2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

Sufficient Conditions for Solving Statistical Filtering Problems by Dynamic Programming

Eugene A. Feinberg, Sayaka Ishizawa, Pavlo O. Kasyanov and David N. Kraemer

Abstract— **The paper studies discrete-time statistical filtering problems with the goal to minimize expected total costs. Such problems are usually defined by pairs of stochastic equations and by one-step cost functions. Stochastic equations describe the state and observation processes, and these equations are defined by transition and observation functions. This paper provides sufficient conditions on observation, transition, and one-step cost functions for convergence of value-iteration algorithms for problems with finite and infinite horizons. It is well-known that nonlinear and linear filtering problems can be presented as Partially Observable Markov Decision Processes (POMDPs). The paper applies contemporary results on convergence of value iterations for Markov Decision Processes (MDPs) and for POMDPs to filtering problems. It formulates conditions on observation and transition functions which imply weak continuity of the filter. Weak continuity of the filter means weak continuity of transition probabilities between belief states. The sufficient condition on one-step functions is their K-inf-compactness. The described conditions hold for broad classes of nonlinear filters and for Kalman filters.**

I. Introduction

This paper provides conditions for weak continuity of filters for discrete-time nonlinear control systems defined by stochastic equations. It is known that discrete-time problems with incomplete information can be reduced to belief-MDPs, whose states are posterior probabilities of the original states. These probabilities are defined by filtering equations, which are equivalent to transition probabilities between belief states in belief-MDPs [14, pp. 86,87]. Karman's filter is an algorithm implementing filtering equations for linear Gaussinan problems. Thus the transition probability for a belief-MDP is a filter, and weak continuity of a filter means weak continuity of the transition probability for the belief-MDP. Weak continuity of a filter and the appropriate conditions on one-step costs imply the existence of optimal policies, validity of optimality equations, and convergence of value iterations for problems with expected total costs. The analysis in this paper is based on recently developed sufficient conditions for weak continuity of filters in stochastic control problems with incomplete information defined by transition and observation kernels; [11], [12], [13] and [15].

Controlled models with incomplete information can be defined either by stochastic equations or by transition and observation kernels. These two approaches are equivalent; see Aumann's lemma (Theorem 2 below). The natural way to find optimal policies for POMDPs is to reduce the initial problem to a belief-MDP, with states being posterior probability distributions of the original states. If an optimal policy for the belied-MDP is found, then it can be used to construct an optimal policy for the initial problem that minimizes the expected total discounted costs [4], [6], [14]. References [1], [2], [6], and [23] introduced this approach for problems with finite state, observation, and action spaces. Reference [20] extended it to problems with countable state spaces, and references [18] and [25] developed the theory for Borel state spaces. In an important well-studied and relatively simple case of linear systems with Gaussian noise, the corresponding stochastic processes are Gaussian, and they are defined by mean vectors and covariance matrices, and the Kalman filter completely characterizes transition probabilities.

In general, the existence of optimal policies for statistical filtering problems is a nontrivial question even in the case of expected total costs. For MDPs with expected total costs the existence of optimal policies and the validity of other important properties, including the existence of solutions to optimality equations and convergence of value iterations, follow from continuity properties of the transition kernel and one-step costs. For MDPs with compact action sets, the corresponding properties are weak continuity of transition kernels and lower semicontinuity of one-step costs; see [21], [22]. For MDPs with possibly noncompact action sets, the corresponding properties are weak continuity of transition kernels and $\mathbb{K}$-inf-compactness of one-step costs; see [10], [11] for details. There is also a parallel theory for MDPs with setwise continuous transition probabilities, including [21], [22] and [9], but so far it has not found applications to problems with incomplete observations.

While lower semicontinuity and $\mathbb{K}$-inf-compactness of one-step cost functions are preserved by the reduction of a POMDP to the belief-MDP, weak continuity of the transition kernel may not be preserved [11, Theorem 3.3, Lemma 2.1, and Examples 4.2 and 4.3]. For a long time, sufficient conditions for weak continuity of filters were unavailable.

Feinberg's research was partially supported by SUNY System Administration under SUNY Research Seed Grant Award #231087 and by ONR Grant N000142412608. Kasyanov's research was partially supported by ONR Grant N000142412646.

Eugene A. Feinberg is with the Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, NY 11794-3600, USA eugene.feinberg@stonybrook.edu

Sayaka Ishizawa is with the Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, NY 11794-3600, USA sayaka.ishizawa@alumni.stonybrook.edu

Pavlo O. Kasyanov is with Institute for Applied System Analysis, National Technical University of Ukraine "Igor Sikorsky Kyiv Polytechnic Institute", Peremogy ave., 37, build, 35, 03056, Kyiv, Ukraine kasyanov@i.ua

David N. Kraemer is with the Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, NY 11794-3600, USA david.kraemer@stonybrook.edu

Monographs [14, pp. 90-93] and [19, Section 2] introduced some particular conditions assuming, among other conditions, weak continuity of transition kernels and continuity in total variation of the observation kernels. Reference [11] proved that these two conditions imply weak continuity of filters. Reference [15] provides another proof of this fact and proved that continuity of the transition kernel is sufficient if the observation kernel does not depend on observations. References [12], [13] introduced the notion of semi-uniform Feller transition probabilities and showed that this property is preserved when a problem with incomplete observations is reduced to a belief-MDP, and this is also true for more general problems with incomplete information than POMDPs. This property implies weak continuity of the filter. In particular, these facts provide another proof that weak continuity of the transition kernel and continuity of the observation kernels imply weak continuity of the filter. They also imply that, if the transition kernel is continuous in total variation and the observation kernel is continuous in total variation in the control parameter, then the filter is weakly continuous.

