

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Techakesari, O;Ford, JJ;Nesic, D

Title:

Practical Stability of Approximating Discrete-Time Filters with Respect to Model Mismatch Using Relative Entropy Concepts

Date:

2011-01-01

Citation:

Techakesari, O., Ford, J. J. & Nesic, D. (2011). Practical Stability of Approximating Discrete-Time Filters with Respect to Model Mismatch Using Relative Entropy Concepts. 2011 50TH IEEE CONFERENCE ON DECISION AND CONTROL AND EUROPEAN CONTROL CONFERENCE (CDC-ECC), pp.7888-7894. IEEE. <https://doi.org/10.1109/CDC.2011.6160225>.

Persistent Link:

<https://hdl.handle.net/11343/299732>

Practical Stability of Approximating Discrete-Time Filters with Respect to Model Mismatch Using Relative Entropy Concepts

Onvaree Techakesari, Jason J. Ford, and Dragan Nešić

Abstract—This paper establishes practical stability results for an important range of approximate discrete-time filtering problems involving mismatch between the true system and the approximating filter model. Using local consistency assumption, the practical stability established is in the sense of an asymptotic bound on the amount of bias introduced by the model approximation. Significantly, these practical stability results do not require the approximating model to be of the same model type as the true system. Our analysis applies to a wide range of estimation problems and justifies the common practice of approximating intractable infinite dimensional nonlinear filters by simpler computationally tractable filters.

I. INTRODUCTION

Many filtering problems involve estimation of system quantities from noisy measurements in situations where the exact (or true) model of the system is either unknown or is more complicated than can be handled using standard techniques. In these types of problems, tractable filters are often proposed on an *ad hoc* basis of an approximating system that reasonably represents the true dynamics. For example, using this informal idea, approximating filters based on hidden Markov model (HMM) and Kalman filter have been exploited in a wide range of signal and image processing applications, see [1], [2], [3].

Despite the successful application of approximate filters in a large number of applications, conditions that ensure reasonable filter behaviour have not yet been completely established in many situations. Specifically, only a small number of stability type results involving error in model dynamics have been presented. These include stability with respect to model mismatch for Kalman filters [4], [5] and particle filters [6], [7], [8]. Further, some convergence and stochastic stability type results for extended Kalman filters are presented in [9], [10].

One apparent method for developing a reasonable model approximation (and hence an approximate filter) is through the application of classical data-based model inference techniques to infer the approximate model, within a useful class, that best matches the sample measurement and state sequences generated by the true model. Data-based model inference is a classical signal processing problem that has been solved, over many decades, using a variety of techniques, including: information theory and entropy based

techniques (such as Akaike's information criterion [11]); maximum-likelihood based techniques [12] (such as the EM algorithm [13], [14]); and prediction error based techniques [15], [16]. However, the use of data inference to identify an approximating model does not immediately imply acceptable filtering performance will be achieved.

Another alternative avenue for attempting mismatch filter design is through the application of robust filter design techniques such as H^∞ , risk-sensitive or minimax type approaches [17], [18]. These robust techniques seek a (compromise) filter design that offers conservative filter performance under some description of system uncertainty (or model mismatch). It is worth emphasizing that relative entropy descriptions of the system uncertainty (or model error) have been used in a number of recent robust filter design approaches [17], [18].

In other related work, relative entropy concepts have been identified as an important design criterion in a range of model approximation problems [3], [14], [17], [18], [21], [22]. In specific modeling situations, such as HMMs, it has been recently shown that the relative entropy rate (RER) between the joint state and measurement processes of two HMMs allows model parameters to be related to the relative likelihood that two compared HMMs produce a particular output sequence [21], [22]. Importantly, in these types of problems, RER concepts allow HMM parameter values to be directly interpreted in terms of expected filter performance and hence, RER concepts can be exploited in the design of HMMs [3]. These results motivate consideration of relative entropy concepts in general filtering problems.

In this paper, we establish practical stability of general approximating filters with respect to modeling errors under some mild assumptions, including one-step (or local) consistency and forgetting properties. Moreover, we show how relative entropy concepts can be exploited to establish the required local consistency conditions without reference to specific property of the filtering equations. The results of this paper are established using the local consistency techniques that have previously been used to establish semi-global practical stability results for discretisation of nonlinear controllers [23].

This paper is structured as follows: In Section II, we introduce our nominal dynamics, our information state concepts, our modeling approximations, our assumptions, and the concepts of relative entropy rate. In Section III, we then establish some important consistency results and the main practical stability results of this paper. In Section IV, we illustrate our results in the case of HMM approximation.

O. Techakesari and J.J. Ford are with the School of Engineering, Queensland University of Technology, Brisbane QLD 4001, Australia. o.techakesari@qut.edu.au, j2.ford@qut.edu.au

D. Nešić is with the Department of Electrical and Electronic Engineering, The University of Melbourne, Parkville VIC 3010, Australia. dnesic@unimelb.edu.au

Some conclusions are presented in Section V.

II. PROBLEM FORMULATION

A. Dynamics

For the time step $k \geq 0$, we will consider the following state process $x_k \in \mathbb{R}^n$ and measurement process $y_k \in \mathbb{R}^m$,

$$\begin{aligned} x_{k+1} &= f(x_k) + v_k \\ y_k &= c(x_k) + w_k \end{aligned} \quad (1)$$

where x_0 has *a priori* distribution σ_0 , $f(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^n$, and $c(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^m$. Here, $v_k \in \mathbb{R}^n$ and $w_k \in \mathbb{R}^m$ are sequences of independent and identically distributed *i.i.d.* random variables with densities $\phi_v(\cdot)$ and $\phi_w(\cdot)$, respectively. The random variables v_k , w_k , and x_0 are assumed to be mutually independent for all k . Let λ^e denote the exact (or true) model of the x_k and y_k processes. We will use the shorthand $y_{[\ell, m]}$ to denote the measurement sequences $\{y_\ell, \dots, y_m\}$. We likewise define $x_{[\ell, m]}$.

Throughout the rest of this paper, we will consider all processes to be defined on a probability space $(\Omega, \mathcal{F}, P)$ where Ω is defined to consist of all infinite sequences $\{x_0, \dots, x_k, \dots; y_1, \dots, y_k, \dots\}$ (with elements $\omega \in \Omega$), $\mathcal{F}$ is defined to be a σ -algebra generated by these sequences, and P will be a probability measure given by Kolmogorov extension theorem applied to these sequences [24]. Finally, we will let $\mathcal{Y}_{[1, k]}$ denote the complete filtration generated by the sequence $y_{[1, k]}$, see [15, p. 18].

