 2) all functions with a compact domain and a finite co-domain.

**Lemma 4.** *Let  $\mathcal{X}$  and  $\mathcal{Y}$  be two compact subsets of a metric space  $(\mathcal{S}, \eta)$ . Then, any bi-Lipschitz function  $f : \mathcal{X} \rightarrow \mathcal{Y}$  is  $L$ -invertible.*

*Proof.* We begin by considering the pre-image set for any  $y \in \mathcal{Y}$  under the function  $f$ . Note that the function  $f$  is continuous because it is bi-Lipschitz, and the singleton  $\{y\}$  is a compact subset of a metric space. Consequently, the pre-image  $f^{-1}(y)$  is a bounded subset of  $\mathcal{X}$ . Next, let  $\mathcal{B}(\mathcal{X})$  denote the set of all bounded subsets of  $\mathcal{X}$ . Given the first result, we can consider a set-valued mapping  $\mathcal{H} : \mathcal{Y} \rightarrow \mathcal{B}(\mathcal{X})$  which returns the pre-image for each  $y \in \mathcal{Y}$ . Then, for any  $y^1, y^2 \in \mathcal{Y}$ , using the definition of the Hausdorff distance in (1):

$$\begin{aligned} \mathcal{H}(f^{-1}(y^1), f^{-1}(y^2)) &= \max \left\{ \sup_{x^1 \in f^{-1}(y^1)} \inf_{x^2 \in f^{-1}(y^2)} \right. \\ &\quad \left. \eta(x^1, x^2), \sup_{x^2 \in f^{-1}(y^2)} \inf_{x^1 \in f^{-1}(y^1)} \eta(x^1, x^2) \right\}. \end{aligned} \quad (55)$$

In the RHS of (55), the bi-Lipschitz property of  $f$  implies that there exist constants  $\underline{L}_f, \bar{L}_f \in \mathbb{R}_{>0}$  such that  $\underline{L}_f \eta(x^1, x^2) \leq |f(x^1) - f(x^2)| \leq \bar{L}_f \eta(x^1, x^2)$ , for all  $x^1, x^2 \in \mathcal{X}$ . Thus, for all  $x^1 \in g^{-1}(y^1)$  and  $x^2 \in g^{-1}(y^2)$ , we write that

$$\eta(x^1, x^2) \leq \underline{L}_f^{-1} \cdot \eta(y^1, y^2). \quad (56)$$

The proof is complete by substituting (56) into (55) and defining the constant  $L_{f^{-1}} := \underline{L}_f^{-1}$ .  $\square$

**Lemma 5.** *Let  $\mathcal{X}$  be a compact subset and  $\mathcal{Y}$  be a finite subset of  $(\mathcal{S}, \eta)$ . Then, any function  $f : \mathcal{X} \rightarrow \mathcal{Y}$  is  $L$ -invertible.*

*Proof.* Let  $\|\mathcal{Y}\| > 0$  denote the minimum distance between two distinct elements in the finite, non-empty set  $\mathcal{Y}$ . Then, for any  $y^1, y^2 \in \mathcal{Y}$  such that  $y^1 \neq y^2$ ,  $\frac{\mathcal{H}(f^{-1}(y^1), f^{-1}(y^2))}{\eta(y^1, y^2)} \leq \sup_{y^1, y^2 \in \mathcal{Y}} \frac{\mathcal{H}(f^{-1}(y^1), f^{-1}(y^2))}{\|\mathcal{Y}\|} =: L_{f^{-1}}$ , where  $L_{f^{-1}} \in \mathbb{R}_{\geq 0}$  is guaranteed to be finite because the set  $\mathcal{X}$  is bounded and thus, so is the numerator. Thus, the function  $f$  is  $L$ -invertible as defined in (55).  $\square$

## APPENDIX B – PRELIMINARY RESULTS

In this subsection, we derive results necessary to prove the properties of the approximate DP in Subsection IV-A.

**Lemma 6.** *Consider three bounded subsets  $\mathcal{X}$ ,  $\mathcal{Y}$  and  $\mathcal{Z}$  of a metric space  $(\mathcal{S}, \eta)$ . Let  $X \in \mathcal{X}$ ,  $Y \in \mathcal{Y}$  and  $Z \in \mathcal{Z}$  be uncertain variables satisfying  $Y = g(X)$ , where  $g : \mathcal{X} \rightarrow \mathcal{Y}$  is  $L$ -invertible, and  $Z = h(X)$ , where  $h : \mathcal{X} \rightarrow \mathcal{Z}$  is Lipschitz. Then, there exists an  $L_{Z|Y} \in \mathbb{R}_{\geq 0}$  such that:*

$$\mathcal{H}([[Z|y^1]], [[Z|y^2]]) \leq L_{Z|Y} \cdot \eta(y^1, y^2), \quad \forall y^1, y^2 \in [[Y]]. \quad (57)$$

*Proof.* We prove the result by constructing a feasible constant  $L_{Z|Y} \in \mathbb{R}_{\geq 0}$  which ensures that (57) is satisfied for all

$y^1, y^2 \in [[Y]]$ . We begin by using the definition of the Hausdorff distance in (1) to expand the LHS of (57) as

$$\mathcal{H}([[Z|y^1]], [[Z|y^2]]) = \max \left\{ \sup_{x^1 \in g^{-1}(y^1)} \inf_{x^2 \in g^{-1}(y^2)} \eta(h(x^1), h(x^2)), \sup_{x^2 \in g^{-1}(y^2)} \inf_{x^1 \in g^{-1}(y^1)} \eta(h(x^1), h(x^2)) \right\}, \quad (58)$$

where, note that  $[[Z|y]] = \{z \in \mathcal{Z} \mid z = h(x), \forall x \in g^{-1}(y)\}$  for any realization  $y \in [[Y]]$ . Next, recall that  $h$  is Lipschitz continuous with a constant  $L_h \in \mathbb{R}_{\geq 0}$ . Substituting this property into the RHS of (58), we write that  $\mathcal{H}([[Z|y^1]], [[Z|y^2]]) \leq L_h \cdot \max \{ \sup_{x^1 \in g^{-1}(y^1)} \inf_{x^2 \in g^{-1}(y^2)} \eta(x^1, x^2), \sup_{x^2 \in g^{-1}(y^2)} \inf_{x^1 \in g^{-1}(y^1)} \eta(x^1, x^2) \} = L_h \cdot \mathcal{H}(g^{-1}(y^1), g^{-1}(y^2)) = L_h \cdot L_{g^{-1}} \cdot \eta(y^1, y^2)$ , where, in the second equality, we use the L-invertible property of  $g$ . Then, the result follows by selecting  $L_{Z|Y} := L_h \cdot L_{g^{-1}}$ .  $\square$