This paper describes sufficient conditions on the transition and observation stochastic equations defining the dynamics of the system for weak continuity of the filter. Such conditions were considered in [11, Section 8.1] for problems with real-valued states, observations, and noises. This paper presents stronger results for problems with more general state, observation, and noise spaces. Relevant equations have been studied in [14], [16], [5] and in publications dealing with particular problems and applications.

This paper is organized as follows. Section II introduces mathematical notation and definitions. Section III formulates the main problem, introduces assumptions, and states Theorem 1, which is the main result of this paper. Section IV explains the equivalence between processes defined by stochastic equations and by stochastic kernels (Theorem 2 which was originally introduced by Aumann [3, Lemma F]) and provides criteria and sufficient conditions for weak continuity and for continuity in total variation of stochastic kernels defined by equations (Theorems 3 and 4). Section V discusses some applications.

The assumptions of this paper are satisfied for broad classes of models including models with additive and multiplicative noises. In particular, as described in [8], they are satisfied for linear state space models and for inventory control models.

II. Notation and Background Definitions

For a metric space $\mathbb{S}$, we denote by $\rho_{\mathbb{S}}$ its metric and denote by $\mathcal{B}(\mathbb{S})$ its Borel σ-algebra; i.e., the σ-algebra generated by the open subsets of $\mathbb{S}$. For $S \in \mathcal{B}(\mathbb{S})$ we usually consider the metric space (S, ρ_S), where $\rho_S(s_1, s_2) = \rho_{\mathbb{S}}(s_1, s_2)$ for $s_1, s_2 \in S$. For the Borel σ-algebra $\mathcal{B}(S)$ on S, the equality $\mathcal{B}(S) = \{B \in \mathcal{B}(\mathbb{S}) : B \subset S\}$ holds.

The set of probability measures on the measurable space $(\mathbb{S}, \mathcal{B}(\mathbb{S}))$ is denoted by $\mathbb{P}(\mathbb{S})$. The set $\mathbb{P}(\mathbb{S})$ is endowed with the topology of weak convergence of probability measures, i.e., $p^{(k)} \to p$ in $\mathbb{P}(\mathbb{S})$ if $\int_{\mathbb{S}} f(s)\ p^{(k)}(ds) \to \int_{\mathbb{S}} f(s)\ p(ds)$ for all bounded continuous functions $f : \mathbb{S} \to \mathbb{R}$. Let $\mathbb{S}_1$ and $\mathbb{S}_2$ be Borel spaces; i.e., Borel subsets of Polish (separable, complete metric) spaces. We recall that a stochastic kernel on $\mathbb{S}_1$ given $\mathbb{S}_2$ is a function $\kappa : \mathcal{B}(\mathbb{S}_1) \times \mathbb{S}_2 \to [0, 1]$, written $\kappa(B|s_2)$, such that

1) for each $B \in \mathcal{B}(\mathbb{S}_1)$, the map $s_2 \mapsto \kappa(B|s_2)$ is a Borel measurable function,
2) for each $s_2 \in \mathbb{S}_2$, the map $B \mapsto \kappa(B|s_2)$ is a Borel probability measure on $\mathbb{S}_1$.

A stochastic kernel κ on $\mathbb{S}_1$ given $\mathbb{S}_2$ is weakly continuous if the probability measures $\kappa(\cdot|s_2')$ converge weakly to $\kappa(\cdot|s_2)$ in $\mathbb{P}(\mathbb{S}_1)$ as $s_2' \to s_2$, and continuous in total variation if $\sup_{B\in\mathcal{B}(\mathbb{S}_1)} |\kappa(B|s_2') - \kappa(B|s_2)| \to 0$ as $s_2' \to s_2$. Continuity in total variation implies weak continuity.

In this paper, the variables d, m, and n always refer to positive integers. Given two measures p_1, p_2 on the same measurable space we write $p_2 \ll p_1$ if p_2 is absolutely continuous with respect to p_1. Lebesgue measure on $\mathbb{R}^n$ is denoted $\lambda^{[n]}$. Let $D_x g = \frac{\partial g}{\partial x}$ denote the Jacobian of a differentiable function $g : \mathbb{R}^n \to \mathbb{R}^n$. When we consider the sufficiently smooth function $g(s_2, x)$ on $\mathbb{S}_2 \times \mathbb{R}^n$ we denote its Jacobian in x either by $D_x\phi(s_2, x)$ or by $D_x\phi_{s_2}(x)$.

A function $f : \mathbb{S} \to [-\infty, +\infty]$ is lower semicontinuous at $s \in \mathbb{S}$ if $\liminf_{s'\to s} f(s') \geq f(s)$. If f is lower semicontinuous at each $s \in \mathbb{S}$, then f is lower semicontinuous on $\mathbb{S}$. The function f is inf-compact if for all $\gamma \in \mathbb{R}$, the sublevel set $\{s \in \mathbb{S} : f(s) \leq \gamma\}$ is compact. A function $f : \mathbb{S}_1 \times \mathbb{S}_2 \to [-\infty, +\infty]$ is $\mathbb{K}$-inf-compact on $\mathbb{S}_1 \times \mathbb{S}_2$ if for all nonempty compact sets $C \subset \mathbb{S}_1$ and for all $\gamma \in \mathbb{R}$, the sublevel sets $\{(s_1, s_2) \in C \times \mathbb{S}_2 : f(s_1, s_2) \leq \gamma\}$ are compact.

For example, for spaces $\mathbb{X} = \mathbb{R}^d$, $\mathbb{A} = \mathbb{R}^\ell$, the following two functions $c : \mathbb{X} \times \mathbb{A} \to \mathbb{R}$ are $\mathbb{K}$-inf-compact:

(i) $c(x, a) := x^T\mathbf{X}x + a^T\mathbf{A}a$, where $\mathbf{X} \in \mathbb{R}^{d\times d}$ is positive semidefinite and $\mathbf{A} \in \mathbb{R}^{\ell\times\ell}$ is positive definite;

(ii) $c(x, a) := \rho_{\mathbb{R}^k}(\mathbf{X}x, \mathbf{A}a)$, where $\mathbf{X} \in \mathbb{R}^{k\times d}$ and $\mathbf{A} \in \mathbb{R}^{k\times\ell}$ is nonsingular.

