

Asymptotic Optimality of Finite Model Approximations for Partially Observed Markov Decision Processes With Discounted Cost

Naci Saldi , Serdar Yüksel , and Tamás Linder 

Abstract—We consider finite model approximations of discrete-time partially observed Markov decision processes (POMDPs) under the discounted cost criterion. After converting the original partially observed stochastic control problem to a fully observed one on the belief space, the finite models are obtained through the uniform quantization of the state and action spaces of the belief space Markov decision process (MDP). Under mild assumptions on the components of the original model, it is established that the policies obtained from these finite models are nearly optimal for the belief space MDP, and so, for the original partially observed problem. The assumptions essentially require that the belief space MDP satisfies a mild weak continuity condition. We provide an example and introduce explicit approximation procedures for the quantization of the set of probability measures on the state space of POMDP (i.e., belief space).

Index Terms—Approximations, Markov decision processes, non-linear filtering, quantization, stochastic control.

I. INTRODUCTION

IN PARTIALLY observed Markov decision processes (POMDP) theory, existence of optimal policies in general, have been established via converting the original partially observed stochastic control problem to a fully observed one on the belief space, leading to a belief-Markov decision process (MDP). However, computing an optimal policy for this fully observed model, and so for the original POMDP, using well known dynamic programming algorithms is challenging even if the original system has finite state and action spaces, since the state space of the fully observed model is always uncountable. One way to overcome this difficulty is to compute an

approximately optimal policy instead of a true optimal policy by constructing a reduced model for the fully observed system for which one can apply well known algorithms, such as policy iteration, value iteration, Q -learning, etc. to obtain the optimal policy.

In MDP theory, various methods have been developed to compute near optimal policies by reducing the original problem into a simpler one. A partial list of these techniques is as follows: Approximate dynamic programming, approximate value or policy iteration, simulation-based techniques, neuro-dynamic programming (or reinforcement learning), state aggregation, etc. We refer the reader to [1]–[9] and references therein. However, existing works mostly study systems with discrete (i.e., finite or countable) state and action spaces ([1], [5], [6], [10]), or those that consider general state and action spaces (see, e.g., [7]–[9], [11], [12]) assume in general Lipschitz-type continuity conditions on the transition probability and the one-stage cost function in order to provide a rate of convergence analysis for the approximation error. However, for the fully observed reduction of POMDP, a Lipschitz-type regularity condition on the transition probability is in general prohibitive. Indeed, demonstrating even the arguably most relaxed regularity condition on the transition probability (i.e., weak continuity in state-action variable), is a challenging problem as was recently demonstrated in [13] for general state and action spaces (see also [14], [15] for a control-free setup). Therefore, results developed in the prior literature cannot in general be applied to compute approximately optimal policies for fully observed reduction of POMDP, and so, for the original POMDP.

In [16], [17] we investigated finite action and state approximations of fully observed stochastic control problems with general state and action spaces under the discounted cost and average cost optimality criteria. For the discounted cost case, we showed that the optimal policies obtained from these finite models asymptotically achieve the optimal cost for the original problem under the weak continuity assumption on the transition probability. Here, we apply and properly generalize the results in these papers to obtain approximation results for fully observed reduction of POMDPs, and so, for POMDPs. The versatility of approximation resulting under weak continuity conditions become particularly evident while investigating the applicability of these results to the partially observed case.

Manuscript received March 15, 2018; revised July 30, 2018; accepted March 3, 2019. Date of publication March 25, 2019; date of current version December 27, 2019. This work was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada. Recommended by Associate Editor E. Zhou. (Corresponding author: Serdar Yüksel.)

N. Saldi is with the Department of Natural and Mathematical Sciences, Ozyegin University, Cekmekoy, 34794 Istanbul, Turkey (e-mail: naci.saldi@ozyegin.edu.tr).

S. Yüksel and T. Linder are with the Department of Mathematics and Statistics, Queen's University, Kingston, ON K7L 3N6, Canada (e-mail: yuksel@mast.queensu.ca; linder@mast.queensu.ca).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TAC.2019.2907172

In the literature there exist various, mostly numerical and computational, results for obtaining approximately optimal policies for POMDPs. In the following, we list a number of such related results and comparisons with our paper: 1) A computational algorithm, utilizing structural convexity properties of the value function of belief-MDPs, for the solutions of POMDPs when the state space is continuous and action and measurements are discrete, and with further extensions to continuous action and measurements is developed in [18]. An algorithm which may be regarded as a quantization of the belief space is provided in [19]. However, no rigorous convergence results are given regarding this computational algorithm. 2) Quantization-based algorithms for the belief state, where the state, measurement, and the action sets are finite is presented in [20] and [21]. 3) An explicit quantization method for the set of probability measures containing the belief states, where the state space, unlike in many other contributions in the literature; is continuous and is provided in [22] and [23]. The quantization is done through the approximations as measured by Kullback–Leibler divergence: Kullback–Leibler divergence (or relative entropy) is a very strong pseudodistance measure which is even stronger than total variation (by Pinsker’s inequality [24]), which in turn is stronger than weak convergence. In particular, being able to quantize the space of probability measures with finitely many balls as defined by such a distance measure requires very strict assumptions on the allowable beliefs and it in particular requires, typical equicontinuity conditions (see e.g. [25, Lemma 4.3]). 4) In [26], the authors consider the near optimality of finite-state controllers that are finite-state probabilistic automata taking observations as inputs and producing controls as the outputs. A special case for these type of controllers are the ones that only use finite observation history. A similar finite memory approximation is developed in [27]. 5) In [28] the authors establish finite state approximation schemes for the belief-MDPs under both discounted cost and average cost criteria using concavity properties of the corresponding value function and show that approximate costs can be used as lower bounds for the optimal cost function. A similar finite state approximation is considered in [29], [30] using concavity and convexity properties of the value function for the discounted cost criterion. We refer the reader to the survey papers [31], [32], and the book [33] for further algorithmic and computational procedures for approximating POMDPs.

Contributions of the paper: 1) We show that finite models asymptotically approximate the original POMDP in the sense that the true costs of the policies obtained from these finite models converge to the optimal cost of the original model. The finite models are constructed by discretizing both the state and action spaces of the equivalent fully observed belief space formulation of the POMDP. We establish the result for models with general state and action spaces under mild conditions on the system components. 2) We provide systematic procedures for the quantization of the set of probability measures on the state space of POMDPs which is the state space of belief-MDPs. By choosing an appropriate metric for the weak convergence

topology, any probability measure on the state space can be approximated by the set of probability measures with fixed finite support. Then, one can use the existing methods in the literature for discretizing the probability simplex to obtain finite-state approximations of the belief space. 3) Our rigorous results can be used to justify the novel quantization techniques presented in [19]–[21] as well as a more relaxed version for the results presented in [22] and [23]. In particular, there do not exist approximation results of the generality presented in our paper with asymptotic performance guarantees. We show that provided that the belief space is quantized according to balls generated through metrics that metrize the weak convergence topology (such as Prokhorov, bounded-Lipschitz, or stronger ones, such as the Wasserstein metric), and provided that the action sets are quantized in a uniform fashion, under very weak conditions on the controlled Markov chain (namely the weak continuity of the kernel and total variation continuity of the measurement channel), asymptotic optimality is guaranteed. 4) Our approach also highlights the difficulties of obtaining explicit rates of convergence results for approximation methods for POMDPs. Even in fully observed models, for obtaining explicit rates of convergence, one needs strong continuity conditions of the Lipschitz type, e.g.; [17, Th. 5.1]. As Theorem 1 shows, this is impossible even under quite strong conditions for POMDPs.

The rest of the paper is organized as follows. In Section II we introduce the partially observed stochastic control model and construct the belief space formulation. In Section III we establish the continuity properties that are satisfied by the transition probability of the belief space MDP. In Section IV we construct the finite model approximations and state approximation results. In Section V we provide explicit methods to quantize the set of probability measures on the state space. In Section VI we illustrate our results by considering a numerical example. Section VII concludes the paper.

II. PARTIALLY OBSERVED MARKOV DECISION PROCESSES

A POMDP has the following components: 1) State space \mathbf{X} , action space \mathbf{A} , and observation space \mathbf{Y} , all Borel spaces; 2) $p(\cdot | x, a)$ is the transition probability of the next state given the current state-action pair is (x, a) ; 3) $r(\cdot | x)$ is the observation channel giving the probability of the current observation given the current state variable x ; and 4) the one-stage cost function $c : \mathbf{X} \times \mathbf{A} \rightarrow [0, \infty)$.

To complete the description of the partially observed control model, we must specify how the controller designs its control law at each time step. To this end, define the history spaces $\mathbf{H}_0 = \mathbf{Y}$ and $\mathbf{H}_t = (\mathbf{Y} \times \mathbf{A})^t \times \mathbf{Y}$, $t = 1, 2, \dots$ endowed with their product Borel σ -algebras generated by $\mathcal{B}(\mathbf{Y})$ and $\mathcal{B}(\mathbf{A})$. A policy $\pi = \{\pi_t\}$ is a sequence of stochastic kernels on \mathbf{A} given \mathbf{H}_t . We denote by Π the set of all policies.