**Lemma 7.** Consider a bounded set  $\mathcal{X}$  and two functions  $f: \mathcal{X} \rightarrow \mathbb{R}$  and  $g: \mathcal{X} \rightarrow \mathbb{R}$ . Then,

$$|\sup_{x \in \mathcal{X}} f(x) - \sup_{x \in \mathcal{X}} g(x)| \leq \sup_{x \in \mathcal{X}} |f(x) - g(x)|, \quad (59)$$

$$|\inf_{x \in \mathcal{X}} f(x) - \inf_{x \in \mathcal{X}} g(x)| \leq \sup_{x \in \mathcal{X}} |f(x) - g(x)|. \quad (60)$$

*Proof.* We omit the proof due to space limitations. Then, (60) follows from similar arguments as (59).  $\square$

**Lemma 8.** For any four scalars  $a, b, c, d \in \mathbb{R}$ ,

$$|\max\{a, b\} - \max\{c, d\}| \leq \max\{|a - c|, |b - d|\}. \quad (61)$$

*Proof.* We omit the proof due to space limitations.  $\square$

**Lemma 9.** Consider two nonempty bounded subsets  $\mathcal{A}, \mathcal{B}$  of a metric space  $(\mathcal{X}, \eta)$ . Let  $f: \mathcal{X} \rightarrow \mathbb{R}$  be a bounded continuous function with a Lipschitz constant  $L_f \in \mathbb{R}_{\geq 0}$  on  $\mathcal{X}$ . Then,

$$|\sup_{a \in \mathcal{A}} f(a) - \sup_{b \in \mathcal{B}} f(b)| \leq L_f \cdot \mathcal{H}(\mathcal{A}, \mathcal{B}). \quad (62)$$

*Proof.* We prove this result by considering three cases that are mutually exclusive but cover all the possibilities. Case 1:  $\sup_{a \in \mathcal{A}} f(a) = \sup_{b \in \mathcal{B}} f(b)$ , which implies that  $|\sup_{a \in \mathcal{A}} f(a) - \sup_{b \in \mathcal{B}} f(b)| = 0$ . The result holds directly from the fact that the RHS of (62) is always non-negative. Case 2:  $\sup_{a \in \mathcal{A}} f(a) < \sup_{b \in \mathcal{B}} f(b)$ , which implies  $|\sup_{a \in \mathcal{A}} f(a) - \sup_{b \in \mathcal{B}} f(b)| = \sup_{a \in \mathcal{A}} f(a) - \sup_{b \in \mathcal{B}} f(b)$ . We define the non-empty set  $\mathcal{A}^1(\beta) := \{a \in \mathcal{A} \mid f(a) + \beta \geq \sup_{b \in \mathcal{B}} f(b)\}$  for any infinitesimal  $\beta > 0$ . Then,  $\sup_{a \in \mathcal{A}} f(a) - \sup_{b \in \mathcal{B}} f(b) \leq \sup_{a \in \mathcal{A}^1(\beta)} f(a) + \beta - \sup_{b \in \mathcal{B}} f(b) \leq \sup_{a \in \mathcal{A}^1(\beta)} \inf_{b \in \mathcal{B}} (f(a) - f(b)) + \beta \leq \sup_{a \in \mathcal{A}} \inf_{b \in \mathcal{B}} |f(a) - f(b)| + \beta \leq L_f \cdot \sup_{a \in \mathcal{A}} \inf_{b \in \mathcal{B}} \eta(a, b) + \beta$  for all  $\beta > 0$ . This implies that  $|\sup_{a \in \mathcal{A}} f(a) - \sup_{b \in \mathcal{B}} f(b)| \leq L_f \cdot \sup_{a \in \mathcal{A}} \inf_{b \in \mathcal{B}} \eta(a, b) \leq L_f \cdot \mathcal{H}(\mathcal{A}, \mathcal{B})$ , where, in the second inequality, we invoke the definition of the Hausdorff distance in (1) to complete the proof. Case 3:  $\sup_{a \in \mathcal{A}} f(a) < \sup_{b \in \mathcal{B}} f(b)$  and we can prove the result using the same sequence of arguments as case 1.  $\square$

Using Lemma 9, we can also establish the following property. Consider two bounded sets  $\mathcal{Y}, \mathcal{Z} \subset \mathbb{R}^n$ ,  $n \in \mathbb{N}$ . For two uncertain variables  $Y \in \mathcal{Y}$  and  $Z \in \mathcal{Z}$ , let  $[[Z|y]]$  satisfy

$\mathcal{H}([[Z|y^1]], [[Z|y^2]]) \leq L_{Z|Y} \cdot \eta(y^1, y^2)$  for all realizations  $y^1, y^2 \in \mathcal{Y}$  of  $Y$ . Then, for a continuous function  $f: \mathcal{Z} \rightarrow \mathbb{R}_{\geq 0}$ , we use (62) to state for all  $y^1, y^2 \in [[Y]]$ :

$$\left| \sup_{z^1 \in [[Z|y_1]]} f(z^1) - \sup_{z^2 \in [[Z|y_2]]} f(z^2) \right| \leq L_{Z|Y} \cdot L_f \cdot \eta(y^1, y^2). \quad (63)$$

## APPENDIX C – APPROXIMATION BOUNDS FOR PERFECTLY OBSERVED SYSTEMS

In this appendix, we derive the values of  $\epsilon_t$  and  $\delta_t$  for all  $t = 0, \dots, T$  when an approximate information state is constructed using state quantization for a perfectly observed system, as described in Subsection IV-C. We first state a property of the Hausdorff distance, which we will use in our derivation.