In filtering, we are often interested in the conditional mean estimate of x_k given the measurements $y_{[1, k]}$ and the *a priori* distribution σ_0 , which can be defined, when it exists, as:

$$\hat{x}_{k|[1, k], \sigma_0}^e \triangleq E[x_k | \mathcal{Y}_{[1, k]}, \sigma_0] \quad (2)$$

for all $k > 0$, where $E[\cdot]$ denotes the expectation operation corresponding to P . Similarly, $\hat{x}_{k|[\ell, m], \sigma_{\ell-1}}^e$ will denote the conditional mean estimate at time k , given the measurement sequences $y_{[\ell, m]}$ and the distribution $\sigma_{\ell-1}$ of $x_{\ell-1}$ at time $\ell - 1$.

Unfortunately, in many situations, it may not be possible to implement a filter that produces $\hat{x}_{k|[1, k], \sigma_0}^e$ (for example, such a filter may be computationally intractable). In this paper, we are interested in the performance of sub-optimal or approximate filters that provide approximate estimates for our system state, x_k .

B. Normalised Information State

We now introduce some information state concepts that describe our estimation operations. Consider the space $L^\infty(\mathbb{R}^n)$ which includes $L^1(\mathbb{R}^n)$; see [25] for an introduction into vector space concepts. We will introduce the $\langle \cdot, \cdot \rangle$ notation to denote the operation of $\xi(\cdot) \in L^1(\mathbb{R}^n)$ and $\gamma(\cdot) \in L^\infty(\mathbb{R}^n)$ as $\langle \xi, \gamma \rangle \triangleq \int_{\mathbb{R}^n} \xi(x) \gamma(x) dx$. The $\langle \cdot, \cdot \rangle$ operation is a linear operation in the sense that $\langle \xi + \psi, \gamma \rangle = \langle \xi, \gamma \rangle + \langle \psi, \gamma \rangle$ where $\xi, \psi \in L^1(\mathbb{R}^n)$ and $\gamma \in L^\infty(\mathbb{R}^n)$. We will also introduce the L_1 norm on information state [25]:

$$\|\xi(\cdot)\|_1 \triangleq \int_{x \in \mathbb{R}^n} |\xi(x)| dx. \quad (3)$$

Let $\bar{L}^1(\mathbb{R}^n) \subset L^1(\mathbb{R}^n)$ denote functions in $L^1(\mathbb{R}^n)$ that have L_1 norm equal to 1 in that $\bar{L}^1(\mathbb{R}^n) \triangleq \{\xi(\cdot) : \xi(\cdot) \in L^1(\mathbb{R}^n) \text{ and } \|\xi(\cdot)\|_1 = 1\}$. We can now define a normalised information state process $\sigma_k^e(\cdot) \in \bar{L}^1(\mathbb{R}^n) : \mathbb{R}^n \rightarrow \mathbb{R}$, based on the true model, by

$$\langle \sigma_k^e, \gamma \rangle = E[\gamma(x_k) | \mathcal{Y}_{[1, k]}, \sigma_0] \quad (4)$$

for all $k > 0$, and all test functions $\gamma(\cdot) \in L^\infty(\mathbb{R}^n)$, where $\sigma_0 \in \bar{L}^1(\mathbb{R}^n)$ is the *a priori* distribution of x_0 . This definition highlights that the normalised information state $\sigma_k^e(\cdot)$ can be interpreted as a conditional probability density function of x_k given measurement sequences $y_{[1, k]}$ and *a priori* distribution σ_0 . In particular, when it exists, we can write our conditional mean estimate as

$$\hat{x}_{k|[1, k], \sigma_0}^e = \int_{x \in \mathbb{R}^n} \sigma_k^e(x) x dx \quad (5)$$

We also consider an unnormalised information state $\sigma_k^{e|u}(\cdot) \in L^1(\mathbb{R}^n)$ which provides a method of calculating $\sigma_k^e(\cdot)$. For all $k > 0$, the unnormalised information state is given by [15, Ch.5]

$$\sigma_k^{e|u}(x) = \frac{\phi_w(y - c(x))}{\phi_w(y)} \int_{\mathbb{R}^n} \phi_v(x - f(z)) \sigma_{k-1}^e(z) dz \quad (6)$$

for $x \in \mathbb{R}^n$, where $\sigma_0^e \in L^1(\mathbb{R}^n) = \sigma_0$. The normalised information state $\sigma_k^e(\cdot) \in \bar{L}^1(\mathbb{R}^n)$ can then be written as

$$\sigma_k^e(\cdot) = N_k^{-1} \sigma_k^{e|u}(\cdot) \quad (7)$$

where $N_k = \|\sigma_k^{e|u}(\cdot)\|_1$ is a normalisation factor. We highlight that (6) and (7) together evolve $\sigma_{k-1}^e(\cdot) \in \bar{L}^1(\mathbb{R}^n)$ to produce $\sigma_k^e(\cdot) \in \bar{L}^1(\mathbb{R}^n)$. Here, when required to highlight the initial condition, we will write $\sigma_{k|[1, k], \sigma_0}^e(\cdot)$ to denote the normalised information state $\sigma_k^e(\cdot)$ after evolution by measurements $y_{[1, k]}$ from initial distribution σ_0 at time $k = 0$ (and sometimes further shortened to $\sigma_k^e(\sigma_0)$, especially when used in sub-scripts of other quantities). Similarly, $\sigma_{k|[\ell+1, k], \sigma_\ell}^e(\cdot)$ will denote $\sigma_k^e(\cdot)$ after evolution by measurements $y_{[\ell+1, k]}$ from distribution $\sigma_\ell^e(\cdot)$ at time $k = \ell$. Importantly, the distributive nature of the information state recursions means that $\sigma_{k|[\ell+1, k], \sigma_\ell}^e(\sigma_0)(\cdot) = \sigma_{k|[1, k], \sigma_0}^e(\cdot)$. We highlight that, although not explicitly shown in our notation, all these information state quantities are also $\mathcal{Y}_{[1, k]}$ -measurable random variables.

C. Parameterised Class of Approximating Models

Let $h > 0$ parameterise a class of approximating models (for example, h might be a spatial discretisation size). For each h , let us consider the following approximating model of x_k and y_k (for time step $k \geq 0$):

$$\begin{aligned} x_{k+1} &= f^h(x_k) + v_k^h \\ y_k &= c^h(x_k) + w_k^h \end{aligned} \quad (8)$$

where x_0 has *a priori* distribution σ_0^h , $f^h(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^n$, and $c^h(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^m$. Here, $v_k^h \in \mathbb{R}^n$ and $w_k^h \in \mathbb{R}^m$ are *i.i.d.* random variables with densities $\phi_v^h(\cdot)$ and $\phi_w^h(\cdot)$, respectively, and v_k^h , w_k^h , and x_0 are assumed to be mutually

independent. We will use λ^h to denote this model of x_k and y_k processes. Corresponding to each approximating model λ^h , we introduce a new probability measure P^h which allows us to relate the true model and these approximating models on the common measure space $(\Omega, \mathcal{F})$; such a measure can be defined through the Kolmogorov extension theorem, see [24].

