

Risk-sensitive filtering and smoothing for hidden Markov models [☆]

Subhrakanti Dey ^{*}, John B. Moore

*Cooperative Research Centre for Robust and Adaptive Systems, Department of Systems Engineering,
Research School of Information Sciences and Engineering, Australian National University,
Canberra ACT 0200, Australia*

Received 6 August 1994; revised 29 October 1994

Abstract

In this paper, we address the problem of risk-sensitive filtering and smoothing for discrete-time Hidden Markov Models (HMM) with finite-discrete states. The objective of risk-sensitive filtering is to minimise the expectation of the exponential of the squared estimation error weighted by a risk-sensitive parameter. We use the so-called Reference Probability Method in solving this problem. We achieve finite-dimensional linear recursions in the information state, and thereby the state estimate that minimises the risk-sensitive cost index. Also, fixed-interval smoothing results are derived. We show that L_2 or risk-neutral filtering for HMMs can be extracted as a limiting case of the risk-sensitive filtering problem when the risk-sensitive parameter approaches zero.

Keywords: Hidden Markov model; Risk-sensitive filtering; Information state; Fixed-interval smoothing

1. Introduction

Risk-sensitive filtering involves minimisation of the expectation of an exponential in quadratic cost criteria. As opposed to L_2 filtering, (termed as *risk-neutral filtering* in [6]), which achieves the minimisation of a quadratic error criteria, risk-sensitive filtering robustifies the filter against plant and noise uncertainties by penalising all the higher-order moments of the estimation error energy. It also allows a trade-off between optimal filtering for the nominal model case and the average noise situation, and robustness to worst-case noise and model uncertainty by weighting the index of the exponential by a risk-sensitive parameter.

The risk-sensitive filtering problem has been addressed for linear Gauss–Markov signal models in [14]. In a companion paper [6] to the present one, the problem has been solved for a general class of discrete-time nonlinear state-space signal models via the so-called reference probability method and the linear Gauss–Markov signal model has been treated as a special case. It has been seen that risk-sensitive filters are closely related to H_∞ filters [11]. Also, related risk-sensitive control problems are abundant in literature [1, 3, 8, 15].

The problem of extracting finite-state homogeneous Markov chains hidden in white Gaussian noise has been studied as an off-line estimation problem using the well-known Expectation Maximisation (EM)

[☆] The authors wish to acknowledge the funding of the activities of the Cooperative Research Centre for Robust and Adaptive Systems by the Australian Commonwealth Government under the Cooperative Research Centres Program.

^{*} Corresponding author.

algorithm [2, 5, 12]. On-line estimation schemes for Hidden Markov Models (HMM) have been given in [4, 9, 10]. In all these estimation schemes, the so-called “forward variable” [12] is the true filtered estimate which is also a conditional expectation of the state at a certain point of time given the observations up and until that point. The smoothed estimate of the state is obtained as a maximum-likelihood estimate based on a fixed set of observations. These filtering schemes are essentially related to risk-neutral filtering for HMMs.

In this paper, we address the problem of risk-sensitive filtering and smoothing for discrete-time Hidden Markov Models with finite-discrete states. We derive information-state filters which are linear and finite-dimensional. The optimising state estimate is given as the minimising argument of a finite-dimensional sum. Also, the backward filters and unnormalised smoothed conditional probability measures are derived. The derivation techniques are based on a reference probability method which has been developed in [7] and used in [1, 3, 6].

In Section 2, we describe the Hidden Markov Model, formally define the risk-sensitive filtering problem, and then deal with the change of measure and reformulation of the problem in the new probability measure to achieve the filtering and smoothing results. In Section 3, we establish the connection between risk-sensitive and risk-neutral filtering and Section 4 presents some concluding remarks.

2. Hidden Markov models

2.1. State space model

Let \mathcal{X}_k be a discrete-time homogeneous, first-order Markov process belonging to a finite-discrete set. Define $\mathcal{E} \triangleq \{e_1, e_2, \dots, e_N\}$ where $e_i = (0, \dots, 0, 1, 0, \dots, 0)' \in \mathbb{R}^N$ with 1 in the i th position. Without loss of generality, we can assume that $\mathcal{X}_k \in \mathcal{E}$ [13]. We consider this process to be defined on the probability space $(\Omega, \mathcal{F}, \mathcal{P})$ with $\mathcal{F}_k^0 = \sigma\{\mathcal{X}_0, \dots, \mathcal{X}_k\}$ and complete filtration $\{\mathcal{F}_k\}$. The state-space model is then defined by

$$\mathcal{X}_{k+1} = A' \mathcal{X}_k + \mathcal{W}_{k+1}, \quad y_k = C(\mathcal{X}_k) + v_k, \quad (1)$$

where $\mathcal{W}_k, k \in \mathbb{N}$ is a sequence of \mathcal{F}_k -martingale increments and hence $E[\mathcal{W}_{k+1} | \mathcal{F}_k] = 0$. Also y_k is continuous valued belonging to \mathbb{R}^p and $v_k \in \mathbb{R}^p, k \in \mathbb{N}$ is i.i.d with a strictly positive density function ϕ_k .