**Lemma 10.** Let  $\mathcal{X}$  be a metric space with compact subsets  $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D} \subset \mathcal{X}$ . Then, it holds that

$$\mathcal{H}(\mathcal{A} \cup \mathcal{B}, \mathcal{C} \cup \mathcal{D}) \leq \max \{ \mathcal{H}(\mathcal{A}, \mathcal{C}), \mathcal{H}(\mathcal{B}, \mathcal{D}) \}. \quad (64)$$

*Proof.* The proof is given in [81, Theorem 1.12.15].  $\square$

Next, we state and prove the main result of this appendix.

**Theorem 5.** Consider a perfectly observed system, i.e.,  $Y_t = X_t$ , for all  $t = 0, \dots, T$ . Let  $\mu_t: \mathcal{X}_t \rightarrow \hat{\mathcal{X}}_t$  such that  $\max_{x_t \in \mathcal{X}_t} \eta(x_t, \mu_t(x_t)) \leq \gamma_t$  at each  $t$ . Then,  $\hat{\Pi}_t = \mu_t(X_t)$  is an approximate information state which satisfies (22) with  $\epsilon_t = 2L_{d_t} \cdot \gamma_t$  and (23) with  $\delta_t = 2\gamma_{t+1} + 2L_{f_t} \cdot \gamma_t$  for all  $t$ , where  $\gamma_{T+1} = 0$ , and where  $L_{d_t}$  and  $L_{f_t}$  are Lipschitz constants for  $d_t$  and  $f_t$ , respectively.

*Proof.* For all  $t = 0, \dots, T$ , let  $m_t = (x_{0:t}, u_{0:t-1})$  be the realization of  $M_t$  and let the approximate information state be  $\hat{x}_t = \mu_t(x_t)$ . We first derive the value of  $\epsilon_t$  in the RHS of (22). At time  $t$ , can expand the conditional ranges to write that  $[[X_t|m_t]] = [[X_t|x_t]] = \{x_t\}$  and  $[[X_t|\hat{x}_t]] = \{x_t \in \mathcal{X} \mid \eta(x_t, \hat{x}_t) \leq \gamma_t\}$ . On substituting these into the LHS of (22), we state that  $\left| \sup_{c_t \in [[C_t|m_t, u_t]]} c_t - \sup_{c_t \in [[C_t|\mu_t(x_t), u_t]]} c_t \right| = \left| d_t(x_t, u_t) - \sup_{\bar{x}_t \in [[X_t|\mu_t(x_t)]]} d_t(\bar{x}_t, u_t) \right| \leq \sup_{\bar{x}_t \in [[X_t|\mu_t(x_t)]]} |d_t(x_t, u_t) - d_t(\bar{x}_t, u_t)| \leq L_{d_t} \cdot \sup_{\bar{x}_t \in [[X_t|\mu_t(x_t)]]} \eta(x_t, \bar{x}_t) \leq L_{d_t} \cdot (\eta(x_t, \mu_t(x_t)) + \sup_{\bar{x}_t \in [[X_t|\mu_t(x_t)]]} \eta(\mu_t(x_t), \bar{x}_t)) \leq 2L_{d_t} \cdot \gamma_t =: \epsilon_t$ , where, in the third inequality, we use the triangle inequality. Next, to derive the value of  $\delta_t$ , we expand the LHS of (23) as

$$\begin{aligned} & \mathcal{H}([[\hat{X}_{t+1}|x_t, u_t]], [[\hat{X}_{t+1}|\mu_t(x_t), u_t]]) \\ &= \mathcal{H}(\{\mu_{t+1}(f_t(x_t, u_t, w_t)) \mid w_t \in \mathcal{W}_t\}, \\ & \quad \{\mu_{t+1}(f_t(\bar{x}_t, u_t, w_t)) \mid \bar{x}_t \in [[X_t|\mu_t(x_t)]], w_t \in \mathcal{W}_t\}) \\ &\leq \sup_{w_t \in \mathcal{W}_t} \mathcal{H}(\{\mu_{t+1}(f_t(x_t, u_t, w_t))\}, \\ & \quad \{\mu_{t+1}(f_t(\bar{x}_t, u_t, w_t)) \mid \bar{x}_t \in [[X_t|\mu_t(x_t)]]\}), \end{aligned} \quad (65)$$

where, we use (64) and  $\{\mu_{t+1}(f_t(x_t, u_t, w_t)) \mid w_t \in \mathcal{W}_t\} = \cup_{w_t \in \mathcal{W}_t} \{\mu_{t+1}(f_t(\bar{x}_t, u_t, w_t))\}$ . Using (64) in the RHS of (65),  $\mathcal{H}([[\hat{X}_{t+1}|x_t, u_t]], [[\hat{X}_{t+1}|\mu_t(x_t), u_t]]) \leq \sup_{w_t \in \mathcal{W}_t, \bar{x}_t \in [[X_t|\mu_t(x_t)]]} \eta(\mu_{t+1}(f_t(x_t, u_t, w_t)), \mu_{t+1}(f_t(\bar{x}_t, u_t, w_t))) \leq \sup_{w_t \in \mathcal{W}_t, \bar{x}_t \in [[X_t|\mu_t(x_t)]]} (\eta(\mu_{t+1}(f_t(x_t, u_t, w_t)), f_t(x_t, u_t, w_t)) + \eta(f_t(x_t, u_t, w_t), f_t(\bar{x}_t, u_t, w_t)) + \eta(f_t(\bar{x}_t, u_t, w_t), \mu_{t+1}(f_t(\bar{x}_t, u_t, w_t)))) \leq \gamma_{t+1} + 2L_{f_t} \gamma_t + \gamma_{t+1} =: \delta_t$ .  $\square$

## APPENDIX D – APPROXIMATION BOUNDS FOR PARTIALLY OBSERVED SYSTEMS

In this appendix, we derive the values of  $\epsilon_t$  and  $\delta_t$  for all  $t = 0, \dots, T$ , when an approximate information state is constructed using state quantization for a partially observed system, as described in Subsection IV-C.