We denote by $\mathbb{X}$, $\mathbb{Y}$, and $\mathbb{A}$ the state, observation, and action spaces, respectively. We denote by $\mathcal{X}$ and $\mathcal{H}$ the spaces of state and observation noises. In general, $\mathbb{X}$, $\mathbb{Y}$, $\mathbb{A}$, $\mathcal{X}$, and $\mathcal{H}$ are assumed to be Borel spaces, that is, they are Borel subsets of complete separable metric spaces. In some results there are additional assumptions that some of these spaces are Euclidian. The discrete time parameter is $t = 0, 1, \ldots$, and x_t, y_t, a_t, ξ_t, and η_t denote the state, observation, control, state noise, and observation noise at time t respectively, where $x_t \in \mathbb{X}$, $y_t \in \mathbb{Y}$, $a_t \in \mathbb{A}$, $\xi_t \in \mathcal{X}$, and $\eta_t \in \mathcal{H}$. These variables are defined in (4).

For a Borel space Ω, for a Borel measurable function $\phi : \mathbb{S}_2 \times \Omega \to \mathbb{S}_1$, and for a probability measure $p \in \mathbb{P}(\Omega)$, let us define the stochastic kernel κ on $\mathbb{S}_1$ given $s_2 \in \mathbb{S}_2$,

$$\kappa(B|s_2) := \int_\Omega \mathbf{1}\{\phi(s_2, \omega) \in B\}\ p(d\omega), \quad B \in \mathcal{B}(\mathbb{S}_1). \quad (1)$$

The following properties of ϕ will be used throughout the paper.

Definition 2.1: (Continuity in distribution, total variation, and probability). Let $\mathbb{S}_1$, $\mathbb{S}_2$, and Ω be Borel spaces, and let $p \in \mathbb{P}(\Omega)$ be a probability measure on $(\Omega, \mathcal{B}(\Omega))$. A Borel function $\phi : \mathbb{S}_2 \times \Omega \to \mathbb{S}_1$ is continuous

(i) in distribution p (weakly continuous) if the function $s_2 \mapsto \int_\Omega f(\phi(s_2,\omega))\, p(d\omega)$ is continuous on $\mathbb{S}_2$ for every bounded continuous function $f : \mathbb{S}_1 \to \mathbb{R}$;

(ii) in total variation with respect to (wrt) p if for each $s_2 \in \mathbb{S}_2$,

$$\lim_{s_2' \to s_2} \sup_{B \in \mathcal{B}(\mathbb{S}_1)} \Big| \int_\Omega \mathbf{1}\{\phi(s_2', \omega) \in B\} - \mathbf{1}\{\phi(s_2,\omega) \in B\}\, p(d\omega) \Big| = 0; \tag{2}$$

(iii) in probability p if $\phi(s_2', \,\cdot\,) \xrightarrow{p} \phi(s_2, \,\cdot\,)$ as $s_2' \to s_2$ for each $s_2 \in \mathbb{S}_2$, that is, for each $s_2 \in \mathbb{S}_2$ and each $\varepsilon > 0$,

$$\lim_{s_2' \to s_2} p(\{\omega \in \Omega \,:\, \rho_{\mathbb{S}_1}(\phi(s_2', \omega), \phi(s_2, \omega)) \geq \varepsilon\}) = 0. \tag{3}$$

We note that continuity in probability is called stochastic continuity [24, p. 30].

If ϕ is continuous in total variation wrt p, then it is stochastically continuous. It is well-known that continuity in probability implies continuity in distribution, while the opposite statement is false [24, p. 30].

For a function $\phi : \mathbb{S}_2 \times \Omega \to \mathbb{S}_1$, we will frequently assume that $\mathbb{S}_1 = \mathbb{R}^n$, Ω is an open subset of $\mathbb{R}^n$, and that the function ϕ satisfies the following condition.

Diffeomorphic Condition. For the metric space $\mathbb{S}_2$, open set $\Omega \subset \mathbb{R}^n$, and a function $\phi : \mathbb{S}_2 \times \mathbb{R}^n \to \mathbb{R}^n$ the following statements hold:

(i) ϕ is continuous on $\mathbb{S}_2 \times \Omega$;

(ii) $D_\omega \phi$ exists for all $s_2 \in \mathbb{S}_2$ and $\omega \in \Omega$;

(iii) the matrix $D_\omega \phi(s_2, \omega)$ is nonsingular for all $s_2 \in \mathbb{S}_2$ and $\omega \in \Omega$;

(iv) the function $(s_2, \omega) \mapsto D_\omega \phi(s_2, \omega)$ is continuous on $\mathbb{S}_2 \times \Omega$;

(v) for each $s_2 \in \mathbb{S}_2$ the function $\omega \mapsto \phi(s_2, \omega)$ is a one-to-one mapping of Ω onto $\phi(s_2, \Omega)$.