According to the Ionescu–Tulcea theorem [34], an initial distribution μ on \mathbf{X} and a policy π define a unique probability measure P_μ^π on $\mathbf{H}_\infty \times \mathbf{X}^\infty$. The expectation with respect to P_μ^π is denoted by \mathbb{E}_μ^π . For any initial distribu-

tion μ and policy π , we can think of the POMDP as a stochastic process $\{x_t, y_t, a_t\}_{t \geq 0}$ defined on the probability space $(\Omega, \mathcal{F}, P_\mu^\pi)$, where $\Omega = \mathbf{H}_\infty \times \mathbf{X}^\infty$, the x_t are \mathbf{X} -valued random variables, the y_t are \mathbf{Y} -valued random variables, the a_t are \mathbf{A} -valued random variables, and they satisfy for all $t \geq 1$

$$P_\mu^\pi(x_0 \in \cdot) = \mu(\cdot)$$

$$\begin{aligned} P_\mu^\pi(x_t \in \cdot | x_{\{0,t-1\}}, y_{\{0,t-1\}}, a_{\{0,t-1\}}) \\ = P_\mu^\pi(x_t \in \cdot | x_{t-1}, a_{t-1}) = p(\cdot | x_{t-1}, a_{t-1}) \end{aligned}$$

$$\begin{aligned} P_\mu^\pi(y_t \in \cdot | x_{\{0,t\}}, y_{\{0,t-1\}}, a_{\{0,t-1\}}) \\ = P_\mu^\pi(y_t \in \cdot | x_t) = r(\cdot | x_t) \end{aligned}$$

$$P_\mu^\pi(a_t \in \cdot | x_{\{0,t\}}, y_{\{0,t\}}, a_{\{0,t-1\}}) = \pi_t(\cdot | y_{\{0,t\}}, a_{\{0,t-1\}})$$

where $x_{\{0,t\}} = (x_0, \dots, x_t)$, $y_{\{0,t\}} = (y_0, \dots, y_t)$, and $a_{\{0,t\}} = (a_0, \dots, a_t)$. We denote by $J(\pi, \mu)$ the discounted cost function of the policy $\pi \in \Pi$ with initial distribution μ , which is given by

$$J(\pi, \mu) := \mathbb{E}_\mu^\pi \left[\sum_{t=0}^{\infty} \beta^t c(x_t, a_t) \right]$$

where $\beta \in (0, 1)$ is the discount factor.

With this notation, the discounted value function of the control problem is defined as

$$J^*(\mu) := \inf_{\pi \in \Pi} J(\pi, \mu).$$

A policy π^* is said to be optimal if $J(\pi^*, \mu) = J^*(\mu)$.

In POMDPs, since the information available to the decision maker is a noisy version of the state, one cannot apply the dynamic programming principle directly as the one-stage cost function depends on the exact state information. A canonical way to overcome this difficulty is converting the original partially observed control problem into a fully observed one by taking the posterior state distributions $\Pr\{x_t \in \cdot | y_0, \dots, y_t, a_0, \dots, a_{t-1}\} \in \mathcal{P}(\mathbf{X})$ ("beliefs" of the observer) as the new state variable and the function $\tilde{c}: \mathcal{P}(\mathbf{X}) \times \mathbf{A} \rightarrow [0, \infty)$, defined by

$$\tilde{c}(z, a) := \int_{\mathbf{X}} c(x, a) z(dx)$$

as the new one-stage cost function. This fully observed MDP is called the belief-MDP and is equivalent to the original POMDP in the sense that for any optimal policy for the belief-MDP, one can construct a policy for the belief-MDP which is optimal. Therefore, finite-model approximation results developed for fully observed MDPs can be applied to the belief-MDP, and so to the POMDP.

III. CONTINUITY PROPERTIES OF BELIEF-MDPS

In this section, we first discuss the continuity properties that are satisfied by or prohibitive for the transition probability of the belief-MDP. Then, we derive the conditions satisfied by the components of the belief-MDP.

A. On the Convergence of Probability Measures

Let \mathbf{E} be a Borel space and let $\mathcal{P}(\mathbf{E})$ denote the family of all probability measure on $(\mathbf{E}, \mathcal{B}(\mathbf{E}))$. A sequence $\{\mu_n\}$ is said to converge to $\mu \in \mathcal{P}(\mathbf{E})$ weakly (resp., setwise) if

$$\int_{\mathbf{E}} g(e) \mu_n(de) \rightarrow \int_{\mathbf{E}} g(e) \mu(de)$$

for all continuous and bounded real function g (resp., for all measurable and bounded real function g).

For any $\mu, \nu \in \mathcal{P}(\mathbf{E})$, the total variation norm is given by

$$\begin{aligned} \|\mu - \nu\|_{TV} &:= 2 \sup_{B \in \mathcal{B}(\mathbf{E})} |\mu(B) - \nu(B)| \\ &= \sup_{f: \|f\|_\infty \leq 1} \left| \int_{\mathbf{E}} f(e) \mu(de) - \int_{\mathbf{E}} f(e) \nu(de) \right| \end{aligned}$$

where the supremum is over all measurable real f such that $\|f\|_\infty = \sup_{e \in \mathbf{E}} |f(e)| \leq 1$. A sequence $\{\mu_n\}$ is said to converge to $\mu \in \mathcal{P}(\mathbf{E})$ in total variation if $\|\mu_n - \mu\|_{TV} \rightarrow 0$. As it is clear from the definitions, total variation convergence implies setwise convergence, which in turn implies weak convergence.

The total variation metric leads to a stringent notion for convergence. For example a sequence of discrete probability measures on a finite-dimensional Euclidean space never converges in total variation to a probability measure which admits a density function with respect to the Lebesgue measure. Setwise convergence also induces a topology which is not easy to work with since the space under this convergence is not metrizable [35, p. 59]. However, the space of probability measures on a Borel space endowed with the topology of weak convergence is itself a Borel space [36]. The bounded-Lipschitz metric ρ_{BL} [37, p.109], for example, can be used to metrize this space

$$\rho_{BL}(\mu, \nu) := \sup_{\|f\|_{BL} \leq 1} \left| \int_{\mathbf{E}} f(e) \mu(de) - \int_{\mathbf{E}} f(e) \nu(de) \right| \quad (1)$$

where

$$\|f\|_{BL} := \|f\|_\infty + \sup_{e \neq e'} \frac{f(e) - f(e')}{d_{\mathbf{E}}(e, e')}$$

and $d_{\mathbf{E}}$ is the metric on \mathbf{E} . Finally, the Wasserstein metric of order one, W_1 , can also be used for compact \mathbf{E} (see [37, Th. 6.9])

$$W_1(\mu, \nu) = \inf_{\eta \in \mathcal{H}(\mu, \nu)} \int_{\mathbf{E} \times \mathbf{E}} d_{\mathbf{E}}(e, e') \eta(de, de')$$

where $\mathcal{H}(\mu, \nu)$ denotes the set of probability measures on $\mathbf{E} \times \mathbf{E}$ with first marginal μ and second marginal ν . Indeed, W_1 can also be used as an upper bound to ρ_{BL} for noncompact \mathbf{E} since W_1 can equivalently be written as [37, Remark 6.5]

$$W_1(\mu, \nu) := \sup_{\|f\|_{Lip} \leq 1} \left| \int_{\mathbf{E}} f(e) \mu(de) - \int_{\mathbf{E}} f(e) \nu(de) \right|$$

where

$$\|f\|_{Lip} := \sup_{e \neq e'} \frac{f(e) - f(e')}{d_{\mathbf{E}}(e, e')}.$$

Comparing this with (1), it follows that:

$$\rho_{BL} \leq W_1. \quad (2)$$

This observation will be utilized for the quantization algorithms on the set of probability measures later in the paper.

B. Belief Space Formulation of POMDPs

As indicated in Section II, any POMDP can be reduced to a (completely observable) MDP [38], [39], whose states are the posterior state distributions or *beliefs* of the observer; that is, the state at time t is

$$z_t := \Pr\{x_t \in \cdot | y_0, \dots, y_t, a_0, \dots, a_{t-1}\} \in \mathcal{P}(\mathbf{X}).$$

In this section, we construct the components of this belief-MDP under some assumptions on the components of the POMDP. Later, we establish the conditions satisfied by the components of the belief-MDP, under which we can apply approximation results in our earlier work [16], [17] to the belief-MDP, and so, to the original POMDP.

To this end, let $v : \mathbf{X} \rightarrow [0, \infty)$ be a continuous moment function in the sense that there exists an increasing sequence of compact subsets $\{K_n\}_{n \geq 1}$ of \mathbf{X} such that

$$\lim_{n \rightarrow \infty} \inf_{x \in \mathbf{X} \setminus K_n} v(x) = \infty.$$

The following assumptions will be imposed on the components of the POMDP.

Assumption 1:

- a) The one-stage cost function c is continuous and bounded.
- b) We have one of the following:
 - 1) the stochastic kernel $p(\cdot | x, a)$ is weakly continuous in $(x, a) \in \mathbf{X} \times \mathbf{A}$, i.e., if $(x_k, a_k) \rightarrow (x, a)$, then $p(\cdot | x_k, a_k) \rightarrow p(\cdot | x, a)$ weakly. The observation channel $r(\cdot | x)$ is continuous in total variation, i.e., if $x_k \rightarrow x$, then $r(\cdot | x_k) \rightarrow r(\cdot | x)$ in total variation.
 - 2) the stochastic kernel $p(\cdot | x, a)$ is total variation continuous in $(x, a) \in \mathbf{X} \times \mathbf{A}$, i.e., if $(x_k, a_k) \rightarrow (x, a)$, then $p(\cdot | x_k, a_k) \rightarrow p(\cdot | x, a)$ in total variation.
- c) \mathbf{A} is compact.
- d) There exists a constant $\lambda \geq 0$ such that

$$\sup_{a \in \mathbf{A}} \int_{\mathbf{X}} v(y) p(dy | x, a) \leq \lambda v(x).$$

- e) The initial probability measure μ satisfies

$$\int_{\mathbf{X}} v(x) \mu(dx) < \infty.$$

We let

$$\mathcal{P}_v(\mathbf{X}) := \left\{ \mu \in \mathcal{P}(\mathbf{X}) : \int_{\mathbf{X}} v(x) \mu(dx) < \infty \right\}.$$

Note that since the probability law of x_t is in $\mathcal{P}_v(\mathbf{X})$, by Assumption 1-d),e), under any policy we have $\Pr\{x_t \in \cdot | y_0, \dots, y_t, a_0, \dots, a_{t-1}\} \in \mathcal{P}_v(\mathbf{X})$ almost everywhere. Therefore, the belief-MDP has state space $\mathbf{Z} = \mathcal{P}_v(\mathbf{X})$ instead of