**Theorem 6.** Consider a partially observed system with  $Y_t = h_t(X_t, N_t)$  for all  $t = 0, \dots, T$ . Let  $\mu_t : \mathcal{X}_t \rightarrow \hat{\mathcal{X}}_t$  such that  $\sup_{x_t \in \mathcal{X}_t} \eta(x_t, \mu_t(x_t)) \leq \gamma_t$  at each  $t$ . Then,  $\hat{\Pi}_t = \nu_t(\Pi_t)$  is an approximate information state with  $\epsilon_t = 2L_{d_t} \cdot \gamma_t$  and  $\delta_t = 2\gamma_{t+1} + 2L_{\bar{f}_t} \cdot L_{h_{t+1}} \cdot L_{f_t} \cdot \gamma_t$  for all  $t$ , where  $\gamma_{T+1} = 0$ , and where  $L_{d_t}$ ,  $L_{\bar{f}_t}$ ,  $L_{h_{t+1}}$ , and  $L_{f_t}$  are Lipschitz constants for the respective functions in the subscripts.

*Proof.* For all  $t = 0, \dots, T$ , let  $m_t \in [[M_t]]$ ,  $P_t = [[X_t|m_t]] \in \mathcal{P}_t$ , and  $\hat{P}_t = \nu_t(P_t) \in \hat{\mathcal{P}}_t$  be the realizations of the memory  $M_t$ , the conditional range  $\Pi_t$  and the approximate information state  $\hat{\Pi}_t$ , respectively. Note that the conditional range  $P_t$  satisfies (14) and (15) from Definition 1. Next, to derive the value of  $\epsilon_t$ , we write the LHS of (22) using (14) as

$$\begin{aligned} & \left| \sup_{c_t \in [[C_t|m_t, u_t]]} c_t - \sup_{c_t \in [[C_t|\nu_t(P_t), u_t]]} c_t \right| \\ &= \left| \sup_{x_t \in P_t} d_t(x_t, u_t) - \sup_{\bar{x}_t \in [[X_t|\nu_t(P_t)]]} d_t(\bar{x}_t, u_t) \right| \\ &\leq L_{d_t} \cdot \mathcal{H}(P_t, [[X_t|\nu_t(P_t)]]) \\ &\leq L_{d_t} \cdot (\mathcal{H}(P_t, \nu_t(P_t)) + \mathcal{H}(\nu_t(P_t), [[X_t|\nu_t(P_t)]])), \end{aligned} \quad (66)$$

where, in the equality, we use (14); in the first inequality, we use (62) from Lemma 9; and in the second inequality, we use the triangle inequality for the Hausdorff distance. We can expand the first term in the RHS of (66) as  $\mathcal{H}(P_t, \nu_t(P_t)) = \mathcal{H}(P_t, \{\mu_t(x_t) \in \hat{\mathcal{X}}_t \mid x_t \in P_t\}) = \mathcal{H}(\cup_{x_t \in P_t} \{x_t\}, \cup_{x_t \in P_t} \{\mu_t(x_t) \in \hat{\mathcal{X}}_t\}) \leq \sup_{x_t \in P_t} \eta(x_t, \mu_t(x_t)) \leq \gamma_t$ , where we use (64) from Lemma 10 in the first inequality. We can also expand the second term in the RHS of (66) as  $\mathcal{H}(\nu_t(P_t), [[X_t|\nu_t(P_t)]]) = \mathcal{H}(\nu_t(P_t), \{x_t \in \mathcal{X}_t \mid \inf_{\bar{x}_t \in \nu_t(P_t)} \eta(x_t, \bar{x}_t) \leq \gamma_t\}) = \sup_{x_t \in [[X_t|\nu_t(P_t)]]} \inf_{\bar{x}_t \in \nu_t(P_t)} \eta(x_t, \bar{x}_t) \leq \gamma_t$ , where the second equality holds by expanding the Hausdorff distance and noting that  $\nu_t(P_t) \subseteq [[X_t|\nu_t(P_t)]]$ . The proof is complete by substituting the results for both terms in the RHS of (66).

Next, to derive the value of  $\delta_t$ , we note that  $P_t = \sigma_t(m_t)$ . Then, using the triangle inequality in the LHS of (23),

$$\begin{aligned} & \mathcal{H}([[\nu_{t+1}(\Pi_{t+1})|m_t, u_t]], [[\nu_{t+1}(\Pi_{t+1})|\nu_t(\sigma_t(m_t)), u_t]]) \\ &\leq \mathcal{H}([[\nu_{t+1}(\Pi_{t+1})|m_t, u_t]], [[\Pi_{t+1}|m_t, u_t]]) + \\ &\quad \mathcal{H}([[\Pi_{t+1}|m_t, u_t]], [[\Pi_{t+1}\nu_t(\sigma_t(m_t)), u_t]]) + \\ &\quad \mathcal{H}([[\Pi_{t+1}|\nu_t(\sigma_t(m_t)), u_t]], [[\nu_{t+1}(\Pi_{t+1})|\nu_t(\sigma_t(m_t)), u_t]]) \\ &\leq 2\gamma_{t+1} + \mathcal{H}([[\Pi_{t+1}|m_t, u_t]], [[\Pi_{t+1}\nu_t(\sigma_t(m_t)), u_t]]), \end{aligned} \quad (67)$$

where, in the second inequality, we use the fact that  $\mathcal{H}(P_{t+1}, \nu_{t+1}(P_{t+1})) \leq \gamma_{t+1}$ , which was proved above. We can write the second term in the RHS of (67) using (15) from Definition 1 as  $\mathcal{H}([[\Pi_{t+1}|m_t, u_t]], [[\Pi_{t+1}\nu_t(\sigma_t(m_t)), u_t]]) = \mathcal{H}([[\Pi_{t+1}|P_t, u_t]], [[\Pi_{t+1}\nu_t(P_t), u_t]])$ . Furthermore, note that  $[[\Pi_{t+1}|\nu_t(P_t), u_t]] = \{\tilde{P}_{t+1} \in [[\Pi_{t+1}|\tilde{P}_t, u_t]] \mid \tilde{P}_t \in$