For a given $h > 0$, we can also define the conditional mean estimate associated with the approximate model as:

$$\hat{x}_{k|[1,k],\sigma_0^h}^h \triangleq E^h [x_k | \mathcal{Y}_{[1,k]}, \sigma_0^h], \quad (9)$$

where $E^h[\cdot]$ denotes the expectation operation defined by measure P^h .

Similar to the true model, we also define a normalised information state process $\sigma_k^h(\cdot) \in \bar{L}^1(\mathbb{R}^n) : \mathbb{R}^n \rightarrow \mathbb{R}$, for our h -class of models, as

$$\langle \sigma_k^h, \gamma \rangle = E^h [\gamma(x_k) | \mathcal{Y}_{[1,k]}, \sigma_0^h] \quad (10)$$

for all $k > 0$, all $h > 0$ and all test functions $\gamma(\cdot) \in L^\infty(\mathbb{R}^n)$ where $\sigma_0^h \in \bar{L}^1(\mathbb{R}^n)$ is the *a priori* distribution of x_0 . Furthermore, we can define a recursion for the unnormalised information state process $\sigma_k^{h|u}(\cdot) \in L^1(\mathbb{R}^n)$ as $\sigma_k^{h|u}(x) = \frac{\phi_w^h(y - c^h(x))}{\phi_w^h(y)} \int_{\mathbb{R}^n} \phi_v^h(x - f^h(z)) \sigma_{k-1}^h(z) dz$ so that a normalised information state $\sigma_k^h(\cdot) \in \bar{L}^1(\mathbb{R}^n)$ can be written as $\sigma_k^h(\cdot) = \bar{N}_k^{-1} \sigma_k^{h|u}(\cdot)$ where $\bar{N}_k^{-1} = \left\| \sigma_k^{h|u}(\cdot) \right\|_1$. Again, we highlight that $\sigma_{k-1}^h(\cdot) \in \bar{L}^1(\mathbb{R}^n)$ evolves to $\sigma_k^h(\cdot) \in \bar{L}^1(\mathbb{R}^n)$. As above, we can also write $\sigma_{k|[\ell+1,k],\sigma_0^h}^h(\cdot) = \sigma_{k|[\ell+1,k],\sigma_0^h}^h(\cdot)$.

We are interested in the situations where the quality of the approximation improves as $h \rightarrow 0$ (the meaning of this asymptotic behaviour will be discussed in more detail later). Before this discussion, we introduce some relative entropy concepts.

D. The Relative Entropy between Models

Consider two probability measures μ and ν on the measurable space $(\Omega, \mathcal{F})$. The relative entropy $\mathcal{D}(\mu||\nu)$ of μ with respect to ν is defined as [26]

$$\mathcal{D}(\mu||\nu) \triangleq \begin{cases} \int_{\Omega} \left(\log \frac{d\mu}{d\nu} \right) d\mu; & \text{if } \mu \ll \nu \text{ and} \\ & \left| \log \left(\frac{d\mu}{d\nu} \right) \right| \text{ is integrable} \\ +\infty; & \text{otherwise} \end{cases} \quad (11)$$

where $(d\mu/d\nu)$ is the Radon-Nikodym derivative of μ with respect to ν . Here, we use $\mu \ll \nu$ to indicate that μ is absolutely continuous with respect to ν , in the sense that $\mu = 0$ whenever $\nu = 0$. The relative entropy $\mathcal{D}(\mu||\nu)$ provides a pseudo-distance measure between μ and ν (not a true distance because it is non-symmetric and does not satisfy the triangle inequality).

We now introduce some assumptions about our approximating models that help us prove this type of result.

E. Consistency Assumptions on Approximating Models

Let us say that a function ψ is of class- $\mathcal{K}$ if it is continuous, strictly increasing and $\psi(0) = 0$. Moreover, function β will be of class- $\mathcal{KL}$ if $\beta(\cdot, t)$ is of class- $\mathcal{K}$ for each $t \geq 0$ and $\beta(s, \cdot)$ is decreasing to zero for each $s > 0$.

The following assumptions will hold throughout the remainder of this paper.

A1) Let $N \subset \mathbb{R}$ be a bounded set containing the origin.

The class of approximating filters $\sigma_{k|[1,k],\sigma_0}^h$ is asymptotically stable in $h \in (0, H]$ with respect to initial conditions in N if, there exists $\beta(\cdot, \cdot) \in \mathcal{KL}$ such that, for all $h \in (0, H]$, and all $\|\sigma_0 - \bar{\sigma}_0\| \in N$ (with $\sigma_0, \bar{\sigma}_0 \in \bar{L}^1(\mathbb{R}^n)$) and all $k \geq 0$, we have that

$$\left\| \sigma_{k|[1,k],\sigma_0}^h(\cdot) - \sigma_{k|[1,k],\bar{\sigma}_0}^h(\cdot) \right\|_1 \leq \beta(\|\sigma_0 - \bar{\sigma}_0\|_1, k)$$

P -a.s.

A2) The class of approximating filters $\sigma_{k|[k],\sigma_{k-1}}^h$ is Lipschitz continuous in $h \in (0, H]$ with respect to prior information in the sense that

$$\left\| \sigma_{k|[k],\sigma_{k-1}}^h(\cdot) - \sigma_{k|[k],\bar{\sigma}_{k-1}}^h(\cdot) \right\|_1 \leq K \|\sigma_{k-1}(\cdot) - \bar{\sigma}_{k-1}(\cdot)\|_1 \quad P\text{-a.s.}$$

for all $h \in (0, H]$, all $\sigma_{k-1}, \bar{\sigma}_{k-1} \in \bar{L}^1(\mathbb{R}^n)$, and all $k > 0$, where $K > 0$ is a finite constant.

Remark 1: Assumption A1 is an abstract version of the asymptotic stability property with respect to initial conditions (or exponential forgetting of initial conditions) that is often encountered in discussion of filter behaviour (for example, see [4], [5], [28]). As an example, if $\sigma_{k|[1,k],\sigma_0}^h$ corresponds to a class of Kalman filters, then under observability and other mild conditions, exponential forgetting of covariance matrix and conditional mean estimate, with respect to initial conditions, can be shown [4], [5]. Hence, using the definition of the L_1 norm, and various algebraic manipulations, it can be shown that

$$\left\| \sigma_{k|[1,k],\sigma_0}^h - \sigma_{k|[1,k],\bar{\sigma}_0}^h \right\|_1 \leq \beta(\|\sigma_0 - \bar{\sigma}_0\|_1, k) \quad P^h\text{-a.s.}$$

where $\beta(\|\sigma_0 - \bar{\sigma}_0\|_1, k) = \alpha_1 \|\sigma_0 - \bar{\sigma}_0\|_1 e^{-\alpha_2 k}$ for some $\alpha_1 > 0$ and $\alpha_2 > 0$.

