

Recursive Parameter Estimation for Partially Observed Markov Chains*

Robert J. Elliott

Dept. of Statistics & Applied Probability
University of Alberta
Edmonton, Alberta,
Canada T6G 2G1

John B. Moore

Dept. of Systems Engineering
Australian National University
GPO Box 4, Canberra
ACT 2601, Australia

Abstract

Using the reference probability method, a recursive equation is obtained for the unnormalized joint conditional density of a noisily observed Markov chain, and parameters which determine the transition densities and coefficients in the observations.

1 Introduction

By adapting the techniques of our paper, [2], the problem of parameter estimation and filtering for a noisily observed finite state, discrete time Markov chain is considered. Recursive estimates are given for the conditional unnormalized distribution of the chain and parameters. The parameters are treated as additional (constant) signal variables. The case of time varying parameters is covered by the result of [2]. Although the parameters are constant their conditional expectations, given the observations, are random variables and so have conditional, unnormalized distributions. The technique is to introduce an equivalent probability measure, the 'reference probability', under which the signal and observation processes are independent. That is, we consider the discrete time version of the Zakai equation. A related recursive equation can be found in the notes of Brémaud and van Schuppen [1]. However, the existence of an equivalent probability measure is a hypothesis in [1], as is the existence of certain transition kernels. Our change of measure is very natural and is described explicitly. A related idea, not involving dependence on parameters, can be found in Kumar and Varaiya [4].

A recent paper by Krishnamurthy and Moore, [3], obtains locally asymptotic optimal estimates, which are globally optimal in the Gaussian case. These have been implemented in practice; work is currently under way to implement the optimal estimates of the present paper.

¹Research partially supported by NSERC Grant A7964 and the Centre for Robust and Adaptive Control and Boeing (BCAC). The hospitality of the Department of Systems Engineering, Australian National University, is gratefully acknowledged.

2 Dynamics

All processes will be defined on a complete probability space (Ω, F, P) . The discrete time parameter ℓ will take values in $Z^+ = 1, 2, \dots$. Suppose $\{X_\ell\}$, $\ell \in Z$ is a finite state, time homogeneous Markov chain representing the signal process. Without loss of generality, the state space of X can be identified with the set of unit vectors $S = \{e_1, e_2, \dots, e_N\}$, $e_i = (0, \dots, 0, 1, 0, \dots, 0)' \in R^N$, for a suitable integer N . We assume that X_0 is given, or its distribution is known.

We shall suppose the transition probabilities of X depend on a parameter θ^1 which takes values in a measure space $(\Theta^1, \beta^1, \lambda^1)$. The value of θ^1 is unknown and, in this paper, we suppose it is constant. That is, for $1 \leq i, j \leq N$,

$$\begin{aligned} \rho_{ij}(\theta^1) &= P(X_n = e_i \mid X_{n-1} = e_j, \theta^1) \\ &= P(X_1 = e_i \mid X_0 = e_j, \theta^1). \end{aligned}$$

Write $P(\theta^1)$ for the $N \times N$ matrix $(\rho_{ij}(\theta^1))$, $1 \leq i, j \leq N$. Also, $\{F_\ell, \ell \in Z^+\}$ will denote the complete filtration generated by X , that is, for any $n \in Z^+$, F_n is the complete σ -field generated by $X_\ell, \ell \leq n$, and θ .

Lemma 2.1 *The state equation is*

$$X_n = P X_{n-1} + m_n \quad (2.1)$$

where m_n is a (P, F_n) martingale increment.