$[[\Pi_t|\nu_t(P_t)]]\}$  =  $\cup_{\tilde{P}_t \in [[\Pi_t|\nu_t(P_t)]]} [[\Pi_{t+1}|\tilde{P}_t, u_t]]$ . Next, we use (64) from Lemma 10 to write that

$$\begin{aligned} & \mathcal{H}([[\Pi_{t+1}|P_t, u_t]]), [[\Pi_{t+1}|\nu_t(P_t), u_t]]) \\ &\leq \sup_{\tilde{P}_t \in [[\Pi_t|\nu_t(P_t)]]} \mathcal{H}([[\Pi_{t+1}|P_t, u_t]]), [[\Pi_{t+1}|\tilde{P}_t, u_t]]) \\ &\leq L_{\bar{f}_t} \cdot \sup_{\tilde{P}_t \in [[\Pi_t|\nu_t(P_t)]]} \mathcal{H}([[Y_{t+1}|P_t, u_t]]), [[Y_{t+1}|\tilde{P}_t, u_t]]) \\ &\leq L_{\bar{f}_t} \cdot L_{h_{t+1}} \cdot \sup_{\tilde{P}_t \in [[\Pi_t|\nu_t(P_t)]]} \mathcal{H}([[X_{t+1}|P_t, u_t]]), [[X_{t+1}|\tilde{P}_t, u_t]]), \end{aligned} \quad (68)$$

where in the second inequality, we use the same arguments as in Lemma 2 and the third inequality can be proven by substituting  $Y_{t+1} = h_{t+1}(X_{t+1}, V_{t+1})$  into the equation. We can further expand the third term in the RHS of (68) and use (64) from Lemma (10) to write that  $\sup_{\tilde{P}_t \in [[\Pi_t|\nu_t(P_t)]]} \mathcal{H}([[X_{t+1}|P_t, u_t]]), [[X_{t+1}|\tilde{P}_t, u_t]]) \leq \sup_{\tilde{P}_t \in [[\Pi_t|\nu_t(P_t)]]}, w_t \in \mathcal{W}_t \mathcal{H}(\{f_t(x_t, u_t, w_t) \mid x_t \in P_t\}, \{f_t(x_t, u_t, w_t) \mid x_t \in \tilde{P}_t\}) \leq L_{f_t} \cdot \sup_{\tilde{P}_t \in [[\Pi_t|\nu_t(P_t)]]} \mathcal{H}(P_t, \tilde{P}_t) \leq L_{f_t} \cdot \sup_{\tilde{P}_t \in [[\Pi_t|\nu_t(P_t)]]} (\mathcal{H}(P_t, \nu_t(P_t)) + \mathcal{H}(\nu_t(P_t), \tilde{P}_t)) \leq 2L_{f_t} \cdot \gamma_t$ , where, in the third inequality, we use the triangle inequality and in the fourth inequality we use the fact that for all  $\nu_t(\tilde{P}_t) = \nu_t(P_t)$ , for all  $\tilde{P}_t \in [[\Pi_t|\nu_t(P_t)]]$ .  $\square$

## REFERENCES

- [1] K.-D. Kim and P. R. Kumar, "Cyber-physical systems: A perspective at the centennial," *Proceedings of the IEEE*, vol. 100, no. Special Centennial Issue, pp. 1287–1308, 2012.
- [2] A. A. Malikopoulos, L. E. Beaver, and I. V. Chremos, "Optimal time trajectory and coordination for connected and automated vehicles," *Automatica*, vol. 125, no. 109469, 2021.
- [3] A. Dave, I. V. Chremos, and A. A. Malikopoulos, "Social media and misleading information in a democracy: A mechanism design approach," *IEEE Transactions on Automatic Control*, vol. 67, no. 5, pp. 2633–2639, 2022.
- [4] P. R. Kumar and P. P. Varaiya, *Stochastic Systems: Estimation, Identification, and Adaptive Control*. Englewood Cliffs, NJ: Prentice-Hall, 1986.
- [5] A. A. Malikopoulos, "A duality framework for stochastic optimal control of complex systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 10, pp. 2756–2765, 2016.
- [6] M. Ahmadi, N. Jansen, B. Wu, and U. Topcu, "Control theory meets pomdps: A hybrid systems approach," *IEEE Transactions on Automatic Control*, vol. 66, no. 11, pp. 5191–5204, 2020.
- [7] A. Mahajan, N. C. Martins, M. C. Rotkowitz, and S. Yüksel, "Information structures in optimal decentralized control," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pp. 1291–1306, IEEE, 2012.
- [8] A. Dave and A. A. Malikopoulos, "Decentralized stochastic control in partially nested information structures," *IFAC-PapersOnLine*, vol. 52, no. 20, pp. 97–102, 2019.
- [9] A. Dave and A. A. Malikopoulos, "A dynamic program for a team of two agents with nested information," in *2021 IEEE Conference on Decision and Control (CDC)*, pp. 3768–3773, IEEE, 2021.
- [10] A. Dave, N. Venkatesh, and A. A. Malikopoulos, "Decentralized control of two agents with nested accessible information," in *2022 American Control Conference (ACC)*, pp. 3423–3430, IEEE, 2022.
- [11] A. A. Malikopoulos, "On team decision problems with nonclassical information structures," *IEEE Transactions on Automatic Control*, vol. 68, no. 7, pp. 3915–3930, 2023.
- [12] R. K. Mishra, D. Vasal, and S. Vishwanath, "Decentralized multi-agent reinforcement learning with shared actions," in *2021 55th Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–6, IEEE, 2021.
- [13] K. Zhang, Z. Yang, and T. Başar, "Multi-agent reinforcement learning: A selective overview of theories and algorithms," *Handbook of Reinforcement Learning and Control*, pp. 321–384, 2021.