Due to the Markov nature of \mathcal{X}_k , we can write

$$E[\mathcal{X}_{k+1} | \mathcal{F}_k] = E[\mathcal{X}_{k+1} | \mathcal{X}_k] = A' \mathcal{X}_k, \quad (2)$$

where the entries $\{a_{ij}\}$ are defined as $P(\mathcal{X}_{k+1} = e_j | \mathcal{X}_k = e_i)$. Obviously, $a_{ij} > 0, \forall i, j$ and $\sum_{j=1}^N a_{ij} = 1, \forall i$. We also assume that \mathcal{X}_0 or its distribution is known.

2.2. Problem definition

Our problem objective is to find an estimate $\hat{\mathcal{X}}_k$ of \mathcal{X}_k , where $\hat{\mathcal{X}}_k \in \mathcal{E}$, such that the following criteria is satisfied:

$$\hat{\mathcal{X}}_k = \underset{\zeta \in \mathcal{E}}{\operatorname{argmin}} J_k(\zeta), \quad J_k(\zeta) = E[\theta \exp(\theta \Psi_{0,k}(\zeta)) | \mathcal{Y}_k], \quad \forall k = 0, 1, \dots, \quad (3)$$

where $\theta (> 0)$ is the risk-sensitive parameter, $\{\mathcal{Y}_k\}$ is the complete filtration generated by $\sigma\{y_0, \dots, y_k\}$ and

$$\Psi_{0,k}(\zeta) = \hat{\Psi}_{0,k-1} + \frac{1}{2}(\mathcal{X}_k - \zeta)' Q_k (\mathcal{X}_k - \zeta), \quad Q_k \geq 0 \quad \forall k, \quad (4)$$

where

$$\hat{\Psi}_{m,n} \triangleq \frac{1}{2} \sum_{i=m}^n (\mathcal{X}_i - \hat{\mathcal{X}}_i)' Q_i (\mathcal{X}_i - \hat{\mathcal{X}}_i).$$

Remark 2.1. The assumption that $\hat{\mathcal{X}}_k \in \mathcal{E}$ is to make the search space for $\hat{\mathcal{X}}_k$ finite-dimensional. But this is not an absolute necessity and all the subsequent results can be obtained with $\hat{\mathcal{X}}_k \in \mathbb{R}^N$.

Remark 2.2. Note that (3) holds $\forall k = 0, 1, \dots$ which means that at the k th time instant, only $\hat{\mathcal{X}}_k$ is determined and $\hat{\mathcal{X}}_i, \forall i = 0, 1, \dots, k-1$ are unaffected. Therefore, $\hat{\mathcal{X}}_k$ is a true filtered estimate of \mathcal{X}_k and is not subsequently smoothed.

2.3. Change of measure and reformulated cost index

We define a new measure \bar{P} where $\{y_k\}, k \in \mathbb{N}$ is a sequence of *i.i.d* random variable having density function as ϕ_k . Define

$$\bar{\lambda}_k = \frac{\phi_k(y_k - C(\mathcal{X}_k))}{\phi_k(y_k)}, \quad \bar{\Lambda}_k = \prod_{l=0}^k \bar{\lambda}_l.$$

If we set the Radon–Nikodym derivative $dP/d\bar{P}|_{\mathcal{G}_k} = \bar{\Lambda}_k$, then under P , the random variables $v_k, k \in \mathbb{N}$ are *i.i.d* with density functions ϕ_k . Here, $\{\mathcal{G}_k\}$ is the complete filtration generated by $\mathcal{X}_k^s \triangleq (\mathcal{X}_0, \dots, \mathcal{X}_k)$ and $Y_{k-1} \triangleq (y_0, \dots, y_{k-1})$ (see [7]). Using a version of Bayes' Theorem, we have

$$E[\theta \exp(\theta \Psi_{0,k}(\zeta)) | \mathcal{Y}_k] = \frac{\bar{E}[\bar{\Lambda}_k \theta \exp(\theta \Psi_{0,k}(\zeta)) | \mathcal{Y}_k]}{\bar{E}[\bar{\Lambda}_k | \mathcal{Y}_k]}. \quad (5)$$