$\mathcal{P}(\mathbf{X})$, where \mathbf{Z} is equipped with the Borel σ -algebra generated by the topology of weak convergence. The transition probability η of the belief-MDP can be constructed as follows (see also [40]). Let z denote the generic state variable for the belief-MDP. First consider the transition probability on $\mathbf{X} \times \mathbf{Y}$ given $\mathbf{Z} \times \mathbf{A}$

$$R(x \in A, y \in B | z, a) := \int_{\mathbf{X}} \kappa(A, B | x', a) z(dx')$$

where $\kappa(dx, dy | x', a) := r(dy | x) p(dx | x', a)$. Let us disintegrate R as

$$R(dx, dy | z, a) = H(dy | z, a) F(dx | z, a, y).$$

Then, we define the mapping $F : \mathbf{Z} \times \mathbf{A} \times \mathbf{Y} \rightarrow \mathbf{Z}$ as

$$F(z, a, y) = F(\cdot | z, a, y). \quad (3)$$

In the literature, (3) is called the ‘‘nonlinear filtering equation’’ [40]. Note that, for each $t \geq 0$, we indeed have

$$F(z, a, y)(\cdot) = \Pr\{x_{t+1} \in \cdot | z_t = z, a_t = a, y_{t+1} = y\}$$

and

$$H(\cdot | z, a) = \Pr\{Y_{t+1} \in \cdot | z_t = z, a_t = a\}.$$

Then, η can be written as

$$\eta(\cdot | z, a) = \int_{\mathbf{Y}} \delta_{F(z, a, y)}(\cdot) H(dy | z, a)$$

where δ_z denotes the Dirac–delta measure at point z ; that is, $\delta_z(D) = 1$ if $z \in D$ and otherwise it is zero. Recall that the initial point for the belief-MDP is μ ; that is, $z_0 \sim \delta_\mu$, and the one-stage cost function \tilde{c} of the belief-MDP is given by

$$\tilde{c}(z, a) := \int_{\mathbf{X}} c(x, a) z(dx). \quad (4)$$

Hence, the belief-MDP is a fully observed Markov decision process with the components $(\mathbf{Z}, \mathbf{A}, \eta, \tilde{c})$. For the belief-MDP define the history spaces $\tilde{\mathbf{H}}_0 = \mathbf{Z}$ and $\tilde{\mathbf{H}}_t = (\mathbf{Z} \times \mathbf{A})^t \times \mathbf{Z}$, $t = 1, 2, \dots$ and let $\tilde{\Pi}$ denote the set of all policies for the belief-MDP, where the policies are defined in a usual manner. Let $\tilde{J}(\tilde{\pi}, \xi)$ denote the discounted cost function of policy $\tilde{\pi} \in \tilde{\Pi}$ for initial distribution ξ of the belief-MDP.

Notice that any history vector $\tilde{h}_t = (z_0, \dots, z_t, a_0, \dots, a_{t-1})$ of the belief-MDP is a function of the history vector $h_t = (y_0, \dots, y_t, a_0, \dots, a_{t-1})$ of the POMDP. Let us write this relation as $i(h_t) = \tilde{h}_t$. Hence, for a policy $\tilde{\pi} = \{\tilde{\pi}_t\} \in \tilde{\Pi}$, we can define $\pi^{\tilde{\pi}} = \{\pi_t^{\tilde{\pi}}\} \in \Pi$ as

$$\pi_t^{\tilde{\pi}}(\cdot | h_t) := \tilde{\pi}_t(\cdot | i(h_t)). \quad (5)$$

Let us write this as a mapping from $\tilde{\Pi}$ to Π : $\tilde{\Pi} \ni \tilde{\pi} \mapsto i(\tilde{\pi}) = \pi^{\tilde{\pi}} \in \Pi$. It is straightforward to show that the cost functions $\tilde{J}(\tilde{\pi}, \xi)$ and $J(\pi^{\tilde{\pi}}, \mu)$ are the same, where $\xi = \delta_\mu$. One can also prove that (see [38], [39])

$$\inf_{\tilde{\pi} \in \tilde{\Pi}} \tilde{J}(\tilde{\pi}, \xi) = \inf_{\pi \in \Pi} J(\pi, \mu) \quad (6)$$

and furthermore, that if $\tilde{\pi}$ is an optimal policy for the belief-MDP, then $\pi^{\tilde{\pi}}$ is optimal for the POMDP as well. Hence, the

POMDP and the corresponding belief-MDP are equivalent in the sense of cost minimization. Therefore, approximation results developed for MDPs in [16], [17] can be applied to the belief-MDP and so, to the POMDP.

C. Strong and Weak Continuity Properties of the Belief MDP

The stochastic kernel η is said to be *weakly* continuous if $\eta(\cdot | z_k, a_k) \rightarrow \eta(\cdot | z, a)$ weakly, whenever $(z_k, a_k) \rightarrow (z, a)$. The kernel is said to be *strongly* continuous if, for any $z \in \mathbf{Z}$, $\eta(\cdot | z, a_k) \rightarrow \eta(\cdot | z, a)$ setwise, whenever $a_k \rightarrow a$. In general, for the fully observed reduction of a partially observed MDP (POMDP), requiring strong continuity of the transition probability is a strong condition. This is illustrated through the simple example [16, Example 2.1]. The following theorem is a consequence of [41, Th. 3.7], [53, Th. 2], and [16, Example 2.1].

Theorem 1:

- i) Under Assumption 1-(b), the stochastic kernel η for belief-MDP is weakly continuous in (z, a) .
- ii) The stochastic kernel η may not be setwise continuous in action variable a even if the observation channel is continuous in total variation.

Part i) of Theorem 1 implies that the transition probability η of the belief-MDP is weakly continuous under Assumption 1. On the other hand, the continuity of the observation channel in total variation is not enough for the setwise continuity of η .

The above suggests that our earlier results in [16] and [17], which only require weak continuity conditions on the transition kernel of a given MDP, are particularly suitable in developing approximation methods for POMDPs (through their MDP reduction), in both quantizing the action spaces as well as state spaces.

Remark 1: We refer the reader to [41, Th. 3.2(c)], [53, Sec. 4] for more general conditions implying weak continuity of the transition probability η . We also note that, in the uncontrolled setting, [42] and [14] have established similar weak continuity conditions (i.e., the weak-Feller property) of the nonlinear filter process (i.e., the belief process) in continuous time and discrete time, respectively.

Remark 2: Our aim in this paper is to show the asymptotic optimality of finite model approximations under very general conditions. Establishing general regularity conditions that produce an explicit error bound on the approximation as a function of the number of bins in the quantization, and thus obtaining rates of convergence to optimality, is in general prohibitive. To obtain explicit rates of convergence even for fully observed systems, one needs strong continuity conditions of the Lipschitz type on the components of the fully observed reduction of the POMDP, see for example in [17, Th. 5.2]. In view of Theorem 1, this does not seem possible under mild conditions since even setwise continuity does not hold. Indeed, establishing conditions that result in a Lipschitz type regularity behaviour (for the belief-MDP) on the components of the fully observed transition kernel and the measurement models is an interesting future re-

search topic. Theorem 1 suggests that these conditions should be quite restrictive.

Example 1: In this example we consider the following partially observed model

$$\begin{aligned} x_{t+1} &= F(x_t, a_t, v_t) \\ y_t &= H(x_t, w_t), \quad t = 0, 1, 2, \dots \end{aligned} \quad (7)$$

where $\mathbf{X} = \mathbb{R}^n$, $\mathbf{A} \subset \mathbb{R}^m$, and $\mathbf{Y} \subset \mathbb{R}^d$ for some $n, m, d \geq 1$. The noise processes $\{v_t\}$ and $\{w_t\}$ are sequences of independent and identically distributed (i.i.d.) random vectors taking values in $\mathbf{V} = \mathbb{R}^p$ and $\mathbf{W} = \mathbb{R}^l$, respectively, for some $p, l \geq 1$, and they are also independent of each other. In this system, the continuity of F in (x, a) is sufficient to imply the weak continuity of the transition probability p , and no assumptions are needed on the noise process (not even the existence of a density is required). On the other hand, the continuity of the observation channel r in total variation holds, if for any $x \in \mathbf{X}$, the probability measure $r(\cdot | x)$ has a density $g(y, x)$, which is continuous in x , with respect to some reference probability measure m on \mathbf{Y} . This follows from Scheffé's theorem (see, e.g., [43, Th. 16.2]). For instance, this density condition holds for the following type of models:

- 1) In the first model, we have $\mathbf{Y} = \mathbf{W} = \mathbb{R}^d$, $H(x, w) = H(x) + w$, H is continuous, and w has a continuous density g_w with respect to Lebesgue measure.
- 2) In the second case, \mathbf{Y} is countable and $r(y|x)$ is continuous in x for all $y \in \mathbf{Y}$. Therefore, the transition probability η of the belief space MDP, corresponding to the model in (7), is weakly continuous.

Next, we derive conditions satisfied by the components of the belief-MDP under Assumption 1. Note first that $\mathbf{Z} = \bigcup_{m \geq 1} F_m$, where

$$F_m := \left\{ \mu \in \mathcal{P}_v(\mathbf{X}) : \int_{\mathbf{X}} v(x) \mu(dx) \leq m \right\}.$$

Since v is a moment function, each F_m is tight [34, Prop. E.8]. Moreover, each F_m is also closed since v is continuous. Therefore, each F_m is compact with respect to the weak topology. This implies that \mathbf{Z} is a σ -compact Borel space. Note that by [44, Prop. 7.30], the one-stage cost function \tilde{c} of the belief-MDP, which is defined in (4), is in $C_b(\mathbf{Z} \times \mathbf{A})$ under Assumption 1-a). Therefore, the belief-MDP satisfies the following conditions under Assumption 1, which we formally state as a separate assumption.

Assumption 2:

- i) The one-stage cost function \tilde{c} is bounded and continuous.
- ii) The stochastic kernel η is weakly continuous.
- iii) \mathbf{A} is compact and \mathbf{Z} is σ -compact.