[14] A. A. Malikopoulos, "Separation of learning and control for cyber-physical systems," *Automatica*, vol. 151, no. 110912, 2023.

[15] H. Kao and V. Subramanian, "Common information based approximate state representations in multi-agent reinforcement learning," in *International Conference on Artificial Intelligence and Statistics*, pp. 6947–6967, PMLR, 2022.

[16] A. A. Malikopoulos, "Combining learning and control in linear systems," in *22nd European Control Conference (ECC)*, 2024 (to appear) (arxiv: 2310.14409).

[17] S. Manner, D. Simester, P. Sun, and J. N. Tsitsiklis, "Bias and variance approximation in value function estimates," *Management Science*, vol. 53, no. 2, pp. 308–322, 2007.

[18] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.

[19] T. Ba  r   and P. Bernhard, *H-infinity optimal control and related minimax design problems: a dynamic game approach*. Springer Science & Business Media, 2008.

[20] M. Rasouli, E. Miehling, and D. Teneketzis, "A scalable decomposition method for the dynamic defense of cyber networks," in *Game Theory for Security and Risk Management*, pp. 75–98, Springer, 2018.

[21] Y. Shoukry, J. Araujo, P. Tabuada, M. Srivastava, and K. H. Johansson, "Minimax control for cyber-physical systems under network packet scheduling attacks," in *Proceedings of the 2nd ACM international conference on High confidence networked systems*, pp. 93–100, 2013.

[22] M. Giuliani, J. Lamontagne, P. Reed, and A. Castelletti, "A state-of-the-art review of optimal reservoir control for managing conflicting demands in a changing world," *Water Resources Research*, vol. 57, no. 12, p. e2021WR029927, 2021.

[23] Q. Zhu and T. Ba  r  , "Robust and resilient control design for cyber-physical systems with an application to power systems," in *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pp. 4066–4071, IEEE, 2011.

[24] K. J.   str  m and P. R. Kumar, "Control: A perspective," *Autom.*, vol. 50, no. 1, pp. 3–43, 2014.

[25] N. Venkatesh, A. Dave, I. Faros, and A. A. Malikopoulos, "Stochastic control with distributionally robust constraints for cyber-physical systems vulnerable to attacks," in *review* (arXiv:2311.03666), 2023.

[26] D. Hadfield-Menell, S. J. Russell, P. Abbeel, and A. Dragan, "Cooperative inverse reinforcement learning," *Advances in neural information processing systems*, vol. 29, 2016.

[27] I. Faros, A. Dave, and A. A. Malikopoulos, "A q-learning approach for adherence-aware recommendations," *IEEE Control Systems Letters*, vol. 7, pp. 3645–3650, 2023.

[28] M. Fatemi, T. W. Killian, J. Subramanian, and M. Ghassemi, "Medical dead-ends and learning to identify high-risk states and treatments," *Advances in Neural Information Processing Systems*, vol. 34, pp. 4856–4870, 2021.

[29] B. Kiumarsi, K. G. Vamvoudakis, H. Modares, and F. L. Lewis, "Optimal and autonomous control using reinforcement learning: A survey," *IEEE transactions on neural networks and learning systems*, vol. 29, no. 6, pp. 2042–2062, 2017.

[30] B. Recht, "A tour of reinforcement learning: The view from continuous control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, pp. 253–279, 2019.

[31] D. P. Bertsekas, "Dynamic programming and suboptimal control: A survey from adp to mpc," *European Journal of Control*, vol. 11, no. 4–5, pp. 310–334, 2005.

[32] D. P. Bertsekas, *Control of uncertain systems with a set-membership description of the uncertainty*. PhD thesis, Massachusetts Institute of Technology, 1971.

[33] O. Hern  ndez-Lerma and J. B. Lasserre, *Discrete-time Markov control processes: basic optimality criteria*, vol. 30. Springer Science & Business Media, 2012.

[34] J. Moon and T. Ba  r  , "Minimax control over unreliable communication channels," *Automatica*, vol. 59, pp. 182–193, 2015.

[35] D. Bertsekas, *Dynamic programming and optimal control: Volume I*, vol. 1. Athena scientific, 2012.

[36] C. H. Papadimitriou and J. N. Tsitsiklis, "The complexity of markov decision processes," *Mathematics of operations research*, vol. 12, no. 3, pp. 441–450, 1987.

[37] P. Bernhard, "A separation theorem for expected value and feared value discrete time control," *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 1, pp. 191–206, 1996.

[38] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

[39] A. Dave and A. A. Malikopoulos, "Structural results for decentralized stochastic control with a word-of-mouth communication," in *2020 American Control Conference (ACC)*, pp. 2796–2801, IEEE, 2020.

[40] J. Subramanian and A. Mahajan, "Approximate information state for partially observed systems," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 1629–1636, IEEE, 2019.

[41] Y. Cong, X. Wang, and X. Zhou, "Rethinking the mathematical framework and optimality of set-membership filtering," *IEEE Transactions on Automatic Control*, vol. 67, no. 5, pp. 2544–2551, 2021.

[42] D. Bertsekas and I. Rhodes, "Sufficiently informative functions and the minimax feedback control of uncertain dynamic systems," *IEEE Transactions on Automatic Control*, vol. 18, no. 2, pp. 117–124, 1973.

[43] M. Gagrani and A. Nayyar, "Decentralized minimax control problems with partial history sharing," in *2017 American Control Conference (ACC)*, pp. 3373–3379, IEEE, 2017.