Proof:

$$\begin{aligned} E[m_n \mid F_{n-1}] &= E[X_n - P X_{n-1} \mid X_{n-1}] \\ &= P X_{n-1} - P X_{n-1} = 0. \end{aligned}$$

□

We suppose the chain X is not observed directly; rather there is an observation process $\{y_\ell\}$, $\ell \in Z^+$, which, for simplicity, we suppose is real valued. The

extension to vector observations is straightforward. The real observations process y has the form

$$y_n = g^{\theta^2}(X_n) + b_n. \quad (2.2)$$

Here θ^2 is a second, constant, unknown parameter taking values in a measure space $(\Theta^2, \beta^2, \lambda^2)$, and the b_ℓ , $\ell \in Z^+$, are real, i.i.d. random variables with a non-zero (positive) density ϕ . The extension to the situation where the b_ℓ , $\ell \in Z^+$, are independent but have possibly different non-zero density functions ϕ_ℓ , is immediate. Because X_n is always one of the unit vectors e_i , $1 \leq i \leq N$, for any $\theta^2 \in \Theta^2$ the function $g^{\theta^2}(\cdot)$ is determined by a vector:

$$g(\theta^2) = (g_1(\theta^2), g_2(\theta^2), \dots, g_N(\theta^2)) \in R^N$$

and

$$g^{\theta^2}(X_n) = \langle g(\theta^2), X_n \rangle$$

where $\langle \cdot, \cdot \rangle$ denotes their inner product in R^N .

Notation 2.2 With $G_\ell^\circ = \sigma\{X_0, X_1, \dots, X_\ell, y_1, \dots, y_{\ell-1}\}$, then $\{G_\ell\}$, $\ell \in Z^+$, will denote the complete filtration generated by G_ℓ° ; $\{\mathcal{Y}_\ell\}$, $\ell \in Z^+$, will denote the complete filtration generated by \mathcal{Y}_ℓ° where $\mathcal{Y}_\ell^\circ = \sigma\{y_0, y_1, \dots, y_\ell\}$.

3 Changes of Measure

On (Ω, F, P) our observation process, therefore, has the form

$$y_n = \langle g(\theta^2), X_n \rangle + b_n \quad (3.1)$$

where the b_n are i.i.d. with density $\phi(b) > 0, \forall b \in R$.

We now wish to define a new probability measure \bar{P} by introducing on the ' n^{th} factor', i.e. on the probability space of b_n , a density

$$\gamma_n(X_n, b_n) := \frac{\phi(\langle g(\theta^2), X_n \rangle + b_n)}{\phi(b_n)}$$

Write

$$\Lambda_n = \prod_{\ell=1}^n \gamma_\ell$$

and define a new probability measure \bar{P} on $(\Omega, \bigcup_{\ell=1}^\infty G_\ell)$ by setting the restriction of the Radon-Nikodym derivative $d\bar{P}/dP$ to G_n equal to Λ_n .

The existence of \bar{P} follows from the theorem of Kolmogorov.

Theorem 3.1 Under \bar{P} the terms of the sequence of random variables $\{y_\ell\}$, $\ell \in Z^+$, are i.i.d. with density function ϕ .

Proof: Applying Bayes' formula, and with $I(\cdot)$ denoting the indicator function,

$$\begin{aligned} \bar{P}(y_n \leq t | G_n) &= \bar{E}[I(y_n \leq t) | G_n] \\ &= \frac{E[\Lambda_n I(y_n \leq t) | G_n]}{E[\Lambda_n | G_n]} \\ &= \frac{\Lambda_{n-1} E[\gamma_n I(y_n \leq t) | G_n]}{\Lambda_{n-1} E[\gamma_n | G_n]}. \end{aligned}$$

Now

$$\begin{aligned} E[\gamma_n | G_n] &= \int_{-\infty}^{\infty} \frac{\phi(\langle g(\theta^2), X_n \rangle + b_n)}{\phi(b_n)} \cdot \phi(b_n) db_n \\ &= \int_{-\infty}^{\infty} \phi(\langle g(\theta^2), X_n \rangle + b_n) db_n \\ &= 1, \end{aligned}$$

so

$$\begin{aligned} \bar{P}(y_n \leq t | G_n) &= \int_{-\infty}^{\infty} \frac{\phi(y_n)}{\phi(b_n)} I(y_n \leq t) \phi(b_n) db_n \\ &= \int_{-\infty}^t \phi(y_n) dy_n \\ &= \bar{P}(y_n \leq t). \end{aligned}$$