IV. FINITE MODEL APPROXIMATIONS

A. Finite-Action Approximation

In this section, we consider finite-action approximation of the belief-MDP and so, the POMDP. For these equivalent models, we obtain an approximate finite-action model as follows. Let

d_A denote the metric on \mathbf{A} . Since \mathbf{A} is assumed compact and thus totally bounded, there exists a sequence of finite sets $\Lambda_n = \{a_{n,1}, \dots, a_{n,k_n}\} \subset \mathbf{A}$ such that for each n ,

$$\min_{i \in \{1, \dots, k_n\}} d_A(a, a_{n,i}) < 1/n \text{ for all } a \in \mathbf{A}.$$

In other words, Λ_n is a $1/n$ -net in \mathbf{A} . The sequence $\{\Lambda_n\}_{n \geq 1}$ is used by the finite-action model to approximate the belief-MDP and the POMDP.

In [16], for MDPs with Borel state and action spaces, we studied the problem of approximating an uncountable action set with a finite one and had established the asymptotic optimality of finite action models for such fully observed MDPs that satisfy a number of technical conditions, in particular, the weak continuity condition of the transition kernel in state and action variables. Given the belief-MDP reduction, [16, Th. 3.2] implies the following result.

Theorem 2: Suppose Assumption 1 (and thus Assumption 2) holds for the POMDP. Then we have

$$\lim_{n \rightarrow \infty} |\tilde{J}_n^*(z) - \tilde{J}^*(z)| = 0 \text{ for all } z \in \mathbf{Z}$$

where \tilde{J}_n^* is the discounted cost value function of the belief-MDP $_n$ with the components $(\mathbf{Z}, \Lambda_n, \eta, \tilde{c})$ and \tilde{J}^* is the discounted cost value function of the belief-MDP with components $(\mathbf{Z}, \mathbf{A}, \eta, \tilde{c})$.

The significance of Theorem 2 is reinforced by the following observation. If we let $\Pi(\Lambda_n)$ to be the set of deterministic policies of the POMDP taking values in Λ_n , then the theorem implies that for any given $\varepsilon > 0$ there exists $n \geq 1$ and $\pi_\varepsilon^* \in \Pi(\Lambda_n)$ such that

$$J(\pi_\varepsilon^*, \mu) < \min_{\pi \in \Pi} J(\pi, \mu) + \varepsilon$$

where $\pi_\varepsilon^* = \pi_{\tilde{f}_n^*}$ [see (5)] and \tilde{f}_n^* is the optimal deterministic stationary policy for the belief-MDP $_n$.

B. Finite-State Approximation

The finite-state model for the belief-MDP is obtained as in [17], by quantizing the set of probability measures $\mathbf{Z} = \mathcal{P}_v(\mathbf{X})$; that is, for each m , we quantize compact set F_m similar to the quantization of \mathbf{A} and represent the rest of the points $\mathbf{Z} \setminus F_m$ by some pseudostate.

If \mathbf{Z} is compact, the index m can be fixed to $m = 1$ in the following:

We let d_Z denote a metric on \mathbf{Z} which metrizes the weak topology. For each $m \geq 1$, since F_m is compact and thus totally bounded, there exists a sequence $(\{z_{n,i}^{(m)}\}_{i=1}^{k_n^{(m)}})_{n \geq 1}$ of finite grids in F_m such that for all $n \geq 1$

$$\min_{i \in \{1, \dots, k_n^{(m)}\}} d_Z(z, z_{n,i}^{(m)}) < 1/n \text{ for all } z \in F_m. \quad (8)$$

Let $\{\mathcal{S}_{n,i}^{(m)}\}_{i=1}^{k_n^{(m)}}$ be a partition of F_m such that $z_{n,i}^{(m)} \in \mathcal{S}_{n,i}^{(m)}$ and

$$\max_{z \in \mathcal{S}_{n,i}^{(m)}} d_Z(z, z_{n,i}^{(m)}) < 1/n \quad (9)$$

for all $i = 1, \dots, k_n^{(m)}$. Choose any $z_{n,k_n^{(m)}+1}^{(m)} \in \mathbf{Z} \setminus F_m$ which is a so-called pseudostate and set $\mathcal{S}_{n,k_n^{(m)}+1}^{(m)} = \mathbf{Z} \setminus F_m$. Let $\mathbf{Z}_n^{(m)} := \{z_{n,1}^{(m)}, \dots, z_{n,k_n^{(m)}}^{(m)}, z_{n,k_n^{(m)}+1}^{(m)}\}$ and define function $Q_n^{(m)} : \mathbf{Z} \rightarrow \mathbf{Z}_n^{(m)}$ by

$$Q_n^{(m)}(z) = z_{n,i}^{(m)} \text{ when } z \in \mathcal{S}_{n,i}^{(m)}.$$

Here $Q_n^{(m)}(z)$ maps z to the representative element of the partition it belongs to.

Remark 3:

- 1) Note that given $\{z_{n,i}^{(m)}\}_{i=1}^{k_n^{(m)}} \subset F_m$ that satisfies (8), one way to obtain the corresponding partition $\{\mathcal{S}_{n,i}^{(m)}\}_{i=1}^{k_n^{(m)}}$ of F_m satisfying (9) as follows. Let us define function $Q_{\text{near}} : F_m \rightarrow \{z_{n,1}^{(m)}, \dots, z_{n,k_n^{(m)}}^{(m)}\}$ as

$$Q_{\text{near}}(z) = \arg \min_{z_{n,i}^{(m)}} d_Z(z, z_{n,i}^{(m)})$$

where ties are broken so that Q_{near} is measurable. In the literature, Q_{near} is often called a nearest neighbor quantizer with respect to ‘‘distortion measure’’ d_Z [45]. Then, Q_{near} induces a partition $\{\mathcal{S}_{n,i}^{(m)}\}_{i=1}^{k_n^{(m)}}$ of the space F_m given by

$$\mathcal{S}_{n,i}^{(m)} = \{z \in F_m : Q_n^{(m)}(z) = z_{n,i}^{(m)}\}$$

which satisfies (9). Although one can construct, in theory, the partition using nearest neighbor sense, it is computationally difficult to find these regions when the original state space \mathbf{X} is uncountable.

- 2) The index n indicates the resolution of the quantizer that is applied to discretize the compact set F_m and index m emphasizes the size of the compact set F_m for which quantization is applied.

Let $\{\nu_n^{(m)}\}$ be a sequence of probability measures on \mathbf{Z} satisfying

$$\nu_n^{(m)}(\mathcal{S}_{n,i}^{(m)}) > 0 \text{ for all } i, n, m. \quad (10)$$

One possible choice for $\nu_n^{(m)}$ is

$$\nu_n^{(m)}(\cdot) = \sum_{i=1}^{k_n^{(m)}+1} \delta_{z_{n,i}^{(m)}}(\cdot).$$

We let $\nu_{n,i}^{(m)}$ be the restriction of $\nu_n^{(m)}$ to $\mathcal{S}_{n,i}^{(m)}$ defined by

$$\nu_{n,i}^{(m)}(\cdot) := \frac{\nu_n^{(m)}(\cdot)}{\nu_n^{(m)}(\mathcal{S}_{n,i}^{(m)})}.$$

The measures $\nu_{n,i}^{(m)}$ will be used to define a sequence of finite-state belief MDPs, denoted as $\text{MDP}_n^{(m)}$, which approximate the belief-MDP. To this end, for each n and m define the one-stage cost function $c_n^{(m)} : \mathbf{Z}_n^{(m)} \times \mathbf{A} \rightarrow [0, \infty)$ and the transition

probability $p_n^{(m)}$ on $Z_n^{(m)}$ given $Z_n^{(m)} \times \mathbf{A}$ by

$$c_n^{(m)}(z_{n,i}^{(m)}, a) := \int_{S_{n,i}^{(m)}} c(z, a) \nu_{n,i}^{(m)}(dz)$$

$$p_n^{(m)}(\cdot | z_{n,i}^{(m)}, a) := \int_{S_{n,i}^{(m)}} Q_n^{(m)} * p(\cdot | z, a) \nu_{n,i}^{(m)}(dz) \quad (11)$$

where $Q_n^{(m)} * p(\cdot | z, a) \in \mathcal{P}(Z_n^{(m)})$ is the pushforward of the measure $p(\cdot | z, a)$ with respect to $Q_n^{(m)}$; that is

$$Q_n^{(m)} * p(y | z, a) = p\left(\left\{z \in Z : Q_n^{(m)}(z) = y\right\} | z, a\right)$$

for all $y \in Z_n^{(m)}$. For each n and m , we define $\text{MDP}_n^{(m)}$ as a Markov decision process with the following components: $Z_n^{(m)}$ is the state space, \mathbf{A} is the action space, $p_n^{(m)}$ is the transition probability, and $c_n^{(m)}$ is the one-stage cost function.

Given the belief-MDP, [17, Th. 3.2]¹ implies the following.

Theorem 3: Suppose Assumption 1 (and thus Assumption 2) holds for the POMDP. Then we have

$$\lim_{n,m \rightarrow \infty} \left| \tilde{J}(f_n^{(m)}, \mu) - \tilde{J}^*(\mu) \right| = 0$$

where $f_n^{(m)}$ is obtained by extending the optimal policy of the $\text{MDP}_n^{(m)}$ to Z . Hence, by the equivalence of POMDPs and belief-MDPs, we also have

$$\lim_{n,m \rightarrow \infty} \left| J(\pi_n^{f_n^{(m)}}, \mu) - J^*(\mu) \right| = 0.$$

Theorem 3 implies that to find a near optimal policy for the POMDP, it is sufficient to compute an optimal policy for the finite-state belief-MDP with sufficiently many states, extend this policy to the original state space of the belief-MDP, and then construct the corresponding policy for the POMDP. Next, we discuss explicit methods to quantize the set of probability measures on \mathbf{X} , that is, the belief-space Z .

V. QUANTIZATION OF THE BELIEF-SPACE

An explicit construction for an application requires a properly defined metric on Z . As stated in Section III-A, one can metrize the set of probability measures defined on a Borel space under the weak topology using various distance measures. Building on this fact, in the following we present explicit methods for the quantization of Z for the cases where \mathbf{X} is finite, a compact subset of a finite-dimensional Euclidean space, or the finite-dimensional Euclidean space itself, and $p(\cdot | a)$ is independent of the state variable x .