Hence, we work under \bar{P} where the modified problem objective is to determine $\hat{\mathcal{X}}_k (\in \mathcal{E})$ such that

$$\hat{\mathcal{X}}_k = \underset{\zeta \in \mathcal{E}}{\operatorname{argmin}} \bar{E}[\bar{\Lambda}_k \theta \exp(\theta \Psi_{0,k}(\zeta)) | \mathcal{Y}_k]. \quad (6)$$

Remark 2.3. In our HMM above, we have taken the measurement noise v_k to be distributed independently of the state. It is possible to have a model where the observation is expressed as

$$y_k = C(\mathcal{X}_k) + \sigma(\mathcal{X}_k) v_k,$$

where $\sigma(\mathcal{X}_k) = \langle \sigma, \mathcal{X}_k \rangle$ since \mathcal{X}_k is a unit vector and $\sigma = (\sigma_1, \dots, \sigma_N)'$. Such a case can be handled by defining $\bar{\lambda}_k = \phi_k(v_k) / \langle \sigma, \mathcal{X}_k \rangle \phi_k(y_k)$. For details, see [7].

2.4. Recursive estimates

Definition 2.1. Define the measure $\alpha_k(e_j)$ to be the unnormalised information state such that

$$\alpha_k(e_j) = \bar{E}[\bar{\Lambda}_{k-1} \theta \exp(\theta \hat{\Psi}_{0,k-1}) \langle \mathcal{X}_k, e_j \rangle | \mathcal{Y}_{k-1}], \quad (7)$$

where $\langle x, y \rangle$ denotes the inner product of two vectors x, y .

Remark 2.4. Note that $\alpha_k(e_j)$ can be interpreted as an information state of an augmented plant where the state includes the actual state of the system and part of the risk-sensitive cost. For details, see [11].

Lemma 2.1. The information state $\alpha_k = (\alpha_k(e_1), \dots, \alpha_k(e_N))'$ obeys the following recursion:

$$\alpha_{k+1} = A' \mathcal{D}'_k \mathcal{B}'_k \alpha_k, \quad (8)$$

where

$$\mathcal{B}_k = \text{diag} \left\{ \frac{\phi_k(y_k - C(e_1))}{\phi_k(y_k)}, \dots, \frac{\phi_k(y_k - C(e_N))}{\phi_k(y_k)} \right\},$$

$$\mathcal{D}_k = \text{diag} \left\{ \exp \left(\frac{\theta}{2} (e_1 - \hat{X}_k)' Q_k (e_1 - \hat{X}_k) \right), \dots, \exp \left(\frac{\theta}{2} (e_N - \hat{X}_k)' Q_k (e_N - \hat{X}_k) \right) \right\}.$$

Proof.

$$\begin{aligned} \alpha_{k+1}(e_i) &= \bar{E}[\bar{\Lambda}_k \theta \exp(\theta \hat{\Psi}_{0,k}) (X_{k+1}, e_i) | \mathcal{Y}_k] \\ &= \bar{E} \left[\frac{\phi_k(y_k - C(X_k))}{\phi_k(y_k)} \exp \left(\frac{\theta}{2} (X_k - \hat{X}_k)' Q_k (X_k - \hat{X}_k) \right) \bar{\Lambda}_{k-1} \theta \exp(\theta \hat{\Psi}_{0,k-1}) X_k' A e_i | \mathcal{Y}_k \right] \\ &= \sum_{j=1}^N \left[\frac{\phi_k(y_k - C(e_j))}{\phi_k(y_k)} \exp \left(\frac{\theta}{2} (e_j - \hat{X}_k)' Q_k (e_j - \hat{X}_k) \right) a_{ji} \alpha_k(e_j) \right] \end{aligned} \quad (9)$$

(using (2) and Definition 2.1).

Writing in the matrix notation completes the proof. \square

Remark 2.5. Note here that the information state filter is linear and finite-dimensional.

Theorem 2.1. The optimizing estimate \hat{X}_k is given by

$$\begin{aligned} \hat{X}_k &= e_{m^*}, \quad m^* = \underset{m}{\text{argmin}} \sum_{j=1}^N \frac{\phi_k(y_k - C(e_j))}{\phi_k(y_k)} \exp \left(\frac{\theta}{2} (e_j - e_m)' Q_k (e_j - e_m) \right) \alpha_k(e_j), \\ m &\in \{1, 2, \dots, N\}. \end{aligned} \quad (10)$$

Proof.