[44] A. Dave, N. Venkatesh, and A. A. Malikopoulos, "On decentralized minimax control with nested subsystems," in *2022 American Control Conference (ACC)*, pp. 3437–3444, IEEE, 2022.

[45] R. Moiti  , M. Quincampoix, and V. M. Veliov, "Optimal control of discrete-time uncertain systems with imperfect measurement," *IEEE transactions on automatic control*, vol. 47, no. 11, pp. 1909–1914, 2002.

[46] C. Piccardi, "Infinite-horizon minimax control with pointwise cost functional," *Journal of Optimization Theory and Applications*, vol. 78, no. 2, pp. 317–336, 1993.

[47] P. Bernhard, "Minimax - or feared value -  $L_1 / L_\infty$  control," *Theoretical computer science*, vol. 293, no. 1, pp. 25–44, 2003.

[48] M. Gagrani, Y. Ouyang, M. Rasouli, and A. Nayyar, "Worst-case guarantees for remote estimation of an uncertain source," *IEEE Transactions on Automatic Control*, vol. 66, no. 4, pp. 1794–1801, 2020.

[49] M. R. James, J. S. Baras, and R. J. Elliott, "Risk-sensitive control and dynamic games for partially observed discrete-time nonlinear systems," *IEEE transactions on automatic control*, vol. 39, no. 4, pp. 780–792, 1994.

[50] P. Bernhard, "Max-plus algebra and mathematical fear in dynamic optimization," *Set-Valued Analysis*, vol. 8, no. 1, pp. 71–84, 2000.

[51] S. P. Coraluppi and S. I. Marcus, "Risk-sensitive and minimax control of discrete-time, finite-state markov decision processes," *Automatica*, vol. 35, no. 2, pp. 301–309, 1999.

[52] A. Dave, I. Faros, N. Venkatesh, and A. A. Malikopoulos, "Worst-case control and learning using partial observations over an infinite time horizon," in *Proceedings of the 62nd IEEE Conference on Decision and Control (CDC)*, pp. 6014–6019, IEEE, 2023.

[53] A. Dave, N. Venkatesh, and A. A. Malikopoulos, "Approximate information states for worst-case control of uncertain systems," in *Proceedings of the 61th IEEE Conference on Decision and Control (CDC)*, pp. 4945–4950, 2022.

[54] N. B  auerle and U. Rieder, "Partially observable risk-sensitive markov decision processes," *Mathematics of Operations Research*, vol. 42, no. 4, pp. 1180–1196, 2017.

[55] T. Osogami, "Robust partially observable markov decision process," in *International Conference on Machine Learning*, pp. 106–115, PMLR, 2015.

[56] A. Girard, "Controller synthesis for safety and reachability via approximate bisimulation," *Automatica*, vol. 48, no. 5, pp. 947–953, 2012.

[57] Y. Tazaki and J.-i. Imura, "Finite abstractions of discrete-time linear systems and its application to optimal control," *IFAC Proceedings Volumes*, vol. 41, no. 2, pp. 10201–10206, 2008.

[58] F. de Roo and M. Mazo, "On symbolic optimal control via approximate simulation relations," in *52nd IEEE Conference on Decision and Control*, pp. 3205–3210, IEEE, 2013.

[59] G. Reissig and M. Rungger, "Symbolic optimal control," *IEEE Transactions on Automatic Control*, vol. 64, no. 6, pp. 2224–2239, 2018.

[60] E. Gallesteay, M. James, and W. McEneaney, "Max-plus methods in partially observed h/sup/spl infin/infin/control," in *Proceedings of the 38th IEEE Conference on Decision and Control (Cat. No. 99CH36304)*, vol. 3, pp. 3011–3016, IEEE, 1999.

[61] N. Saldi, T. Linder, and S. Y  kse  , "Asymptotic optimality and rates of convergence of quantized stationary policies in stochastic control," *IEEE Transactions on Automatic Control*, vol. 60, no. 2, pp. 553–558, 2014.

[62] S. Meyn, *Control Systems and Reinforcement Learning*. Cambridge University Press, 2022.

[63] S. Levine, A. Kumar, G. Tucker, and J. Fu, "Offline reinforcement learning: Tutorial, review, and perspectives on open problems," *arXiv preprint arXiv:2005.01643*, 2020.