Remark 2: We highlight that asymptotic stability properties with respect to initial conditions (such as those that could be shown when using Kalman filter approximations) would generally be established P^h -a.s. (not P -a.s. as expressed in Assumption A1). However, the mild additional condition that $P^h \gg P$ can be used on results that hold P^h -a.s. to imply that they also hold P -a.s.. Note that $P^h \gg P$ seems to be a natural prerequisite for approximation.

Remark 3: In many situations, the bounded set N appeared in A1 includes all $\sigma_0, \bar{\sigma}_0 \in \bar{L}^1(\mathbb{R}^n)$ and hence, A1 will imply A2.

In the next section, we will use the relative entropy concepts and our consistency assumptions to establish multi-step consistency of estimators before establishing our main practical stability results.

III. PRACTICAL STABILITY OF APPROXIMATING FILTERS

Let us first introduce an important definition.

Definition 3.1: The class of approximating filters $\sigma_{k|[k],\sigma_{k-1}}^h$ is said to be one-step or locally consistent with the true filter $\sigma_{k|[k],\sigma_{k-1}}^e(\cdot)$ in $h \in (0, H]$ if for each $\rho > 0$, there is a $h \in (0, H]$, we have that

$$\left\| \sigma_{k|[k],\sigma_{k-1}}^e(\cdot) - \sigma_{k|[k],\sigma_{k-1}}^h(\cdot) \right\|_1 \leq \rho \quad P\text{-a.s.} \quad (12)$$

for all initial conditions $\sigma_{k-1} \in \bar{L}^1(\mathbb{R}^n)$

Lemma 3.1: Consider a state process $x_{[0,k]}$ and a measurement process $y_{[0,k]}$ generated by λ^e . The L_1 norm of the approximating filter $\sigma_{k|[k],\sigma_{k-1}}^h$ is bounded by the relative entropy of conditional probability density functions in the sense that

$$\left\| \sigma_{k|[k],\sigma_{k-1}}^e(\cdot) - \sigma_{k|[k],\sigma_{k-1}}^h(\cdot) \right\|_1 \leq B \sqrt{\mathcal{D} \left(\sigma_{k|[k],\sigma_{k-1}}^e(\cdot) \left\| \sigma_{k|[k],\sigma_{k-1}}^h(\cdot) \right\| \right)} \quad P\text{-a.s.} \quad (13)$$

for all initial conditions $\sigma_{k-1} \in \bar{L}^1(\mathbb{R}^n)$ and all $k \geq 0$, where $B > 0$ is a finite constant.

Proof: From [26, Lemma 11.6.1], we obtain

$$\left(\left\| \sigma_{k|[k],\sigma_{k-1}}^e(\cdot) - \sigma_{k|[k],\sigma_{k-1}}^h(\cdot) \right\|_1 \right)^2 \leq \bar{B} \mathcal{D} \left(\sigma_{k|[k],\sigma_{k-1}}^e(\cdot) \left\| \sigma_{k|[k],\sigma_{k-1}}^h(\cdot) \right\| \right) \quad P\text{-a.s.} \quad (14)$$

for some positive finite constant $\bar{B}$ (which is independent of h). The lemma statement then follows under the square root operation. ■

Lemma 3.1 establishes that relative entropy provides a method for establishing one-step consistency for a class of estimators (as long as $\mathcal{D} \left(\sigma_{k|[k],\sigma_{k-1}}^e(\cdot) \left\| \sigma_{k|[k],\sigma_{k-1}}^h(\cdot) \right\| \right) \rightarrow 0$ as $h \rightarrow 0$). We now introduce the concept of multi-step consistency.

Definition 3.2: The class of approximating filters $\sigma_{k|[1,k],\sigma_0}^h(\cdot)$ is said to be multi-step consistent with the true estimator $\sigma_{k|[1,k],\sigma_0}^e(\cdot)$ in $h \in (0, H]$ if for each finite $L \geq 2$ and each $\eta(L) > 0$, there is a $h \in (0, H]$ such that, for all initial conditions $\sigma_0 \in \bar{L}^1(\mathbb{R}^n)$ and all $k \in [1, L]$ we have that

$$\left\| \sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0}^h(\cdot) \right\|_1 \leq \eta(L) \quad P\text{-a.s.} \quad (15)$$

We now establish conditions under which this multi-step consistency condition holds.

Lemma 3.2: Consider a state process $x_{[0,k]}$ and a measurement process $y_{[0,k]}$ generated by λ^e . Consider a class of approximating filters $\sigma_{k|[1,k],\sigma_0}^h(\cdot)$. Assume that there is a finite $H > 0$ such that A2 and the one-step consistency condition hold in $h \in (0, H]$, then the class of approximating filters is multi-step consistent with the true filter $\sigma_{k|[1,k],\sigma_0}^e$ in $h \in (0, H]$.

Proof: From our definition of one-step consistency, we have that for each $\rho > 0$, there is a $H > 0$ such that, at time $k = 1$,

$$\left\| \sigma_{1|[1],\sigma_0}^e(\cdot) - \sigma_{1|[1],\sigma_0}^h(\cdot) \right\|_1 \leq \rho \quad P\text{-a.s.} \quad (16)$$

for all $h \in (0, H]$ and all initial conditions $\sigma_0 \in \bar{L}^1(\mathbb{R}^n)$.

At time $k = 2$, we have that

$$\begin{aligned} & \left\| \sigma_{2|[1,2],\sigma_0}^e(\cdot) - \sigma_{2|[1,2],\sigma_0}^h(\cdot) \right\|_1 \\ &= \left\| \sigma_{2|[2],\sigma_1^e(\sigma_0)}^e(\cdot) - \sigma_{2|[2],\sigma_1^h(\sigma_0)}^h(\cdot) \right\|_1 \\ &\leq \left\| \sigma_{2|[2],\sigma_1^e(\sigma_0)}^e(\cdot) - \sigma_{2|[2],\sigma_1^e(\sigma_0)}^h(\cdot) \right\|_1 \\ &\quad + \left\| \sigma_{2|[2],\sigma_1^e(\sigma_0)}^h(\cdot) - \sigma_{2|[2],\sigma_1^h(\sigma_0)}^h(\cdot) \right\|_1 \\ &\leq \rho + K \left\| \sigma_1^e(\sigma_0) - \sigma_1^h(\sigma_0) \right\|_1 \quad P\text{-a.s.} \\ &= \rho + K \left\| \sigma_{1|[1],\sigma_0}^e(\cdot) - \sigma_{1|[1],\sigma_0}^h(\cdot) \right\|_1 \quad P\text{-a.s.} \\ &\leq \rho + K\rho = (1+K)\rho \quad P\text{-a.s.} \end{aligned} \quad (17)$$

The 2nd step comes from Minkowski's inequality [24, p. 242]. In the 3rd step, we have applied one-step consistency assumption that $\left\| \sigma_{2|[2],\sigma_1^e(\sigma_0)}^e(\cdot) - \sigma_{2|[2],\sigma_1^e(\sigma_0)}^h(\cdot) \right\|_1 \leq \rho$ and using Assumption A2, we have that $\left\| \sigma_{2|[2],\sigma_1^e(\sigma_0)}^h(\cdot) - \sigma_{2|[2],\sigma_1^h(\sigma_0)}^h(\cdot) \right\|_1 \leq K \left\| \sigma_1^e(\sigma_0) - \sigma_1^h(\sigma_0) \right\|_1$. In the 4th step, we have used that $\sigma_1^e(\sigma_0)$ is shorthand for $\sigma_{1|[1],\sigma_0}^e(\cdot)$, etc.