That is, under \bar{P} the y_ℓ are i.i.d. with density ϕ . \square

Remarks 3.2 Conversely, what we wish to do is to suppose we start with a probability measure \bar{P} on $(\Omega, \bigcup_{\ell=1}^\infty G_\ell)$ such that under \bar{P} :

- $\{X_\ell\}$, $\ell \in Z^+$, is a Markov chain with a transition matrix $P(\theta^1)$, for some $\theta^1 \in \Theta^1$, so that

$$X_n = P(\theta^1)X_{n-1} + m_n$$

where

$$\bar{E}[m_n | F_{n-1}, \theta^1] = 0,$$

and

- $\{y_\ell\}$, $\ell \in Z^+$, is a sequence of i.i.d. real random variables with density ϕ , $\phi(b) > 0$.

Note, in particular, that under $P(\theta^1)$ the y_ℓ are independent of X .

We wish to construct a probability P such that under P

$$b_\ell := y_\ell - \langle g(\theta^2), X_\ell \rangle$$

is a sequence of i.i.d. random variables with density ϕ . That is, under P the signal and observation process are related as in Section 2. To construct P starting from \bar{P} we must proceed in an inverse fashion, first defining

$$\bar{\gamma}_n(X_n, y_n) = \frac{\phi(b_n)}{\phi(y_n)}$$

and

$$\bar{\Lambda}_n = \prod_{\ell=1}^n \bar{\gamma}_\ell.$$

P is defined by putting the restriction to G_n of the Radon-Nikodym derivative $dP/d\bar{P}$ equal to

$$\frac{dP}{d\bar{P}} \Big|_{G_n} = \bar{\Lambda}_n.$$

A calculation similar to Theorem 3.1 then shows that

$$\begin{aligned} P(b_n \leq t | G_n) &= P(b_n \leq t) \\ &= \int_{-\infty}^t \phi(b_n) db_n. \end{aligned}$$

4 The Recursive Densities

We shall work under probability measure \bar{P} , so that $\{y_\ell\}$, $\ell \in Z^+$, is a sequence of i.i.d. random variables with density ϕ , $\phi(y) > 0$, and $\{X_\ell\}$, $\ell \in Z^+$, is an independent Markov chain.

Notation 4.1 If ϕ_ℓ , $\ell \in Z^+$, is any measurable sequence we shall write

$$\sigma_\ell(\phi_\ell) = \bar{E}[\bar{\Lambda}_\ell \phi_\ell | \mathcal{Y}_\ell].$$

Then $\sigma_\ell(\phi_\ell)$ is an unnormalized conditional expectation of ϕ_ℓ given \mathcal{Y}_ℓ . By a version of Bayes' theorem:

$$\begin{aligned} E[\phi_\ell | \mathcal{Y}_\ell] &= \frac{\bar{E}[\bar{\Lambda}_\ell \phi_\ell | \mathcal{Y}_\ell]}{\bar{E}[\bar{\Lambda}_\ell | \mathcal{Y}_\ell]} \\ &= \frac{\sigma_\ell(\phi_\ell)}{\sigma_\ell(1)}. \end{aligned} \quad (4.1)$$

Equation (4.1) indicates why the unnormalized form $\sigma_\ell(\phi_\ell)$ is investigated. In our model, however, not only ϕ is to be estimated but also the parameters θ^1 and θ^2 .

Notation 4.2 Write $q_n^i(\theta^1, \theta^2)$, $n \in Z^+$, for the unnormalized conditional density such that

$$\begin{aligned} \bar{E}[\bar{\Lambda}_n I(X_n = e_i) I(\theta^1 \in d\theta^1) I(\theta^2 \in d\theta^2) | \mathcal{Y}_n] \\ = q_n^i(\theta^1, \theta^2) d\theta^1 d\theta^2. \end{aligned}$$

The existence of $q_n^i(\theta^1, \theta^2)$ will be discussed below.