$$\begin{aligned} &\bar{E}[\bar{\Lambda}_k \theta \exp(\theta \Psi_{0,k}(\zeta)) | \mathcal{Y}_k] \\ &= \bar{E} \left[\frac{\phi_k(y_k - C(X_k))}{\phi_k(y_k)} \exp \left(\frac{\theta}{2} (X_k - \zeta)' Q_k (X_k - \zeta) \right) \bar{\Lambda}_{k-1} \theta \exp(\theta \hat{\Psi}_{0,k-1}) | \mathcal{Y}_k \right] \\ &= \sum_{j=1}^N \frac{\phi_k(y_k - C(e_j))}{\phi_k(y_k)} \exp \left(\frac{\theta}{2} (e_j - \zeta)' Q_k (e_j - \zeta) \right) \alpha_k(e_j) \end{aligned} \quad (11)$$

(using Definition 2.1). Using (6), the proof is completed. \square

Remark 2.6. Risk-sensitive filtering for HMMs with discrete measurements can also be carried out in a similar fashion but with a different definition of $\bar{\lambda}_k$ instead of the one introduced in Section 2.3. For details, see Ch. 2 of [7].

2.5. Smoothing

In this section, we present the fixed-interval smoothing results for the Hidden Markov Model. We will find that the unnormalised smoothed conditional probability $\gamma_{k,T}(e_j)$ can be expressed as a product of the information state $\alpha_k(e_j)$ and the backward recursive measure $\beta_{k,T}(e_j)$, $\forall j \in \{1, 2, \dots, N\}$. In fact, this is similar to the result obtained in [12] for risk-neutral HMM filtering where the smoothed estimate is obtained by the “forward-backward” procedure. We define $\gamma_{k,T}(e_j)$ and $\beta_{k,T}(e_j)$ below.

Definition 2.2.

$$\beta_{k,T}(e_j) = \bar{E}[\bar{\Lambda}_{k,T} \exp(\theta \hat{\Psi}_{k,T}) | \mathcal{X}_k = e_j, \mathcal{Y}_T],$$

$$\gamma_{k,T}(e_j) = \bar{E}[\bar{\Lambda}_T \theta \exp(\theta \hat{\Psi}_{0,T}) \langle \mathcal{X}_k, e_j \rangle | \mathcal{Y}_T], \quad \forall j \in \{1, 2, \dots, N\},$$

where $\bar{\Lambda}_{m,n} = \prod_{k=m}^n \bar{\lambda}_k$.

We state the following lemma and theorem without proof.

Lemma 2.2. *The backward recursive measure $\beta_{k,T}(e_j)$ obeys the following backward recursion:*

$$\beta_{k,T}(e_j) = \frac{\phi_k(y_k - C(e_j))}{\phi_k(y_k)} \exp\left(\frac{\theta}{2}(e_j - \hat{X}_k)' Q_k (e_j - \hat{X}_k)\right) \sum_{i=1}^N a_{ji} \beta_{k+1,T}(e_i), \quad (12)$$

where

$$\beta_{T,T}(e_j) = \frac{\phi_T(y_T - C(e_j))}{\phi_T(y_T)} \exp\left(\frac{\theta}{2}(e_j - \hat{X}_T)' Q_T (e_j - \hat{X}_T)\right), \quad \forall j \in \{1, 2, \dots, N\}.$$

Theorem 2.2. *The unnormalised smoothed conditional probability $\gamma_{k,T}(e_j)$ can be expressed as*

$$\gamma_{k,T}(e_j) = \alpha_k(e_j) \beta_{k,T}(e_j), \quad \forall j \in \{1, 2, \dots, N\}. \quad (13)$$

3. Limiting results

In this section, we consider the case when the risk-sensitive parameter θ approaches 0. It is known that for linear and nonlinear state-space signal models with continuous-range states, the risk-neutral (or L_2) filtering problem is recovered as a special case of risk-sensitive filtering when $\theta \rightarrow 0$ [6]. Since, in the case of L_2 filtering, we know that the conditional mean estimate is the minimum variance estimate, we can define the unnormalised information state $\alpha_k(e_j)$ as

$$\alpha_k(e_j) = \bar{E}[\bar{\Lambda}_k \theta \exp(\theta \hat{\Psi}_{0,k-1}) \langle \mathcal{X}_k, e_j \rangle | \mathcal{Y}_k]$$

instead of using Definition 2.1. It is not difficult to show that with this definition, $\alpha_k(e_j)$ will obey the following recursion:

$$\alpha_{k+1}(e_j) = \frac{\phi_{k+1}(y_{k+1} - C(e_j))}{\phi_{k+1}(y_{k+1})} \exp\left(\frac{\theta}{2}(e_j - \hat{X}_k)' Q_k (e_j - \hat{X}_k)\right) \sum_{i=1}^N a_{ij} \alpha_k(e_i).$$