A. Construction with Finite X

If the state space is finite with $|\mathbf{X}| = m$, then $Z = \mathcal{P}_v(\mathbf{X}) = \mathcal{P}(\mathbf{X})$, and Z is a simplex in \mathbb{R}^m . In this case, Euclidean distance can be used to metrize Z . Indeed, one can make use of the

¹Although Theorem 3.2 in [17] is proved under the assumption that the state space is locally compact, a careful examination of the proof reveals that σ -compactness of the state space is also sufficient to establish the result.

algorithm in [46] (see also [47]) to quantize Z in a nearest neighbor manner. To this end, for each $n \geq 1$, define

$$Z_n := \left\{ (p_1, \dots, p_m) \in \mathbb{Q}^m : p_i = \frac{k_i}{n}, \sum_{i=1}^m k_i = n \right\} \quad (12)$$

where \mathbb{Q} is the set of rational numbers and $n, k_1, \dots, k_m \in \mathbb{Z}_+$. The set Z_n is called *type lattice* by analogy with the concept of *types* in information theory [48, Ch. 12]. Then, the algorithm that computes the nearest neighbor levels can be described as follows:

Algorithm: Given $z \in Z$, find nearest $y \in Z_n$:

1) Compute values ($i = 1, \dots, m$)

$$k'_i = \left\lfloor nz_i + \frac{1}{2} \right\rfloor \text{ and } n' = \sum_{i=1}^m k'_i.$$

2) If $n' = n$ the nearest y is given by $(\frac{k'_1}{n}, \dots, \frac{k'_m}{n})$. Otherwise, compute the errors

$$\delta_i = k'_i - nz_i$$

and sort them

$$-\frac{1}{2} \leq \delta_{i_1} \leq \delta_{i_2} \leq \dots \leq \delta_{i_m} \leq \frac{1}{2}.$$

3) Let $\Delta = n' - n$. If $\Delta > 0$, set

$$k_{i_j} = \begin{cases} k'_{i_j} & \text{if } j = 1, \dots, m - \Delta - 1 \\ k'_{i_j} - 1 & \text{if } j = m - \Delta, \dots, m. \end{cases}$$

If $\Delta < 0$, set

$$k_{i_j} = \begin{cases} k'_{i_j} + 1 & \text{if } j = 1, \dots, |\Delta| \\ k'_{i_j} & \text{if } j = |\Delta| + 1, \dots, m. \end{cases}$$

Then, the nearest y is given by $(\frac{k_1}{n}, \dots, \frac{k_m}{n})$.

One can also compute the maximum radius of the quantization regions for this algorithm. To this end, let d_∞ and d_p denote respectively the metrics induced by L_∞ and L_p ($p \geq 1$) norms on \mathbb{R}^m , which metrizes the weak topology on Z . Then, we have [46, Prop. 2]

$$b_\infty := \max_{z \in Z} \min_{y \in Z_n} d_\infty(z, y) = \frac{1}{n} \left(1 - \frac{1}{m} \right)$$

$$b_2 := \max_{z \in Z} \min_{y \in Z_n} d_2(z, y) = \frac{1}{n} \sqrt{\frac{a(m-a)}{m}}$$

$$b_1 := \max_{z \in Z} \min_{y \in Z_n} d_1(z, y) = \frac{1}{n} \frac{2a(m-a)}{m}$$

where $a = \lfloor m/2 \rfloor$. Hence, for each $n \geq 1$, the set Z_n is an b_j -net in Z with respect to d_j metric, where $j \in \{\infty, 2, 1\}$.

B. Construction With Compact X

The analysis in the previous subsection shows that a finitely supported measure can be approximated through *type lattices*. Thus, if compactly supported probability measures can be approximated with those having finite support, the analysis in Section V-A yields approximately optimal policies. In the following, we assume that \mathbf{X} is a compact subset of \mathbb{R}^d for some

$d \geq 1$. Then $\mathbf{Z} := \mathcal{P}_v(\mathbf{X}) = \mathcal{P}(\mathbf{X})$ is also compact (under the weak convergence topology) and can be metrized using the Wasserstein metric W_1 (here, in defining W_1 , we use the metric on \mathbf{X} induced by the Euclidean norm $\|\cdot\|$) as discussed in Section III-A.

For each $n \geq 1$, let Q_n be some lattice quantizer [45] on \mathbf{X} such that $\|x - Q_n(x)\| < 1/n$ for all $x \in \mathbf{X}$. Set $\mathbf{X}_n = Q_n(\mathbf{X})$, i.e., the output levels of Q_n (note that \mathbf{X}_n is finite since \mathbf{X} is compact). Then, one can approximate any probability measure in \mathbf{Z} with probability measures in

$$\mathcal{P}(\mathbf{X}_n) := \{\mu \in \mathcal{P}(\mathbf{X}) : \mu(\mathbf{X}_n) = 1\}.$$

Indeed, for any $\mu \in \mathbf{Z}$, we have [49, Th. 2.6]

$$\begin{aligned} \inf_{\mu' \in \mathcal{P}(\mathbf{X}_n)} W_1(\mu, \mu') &\leq \inf_{Q: \mathbf{X} \rightarrow \mathbf{X}_n} \int_{\mathbf{X}} \|x - Q(x)\| \mu(dx) \\ &\leq \int_{\mathbf{X}} \|x - Q_n(x)\| \mu(dx) \leq \frac{1}{n}. \end{aligned}$$

Once this is obtained, we can further approximate the probability measure induced by Q_n via the algorithm introduced in Section V-A with asymptotic performance guarantees. Thus, through a sequence of type lattices \mathbf{Z}_{m_n} as given in (12) with a successively refined support set so that $\mathbf{X}_n \subset \mathbf{X}_{n+1}$ for $n \in \mathbb{N}$ with $m_n = |\mathbf{X}_n|$, one can quantize \mathbf{Z} to obtain a sequence of finite state-action MDPs through (11) leading to Theorem 3.

For some related properties of approximations of probability measures with those with finite support, and the relation to optimal quantization, we refer the reader to [49].

C. Construction With Non-Compact \mathbf{X}

Here we assume that $\mathbf{X} = \mathbb{R}^d$ for some $d \geq 1$ and that Assumption 1 holds for $v(x) = \|x\|^2$. In this case, $\mathbf{Z} := \mathcal{P}_v(\mathbf{X})$ becomes the set of probability measures with finite second moment and F_m is the set of probability measures with finite second moments bounded by m . We endow here \mathbf{Z} with the bounded-Lipschitz metric ρ_{BL} , which metrizes weak convergence (see Section III-A).

We first describe the discretization procedure for F_m . For each $n \geq 1$, set $K_n := [-n, n]^d$ and let q_n denote a lattice quantizer on K_n satisfying

$$\sup_{x \in K} \|x - q_n(x)\| < 1/n.$$

Let \mathbf{X}_n denote the set of output levels of q_n ; that is, $\mathbf{X}_n = q_n(K_n)$. Define

$$Q_n(x) = \begin{cases} q_n(x) & \text{if } x \in K \\ 0 & \text{if } x \in K^c. \end{cases}$$

Then, any measure in F_m can be approximated by probability measures in

$$\mathcal{P}(\mathbf{X}_n) := \{\mu \in \mathcal{P}(\mathbf{X}) : \mu(\mathbf{X}_n) = 1\}.$$

Indeed, for any $\mu \in F_m$, we have

$$\begin{aligned} \inf_{\mu' \in \mathcal{P}(\mathbf{X}_n)} \rho_{BL}(\mu, \mu') &\leq \inf_{\mu' \in \mathcal{P}(\mathbf{X}_n)} W_1(\mu, \mu') \\ &\leq \inf_{Q: \mathbf{X} \rightarrow \mathbf{X}_n} \int_{\mathbf{X}} \|x - Q(x)\| \mu(dx) \\ &\leq \int_{\mathbf{X}} \|x - Q_n(x)\| \mu(dx) \\ &= \int_K \|x - Q_n(x)\| \mu(dx) + \int_{K^c} \|x\| \mu(dx) \\ &\leq \frac{1}{n} + \int_{\{\|x\| > n\}} \|x\|^2 \mu(dx) \frac{1}{n} \\ &\leq \frac{(1+m)}{n}. \end{aligned} \tag{13}$$

In the derivation above, (13) follows from (2). Thus, μ in F_m can be approximated by the $\mu_n \in \mathcal{P}(\mathbf{X}_n)$, which is induced by the quantizer Q_n , with a bound $\rho_{BL}(\mu, \mu_n) \leq (1+m)/n$. Then, similar to Section V-B, we can further approximate probability measure μ_n via the algorithm introduced in Section V-A with again asymptotic performance guarantees by Theorem 3. Thus, analogous to compact case, using a sequence of type lattices \mathbf{Z}_{m_n} as given in (12) with a successively refined support set $\mathbf{X}_n \subset \mathbf{X}_{n+1}$ for $n \in \mathbb{N}$ with $m_n = |\mathbf{X}_n|$, one can quantize \mathbf{Z} to obtain a sequence of finite state-action MDPs through (11).

Remark 4: Note that the complexity of the quantization algorithm scales with the dimension of the space on which it is applied. Therefore, quantizing a belief space when the underlying state space is continuous (compact or noncompact) is in general computationally very demanding since, in this case, the quantized space is infinite dimensional. One way to overcome this issue is to (approximately or exactly) parameterize the reachable beliefs with a low-dimensional parameter as will be discussed in detail in the next section.

D. Construction for Special Models Leading to Quantized Beliefs With Continuous Support

So far, we have obtained quantized beliefs where each such quantized belief measure was supported on a finite set. For some applications, this may not be efficient and it may be more desirable to quantize the measurement space appropriately. For some further applications, a parametric representation of the set of reachable beliefs may be present and the construction of bins may be more immediate through quantizing the parameters in a parametric class. What is essential in such models is that the bins designed to construct the finite belief-MDP correspond to balls which are small under the metrics that metrize the weak convergence as discussed in Section III-A. Therefore, although the belief space is originally infinite dimensional in such models, because of the special parametrization of the reachable beliefs, the complexity of the finite model approximation algorithm can be reduced to the complexity of the quantization of the parameter space. This drastically decreases the complexity of the algorithm as can be seen in the numerical example in Section VI-B. For instance, if we obtain the discretized states

by quantizing the observation space (which can be done if the transition probability only depends on the actions as discussed in the next section), then the complexity of the finite model approximation algorithm scales with the dimension of the observation space. Hence, exact or approximate parametrization of reachable beliefs can drastically reduce the complexity of the algorithm. This is a very efficient way to overcome the *curse-of-dimensionality* in approximation of POMDPs.