Remark 1: Let us consider the notation $\phi_{s_2}(\omega) := \phi(s_2, \omega)$. The Diffeomorphic Condition implies that for each $s_2 \in \mathbb{S}_2$ the inverse function $\phi_{s_2}^{-1}(s_1)$ exists for all $s_1 \in \phi_{s_2}(\mathbb{R}^n)$, and this function is continuously differentiable on $\phi_{s_2}(\mathbb{R}^n)$. Diffeomorphic Condition and the inverse function rule imply $\det D_{s_1} \phi_{s_2}^{-1}(s_1) \neq 0$ for all all $s_2 \in \mathbb{S}_2$ and for all $s_1 \in \phi_{s_2}(\Omega)$, and for each compact set $\mathbb{K} \subset \Omega$ be the mappings $\phi_{s_2}^{-1}(s_1) \;:\; (\phi(\mathbb{S}_2 \times \mathbb{K}), \rho_{\mathbb{S}_2 \times \mathbb{R}^n}) \to (\mathbb{K}, \rho_{\mathbb{R}^n})$ and $\det D_{s_1} \phi_{s_2}^{-1}(s_1) \;:\; (\phi(\mathbb{S}_2 \times \mathbb{K}), \rho_{\mathbb{S}_2 \times \mathbb{R}^n}) \to (\mathbb{R}, \rho_{\mathbb{R}})$ are continuous. We also note that, in the one-dimensional case Diffeomorphic Condition (v) follows from Diffeomorphic Conditions (i-iv) since each function $\phi_{s_2}(\omega)$ is continuous and strictly monotonic. If $n > 1$, finding broad sufficient conditions for injectivity of the mapping $\omega \mapsto \phi_{s_2}(\omega)$ over the entire set $\mathbb{R}^n$ is a challenging mathematical problem. However, as shown in Section V below, in the case of additive or multiplicative noises, Diffeomorphic Condition (v) is satisfied if Diffeomorphic Conditions (i-iv) are satisfied.

III. Discrete-Time Partially Observable Controlled Markov Models

Let $\mathbb{X}$, $\mathbb{Y}$, and $\mathbb{A}$ be Borel spaces (Borel subsets of complete separable metric spaces) of states, observations, and actions, respectively, and $\mathcal{X}, \mathcal{H}$ be Borel spaces of state and observation noises. We consider a discrete-time control system with dynamics and observations defined for time $t = 0, 1, \ldots$ by stochastic equations

$$x_{t+1} = F(x_t, a_t, \xi_t), \quad x_t \in \mathbb{X},\ a_t \in \mathbb{A},\ \xi_t \in \mathcal{X}, \tag{4a}$$

$$y_{t+1} = G(a_t, x_{t+1}, \eta_{t+1}), \quad y_{t+1} \in \mathbb{Y},\ \eta_{t+1} \in \mathcal{H}, \tag{4b}$$

where $F : \mathbb{X} \times \mathbb{A} \times \mathcal{X} \to \mathbb{X}$ is the Borel-measurable functions representing the transition dynamics of the system, and $G : \mathbb{A} \times \mathbb{X} \times \mathcal{H} \to \mathbb{Y}$ is the Borel-measurable function representing the observations. In addition, there is the probability distribution $p_0 \in \mathbb{P}(\mathbb{X})$ of the initial state x_0, and $y_0 = G_0(x_0, \eta_0)$, where $G_0 \,:\, \mathbb{X} \times \mathcal{H} \to \mathbb{Y}$ a Borel-measurable function defining the initial observation, and $\{\xi_t\}_{t=0}^\infty$ and $\{\eta_t\}_{t=0}^\infty$ are sequences of independent and identically distributed (iid) random variables with distributions $\mu \in \mathbb{P}(\mathcal{X})$ and $\nu \in \mathbb{P}(\mathcal{H})$, respectively. It is also assumed that the sequence $\{x_0, \xi_0, \eta_0, \xi_1, \eta_1, \ldots\}$ is mutually independent. The process evolves as follows. The initial hidden state is taken $x_0 \sim p_0$, and the initial observation $y_0 = G(x_0, \eta_0)$ is made. For each $t = 0, 1, \ldots,$ the decision maker observes y_t and selects the action $a_t \in \mathbb{A}$. The next state x_{t+1} and observation y_{t+1} are determined by the equations (4). At each epoch $t = 0, 1, \ldots$, in addition to the observation y_t, the decision maker knows the initial state distribution p_0, the previous observations $y_0, y_1, \ldots, y_{t-1}$, and previously chosen actions $a_0, a_1, \ldots, a_{t-1}$. Equations (4) are slightly more general that equations for filtering problems because the function G in (4b) can depend on controls a_t.

The evolution of a system defined by stochastic equations such as (4) can also be represented in terms of stochastic kernels, and the corresponding model is a POMDP. For POMDPs, transitions of states are defined by a stochastic kernel $\mathcal{T}$ on $\mathbb{X}$ given $\mathbb{X} \times \mathbb{A}$ called the transition kernel or transition probabilities, and observations are defined by a stochastic kernel Q on $\mathbb{Y}$ given $\mathbb{A} \times \mathbb{X}$ called the observation kernel or observation probabilities. For POMDPs, the state x_{t+1} is defined by the distribution $\mathcal{T}(\,\cdot\,|x_t, a_t)$ instead of by equality (4a), and the observation y_{t+1} is defined by the distribution $Q(\,\cdot\,|a_t, x_{t+1})$ instead of by equality (4b). In addition, the initial distribution of observation y_0 is defined by a stochastic kernel Q_0 on Y given X, that is, y_0 is defined by the distribution $Q(\,\cdot\,|x_0)$; see [14, Chapter 4] or [11] for details.