It follows immediately that as $\theta \rightarrow 0$, we have

$$\alpha_{k+1}(e_j) = \frac{\phi_{k+1}(y_{k+1} - C(e_j))}{\phi_{k+1}(y_{k+1})} \sum_{i=1}^N a_{ij} \alpha_k(e_i),$$

which is the well-known risk-neutral recursive filter for HMMs. A variation of this recursion in the so-called forward variable (analogous to $\alpha_k(e_j)$ here) appears in [12]. The MAP estimate is defined as $\hat{X}_k = e_{m^*}$, $m^* = \operatorname{argmax}_m \alpha_k(e_m)$. This implies that the standard filtering equations for HMMs can be obtained as a special case of the risk-sensitive filtering equations when $\theta \rightarrow 0$.

4. Conclusion

The problem of discrete-time filtering and smoothing for Hidden Markov Models with finite-discrete states with an exponential of quadratic cost criteria, termed risk-sensitive filtering in [6] is addressed in this paper

using the reference probability method. A new probability measure is defined where observations are *i.i.d* and the reformulated cost-criteria is minimised to give filtering and smoothing results for HMMs. Finite-dimensional linear recursions are obtained in the information state. Closed-form results for the optimising state estimate and unnormalised smoothed conditional probability measure are given and connection between risk-sensitive filtering and risk-neutral filtering for HMMs has been obtained as a limit result when θ approaches 0. Simulation studies (not reported here due to lack of space) have confirmed that risk-sensitive filters are more robust to uncertain noise environments, specially to coloured and occasionally high noise. They show that there is a certain range of θ where the risk-sensitive filter outperforms the L_2 or risk-neutral filter in terms of desirable robustness properties, but increasing θ beyond this range actually degrades the performance of the risk-sensitive filter.

References

- [1] L. Aggoun, A. Bensoussan, R.J. Elliott and J.B. Moore, Finite-dimensional quasi-linear risk-sensitive control, *Systems Control Lett.* **25** (1995) 151–157.
- [2] L.E. Baum and T. Petrie, Statistical inference for probabilistic functions of finite state Markov chains, *Ann. Math. Statist.* **37** (1966) 1554–1563.
- [3] I.B. Collings, M.R. James and J.B. Moore, An information-state approach to risk-sensitive tracking problems, submitted to: *J. Math. Systems, Estimation, Control* (1994).
- [4] I.B. Collings, V. Krishnamurthy and J.B. Moore, On-line identification of hidden Markov models via recursive prediction error techniques, to appear in: *IEEE Trans. Signal Process.* (1994).
- [5] A.P. Dempster, N.M. Laird and D.B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, *J. Royal Statist. Soc.* **6** (1977) 1–38.
- [6] S. Dey and J.B. Moore, Risk-sensitive filtering and smoothing via reference probability methods, submitted to: *IEEE Trans. Automat. Control* (1994).
- [7] R.J. Elliott, L. Aggoun and J.B. Moore, *Hidden Markov Models: Estimation and Control* (Springer, Berlin, Heidelberg, New York, 1994).
- [8] M.R. James, J.S. Baras and R.J. Elliott, Risk-sensitive control and dynamic games for partially observed discrete-time systems, *IEEE Trans. Automat. Control* **39** (1994) 780–792.
- [9] V. Krishnamurthy, S. Dey and J.P. LeBlanc, Blind equalization of IIR channels using hidden Markov models and extended least squares, accepted for publication in: *IEEE Trans. Signal Process.* (1995).
- [10] V. Krishnamurthy and J.B. Moore, On-line estimation of hidden Markov model parameters based on the Kulback–Leibler information measure, *IEEE Trans. Signal Process.* **41** (1993) 2557–2573.
- [11] J.B. Moore, R.J. Elliott and S. Dey, Risk-sensitive generalizations of minimum variance estimation and control, submitted to: *J. Math. Systems, Estimation Control* (1994).
- [12] L.R. Rabiner, A tutorial on hidden Markov models and selected applications in speech recognition, *Proc. IEEE* **77** (1989) 257–285.
- [13] A. Segall, Recursive estimation from discrete-time point processes, *IEEE Trans. Inform. Theory* **22** (1976) 422–431.
- [14] J.L. Speyer, C. Fan and R.N. Banavar, Optimal stochastic estimation with exponential cost criteria, *Proc. 31st Conf. on Decision and Control*, Vol. 2 (1992) 2293–2298.
- [15] P. Whittle, Risk-sensitive linear/quadratic/Gaussian control, *Adv. Appl. Probab.* **13** (1981) 746–777.