Now at time $k = 3$, we obtain

$$\begin{aligned} & \left\| \sigma_{3|[1,3],\sigma_0}^e(\cdot) - \sigma_{3|[1,3],\sigma_0}^h(\cdot) \right\|_1 \\ &\leq \left\| \sigma_{3|[3],\sigma_2^e(\sigma_0)}^e(\cdot) - \sigma_{3|[3],\sigma_2^e(\sigma_0)}^h(\cdot) \right\|_1 \\ &\quad + \left\| \sigma_{3|[3],\sigma_2^e(\sigma_0)}^h(\cdot) - \sigma_{3|[3],\sigma_2^h(\sigma_0)}^h(\cdot) \right\|_1 \\ &\leq \rho + K \left\| \sigma_2^e(\sigma_0) - \sigma_2^h(\sigma_0) \right\|_1 \quad P\text{-a.s.} \\ &= \rho + K \left\| \sigma_{2|[1,2],\sigma_0}^e(\cdot) - \sigma_{2|[1,2],\sigma_0}^h(\cdot) \right\|_1 \quad P\text{-a.s.} \\ &\leq \rho + K(1+K)\rho = (1+K+K^2)\rho \quad P\text{-a.s.} \end{aligned} \quad (18)$$

By induction, for each $L \geq 2$ and each $\eta(L) > 0$, from our one-step consistency we can select a $\rho > 0$, and hence $h \in (0, H]$, such that

$$\begin{aligned} \left\| \sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0}^h(\cdot) \right\|_1 &\leq \rho \sum_{i=0}^{L-1} K^i \quad P\text{-a.s.} \\ &\leq \eta(L) \quad P\text{-a.s.} \end{aligned} \quad (19)$$

for all initial conditions $\sigma_0 \in \bar{L}^1(\mathbb{R}^n)$. This establishes the lemma statement. ■

We now establish the main result of this paper.

Theorem 3.1: (Practical asymptotic stability with respect to modeling errors) Consider a state process $x_{[0,k]}$ and a measurement process $y_{[0,k]}$ generated by λ^e . Assume that there is a finite $H > 0$ such that A1 holds in $h \in (0, H]$ and in N , and that the class of approximating filters is multi-step consistent with the true filter $\sigma_{k|[1,k],\sigma_0}^e(\cdot)$ in $h \in (0, H]$. Then, for any selected $R > 0$, there is a $h \in (0, H]$ such that the class of approximating filters is practically stable in the presence of modeling errors in the sense that

$$\begin{aligned} & \left\| \sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0}^h(\cdot) \right\|_1 \\ &\leq \beta \left(\left\| \sigma_0 - \sigma_0^h \right\|_1, k \right) + R \quad P\text{-a.s.} \end{aligned} \quad (20)$$

for all initial conditions $\|\sigma_0 - \sigma_0^h\|_1 \in N$ (with $\sigma_0, \sigma_0^h \in \bar{L}^1(\mathbb{R}^n)$).

Proof: For the selected R , it follows from Assumption A1 and multi-step consistency condition that for any choice of L , we can find a h such that:

- 1) $2\eta(L) \subset N$,
- 2) $R > \frac{3}{2}(2\eta(L), 0)$
- 3) $\beta(M, L) \leq \eta(L)$ where $M = \sup_{\xi \in N} |\xi|$, and

Then note that from Assumption A1 with $k = 0$ that $\beta(s, 0) \geq s$. Hence, $R > \frac{3}{2}\beta(2\eta(L), 0)$ ensures that $\eta(L) \leq \frac{1}{3}R$.

Now for the selected L , we have for any $\|\sigma_0 - \sigma_0^h\|_1 \in N$ and any $k \in [1, L]$ that

$$\left\| \sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0^h}^h(\cdot) \right\|_1 \leq \eta(L) \quad P\text{-a.s.} \quad (21)$$

Thus we have, for such $(h, k, \|\sigma_0 - \sigma_0^h\|_1)$,

$$\begin{aligned} & \left\| \sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0^h}^h(\cdot) \right\|_1 \\ & \leq \left\| \sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0}^h(\cdot) \right\|_1 \\ & \quad + \left\| \sigma_{k|[1,k],\sigma_0}^h(\cdot) - \sigma_{k|[1,k],\sigma_0^h}^h(\cdot) \right\|_1 \\ & \leq \eta + \beta(\|\sigma_0 - \sigma_0^h\|_1, k) \quad P\text{-a.s.} \end{aligned} \quad (22)$$

Since $\eta(L) < \frac{1}{3}R$, we have that (20) holds for all $\|\sigma_0 - \sigma_0^h\|_1 \in N$ and for all $k \in [1, L]$. It remains to establish this holds for larger k .

We highlight that from our choice of L , (22), and noting that $N \subset M$, we obtain at time $k = L$ that

$$\left\| \sigma_{L|[1,L],\sigma_0}^e(\cdot) - \sigma_{L|[1,L],\sigma_0^h}^h(\cdot) \right\|_1 \leq 2\eta(L) \quad P\text{-a.s.} \quad (23)$$