The model definitions based on stochastic equations and on POMDPs are equivalent. Indeed, for the functions F and G from (4), the transition kernel $\mathcal{T}$ and observation kernel Q are defined in (1), where the transition kernel $\mathcal{T} = \kappa$ for $\phi := F$, $\mathbb{S}_1 := \mathbb{X}$, $\mathbb{S}_2 := \mathbb{X} \times \mathbb{A}$, $\Omega := \mathcal{X}$ and the observation kernel $Q = \kappa$ for $\phi := G$, $\mathbb{S}_1 := \mathbb{Y}$, $\mathbb{S}_2 := \mathbb{A} \times \mathbb{X}$, and $\Omega := \mathcal{H}$, Of course, the function G_0 also defines the stochastic kernel $Q_0(C|a,x) = \int_{\mathcal{H}} \mathbf{1}\{G_0(x,\eta) \in C\}\nu(d\eta)$, where $C \in \mathcal{B}(\mathbb{Y})$, $a \in \mathbb{A}$, and $x \in \mathbb{X}$. The reduction of a POMDP to the stochastic system defined by equations (4), with $\mathcal{X} = [0,1]$, $\mathcal{H} = [0,1]$, and μ and η are uniform distributions, follows from [3, Lemma F]; see Theorem 2 below. However, we do not assume here that $\mathcal{X} = [0,1]$, $\mathcal{H} = [0,1]$, and μ and η are distributed uniformly, because sometimes it can be more convenient to consider other distributions; see, e.g., Corollary 1.

If decision a is chosen at state x then the one-step cost is $c(x,a)$, where $c : \mathbb{X} \times \mathbb{Y}$ is a bounded below Borel-measurable function. Then expected total discounted costs for a policy π are defined as

$$V^{\pi}_{T,\alpha}(p_0) := \mathbf{E}^{\pi}_{p_0} \sum_{t=0}^{T-1} \alpha^t c(x_t, a_t), \qquad p_0 \in \mathbb{P}(\mathbb{X}), \tag{5}$$

where $T = 0, 1, \dots$ or $T = \infty$ is the fixed finite or infinite planning horizon, and $\beta \in [0,1)$ is the discount factor. If $T < \infty$, than β can be an arbitrary nonnegative number. The goal is to find an optimal policy minimizing expected total discounted costs for a given planning horizon.

A natural approach to solving a POMDP is based on its reduction to a completely observable MDP, whose states are probability distributions of the states of the system, sometimes called belief distributions, and this MDP is sometimes called a filter. Weak continuity of the transition probabilities of the filter, also called weak continuity of the filter, is an important property. If one-step costs c are $\mathbb{K}$-inf-compact and bounded below, then weak continuity of the filter implies the existence of optimal policies for problems with expected total costs, the validity of optimality equations, and convergence of value iterations [11].

In view of space limitations, we do not provide here the classic definitions of the belief-MDP ; see, e.g., [7], [14]. The state space of the belief-MDP is the space $\mathbb{P}(\mathbb{X})$ of probability measures on $\mathbb{X}$ endowed with the topology of weak of convergence of probability measures, and the transition kernel q for the filter is the transition probability from $\mathbb{P}(\mathbb{X}) \times \mathbb{A}$ to $\mathbb{P}(\mathbb{X})$. Weak and setwise continuity assumptions on kernels $\mathcal{T}$ and Q are not sufficient for weak continuity of q; see examples in [11]. The following sufficient conditions are known for weak continuity of filters:

(i) $\mathcal{T}$ is weakly continuous, and Q is continuous in total variation [11];
(ii) $\mathcal{T}$ is continuous in total variation, and Q is continuous in a in total variation [12];

In particular, if Q does not depend on the control parameter, then continuity of $\mathcal{T}$ in total variation implies weak continuity of the filter [15].

As explained in the next section, the following two assumptions on functions F and G imply respectively conditions (i) and (ii) above.

Assumption 1: The following statements hold:

(i) the function $((x,a),\xi) \mapsto F(x,a,\xi)$ is continuous in distribution μ;
(ii) $\mathbb{Y} = \mathbb{R}^n$, $\mathcal{H}$ is an open subset of $\mathbb{R}^m$, $\nu \ll \lambda^{[m]}$, and the function $((a,x),\eta) \mapsto G(a,x,\eta)$ satisfies the Diffeomorphic Condition.

Assumption 2: The following statements hold:

(i) $\mathbb{X}$ is an open subset of $\mathbb{R}^d$, $\mathcal{X} = \mathbb{R}^d$, $\mu \ll \lambda^{[d]}$, and the function $((x,a),\xi) \mapsto F(x,a,\xi)$ satisfies the Diffeomorphic Condition;
(ii) either G does not depend on a, or the following conditions (a) and (b) hold: (a) $\mathbb{Y} = \mathbb{R}^m$, $\mathcal{H}$ is an open subset of $\mathbb{R}^m$, and $\nu \ll \lambda^{[m]}$, and (b) for each $x \in \mathbb{X}$ the function $(a,\eta) \mapsto G(a,x,\eta)$ satisfies the Diffeomorphic Condition.

Theorem 1 describes sufficient conditions on the kernels $\mathcal{T}$ and Q for weak continuity of the filter. The main goal of this paper is to describe sufficient conditions on the functions F and G for weak continuity of the filter. Such results were provided in [11, Section 8.1] for problems with real-valued states, observations, and noises. The following theorem provides stronger results for more general state, observation, action, and noise spaces.

Theorem 1: If either Assumption 1 or Assumption 2 holds, then the filter is weakly continuous, that is, the transition probability q from $\mathbb{P}(\mathbb{X}) \times \mathbb{A}$ to $\mathbb{P}(\mathbb{X})$ for the belief MDP is weakly continuous. Thus, if the one-step cost is $\mathbb{K}$-inf-compact, then the value function for the filter is lower semicontinuous and satisfies the optimality equations for the filter, value iterations converge to optimal values, and optimality equations define optimal policies.

IV. Continuity of Stochastic Kernels Defined by Stochastic Equations

Let $\mathbb{S}_1$, $\mathbb{S}_2$, and Ω again denote Borel spaces, and let us consider the stochastic kernel κ on $\mathbb{S}_1$ given $\mathbb{S}_2$. This section deals with two distinct but related issues concerning κ. We shall first consider the existence of a Borel measurable function $\phi : \mathbb{S}_2 \times \Omega \to \mathbb{S}_1$ and probability measure $p \in \mathbb{P}(\Omega)$ such that representation (1) holds. Theorem 2 is Aumann's lemma [3, Lemma F], which shows that such a representation exists with $\Omega = [0,1]$ and $p = \lambda^{[1]}$. Corollary 1 states that for every natural number n such a representation exists with $\Omega = [0,1]^n$ and $p = \lambda^{[n]}$.