1) Quantized Measures Through Quantized Measurements: For this section, we assume that transition probability $p(\cdot|a)$ is independent of the state variable x , $\mathbf{Y} \subset \mathbb{R}^p$ for some $p \geq 1$, and Assumption 1 holds for some v . In the view of Theorem 2, as a preprocessing setup, we quantize the action space \mathbf{A} , where the finite set \mathbf{A}_q represents the output levels of this quantizer. Hence, in the sequel, we assume that the action space is \mathbf{A}_q .

Since $\kappa(dx, dy|a) := r(dy|x) \otimes p(dx|a)$, we have

$$\begin{aligned} R(x \in A, y \in B|z, a) &= \int_{\mathbf{X}} \kappa(A, B|a) z(dx') \\ &= \kappa(A, B|a) \end{aligned}$$

and so, the disintegration of R becomes

$$R(dx, dy|a) = H(dy|a) \otimes F(dx|a, y).$$

Then, η is given by

$$\eta(\cdot|a) = \int_{\mathbf{Y}} \delta_{F(\cdot|a, y)}(\cdot) H(dy|a).$$

This implies that we can take the following set as the state space \mathbf{Z} of the fully observed model instead of $\mathcal{P}_v(\mathbf{X})$:

$$\mathbf{Z} = \{F(\cdot|a, y) : (a, y) \in \mathbf{A}_q \times \mathbf{Y}\}.$$

We endow \mathbf{Z} with the bounded-Lipschitz metric ρ_{BL} . For each $n \geq 1$, set $L_n := [-n, n]^p$ and let l_n denote a lattice quantizer on L_n satisfying

$$\sup_{y \in L_n} \|y - l_n(y)\| < 1/n.$$

Let \mathbf{Y}_n denote the set of output levels of l_n ; that is, $\mathbf{Y}_n = l_n(L_n)$. Define

$$q_n(y) = \begin{cases} l_n(y) & \text{if } y \in L_n \\ 0 & \text{if } y \in L_n^c. \end{cases}$$

Then, finite set $\mathbf{Z}_n \subset \mathbf{Z}$, which is used to quantize \mathbf{Z} , is given by

$$\mathbf{Z}_n = \{F(\cdot|a, y) : (a, y) \in \mathbf{A}_q \times \mathbf{Y}_n\}$$

and the corresponding quantizer $Q_n : \mathbf{Z} \rightarrow \mathbf{Z}_n$ is defined as follows: Given $z = F(\cdot|a, y)$, we define

$$Q_n(z) = F(\cdot|a, q_n(y)).$$

Note that to use Q_n for constructing finite models, we have to obtain an upper bound on the ρ_{BL} -distance between z and $Q_n(z)$. This can be achieved under various assumptions on the system components. One such assumption is the following: 1) $\mathbf{X} = \mathbb{R}^d$ for some $d \geq 1$, 2) \mathbf{Y} is compact, 3) $p(dx|a) = g_p(x|a)m(dx)$ and $r(dy|x) = g_r(y|x)m(dy)$, 4) g_r is Lipschitz continuous with Lipschitz constant K_r , $g_r > \theta$ for

some $\theta > 0$, and $\sup_{\{(y,x) \in \mathbf{Y} \times \mathbf{X}\}} |g_r(y, x)| =: \|g_r\| < \infty$. Since \mathbf{Y} is compact, there exists $\epsilon(n)$ for each n such that $\epsilon(n) \rightarrow 0$ as $n \rightarrow \infty$ and $\|y - q_n(y)\| \leq \epsilon(n)$ for all $y \in \mathbf{Y}$. Under the above assumptions, we have

$$F(dx|a, y) = f(x|a, y)m(dx)$$

where

$$f(x|a, y) = \frac{g_r(y|x)g_p(x|a)}{\int_{\mathbf{X}} g_r(y|x)g_p(x|a)m(dx)}.$$

Since the bounded-Lipschitz metric ρ_{BL} is upper bounded by the total variation distance, we obtain

$$\begin{aligned} \rho_{BL}(z, Q_n(z)) &\leq \|z - Q_n(z)\|_{TV} \\ &= \int_{\mathbf{X}} |f(x|a, y) - f(x|a, q_n(y))| m(dx) \\ &\leq \frac{2\|g_r\|K_r}{\theta^2} \|y - q_n(y)\| \leq \frac{2\|g_r\|K_r}{\theta^2} \epsilon(n). \end{aligned}$$

Hence, Q_n is a legitimate quantizer for constructing the finite models. Section VI-B exhibits another example where we have such an upper bound.

2) Construction From a Parametrically Represented Class:

For some applications, the set of belief measures can be first approximated by some parametric class of measures, where parameters belong to some low-dimensional space [23], [50], [51]. For instance, in [23], densities of belief measures are projected onto exponential family of densities using the *Kullback–Leibler (KL) divergence*, where it was assumed that projected beliefs are close enough to true beliefs in terms of cost functions. In [51], densities are parameterized by unimodal Gaussian distributions and parameterized MDP are solved through Monte Carlo simulation-based method. In [50], densities are represented by sufficient statistics, and in particular represented by Gaussian distributions, and the parameterized MDP is solved through fitted value iteration algorithm. However, among these works, only the [23] develop rigorous error bounds for their algorithms using the KL divergence and the other works do not specify distance measures to quantify parametric representation approximations.

In these methods, if the parameterized beliefs are good enough to represent true beliefs as it was shown in [23], then the method presented in the earlier sections (of first quantizing the state space, and then quantizing the beliefs on the state space) may not be necessary and one can, by quantizing the parameters for the class of beliefs considered, directly construct the finite belief-MDP. As noted earlier, what is essential in such methods is that the bins designed to construct the finite belief-MDP correspond to balls which are small under the metrics that metrize the weak convergence as discussed in Section III-A. This possible if the projected beliefs are proved to be close to the true beliefs with respect to some metric that generates the weak topology or with respect to some (pseudo) distance which is stronger than weak topology. For instance, since convergence in KL-divergence is stronger than weak convergence, the projected beliefs constructed in [23] indeed satisfies this requirement. Hence, one can apply our results to conclude the

convergence of the reduced model to the original model in [23]. As noted earlier, relative entropy is a very strong pseudodistance measure which is even stronger than total variation (by Pinsker's inequality [24]) and for being able to quantize a set of probability measures with finitely many balls as defined by such a distance measure requires very strict assumptions on the allowable beliefs and it in particular requires, typical equicontinuity conditions (see e.g., [25, Lemma 4.3]). In turn, it is in general necessary to assume that transition probability and observation channel have very strong regularity conditions.

VI. NUMERICAL EXAMPLES

A. Example with Finite X

We consider a machine repair problem in order to illustrate our results numerically for finite state POMDPs. In this model, we have $\mathbf{X} = \mathbf{A} = \mathbf{Y} = \{0, 1\}$ with the following interpretation:

$$x_t = \begin{cases} 1 & \text{machine is working at time } t \\ 0 & \text{machine is not working at time } t \end{cases}$$

$$a_t = \begin{cases} 1 & \text{machine is being repaired at time } t \\ 0 & \text{machine is not being repaired at time } t \end{cases}$$

κ is the probability that the machine repair was successful given an initial "not working" state

$$\Pr\{x_{t+1} = 1 | x_t = 0, a_t = 1\} = \kappa.$$

Finally, the probability that the machine does not break down in one time step is denoted by α

$$\Pr\{x_{t+1} = 0 | x_t = 1, a_t = 0\} = \alpha.$$

The probability that the measured state is not the true state is given by ε ; that is

$$\Pr\{y_t = 0 | x_t = 1\} = \Pr\{y_t = 1 | x_t = 0\} = \varepsilon.$$

The one-stage cost function for this model is given by

$$c(x, a) = \begin{cases} R + E & x = 0 \text{ and } a = 1 \\ E & x = 0 \text{ and } a = 0 \\ 0 & x = 1 \text{ and } a = 0 \\ R & x = 1 \text{ and } a = 1 \end{cases}$$

where R is defined to be the cost of repair and E is the cost incurred by a broken machine. The cost function to be minimized is the discounted cost function with a discount factor β . In order to find the approximately optimal policies, we first construct the belief space formulation of the above model. Note that the state space of the belief space model is the interval $[0, 1]$. Hence, we can use uniform quantization on $[0, 1]$ to obtain the finite model. For the numerical results, we use the following parameters: $\varepsilon = 0.17$, $\kappa = 0.9$, $\alpha = 0.9545$, and $\beta = 0.3$. We selected 20 different values for the number n of grid points to discretize $[0, 1]$: $n = 10, 20, 30, \dots, 200$. The grid points are chosen uniformly. For each n , the finite state models are constructed as in [17, Sec. 2].




Fig. 1. Optimal costs of the finite models when the initial state is $x = 1$.

Fig. 1 shows the graph of the value functions of the finite models corresponding to the different values of n (number of grid points), when the initial state is $x = 1$. The resulting graph suggest convergence of the value functions (to the value function of the original model). Note that the rigorous justification of eventual convergence is given in the paper; see Theorems 2 and 3.

B. Example With Compact X

In this example we consider the following model:

$$x_{t+1} = \exp\{-\theta_1 a_t + v_t\}, \quad t = 0, 1, 2, \dots \quad (15)$$

$$y_t = x_t + \xi_t, \quad t = 0, 1, 2, \dots \quad (16)$$

where $\theta_1 \in \mathbb{R}_+$, x_t is the state at t , and a_t is the action at t . The one-stage "reward" function is $u(x_t - a_t)$, where u is some utility function. In this model, the goal is to maximize the discounted reward. This model is the modified and partially observed version of the population growth model in [34, Sec. 1.3].