Now consider the time interval $k \in [L+1, 2L]$ and let $\bar{k} = k - L$. From Assumption A1, time-invariance of the true system λ^e and the approximating model λ^h , the bounds previously established in (21) and (23), and the fact that a ball of radius $2\eta(L)$ is contained in N , we obtain that, for all $k \in [L+1, 2L]$,

$$\begin{aligned} & \left\| \sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0^h}^h(\cdot) \right\|_1 \\ & = \left\| \sigma_{\bar{k}|[L+1,k],\sigma_L^e(\sigma_0)}^e(\cdot) - \sigma_{\bar{k}|[L+1,k],\sigma_L^h(\sigma_0^h)}^h(\cdot) \right\|_1 \\ & \leq \left\| \sigma_{\bar{k}|[L+1,k],\sigma_L^e(\sigma_0)}^e(\cdot) - \sigma_{\bar{k}|[L+1,k],\sigma_L^e(\sigma_0)}^h(\cdot) \right\|_1 \\ & \quad + \left\| \sigma_{\bar{k}|[L+1,k],\sigma_L^e(\sigma_0)}^h(\cdot) - \sigma_{\bar{k}|[L+1,k],\sigma_L^h(\sigma_0^h)}^h(\cdot) \right\|_1 \\ & \leq \eta(L) + \beta(\|\sigma_L^e(\sigma_0) - \sigma_L^h(\sigma_0^h)\|_1, \bar{k}) \quad P\text{-a.s.} \\ & \leq \eta(L) + \beta(2\eta(L), \bar{k}) \quad P\text{-a.s.} \end{aligned} \quad (24)$$

Here, Assumption A1 is used in the 2nd step. In the 3rd step, we have used that $\sigma_L^e(\sigma_0)$ is shorthand for $\sigma_{L|[1,L],\sigma_0}^e(\cdot)$, etc. The 4th step follows (23).

Now we highlight that, for all $k \in [L+1, 2L]$,

$$\begin{aligned} & \left\| \sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0^h}^h(\cdot) \right\|_1 \\ & \leq \eta(L) + \beta(2\eta(L), 0) \quad P\text{-a.s.} \\ & \leq R/4 + R/2 \quad P\text{-a.s.} \\ & \leq R \quad P\text{-a.s.} \end{aligned} \quad (25)$$

We also note that at the end of the interval $[L+1, 2L]$, we have that

$$\begin{aligned} & \left\| \sigma_{2L|[1,2L],\sigma_0}^e(\cdot) - \sigma_{2L|[1,2L],\sigma_0^h}^h(\cdot) \right\|_1 \\ & \leq \eta(L) + \beta(2\eta(L), L) \\ & \leq 2\eta(L) \quad P\text{-a.s.} \end{aligned} \quad (26)$$

The result then follows by induction. $\blacksquare$

The importance of Theorem 3.1 is that if the approximating filter is asymptotically stable with respect to initial conditions and locally consistent with the true filter, then the error between the true and approximating filters is asymptotically small. This means that approximations can be designed that, asymptotically, have any desired level of relative performance. We stress that we only require the error between the approximating filter and differently initialised versions of the approximating filter itself is bounded by $\beta(\cdot, \cdot)$. This assumption, combined with multi-step consistency property, is used to establish the practical stability result.

Remark 4: The role of relative entropy in the presented practical asymptotic stability result is limited to establishing the useful one-step consistency property and we acknowledge that it may be possible to establish the required one-step consistency using other techniques. However, we highlight that relative entropy concepts allow the important one-step consistency property to be established without appealing to the specific nature of the filtering recursions (in comparison, we note that previous stability results of this type have only been established by appealing to specific features of the recursions involved, for example the stability results for particle filters established in [6], [7], [8]).

A. Models with Sufficiently Informative Observations

Let us now introduce an assumption under which our true and approximation models will be said to have sufficiently informative observations.

- C1) The difference between distributions $\Delta\sigma_k^h(\cdot) = \sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0^h}^h(\cdot)$ has light tails in $h \in (0, H]$ in the sense that, for all $h \in (0, H]$ and all $k \geq 0$, we have that

$$\|\Delta\sigma_k^h(x)\|_1 \geq B \|\Delta\sigma_k^h(x)x\|_1 \quad P\text{-a.s.}$$

for some finite constant B (which is independent of h).

Remark 5: Assumption C1 implies that observations should be sufficiently informative so that the information states corresponding to the true and approximating models sufficiently match outside some compact set. As simple example, C1 automatically holds if w_k has compact support.

Moreover, this condition also holds when w_k has Gaussian density (which tends to “localise” the state values enough for C1 to hold). Although we admit this condition seems difficult to establish without examination of the specific filters involved, we highlight that Assumption C1 seems no more restrictive than the usual type of observation model assumptions that appears in establishment of the asymptotic stability with respect to initial conditions, see [19], [20].

We can now establish an important result related to the conditional mean estimates produced by the approximate filters.

Theorem 3.2: Consider a state process $x_{[0,k]}$ and a measurement process $y_{[0,k]}$ generated by λ^e . Assume that there is a finite $H > 0$ such that A1 holds in $h \in (0, H]$ and in N , that C1 holds in $h \in (0, H]$, and that the class of approximating filters is multi-step consistent with the true filter $\sigma_{k|[1,k],\sigma_0}^e(\cdot)$ in $h \in (0, H]$. Then, for any selected $R > 0$, there is a $h \in (0, H]$ such that the class of approximating filters is practically stable in the presence of modeling errors in the sense that

$$\left| \hat{x}_{k|[1,k],\sigma_0}^e - \hat{x}_{k|[1,k],\sigma_0^h}^h \right| \leq \beta \left(\left\| \sigma_0 - \sigma_0^h \right\|_1, k \right) + R \quad P\text{-a.s.} \quad (27)$$

for all initial conditions $\left\| \sigma_0 - \sigma_0^h \right\|_1 \in N$ (with $\sigma_0, \sigma_0^h \in \bar{L}^1(\mathbb{R}^n)$).

Proof: We first note that

$$\begin{aligned} \hat{x}_{k|[1,k],\sigma_0}^e - \hat{x}_{k|[1,k],\sigma_0^h}^h &= \int_{x \in \mathbb{R}^n} \left(\sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0^h}^h(\cdot) \right) x dx. \end{aligned} \quad (28)$$

By taking the magnitude, we obtain

$$\begin{aligned} \left| \hat{x}_{k|[1,k],\sigma_0}^e - \hat{x}_{k|[1,k],\sigma_0^h}^h \right| &= \left| \int_{x \in \mathbb{R}^n} \left(\sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0^h}^h(\cdot) \right) x dx \right| \\ &\leq \left\| \left(\sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0^h}^h(\cdot) \right) x \right\|_1 \\ &\leq \bar{B} \left\| \left(\sigma_{k|[1,k],\sigma_0}^e(\cdot) - \sigma_{k|[1,k],\sigma_0^h}^h(\cdot) \right) \right\|_1 \end{aligned} \quad (29)$$

for some finite positive constant $\bar{B}$ (which is independent of h).

In the 2nd step, we have used the definition of $\|\cdot\|_1$ operation and the integral property that $\|\xi(\cdot)x\|_1 \geq \left| \int_{x \in \mathbb{R}^n} \xi(x) dx \right|$. The 3rd step follows from Assumption C1. The theorem statement then follows from the result of Theorem 3.1. ■

IV. CASE STUDY

In this section, we illustrate our results in the approximation of a scalar continuous-valued nonlinear system by a hidden Markov model (HMM). For this purpose, we will create a special interpretation of the HMM’s underlying discrete-state process via a spatial “blurred” version of the state process.