We will then consider conditions on ϕ and p that imply different continuity properties of the stochastic kernel κ. Theorem 3 shows that a necessary and sufficient condition for κ to be weakly continuous is for ϕ to be continuous in distribution p. It also shows that stochastic continuity is a sufficient condition for ϕ to imply that κ is weakly continuous. Theorem 4 gives a necessary and sufficient condition for κ to be continuous in total variation; namely, that ϕ is continuous in total variation wrt p. In addition, it shows that, if p is a continuous distribution on $\mathbb{R}^n$ and ϕ satisfies the Diffeomorphic Condition, then κ is continuous in total variation. Finally, we provide an example showing the significance of the assumption that p has a density for continuity of κ in total variation.

The following theorem is due to [3, Lemma F]. It shows that κ always admits representation (1) when $\Omega = [0,1]$ and p is the uniform distribution.

Theorem 2 ([3, Lemma F]): Let $\mathbb{S}_1$ and $\mathbb{S}_2$ be Borel spaces, and let κ be a stochastic kernel on $\mathbb{S}_1$ given $\mathbb{S}_2$. Then there exists a Borel measurable function $\phi : \mathbb{S}_2 \times [0,1] \to \mathbb{S}_1$, where the Borel σ-algebra is considered on the unit interval $[0,1]$, such that

$$\kappa(B|s_2) = \int_0^1 \mathbf{1}\{\phi(s_2,\omega) \in B\}\, d\omega, \quad B \in \mathcal{B}(\mathbb{S}_1). \tag{6}$$

As a remark, we would like to mention a relevant fact, which is not used in this paper. [17, Theorem 1.1] and the isomorphism theorem for Borel spaces imply that it is possible to define a probability measure $\mathfrak{m}$ on the space $\mathfrak{B}(\mathbb{S}_2,\mathbb{S}_1)$ of Borel measurable functions $f : \mathbb{S}_2 \to \mathbb{S}_1$ such that $\kappa(B|s_2) = \int_{\mathfrak{B}(\mathbb{S}_2,\mathbb{S}_1)} \mathbf{1}\{f(s_2) \in B\}\, \mathfrak{m}(df)$.

The following corollary is a generalization of Theorem 2. It shows that the space $[0,1]$ can be replaced by $[0,1]^n$, $n = 1,2,\ldots,$ without loss of generality.

Corollary 1: Let $\mathbb{S}_1$ and $\mathbb{S}_2$ be Borel spaces, and let κ be a stochastic kernel on $\mathbb{S}_1$ given $\mathbb{S}_2$. Then for each natural number n there exists a Borel measurable function $\phi : \mathbb{S}_2 \times [0,1]^n \to \mathbb{S}_1$, where the Borel σ-algebra is considered on the unit box $[0,1]^n$, such that

$$\kappa(B|s_2) = \int_{[0,1]^n} \mathbf{1}\{\phi(s_2,\omega) \in B\}\, d\omega, \quad B \in \mathcal{B}(\mathbb{S}_1). \tag{7}$$

The remainder of this section concerns continuity of the stochastic kernel κ defined in (1). There are well-known sufficient conditions on the function ϕ that imply that κ is weakly continuous. For example, continuity of ϕ is sufficient; [14, p. 92]. As discussed in [11, Section 8.1], it is sufficient for $s_2 \mapsto \phi(s_2,\omega)$ to be continuous for p-a.s. ω. The following theorem describes necessary and sufficient conditions on the Borel measurable function ϕ for the kernel κ to be weakly continuous. It also shows that stochastic continuity is another sufficient condition for weak continuity of κ, which is weaker than p-a.s. continuity of ϕ at each $s_2 \in \mathbb{S}_2$, which in turn is weaker than continuity of ϕ.

Theorem 3 (Weak continuity): Let $p \in \mathbb{P}(\Omega)$ and $\phi : \mathbb{S}_2 \times \Omega \to \mathbb{S}_1$ be Borel measurable, and consider the stochastic kernel κ on $\mathbb{S}_1$ given $\mathbb{S}_2$ defined in (1). The following statements hold:

(a) the function ϕ is continuous in distribution p if and only if the stochastic kernel κ is weakly continuous;
(b) if the function ϕ is continuous in probability p, then κ is weakly continuous.

The next theorem states necessary and sufficient conditions for the stochastic kernel κ to be continuous in total variation. It also shows that for the spaces $\mathbb{S}_1 = \Omega = \mathbb{R}^n$, if p is a continuous distribution and ϕ satisfies the Diffeomorphic Condition, the stochastic kernel κ is continuous in total variation. Theorem 4 is applicable to linear and nonlinear filtering problems including to linear filtering problems with additive and multiplicative noises and to inventory control problems [8].

Theorem 4 (Continuity in total variation): Let $p \in \mathbb{P}(\Omega)$ and $\phi : \mathbb{S}_2 \times \Omega \to \mathbb{S}_1$ be Borel measurable, and consider the stochastic kernel κ on $\mathbb{S}_1$ given $\mathbb{S}_2$ defined in (1). The following statements hold:

(a) the function ϕ is continuous in total variation wrt p if and only if the stochastic kernel κ is continuous in total variation;
(b) if $\mathbb{S}_1 = \mathbb{R}^n$, Ω is an open subset of $\mathbb{R}^n$, $p \ll \lambda^{[n]}$, where $\lambda^{[n]}$ is restricted to Ω, and the function ϕ satisfies the Diffeomorphic Condition, then the stochastic kernel κ is continuous in total variation.