Our main result follows. This provides a recursive, closed form update for $q_n^i(\theta^1, \theta^2)$. As in (4.1), the normalized conditional density

$$\begin{aligned} p_n^i(\theta^1, \theta^2) d\theta^1 d\theta^2 \\ = E[I(X_n = e_i) I(\theta^1 \in d\theta^1) I(\theta^2 \in d\theta^2) | \mathcal{Y}_n] \end{aligned}$$

is given by:

$$p_n^i(\theta^1, \theta^2) = \frac{q_n^i(\theta^1, \theta^2)}{\sum_{j=1}^N \int_{\Theta^1} \int_{\Theta^2} q_n^j(\lambda^1, \lambda^2) d\lambda^1 d\lambda^2}.$$

Theorem 4.3

$$q_n^i(\lambda^1, \lambda^2) = \frac{\phi(y_n - g_i(\lambda^2))}{\phi(y_n)} \sum_{j=1}^N \rho_{ij}(\lambda^1) q_{n-1}^j(\lambda^1, \lambda^2). \quad (4.2)$$

Proof: Suppose $f : \Theta^1 \times \Theta^2 \rightarrow R$ is any measurable function. Then

$$\begin{aligned} \bar{E}[f(\theta^1, \theta^2)(X_n, e_i) \bar{\Lambda}_n | \mathcal{Y}_n] \\ = \bar{E}[f(\theta^1, \theta^2)(X_n, e_i) \\ \times \bar{\Lambda}_{n-1} \phi(y_n - \langle g(\theta^2), X_n \rangle) | \mathcal{Y}_n] \phi(y_n)^{-1} \\ = \bar{E}[f(\theta^1, \theta^2)(P(\theta^1)X_{n-1} + m_n, e_i) \\ \times \bar{\Lambda}_{n-1} \phi(y_n - \langle g(\theta^2), e_i \rangle) | \mathcal{Y}_n] \phi(y_n)^{-1} \\ = \bar{E}[f(\theta^1, \theta^2)(P(\theta^1)X_{n-1}, e_i) \\ \times \bar{\Lambda}_{n-1} \phi(y_n - g_i(\theta^2)) | \mathcal{Y}_n] \phi(y_n)^{-1} \\ = \sum_{j=1}^N \bar{E}[f(\theta^1, \theta^2)(X_{n-1}, e_j) \rho_{ij}(\theta^1) \\ \times \bar{\Lambda}_{n-1} \phi(y_n - g_i(\theta^2)) | \mathcal{Y}_{n-1}] \phi(y_n)^{-1} \\ = \phi(y_n)^{-1} \sum_{j=1}^N \int_{\Theta^1} \int_{\Theta^2} f(\lambda^1, \lambda^2) \rho_{ij}(\lambda^1) \\ \times \phi(y_n - g_i(\lambda^2)) q_{n-1}^j(\lambda^1, \lambda^2) d\lambda^1 d\lambda^2. \end{aligned}$$

Therefore,

$$q_n^i(\lambda^1, \lambda^2) = \frac{\phi(y_n - g_i(\lambda^2))}{\phi(y_n)} \sum_{j=1}^N \rho_{ij}(\lambda^1) q_{n-1}^j(\lambda^1, \lambda^2). \quad \square$$

Remarks 4.4 Suppose $\pi = (\pi_1, \dots, \pi_N)$, $\pi_i = P(X_0 = e_i)$, is the initial distribution for X_0 and $h(\lambda^1, \lambda^2)$, (possibly of the form $h^1(\lambda^1)h^2(\lambda^2)$), is the prior density for (θ^1, θ^2) . Then

$$q_0^i(\lambda^1, \lambda^2) = \pi_i h(\lambda^1, \lambda^2),$$

and the up-dated estimates are obtained by substituting in (4.2) for $n \geq 1$.

If the prior estimates are delta functions, (unit masses at particular values of e_i , θ^1 and θ^2), $q_1^i(\lambda^1, \lambda^2)$ and higher unnormalized conditional distributions can be calculated by a formula analogous to (4.2). However, because no noise or dynamics enters into θ^1 and θ^2 , if delta functions are taken as the prior distributions for θ^1 and θ^2 no up-dating takes place; (this is not the case with the distribution for X). This is to be expected because in the filtering procedure the prior does not represent an initial guess for θ^i given no information, but the best estimate for the distribution of θ^i given the initial information. Care must, therefore, be taken with the choice of the prior for θ^1 and θ^2 and, unless there is reason to choose otherwise, priors should be taken so that they have support on the whole range of θ^1 and θ^2 .