For presentation purposes we limit our example to a scalar example (but this approach can be generalised). Consider

a scalar true model with dynamics $x_k \in \mathbb{R}$ described by (1), where $f(x) = \text{mod}(ax + b, 2b) - b$ for some $a \in \bar{N}$ and $b > 0$, $c(x) = x$, $\phi_w(\cdot)$ is a zero-mean unit-variance Gaussian density, and $\phi_v(\cdot)$ is some density function with support only in the interval $[-1, 1]$ (here $\text{mod}(\cdot, \cdot)$ is the modulus operation). Under these assumptions $x_k \in S_x = [-(b+1), b+1]$ for all k .

The restriction of dynamics to a finite region of state-space is somewhat limiting but is also understandable considering the nature of HMMs (also, admittedly, the restriction to bounded region immediately implies that filtering errors are finite, but our results establish more than this property).

We will now introduce a HMM process which approximates the true system described above.

Let $e_i = [0, \dots, 0, 1, 0, \dots, 0]' \in \mathbb{R}^{\bar{N}}$ denote an indicator vector with 1 in the i th position and zero elsewhere, and let $\bar{N}$ denotes the number of HMM states (to simplify later construction, we will assume $\bar{N}$ is even). At time k , we will let $X_k \in \{e_1, e_2, \dots, e_{\bar{N}}\}$ denote the state of the HMM process. This HMM state process is described by a transition probabilities matrix A with ij th element, $A_{ij} = p(X_{k+1} = e_i | X_k = e_j)$, where $p(\cdot)$ is the probability law describing our HMM state process. The HMM state process is also assumed to have an initial probabilities vector π with i th element, $\pi_i = p(X_0 = e_i)$. The measurement process y_k associated with the HMM state process is described by an output probability matrix $B(y_k)$ given by $B(y_k) = \text{diag}([p(y_k | X_k = e_1), \dots, p(y_k | X_k = e_{\bar{N}})])$, where $\text{diag}(x)$ is the diagonal matrix with x on its diagonal.

The HMM state process is assumed to exist on the following spatial discretisation of S_x . Let G be the spatial grid (with $\bar{N}$ grid points) that approximates S_x such that $G = \{x : x = \pm mh\}$ where $m = 1, \dots, (\bar{N}/2)$, and $h = 2(b+1)/\bar{N}$ is spacing parameter. This allows us to relate each grid point with a HMM state value. We will use $G(e_i)$ to denote the specific location on G corresponding to state value e_i .

We will now introduce a blurred approximating process associated with this HMM state process. Let $C(e_i)$ denote a h -sized cell containing grid location $G(e_i)$. The cell $C(e_i)$ is used to describe the region of S_x represented by the state value e_i . The cells are assumed designed to completely cover S_x in the sense that for all $x \in S_x$, $x \in C(e_i)$ for some e_i . Conversely, let $e(x)$ be the indicator vector denoting the cell containing the value x (ie. the inverse association), that is $x \in C(e(x))$ for all possible values of $x \in S_x$. We will also assume that the boundaries between adjacent cells are not shared. For approximation purposes, we define x_k^a to be a blurred version of X_k , with the properties that, for all $k \geq 0$, $x_k^a \in C(X_k)$, $X_k = e(x_k^a)$ and x_k^a has uniform distribution over the cell $C(X_k)$. Further, we assume that $p(y_k | x_k^a) = p(y_k | X_k = e(x_k^a))$ for all possible values of y_k and x_k^a . We also assume that the grid points $G(e_i)$ are centred in their corresponding cells $C(e_i)$ so that $E[x_k^a | X_k = e(x_k^a)] = G^h(e(x_k^a))$ for all possible values of x_k^a . At time k , the information state associated with this blurred process, given the measurements $y_{[1,k]}$ and an blurred initial condition

corresponding π , can be written as:

$$\sigma_{k|[1,k],\sigma_0}^{a,h}(x) = \frac{1}{h} e(x)' \hat{X}_{k|[1,k],\pi}$$

where $\hat{X}_{k|[1,k],\pi}$ denotes the HMM filter estimate at time k given the measurements $y_{[1,k]}$ and the initial condition π (see [15]). Note that the inner product $e(x)' \hat{X}_{k|[1,k],\pi}$ simply extracts the element of $\hat{X}_{k|[1,k],\pi}$ corresponding to the filtered probability of being in cell $C(x)$.

We now consider the application of Theorem 3.1 and Theorem 3.2 to the approximate filter $\sigma_{k|[1,k],\sigma_0}^{a,h}(x)$ by introducing some assumptions about the approximation model. Let $x_{k+1}^+ = E[x_{k+1} - x_k | x_k]$. Let us choose an irreducible and aperiodic A such that

$$A_{ij} = \frac{1}{2} \max \left[0, \min \left(G(e_i) + \frac{h}{2}, f(G(e_j)) + 1 \right) - \max \left(G(e_i) - \frac{h}{2}, f(G(e_j)) - 1 \right) \right],$$

so that the state dynamics are matched in means and variances (that is, locally consistent as suggested in [27]) in the sense that

- 1) $E[G(X_{k+1}) - G(X_k) | X_k = e(x_k)] = x_{k+1}^+ + \alpha_1 h$ for some $\alpha_1 > 0$ and for all x_k ,
- 2) $E[(G(X_{k+1}) - G(X_k) - x_{k+1}^+)^2 | X_k = e(x_k)] = E[(x_{k+1} - x_k - x_{k+1}^+)^2 | x_k] + \alpha_2 h$ for some $\alpha_2 > 0$ and for all x_k .

Assume the observation model $p(y_k | X_k) = \phi_w(y_k - G(X_k))$. Under these assumptions on HMM parameters (specifically, irreducible and aperiodic A and positive observation density $p(y_k | X_k)$), Theorem 2.2 of [28] shows that corresponding HMM filter is exponential forgetting in the sense that

$$\|\hat{X}_{k|[1,k],\pi} - \hat{X}_{k|[1,k],\pi_0}\|_1 \leq \beta^{HMM}(\|\pi - \pi_0\|_1, k) \text{ } P\text{-a.s.},$$

where π, π_0 are two different initial conditions. Here, $\beta^{HMM}(s, k) = \alpha \bar{\varepsilon}^k \|s\|_1$ where $\bar{\varepsilon} < 1$ and α is a finite constant. Under our definition of x_k^a , the same exponential stability with respect to initial conditions also holds for the blurred process x_k^a , P -a.s.. Also, Assumption A2 holds because HMM conditional mean estimates are linear in previous estimate, see [15].