Let us consider the special case $\mathbb{S}_2 = \mathbb{R}^n$, $\Omega = \mathbb{R}^n$, and assume that the distribution p has a continuous density f. Let us define random variables $Y_{s_2} = \phi(s_2,\xi)$, where $s_2 \in \mathbb{S}_2$ and $\xi \sim p$. Then $Y_{s_2} \sim \kappa(\cdot|s_2)$. Following [16, example 1.3.2(v)], the formula for the density of a function of a random vector implies that the density $f_{Y_{s_2}}$ of Y_{s_2} is

$$f_{Y_{s_2}}(s_1) = f(\phi^{-1}_{s_2}(s_1))|D_{s_1}\phi(s_2.s_1)|^{-1}\mathbf{1}\{s_1 \in \phi(s_2,\Omega)\}.$$

If $\phi(s_2,\Omega) = \mathbb{R}^n$ for all $s_2 \in \mathbb{S}_2$, then continuity of f and the Diffeomorphic Condition condition on ϕ imply that $\phi(s_2',\Omega) \to \phi(s_2,\Omega)$ when $s_2' \to s_2$ and $s_2, s_2' \in \mathbb{S}_2$. Therefore, Scheffé's theorem implies that the kernel $\kappa(\cdot|s_s)$ is continuous in total variation. However, in general convergence $\phi(s_2',\Omega) \to \phi(s_2,\Omega)$ may not take place for all $s_2 \in \mathbb{S}_2$ since indicators are not continuous functions. Lemma 1 was used in [8] to prove Theorem 4(b), which does not assume that $\phi(s_2,\Omega) = \mathbb{R}^n$ for all $s_2 \in \mathbb{S}_2$.

Lemma 1: ([8, Lemma 9.1]) Let $\mathbb{S}_2$ be a metric space, and let $\Omega \subset \mathbb{R}^n$ be open. If a function $\phi : \mathbb{S}_2 \times \mathbb{R}^n \to \mathbb{R}^n$ satisfies the Diffeomorphic Condition. Then, for every compact set $K \subset \Omega$ and $s_2' \to s_2$,

$$\lim_{s_2' \to s_2} \lambda^{[n]}(\phi(s_2',K) \triangle \phi(s_2,K)) = 0. \tag{8}$$

The following example shows that the assumption $p \ll \lambda^{[n]}$ is essential in statement (b) of Theorem 4.

Example 1: Suppose $\mathbb{S}_1 = \mathbb{S}_2 = \Omega = \mathbb{R}$, suppose the distribution p is concentrated at the point 0, that is, $p(B) = \mathbf{1}\{0 \in B\}$ for $B \in \mathcal{B}(\mathbb{R})$, and suppose $\phi(s_2,\omega) = s_2 + \omega$. Then $\kappa(B|s_2) = \mathbf{1}\{s_2 \in B\}$ for $B \in \mathcal{B}(\mathbb{R})$ since $p(0) = 1$. Consider the sequence $s_2^{(k)} = k^{-1}$ for $k = 1,2,\ldots,$ and let $B = \{0\}$. Then $s_2^{(k)} \to 0$, $\kappa(B|0) = 1$, and $\kappa(B|s_2^{(k)}) = 0$ for all $k = 1,2,\ldots$; hence, κ is not continuous in total variation.

V. Nonlinear Filtering

This section deals with the version of equations (4) in which the function G does not depend on the action:

$$x_{t+1} = F(x_t, a_t, \xi_t), \qquad x_t \in \mathbb{X},\ a_t \in \mathbb{A},\ \xi_t \in \mathcal{X}, \tag{9a}$$

$$y_{t+1} = \tilde{G}(x_{t+1}, \eta_{t+1}), \quad y_{t+1} \in \mathbb{Y}, \quad \eta_{t+1} \in \mathcal{H}, \tag{9b}$$

where $\tilde{G} : \mathbb{X} \times \mathcal{H} \to \mathbb{Y}$ is a Borel measurable function. Such equations commonly appear in statistical filtering theory, as discussed below. Model (9) is a simpler model than (4) because it does not include actions in the observation equation. The following theorem is the main result of this section. It follows directly from Theorem 1 applied to the model (9).

Theorem 5 (Nonlinear filtering): Consider the model (9) for Borel measurable functions $F : \mathbb{X} \times \mathbb{A} \times \mathcal{X} \to \mathbb{X}$ and $\tilde{G} : \mathbb{X} \times \mathcal{H} \to \mathbb{Y}$, and for transition disturbances $\{\xi_t\}_{t=0}^{\infty} \overset{\text{iid}}{\sim} \mu$ and observation disturbances $\{\eta_t\}_{t=0}^{\infty} \overset{\text{iid}}{\sim} \nu$. Let us consider the following two two assumptions:

(i) The following statements hold:
   a) the function $((x, a), \xi) \mapsto F(x, a, \xi)$ is continuous in distribution μ;
   b) $\mathbb{Y} = \mathbb{R}^m$, $\mathcal{H}$ is an open subset of $\mathbb{R}^m$, $\nu \ll \lambda^{[m]}$, and the function $(x, \eta) \mapsto \tilde{G}(x, \eta)$ satisfies the Diffeomorphic Condition;

(ii) $\mathbb{X} = \mathbb{R}^d$, $\mathcal{X}$ is an open subset of $\mathbb{R}^d$, $\mu \ll \lambda^{[d]}$, and the function $((x, a), \xi) \mapsto F(x, a, \xi)$ satisfies the Diffeomorphic Condition.

Each of assumptions (i) or (ii) implies that the transition probability for the belief-MDP is weakly continuous. In addition, if the one-step cost function c is $\mathbb{K}$-inf-compact, then for expected total discounted costs, value functions are lower semicontinuous, finite-horizon values converge to the infinite-horizon value, optimality equations hold, and optimality equations define for the filter Markov optimal policies for finite-horizon problems and stationary optimal policies for infinite-horizon problems.