The state and action spaces are $\mathbf{X} = \mathbf{A} = [0, L]$, for some $L \in \mathbb{R}_+$, and the observation space is $\mathbf{Y} = [0, K]$ for some $K \in \mathbb{R}_+$. Since θ_1 is merely a constant, by taking $[0, \frac{L}{\theta_1}]$ as our new action space, instead of dynamics in (15) we can write the dynamics of the state as

$$x_{t+1} = \exp\{-a_t + v_t\}, \quad t = 0, 1, 2, \dots$$

The noise processes $\{v_t\}$ and $\{\xi_t\}$ are sequences of i.i.d. random variables which have common densities g_v supported on $[0, \lambda]$ and g_ξ supported on $[0, \tau]$, respectively. Therefore, the transition probability $p(\cdot | x, a)$ is given by

$$p(D|x, a) = \int_D g_v(\log(v) + a) \frac{1}{v} m(dv)$$

for all $D \in \mathcal{B}(\mathbb{R})$ and the observation kernel $r(\cdot | x)$ is given by

$$r(B|x, a) = \int_B g_\xi(\xi - x) m(d\xi)$$

for all $B \in \mathcal{B}(\mathbb{R})$. To make the model consistent, we must have $\exp\{-a + v\} \in [0, L]$ for all $(a, v) \in [0, L] \times [0, \lambda]$. We assume that g_v and g_ξ are uniform probability density functions; that is, $g_v = \frac{1}{\lambda}$ on $[0, \lambda]$ and $g_\xi = \frac{1}{\tau}$ on $[0, \tau]$. Hence, Assumption 1 holds for this model with $v(x) = 1$.

In the view of Theorem 2, as a preprocessing setup, we quantize the action space \mathbf{A} , where the finite set $\mathbf{A}_q = \{a_1, a_2, \dots, a_q\}$ represents the output levels of this quantizer with $0 < a_1 < a_2 < \dots < a_q$. In the remainder of this example we assume that the action space is \mathbf{A}_q .

We now obtain the stochastic kernels $H(\cdot|z, a)$ and $F(\cdot|z, a, y)$ that describe the transition probability η of the reduced MDP. Indeed, we have $H(dy|z, a) = h(y|a)m(dy)$ where $h(y|a)$ is given by

$$h(y|a) = \int_{\mathbf{X}} g_{\xi}(y-x)g_v(\log(x)+a)\frac{1}{x}m(dx).$$

Similarly, we have $F(dx|z, a, y) = f(dx|a, y)m(dx)$ where $f(x|z, a, y)$ is given by

$$f(x|a, y) = \frac{g_{\xi}(y-x)g_v(\log(x)+a)\frac{1}{x}}{\int_{\mathbf{X}} g_{\xi}(y-x)g_v(\log(x)+a)\frac{1}{x}m(dx)}. \quad (17)$$

Hence, for any (z, a) , the transition probability $\eta(\cdot|z, a)$ has a support on the set of probability measures on \mathbf{X} having densities given by (17). This implies that we can take the following set as the state space \mathbf{Z} of the fully observed model instead of $\mathcal{P}(\mathbf{X})$:

$$\mathbf{Z} = \{f(x|a, y)m(dx) : (a, y) \in \mathbf{A} \times \mathbf{Y} \text{ and } f \text{ as in (17)}\}.$$

For each n , let q_n denote the uniform quantizer on \mathbf{Y} having n output levels; that is, $q_n : \mathbf{Y} \rightarrow \{y_1, \dots, y_n\} =: \mathbf{Y}_n \subset \mathbf{Y}$ where $y_j = (j - \frac{1}{2})\Delta_n$, $j = 1, \dots, n$, and

$$q_n^{-1}(y_j) = \left[y_j - \frac{\Delta_n}{2}, y_j + \frac{\Delta_n}{2} \right)$$

where $\Delta_n = \frac{K}{n}$. We define

$$\mathbf{Z}_n := \{f(x|a, y)m(dx) \in \mathbf{Z} : (a, y) \in \mathbf{A}_q \times \mathbf{Y}_n\}.$$

Then, the quantizer $Q_n : \mathbf{Z} \rightarrow \mathbf{Z}_n$, which is used to construct the finite model, is defined as follows: Given $z = f(x|a, y)m(dx)$, we define $Q_n(z) = f(x|a, q_n(y))m(dx)$.

To be able to use Q_n for constructing finite models, we need to obtain an upper bound on the ρ_{BL} -distance between z and $Q_n(z)$. Indeed, one can prove that

$$\rho_{BL}(z, Q_n(z)) \leq \|z - Q_n(z)\|_{TV} \leq M\Delta_n \quad (18)$$

for some constant M . The utility function u is taken to be quadratic function; i.e., $u(t) = t^2$.

For the numerical results, we use the parameters as: $\lambda = 1$, $\tau = 0.5$, and $\beta = 0.2$. As a preprocessing setup, we uniformly discretize the action space \mathbf{A} by using 20 grid points. Then, we select 99 different values for the number n of grid points to discretize the state space \mathbf{Z} using the quantizer Q_n , where n varies from 29 to 1436.

We use the value iteration algorithm to compute the value functions of the finite models. Since the components of the transition matrices and the cost function matrix can be computed analytically in this example, complexity of constructing these matrices is polynomial in the size of the action space (i.e., $|\mathbf{A}_q|$) and state space (i.e., $|\mathbf{Z}_n|$). In addition, it is known that the computation of an optimal policy using value iteration algorithm is also polynomial in $|\mathbf{A}_q|$, $|\mathbf{Z}_n|$, $\frac{1}{1-\beta}$, and the maximum number

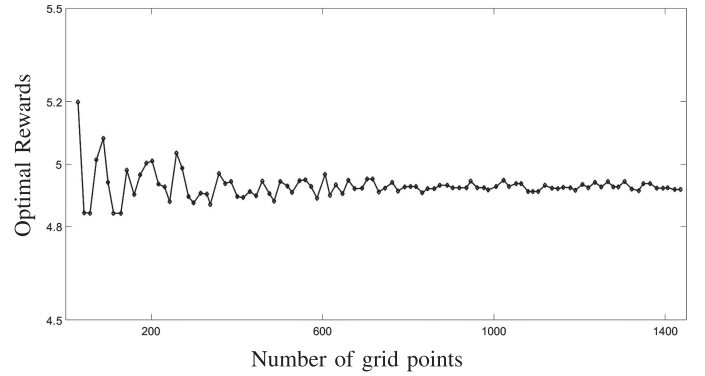


Fig. 2. Optimal rewards of the finite models when the initial state is $x = 2$.

of bits B needed to represent any numerator or denominator of β or one of the components of the transition matrices or cost function matrix [52]. Therefore, the computation of an optimal policy for any finite model with n grid points is polynomial in $|\mathbf{A}_q|$, $|\mathbf{Z}_n|$, $\frac{1}{1-\beta}$, and B .

The simulation was implemented by using MATLAB and it took 411.75 s using an HP EliteDesk 800G2 SFF desktop computer with CPU Intel Core i7-6700 3.4 GHz 4-core. Fig. 2 displays the graph of the value functions corresponding to the different values for the number of grid points when the initial state is $x = 2$. The resulting graph suggests convergence of the value functions; the rigorous justification of convergence is given in Theorems 2 and 3.

VII. CONCLUDING REMARKS

We studied the approximation of discrete-time POMDPs under the discounted cost criterion. An essential observation was that establishing strong continuity properties for the reduced (belief) model is quite difficult for general state and action models, whereas weak continuity can be established under fairly mild conditions on the transition kernel of the original model and the measurement equations. This allowed us to apply our prior approximation results [16], [17], developed under weak continuity conditions, to partially observed models. Extending the analysis to average cost problems is a future task.

ACKNOWLEDGMENT

The authors would like to thank Prof. E. Feinberg for general discussions related to the regularity properties of POMDPs. We acknowledge M. Galal, M. Gaskin, I. Harbell, and D. Kao (all senior students in the Mathematics and Engineering program at Queen's University) for their help with the numerical computations.

REFERENCES

- [1] B. Fox, "Finite-state approximations to denumerable state dynamic programs," *J. Math. Anal. Appl.*, vol. 34, pp. 665–670, 1971.
- [2] W. Whitt, "Approximations of dynamic programs I," *Math. Oper. Res.*, vol. 3, no. 3, pp. 231–243, 1978.
- [3] H. Langen, "Convergence of dynamic programming models," *Math. Oper. Res.*, vol. 6, no. 4, pp. 493–512, Nov. 1981.