Now we note that the bounded nature of S_x implies that the information states $\sigma_k^e(\cdot)$ and $\sigma_k^h(\cdot)$ have compact support (which is independent of h) and hence, Assumption C1 holds because

$$\begin{aligned} \|\Delta \sigma_k^h(x_k) x_k\|_1 &= \int |\Delta \sigma_k^h x_k| dx_k \\ &\leq \int |\Delta \sigma_k^h| |x_k| dx_k \\ &\leq |x_{max}| \int |\Delta \sigma_k^h| dx_k \\ &= B \|\Delta \sigma_k^h\|_1 \end{aligned}$$

where x_{max} is the largest absolute value of x in the support of $\Delta \sigma_k^h$. Hence, if we select our HMM design so that

$\mathcal{D} \left(\sigma_{k|[k],\sigma_{k-1}(\cdot)}^e \left\| \sigma_{k|[k],\sigma_{k-1}(\cdot)}^h \right\| \right) \rightarrow 0$ as $h \rightarrow 0$, then Assumptions A2 can be used to give that the approximating HMM model is multi-step consistent with the true model. Consequently, Assumptions A1, A2, C1 and the multi-step consistency property allow us to apply Theorem 3.1 (and Theorem 3.2) to establish practical stability of the approximating filters with respect to modeling errors. These practical stability results provide the first theoretical justification to the widespread application of HMMs in various approximation problems [1], [2], [3].

V. CONCLUSION

In this paper, we present practical stability with respect to modeling errors results for a range of approximate filtering problems. The results are established using some important filter local consistency and relative entropy concepts. We illustrated the application of our practical stability results in the case of hidden Markov model based approximations.

VI. ACKNOWLEDGMENTS

This research was supported under Australian Research Council's Linkage Projects funding scheme (project number LP100100302) and the Smart Skies Project, which is funded, in part, by the Queensland State Government Smart State Funding Scheme.

REFERENCES

- [1] R. L. Streit and R. F. Barrett, "Frequency line tracking using hidden Markov models," *IEEE Trans. Acoustic, Speech, and Signal Process.*, vol. 38(4), pp. 2659-2676, Apr. 1999.
- [2] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. IEEE*, vol. 77(2), pp. 257-286, Feb. 1989.
- [3] J. Lai and J. J. Ford, "Relative entropy rate based multiple hidden Markov model approximation," *IEEE Trans. Signal Process.*, vol. 58(1), pp. 165-174, Jan. 2010.
- [4] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*, New York: Academic Press, 1970.
- [5] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*, Englewood Cliffs: Prentice-Hall, 1979.
- [6] D. Crisan and K. Heine, "Stability of the discrete time filter in terms of the tails of noise distributions," *J. London Math. Soc.*, vol. 78(2), pp. 441-458, Jul. 2008.
- [7] B. Azimi-Sadjadi and P. S. Krishnaprasad, "Approximate nonlinear filtering and its application in navigation," *Automatica*, vol. 41(6), pp. 945-956, Jun. 2005.
- [8] K. Heine and D. Crisan, "Uniform approximations of discrete-time filters," *Adv. Appl. Prob.*, vol. 40(4), pp. 979-1001, 2008.
- [9] K. Reif, S. Gunther, and E. Yaz, "Stochastic stability of the discrete-time extended Kalman filter," *IEEE Trans. Autom. Control*, vol. 44(4), pp. 714-728, Apr. 1999.
- [10] B. F. La Scala, R. R. Bitmead, and M. R. James, "Conditions for stability of the extended Kalman filter and their application to the frequency tracking problem," *Mathematics of Control, Signals, and Systems*, vol. 8(1), pp. 1-26, 1995.
- [11] H. Akaike, "Stochastic theory of minimal realization," *IEEE Trans. Autom. Control*, vol. 19(6), pp. 667-674, Dec. 1974.
- [12] I. E. Baum and T. Petrie, "Statistical inference for probabilistic functions of finite state Markov chains," *The Annals of Mathematical Statistics*, vol. 37(6), pp. 1554-1563, Dec. 1966.
- [13] G. J. McLachlan and T. Krishnan, *The EM Algorithm and Extensions*, New York: Wiley, 1997.
- [14] Y. Singer and M. K. Warmuth, "Training algorithms for hidden Markov models using entropy based distance functions," *Advances in Neural Information Processing Systems*, vol. 9, pp. 641-647, 1996.
- [15] R. J. Elliott, L. Aggoun, and J. B. Moore, *Hidden Markov Models: Estimation and Control*, Berlin: Springer-Verlag, 1995.

- [16] J. J. Ford and J. B. Moore, "Adaptive estimation of HMM transition probabilities," *IEEE Trans. Signal Process.*, vol. 46(5), pp. 1374-1385, May 1998.
- [17] B. C. Levy and R. Nikoukhah, "Robust least-squares estimation with a relative entropy constraint," *IEEE Trans. Inf. Theory*, vol. 50(1), pp. 89-104, Jan. 2004.
- [18] R. K. Boel, M. R. James, and I. R. Peterson, "Robustness and risk-sensitive filtering," *IEEE Trans. Autom. Control*, vol. 47(3), pp. 451-461, Mar. 2002.
- [19] R. van Handel, "Discrete time nonlinear filters with informative observations are stable," *Elect. Comm. in Probab.*, vol. 13, pp. 562-575, Oct. 2008.
- [20] A. Budhiraja and D. Ocone, "Exponential stability in discrete-time filtering for non-ergodic signals," *Stochastic Processes and their Applications*, vol. 82(2), pp. 245-257, Aug. 1999.
- [21] L. Xie, V. A. Ugrinovskii, and I. R. Petersen, "Probabilistic distances between finite-state finite-alphabet hidden Markov models," *IEEE Trans. Autom. Control*, vol. 50(4), pp. 505-511, Apr. 2005.
- [22] M. Vidyasagar, "Bounds on the Kullback-Leibler divergence rate between hidden Markov models," In *Proc. 46th IEEE Conf. on Decision and Control*, pp. 6160-6165, 2007.
- [23] D. Nesic, A. R. Teel, and P. V. Kokotovic, "Sufficient conditions for stabilization of sampled-data nonlinear systems via discrete-time approximations," *Syst. Control Lett.*, vol. 38(4-5), pp. 259-270, Dec. 1999.
- [24] P. Billingsley, *Probability and Measure*, 3rd ed., New York: Wiley, 1995.
- [25] D. G. Luenberger, *Optimization by Vector Space Methods*, New York: Wiley, 1969.
- [26] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., New York: Wiley, 2005.
- [27] H. J. Kushner and P. Dupuis, *Numerical Methods for Stochastic Control Problems in Continuous Time*, 2nd ed., New York: Springer-Verlag, 2001.
- [28] F. Le Gland and L. Mevel, "Exponential forgetting and geometric ergodicity in hidden Markov models," *Mathematics of Control, Signals, and Systems*, vol. 13(1), pp. 63-93, 2000.