VI. Conclusions

Following the progress in understanding sufficient conditions for the existence of optimal policies and applicability of dynamic programming algorithms for MDPs with infinite state spaces and for POMDPs, this paper provides such conditions for discrete-time stochastic problems with incomplete state observations, when the problem is defined by stochastic equations, and the goal is to optimize expected total discounted costs. For POMDPs the important question is whether its belief-MDP has a weakly continuous transition probability, which defines a filter. The answer to this question depends on continuity properties of transition and observation kernels. These properties of transition and observation kernels are their weak continuity and continuity in total variation. This paper links continuity properties of transition and observation functions to continuity properties of the corresponding stochastic kernels.

Continuity in total variation is the more challenging property than weak continuity, and the paper describes a sufficient condition of continuity of a stochastic kernel in total variation. This condition is called the Diffeomorphic Condition. In general, verification of this condition can be challenging, but it can be easily done for problems with additive and multiplicative noises, and such problems are widely used in practice. The paper does not assume that noises are Gaussian.

As was recently understood in [12], all currently known conditions for weak continuity of filters also imply a stronger property than weak continuity of a filter. If the state space of the problem is augmented by an observation, then the currently known sufficient conditions imply semi-uniform Feller continuity of the filter, which is the stronger property than weak continuity. This observation indicates that there could be more general sufficient conditions for weak continuity of filters than the currently known conditions.

References

[1] Aoki, M. Optimal control of partially observable Markovian systems. *Journal of the Franklin Institute 280*, 5 (1965), 367–386.

[2] Åström, K. J. Optimal control of Markov processes with incomplete state information. *Journal of Mathematical Analysis and Applications 10*, 1 (1965), 174–205.

[3] Aumann, R. J. Mixed and behavior strategies in infinite exstensive games. *Advances in Game Theory 52* (1964), 627–650.

[4] Bertsekas, D., and Shreve, S. E. *Stochastic Optimal Control: the Discrete-Time Case*. Athena Scientific, Belmont, 1996.

[5] Chapman, M. P., Jensen, E., Chan, S. M., and Lessard, L. Information-theoretic multi-time-scale partially observable systems with inspiration from leukemia treatment. *Automatica 163* (2024), 111546.

[6] Dynkin, E. B. Controlled random sequences. *Theory of Probability and Its Applications 10*, 1 (1965), 1–14.

[7] Dynkin, E. B., and Yushkevich, A. A. *Controlled Markov Processes*. Springer-Verlag, New York, 1979.

[8] Feinberg, E. A., Ishizawa, S., Kasyanov, P. O., and N., K. D. Continuity of filters for discrete-time control problems defined by explicit equations. *Arxiv e-prints* (2024).

[9] Feinberg, E. A., and Kasyanov, P. O. MDPs with setwise continuous transition probabilities. *Operations Research Letters 49*, 5 (2021), 734–740.

[10] Feinberg, E. A., Kasyanov, P. O., and Zadoianchuk, N. V. Average cost Markov decision processes with weakly continuous transition probabilities. *Mathematics of Operations Research 37*, 4 (2012), 591–607.

[11] Feinberg, E. A., Kasyanov, P. O., and Zgurovsky, M. Z. Partially observable total-cost Markov decision processes with weakly continuous transition probabilities. *Mathematics of Operations Research 41*, 2 (2016), 656–681.

[12] Feinberg, E. A., Kasyanov, P. O., and Zgurovsky, M. Z. Markov decision processes with incomplete information and semi-uniform Feller transition probabilities. *SIAM Journal on Control and Optimization 60*, 4 (2022), 2488–2513.

[13] Feinberg, E. A., Kasyanov, P. O., and Zgurovsky, M. Z. Semi-uniform Feller stochastic kernels. *Journal of Theoretical Probability 36* (2023), 2262–2283.

[14] Hernández-Lerma, O. *Adaptive Markov Control Processes*. Springer-Verlag, New York, 1989.

[15] Kara, A. D., Saldi, N., and Yüksel, S. Weak Feller property of non-linear filters. *Systems & Control Letters 134* (2019), 104512.

[16] Kara, A. D., and Yüksel, S. Robustness to incorrect system models in stochastic control. *SIAM Journal on Control and Optimization 58* (2020), 1144–1182.

[17] Kifer, Y. *Ergodic Theory of Random Transformations*. Birkhäuser, Boston, 1986.

[18] Rhenius, D. Incomplete information in Markovian decision models. *The Annals of Statistics 2*, 6 (1974), 1327–1334.

[19] Runggaldier, W. J., and Stettner, L. *Approximations of Discrete Time Partially Observed Control Problems*. Applied Mathematics Monographs CNR, Giardini Editori, Pisa, 1994.

[20] Sawaragi, Y., and Yoshikawa, T. Discrete-time Markovian decision processes with incomplete state observation. *The Annals of Mathematical Statistics 41*, 1 (1970), 78–86.

[21] Schäl, M. On continuous dynamic programming with discrete time parameter. *Z. Wahrscheinlichkeitstheorie verw. Gebiete 21* (1972), 279–288.

[22] Schäl, M. Average optimality in dynamic programming with general state space. *Mathematics of Operations Research 18*, 1 (1993), 163–172.

[23] Shiryaev, A. On the theory of decision functions and control of a process of observation based on incomplete information. pp. 678–681. Engl. transl. in *Select. Transl. in Math., Statist. Probab.* 6(1966), 162–188.

[24] Wentzell, A. D. *A Course in the Theory of Stochastic Processes*. McGraw-Hill, New York, 1996.

[25] Yushkevich, A. A. 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 of Probability and Its Applications 21*, 1 (1976), 153–158.