- [4] D. Bertsekas and J. Tsitsiklis, *Neuro-Dynamic Programming*. Belmont, MA, USA, Athena Scientific, 1996.
- [5] R. Ortner, "Pseudometrics for state aggregation in average reward Markov decision processes," in *Algorithmic Learning Theory*, M. Hutter, R.A. Servedio, and E. Takimoto, Eds. Berlin, Germany: Springer-Verlag, 2007.
- [6] D. White, "Finite-state approximations for denumerable state infinite horizon discounted Markov decision processes," *J. Math. Anal. Appl.*, vol. 74, pp. 292–295, 1980.
- [7] D. Bertsekas, "Convergence of discretization procedures in dynamic programming," *IEEE Trans. Autom. Control*, vol. 20, no. 3, pp. 415–419, Jun. 1975.
- [8] F. Dufour and T. Prieto-Rumeau, "Finite linear programming approximations of constrained discounted Markov decision processes," *SIAM J. Control Optim.*, vol. 51, no. 2, pp. 1298–1324, 2013.
- [9] F. Dufour and T. Prieto-Rumeau, "Approximation of average cost Markov decision processes using empirical distributions and concentration inequalities," *Stochastics*, vol. 87, no. 2, pp. 273–307, 2014.
- [10] B. Roy, "Performance loss bounds for approximate value iteration with state aggregation," *Math. Oper. Res.*, vol. 31, no. 2, pp. 234–244, May 2006.
- [11] F. Dufour and T. Prieto-Rumeau, "Approximation of Markov decision processes with general state space," *J. Math. Anal. Appl.*, vol. 388, pp. 1254–1267, 2012.
- [12] C.-S. Chow and J. N. Tsitsiklis, "An optimal one-way multigrid algorithm for discrete-time stochastic control," *IEEE Trans. Autom. Control*, vol. 36, no. 8, pp. 898–914, Aug. 1991.
- [13] E. Feinberg, P. Kasyanov, and N. Zadioanchuk, "Average cost Markov decision processes with weakly continuous transition probabilities," *Math. Oper. Res.*, vol. 37, no. 4, pp. 591–607, Nov. 2012.
- [14] A. Budhiraja, "On invariant measures of discrete time filters in the correlated signal-noise case," *Ann. Appl. Probab.*, vol. 12, no. 3, pp. 1096–1113, 2002.
- [15] D. Crisan and A. Doucet, "A survey of convergence results on particle filtering methods for practitioners," *IEEE Trans. Signal Process.*, vol. 50, no. 3, pp. 736–746, Mar. 2002.
- [16] N. Saldi, S. Yüksel, and T. Linder, "Near optimality of quantized policies in stochastic control under weak continuity conditions," *J. Math. Anal. Appl.*, vol. 435, pp. 321–337, 2016.
- [17] N. Saldi, S. Yüksel, and T. Linder, "On the asymptotic optimality of finite approximations to Markov decision processes with Borel spaces," *Math. Oper. Res.*, vol. 42, no. 4, pp. 945–978, Mar. 2017.
- [18] J. M. Porta, N. Vlassis, M. T. J. Spaan, and P. Poupart, "Point-based value iteration for continuous POMDPs," *J. Mach. Learn. Res.*, vol. 7, pp. 2329–2367, 2006.
- [19] N. Vlassis and M. T. J. Spaan, "Perseus: Randomized point-based value iteration for POMDPs," *J. Artif. Intell. Res.*, vol. 24, pp. 195–220, 2005.
- [20] T. Smith and R. Simmons, "Point-based POMDP algorithms: Improved analysis and implementation," 2012, arXiv:1207.1412.
- [21] J. Pineau, G. Gordon, and S. Thrun, "Anytime point-based approximations for large POMDPs," *J. Artif. Intell. Res.*, vol. 27, pp. 335–380, 2006.
- [22] E. Zhou, M. C. Fu, and S. I. Marcus, "A density projection approach to dimension reduction for continuous-state POMDPs," in *Proc. 47th IEEE Conf. Decis. Control*, 2008, pp. 5576–5581.
- [23] E. Zhou, M. C. Fu, and S. I. Marcus, "Solving continuous-state POMDPs via density projection," *IEEE Trans. Autom. Control*, vol. 55, no. 5, pp. 1101–1116, May 2010.
- [24] R. M. Gray, *Entropy and Information Theory*. New York, NY, USA: Springer-Verlag, 1990.
- [25] S. Yüksel and T. Linder, "Optimization and convergence of observation channels in stochastic control," *SIAM J. Control Optim.*, vol. 50, pp. 864–887, 2012.
- [26] H. Yu and D. Bertsekas, "On near optimality of the set of finite-state controllers for average cost POMDP," *Math. Oper. Res.*, vol. 33, no. 1, pp. 1–11, Feb. 2008.
- [27] C. White and W. Scherer, "Finite-memory suboptimal design for partially observed Markov decision processes," *Oper. Res.*, vol. 42, no. 3, pp. 439–455, 1994.
- [28] H. Yu and D. Bertsekas, "Discretized approximations for POMDP with average cost," in *Proc. Conf. Uncertainty Artif. Intell.*, 2004, pp. 619–627.
- [29] W. Lovejoy, "Computationally feasible bounds for partially observed Markov decision processes," *Oper. Res.*, vol. 39, no. 1, pp. 162–175, 1991.
- [30] R. Zhou and E. Hansen, "An improved grid-based approximation algorithm for POMDPs," in *Proc. Int. J. Conf. Artif. Intell.*, 2001, pp. 707–714.
- [31] W. Lovejoy, "A survey of algorithmic methods for partially observed Markov decision processes," *Ann. Oper. Res.*, vol. 28, pp. 47–66, 1991.
- [32] C. White, "A survey of solution techniques for the partially observed Markov decision process," *Ann. Oper. Res.*, vol. 32, pp. 215–230, 1991.
- [33] V. Krishnamurthy, *Partially Observed Markov Decision Processes: From Filtering to Controlled Sensing*. New York, NY, USA: Cambridge Univ. Press, 2016.
- [34] O. Hernández-Lerma and J. Lasserre, *Discrete-Time Markov Control Processes: Basic Optimality Criteria*. New York, NY, USA: Springer, 1996.
- [35] J. K. Ghosh and R. V. Ramamoorthi, *Bayesian Nonparametrics*. New York, NY, USA: Springer, 2003.
- [36] P. Billingsley, *Convergence of Probability Measures*, 2nd ed. New York, NY, USA: Wiley, 1999.
- [37] C. Villani, *Optimal Transport: Old and New*. Berlin, Germany: Springer, 2009.
- [38] A. Yushkevich, "Reduction of a controlled Markov model with incomplete data to a problem with complete information in the case of Borel state and control spaces," *Theory Prob. Appl.*, vol. 21, pp. 153–158, 1976.
- [39] D. Rhenius, "Incomplete information in Markovian decision models," *Ann. Statist.*, vol. 2, pp. 1327–1334, 1974.
- [40] O. Hernández-Lerma, *Adaptive Markov Control Processes*. New York, NY, USA: Springer-Verlag, 1989.
- [41] E. Feinberg, P. Kasyanov, and M. Zgurovsky, "Partially observable total-cost Markov decision process with weakly continuous transition probabilities," 2014, arXiv:1401.2168.
- [42] A. G. Bhatt, A. Budhiraja, and R. L. Karandikar, "Markov property and ergodicity of the nonlinear filter," *SIAM J. Control Optim.*, vol. 39, no. 3, pp. 928–949, 2000.
- [43] P. Billingsley, *Probability and Measure*, 3rd ed. New Delhi, India: Wiley, 1995.
- [44] D. P. Bertsekas and S. E. Shreve, *Stochastic Optimal Control: The Discrete Time Case*. New York, NY, USA: Academic, 1978.
- [45] G. Gray and D. Neuhoff, "Quantization," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2325–2383, Oct. 1998.
- [46] Y. A. Reznik, "An algorithm for quantization of discrete probability distributions," in *Proc. Data Compression Conf.*, 2011, pp. 333–342.
- [47] G. Böcherer and B. C. Geiger, "Optimal quantization for distribution synthesis," *IEEE Trans. Inf. Theory*, vol. 62, no. 11, pp. 6162–6172, Nov. 2016.
- [48] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley, 1991.
- [49] W. Kreitmeier, "Optimal vector quantization in terms of Wasserstein distance," *J. Multivariate Anal.*, vol. 102, no. 8, pp. 1225–1239, 2011.
- [50] A. Brooks, A. Makarenko, S. Williams, and H. Durrant-Whyte, "Parametric POMDPs for planning in continuous state spaces," *Robot. Auton. Syst.*, vol. 54, no. 11, pp. 887–897, 2006.
- [51] A. Brooks and S. Williams, "A Monte Carlo update for parametric POMDPs," *Robot. Res.*, vol. 66, pp. 213–223, 2011.
- [52] M. Littman, T. Dean, and L. Kaelbling, "On the complexity of solving Markov decision problems," in *Proc. 11th Conf. Uncertainty Artif. Intell.*, 1995, pp. 394–402.
- [53] A. D. Kara and N. Saldi and S. Yüksel, "Weak Feller Property of Non-linear Filters," 2019, arXiv preprint arXiv: 1812.05509.



Naci Saldi (S'14–M'16) received the B.Sc. and M.S. degrees in electrical and electronics engineering from Bilkent University, Ankara, Turkey in 2008 and 2010, respectively and the Ph.D. degree from the Department of Mathematics and Statistics, Queen's University, Kingston, ON, Canada in 2015.

He was a Postdoctoral Researcher with the University of Illinois at Urbana-Champaign, Champaign, IL, USA, before joining the Department of Natural and Mathematical Sciences at Ozyegin University, Çekmeköy, Turkey as an Assistant Professor. His research interests include stochastic and decentralized control, source coding, mean-field games, and applied probability.



Serdar Yüksel (S'02–M'11) received the B.Sc. degree in electrical and electronics engineering from Bilkent University, Ankara, Turkey in 2001; the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Illinois at Urbana-Champaign, Champaign, IL, USA in 2003 and 2006, respectively.

He was a Postdoctoral Researcher at Yale University, New Haven, CT, USA for a year before joining Queen's University, Kingston, ON, Canada, as an Assistant Professor of mathematics and engineering with the Department of Mathematics and Statistics, where he is now an Associate Professor. His research interests include stochastic control, decentralized control, information theory, and applied probability.

He was the recipient of the 2013 CAIMS/PIMS Early Career Award in Applied Mathematics. He is an Associate Editor for the IEEE TRANSACTIONS ON AUTOMATIC CONTROL and *Automatica*.



Prof. Tamás Linder (S'92–M'93–SM'00–F'13) received the M.S. degree in electrical engineering from the Technical University of Budapest, Budapest, Hungary, in 1988, and the Ph.D. degree in electrical engineering from the Hungarian Academy of Sciences, Budapest, Hungary, in 1992.

He was a Postdoctoral Researcher with the University of Hawaii in 1992 and a Visiting Fulbright Scholar with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Champaign, IL, USA during 1993–1994. From 1994 to 1998 he was a faculty member with the Department of Computer Science and Information Theory at the Technical University of Budapest, Hungary. From 1996 to 1998 he was also a Visiting Research Scholar with the Department of Electrical and Computer Engineering, University of California, San Diego, CA, USA. In 1998 he joined Queen's University, Kingston, ON, Canada, where he is now a Professor of mathematics and engineering with the Department of Mathematics and Statistics. His research interests include communications and information theory, source coding and vector quantization, machine learning, and statistical pattern recognition.

Dr. Linder was the recipient of the Premier's Research Excellence Award of the Province of Ontario in 2002 and the Chancellor's Research Award of Queen's University in 2003. He was an Associate Editor for Source Coding of the IEEE TRANSACTIONS ON INFORMATION THEORY during 2003–2004.