- [64] A. A. Malikopoulos, P. Y. Papalambros, and D. N. Assanis, "A Real-Time Computational Learning Model for Sequential Decision-Making Problems Under Uncertainty," *Journal of Dynamic Systems, Measurement, and Control*, vol. 131, 05 2009. 041010.
- [65] T. M. Moerland, J. Broekens, and C. M. Jonker, "Model-based reinforcement learning: A survey," *arXiv preprint arXiv:2006.16712*, 2020.
- [66] J. Ramírez, W. Yu, and A. Perrusquía, "Model-free reinforcement learning from expert demonstrations: a survey," *Artificial Intelligence Review*, vol. 55, no. 4, pp. 3213–3241, 2022.
- [67] J. Morimoto and K. Doya, "Robust reinforcement learning," *Neural computation*, vol. 17, no. 2, pp. 335–359, 2005.
- [68] M. Heger, "Consideration of risk in reinforcement learning," in *Machine Learning Proceedings 1994*, pp. 105–111, Elsevier, 1994.
- [69] G. Jiang, C.-P. Wu, and G. Cybenko, "Minimax-based reinforcement learning with state aggregation," in *Proceedings of the 37th IEEE Conference on Decision and Control (Cat. No. 98CH36171)*, vol. 2, pp. 1236–1241, IEEE, 1998.
- [70] J. García and F. Fernández, "A comprehensive survey on safe reinforcement learning," *Journal of Machine Learning Research*, vol. 16, no. 1, pp. 1437–1480, 2015.
- [71] A. P. Valadbeigi, A. K. Sedigh, and F. L. Lewis, " $H_\infty$  static output-feedback control design for discrete-time systems using reinforcement learning," *IEEE transactions on neural networks and learning systems*, vol. 31, no. 2, pp. 396–406, 2019.
- [72] S. Chakravorty and D. Hyland, "Minimax reinforcement learning," in *AIAA Guidance, Navigation, and Control Conference and Exhibit*, p. 5718, 2003.
- [73] B. Kiumarsi, F. L. Lewis, and Z.-P. Jiang, " $H_\infty$  control of linear discrete-time systems: Off-policy reinforcement learning," *Automatica*, vol. 78, pp. 144–152, 2017.
- [74] N. Agarwal, B. Bullins, E. Hazan, S. Kakade, and K. Singh, "Online control with adversarial disturbances," in *International Conference on Machine Learning*, pp. 111–119, PMLR, 2019.
- [75] P. Gradu, J. Hallman, and E. Hazan, "Non-stochastic control with bandit feedback," *Advances in Neural Information Processing Systems*, vol. 33, pp. 10764–10774, 2020.
- [76] J. Subramanian, A. Sinha, R. Seraj, and A. Mahajan, "Approximate information state for approximate planning and reinforcement learning in partially observed systems," *Journal of Machine Learning Research*, vol. 23, no. 12, pp. 1–83, 2022.
- [77] L. Yang, K. Zhang, A. Amice, Y. Li, and R. Tedrake, "Discrete approximate information states in partially observable environments," in *2022 American Control Conference (ACC)*, pp. 1406–1413, IEEE, 2022.
- [78] T. W. Killian, H. Zhang, J. Subramanian, M. Fatemi, and M. Ghassemi, "An empirical study of representation learning for reinforcement learning in healthcare," in *Machine Learning for Health*, pp. 139–160, PMLR, 2020.
- [79] G. N. Nair, "A nonstochastic information theory for communication and state estimation," *IEEE Transactions on automatic control*, vol. 58, no. 6, pp. 1497–1510, 2013.
- [80] A. Girard and G. J. Pappas, "Approximation metrics for discrete and continuous systems," *IEEE Transactions on Automatic Control*, vol. 52, no. 5, pp. 782–798, 2007.
- [81] M. F. Barnsley, *Superfractals*. Cambridge University Press, 2006.
- [82] G. Didinsky, *Design of minimax controllers for nonlinear systems using cost-to-come methods*. University of Illinois at Urbana-Champaign, 1995.
- [83] D. Bertsekas, "Convergence of discretization procedures in dynamic programming," *IEEE Transactions on Automatic Control*, vol. 20, no. 3, pp. 415–419, 1975.
- [84] D. Karimi and S. E. Salcudean, "Reducing the hausdorff distance in medical image segmentation with convolutional neural networks," *IEEE Transactions on medical imaging*, vol. 39, no. 2, pp. 499–513, 2019.
- [85] O. U. Aydin, A. A. Taha, A. Hilbert, A. A. Khalil, I. Galinovic, J. B. Fiebach, D. Frey, and V. I. Madai, "On the usage of average hausdorff distance for segmentation performance assessment: hidden error when used for ranking," *European radiology experimental*, vol. 5, pp. 1–7, 2021.
- [86] A. Dave, N. Venkatesh, and A. A. Malikopoulos, "On robust control of partially observed uncertain systems with additive costs," in *Proceedings of the 2023 American Control Conference (ACC)*, pp. 4639–4644, 2023.



**Aditya Dave** (S'18 - SM'23) is a postdoctoral researcher in the School of Civil and Environmental Engineering at Cornell University, Ithaca, NY. He received his PhD from the Department of Mechanical Engineering at the University of Delaware in 2023. He completed his B. Tech. in Mechanical Engineering from the Indian Institute of Technology, Bombay, India, in 2016. His current research interests span several areas, including worst-case control, reinforcement learning, decentralized systems, and mechanism design. He is a member of IEEE.



**Nishanth Venkatesh S** (S'21) is PhD student with the Systems Engineering Program at Cornell University, Ithaca, NY. He received a Master's degree in Robotics from the University of Delaware, Newark, DE in 2021. Prior to his Master's, he received a B. Tech. in Mechanical Engineering from the Indian Institute of Technology, Bombay, India, in 2019. His current research interests span several areas, including worst-case control, reinforcement learning, and decentralized systems. He is a student member of IEEE.



**Andreas A. Malikopoulos** (S'06–M'09–SM'17) received a Diploma in mechanical engineering from the National Technical University of Athens (NTUA), Greece, in 2000. He received M.S. and Ph.D. degrees in mechanical engineering at the University of Michigan, Ann Arbor, Michigan, USA, in 2004 and 2008, respectively. He is a professor in the School of Civil and Environmental Engineering at Cornell University and the director of the Information and Decision Science (IDS) Laboratory. Prior to these appointments, he was the Terri Connor Kelly and John Kelly Career Development Professor in the Department of Mechanical Engineering (2017–2023) and the founding Director of the Sociotechnical Systems Center (2019–2023) at the University of Delaware (UD). Before he joined UD, he was the Alvin M. Weinberg Fellow (2010–2017) in the Energy & Transportation Science Division at Oak Ridge National Laboratory (ORNL), the Deputy Director of the Urban Dynamics Institute (2014–2017) at ORNL, and a Senior Researcher in General Motors Global Research & Development (2008–2010). His research spans several fields, including analysis, optimization, and control of cyber-physical systems (CPS); decentralized stochastic systems; stochastic scheduling and resource allocation; and learning in complex systems. His research aims to develop theories and data-driven system approaches at the intersection of learning and control for making CPS able to realize their optimal operation while interacting with their environment. He has been an Associate Editor of the *IEEE Transactions on Intelligent Vehicles* and *IEEE Transactions on Intelligent Transportation Systems* from 2017 through 2020. He is currently an Associate Editor of *Automatica* and *IEEE Transactions on Automatic Control*, and a Senior Editor of *IEEE Transactions on Intelligent Transportation Systems*. He is a member of SIAM, AAAS, and a Fellow of the ASME.