5 Multidimensional Observations

Again, suppose the Markov chain has state space $\{e_1, \dots, e_N\}$ and

$$X_n = P(\theta^1)X_{n-1} + m_n, \quad n \in Z^+,$$

for some (unknown) $\theta^1 \in \Theta^1$.

Consider now the case where the observation process is k dimensional with components:

$$\begin{aligned} y_n^1 &= \langle g^1(\theta_1^2), X_n \rangle + b_n^1 \\ y_n^2 &= \langle g^2(\theta_2^2), X_n \rangle + b_n^2 \\ &\vdots \\ y_n^k &= \langle g^k(\theta_k^2), X_n \rangle + b_n^k, \quad n \in Z^+. \end{aligned}$$

Here, for $1 \leq j \leq k$, $g^j(\theta_j^2) = (g_1^j(\theta_j^2), \dots, g_N^j(\theta_j^2))' \in R^N$, where $(\theta_1^2, \theta_2^2, \dots, \theta_k^2) \in \Theta_1^2 \times \Theta_2^2 \times \dots \times \Theta_k^2$, and each Θ_j^2 is a measurable space. Further, the b_ℓ^j , $1 \leq j \leq k$, $\ell \in Z^+$, are a family of independent random variables with non-zero densities $\phi_j(b_\ell^j)$.

The same techniques then establish the following result:

Theorem 5.1 Write $q_n^i(\theta^1, \theta_1^2, \theta_2^2, \dots, \theta_k^2)$, $n \in Z^+$, for the unnormalized conditional density such that

$$\begin{aligned} &\overline{E}[\Lambda_n I(X_n = e_i) I(\theta^1 \in d\theta^1) I(\theta_1^2 \in d\theta_1^2) \\ &\quad I(\theta_2^2 \in d\theta_2^2) \dots I(\theta_k^2 \in d\theta_k^2) | \mathcal{Y}_n] \\ &= q_n^i(\theta^1, \theta_1^2, \dots, \theta_k^2) d\theta^1 d\theta_1^2 d\theta_2^2 \dots d\theta_k^2. \end{aligned}$$

Then

$$\begin{aligned} &q_n^i(\lambda^1, \lambda_1^2, \dots, \lambda_k^2) \\ &= \frac{\sum_{j=1}^k \phi_j(y_n^j - g_j(\lambda_j^2))}{\sum_{j=1}^k \phi_j(y_n^j)} \\ &\quad \cdot \sum_{j=1}^N \rho_{ij}(\lambda^1) q_{n-1}^j(\lambda^1, \lambda_1^2, \dots, \lambda_k^2). \end{aligned}$$

6 Conclusion

Using a change of measure, a recursive expression has been derived for an unnormalized, joint conditional distribution of the state of a partially observed Markov chain, together with the unknown parameters of the transition matrix and observation process. The recursive expression is in fact linear in the unnormalized distribution.

References

- [1] P.M. Brémaud and J.H. Van Schuppen, "Discrete Time Stochastic Systems," Part II, *Estimation Theory*, 1976.
- [2] R.J. Elliott, "A General Recursive Discrete Time Filter," *Jour. of Appl. Prob.* (to appear).
- [3] V. Krishnamurthy and J.B. Moore, "On-line Hidden Markov Model Signal Processing and Elimination of Deterministic Interferences," *IEEE Trans. on Signal Processing*, to appear.
- [4] P.R. Kumar and P. Varaiya, *Stochastic Systems: Estimation, Identification and Adaptive Control*, Prentice Hall Inc., 1986.
- [5] L.R. Rabiner, "A tutorial on Hidden Markov Models and selected applications in speech recognition," *Proc. IEEE*, Vol. 77, No. 2, pp. 257-285, 1989.