

RISK AND INFORMATION IN THE ESTIMATION OF HIDDEN MARKOV MODELS

Vahid R. Ramezani
Steven I. Marcus

Institute for Systems Research
Department of Electrical and computer Engineering
University of Maryland
College Park, MD 20742, U.S.A.

Michael Fu

R. H. School of Business
and Institute for Systems Research
University of Maryland
College Park, MD 20742, U.S.A.

ABSTRACT

In this paper, we consider the relationship between risk-sensitivity and information. Product estimators are introduced as a generalization of Maximum A Posteriori Probability (MAP) estimator for Hidden Markov Models. We study the relationship between the inclusion of higher order moments, the underlying dynamics and the availability of information. Asymptotic periodicity of these estimators and the relationship between risk and information is studied via simulation.

1 INTRODUCTION

It is well known that the probability distribution of an ergodic Markov chain is asymptotically stationary, independent of the initial probability distribution, and that the stationary distribution is the solution to a fixed point problem (Shiryayev 1984). This probability distribution can be viewed as the information state for an estimation problem arising from the Maximum A Posterior Probability Estimator (MAP) estimation of the Markov chain for which no observation is available.

Risk-sensitive filters (Baras and James 1994; Boel, James, and Petersen 1997; Dey and Moore 1995; James, Baras, and Elliot 1994; Whittle 1990) take into account the "higher order" moments of the estimation error. Roughly speaking, this follows from the analytic property of the exponential $e^x = \sum_{k=0}^{\infty} x^k/k!$ so that if Ψ stands for the sum of the error functions over some interval of time then

$$E[\exp(\gamma\Psi)] = E[1 + \gamma\Psi + (\gamma)^2(\Psi)^2/2 + \dots].$$

Thus, at the expense of the mean error cost, the higher order moments are included in the minimization of the expected cost, reducing the "risk" of large deviations and increasing our "confidence" in the estimator. The parameter $\gamma > 0$ controls the extent to which the higher order moments are included. In particular, the first order approximation, $\gamma \rightarrow 0$,

$E[\exp(\gamma\Psi)] \cong 1 + \gamma E\Psi$, indicates that the original minimization of the sum criterion or the risk-neutral problem is recovered as the small risk limit of the exponential criterion.

Another point of view is that the exponential function has the unique algebraic property of converting the sum into a product. A notion of probability for Markov chains follows from this point of view which due to its connection to risk-sensitive filters, will be termed "risk-sensitive probability (RS-probability)". We consider an estimation problem of the states of a Markov chain in which the cost has a product structure. The asymptotic behavior of RS-probability appears to be periodic.

In Section 2, we view the probability distribution of a Markov chain as the information state of an additive optimization problem. RS-probability for Markov chains are introduced in section 3. We show that its evolution is governed by an operator (denoted by F^γ) which can be viewed as a generalization of the usual linear Markov operator. The behavior of this operator is studied in section 4 via simulation. It appears that RS-probability is asymptotically periodic.

2 PROBABILITY AS AN INFORMATION STATE

In Ramezani and Marcus 2002, Ramezani 2001, we studied the exponential (risk-sensitive) criterion for the estimation of HMM's and introduced risk-sensitive filter banks.

The probability distribution of a Markov chain, knowing only initial distribution, determines the most "likely state" in the sense of MAP. In the context of Hidden Markov Models (HMM), the problem can be viewed as that of "pure prediction"; i.e., an HMM whose states are entirely hidden.

Define a Hidden Markov Model as a five tuple $\langle \mathbf{X}, \mathbf{Y}, \mathbf{X}, \mathcal{A}, Q \rangle$; here \mathcal{A} is the transition matrix, $\mathbf{Y} = \{1, 2, \dots, N_Y\}$ is the set of observations and $\mathbf{X} = \{1, 2, \dots, N_X\}$ is the finite set of (internal) states as well as the set of estimates or decisions. We can associate with each $i \in X$, a unit vector in R_X^N whose i^{th} component is

1. In addition, we have that $Q := [q_{x,y}]$ is the $N_X \times N_Y$ state/observation matrix, i.e., $q_{x,y}$ is the probability of observing y when the state is x . We consider the following information pattern. At decision epoch t , the system is in the (unobservable) state $X_t = i$ and the corresponding observation Y_t is gathered, such that

$$P(Y_t = j | X_t = i) = q_{i,j}. \quad (1)$$

In Equation (1), $q_{i,j}$ can be considered as an element of the "observation matrix". The estimators V_t are functions of observations (Y_0, \dots, Y_t) and are chosen according to some specified criterion. Consider a sequence of finite dimensional random variables X_t and the corresponding observations Y_t defined on the common probability space $(\Omega, \mathcal{M}, \mathbf{P})$. Let \hat{X}_t be a Borel measurable function of the filtration generated by observations up to Y_t denoted by \mathcal{Y}_t . The Maximum A Posteriori Probability (MAP) estimator is defined recursively; given $\hat{X}_0, \dots, \hat{X}_{t-1}$, \hat{X}_t is chosen such that the following sum is minimized:

$$E\left[\sum_{i=0}^t \rho(X_i, \hat{X}_i)\right], \quad (2)$$

where

$$\rho(u, v) = \begin{cases} 0 & \text{if } u = v; \\ 1 & \text{otherwise.} \end{cases} \quad (3)$$

The usual definition of MAP as the argument with the greatest probability given the observation follows from (2) and (3). (H. V. Poor 1994). The solution is well known; we need to define recursively an *information state*

$$\sigma_{t+1} = N_Y \cdot \bar{Q}(Y_{t+1}) \mathcal{A}^T \cdot \sigma_t, \quad (4)$$

where $\bar{Q}(y) := \text{diag}(q_{i,y})$, \mathcal{A}^T denotes the transpose of the matrix \mathcal{A} . σ_0 is set equal to $N_Y \cdot \bar{Q}(Y_0) p_0$, where p_0 is the initial distribution of the state and is assumed to be known.

When no observation is available, it is easy to see that $N_Y \cdot \bar{Q}(Y_t) = I$, where I is the identity matrix. Thus, the information state for the prediction case evolves according to $\sigma_{t+1} = \mathcal{A}^T \cdot \sigma_t$ which when normalized is simply the probability distribution of the chain.

This "prediction" optimization problem for a multiplicative cost will be considered next.

3 RS-PROBABILITY FOR MARKOV CHAINS

With the notation of the previous section, given $\hat{X}_0, \dots, \hat{X}_{t-1}$, define \hat{X}_t recursively as the estimator which minimizes the

exponential (risk-sensitive) cost

$$E[\exp\{\gamma \sum_{i=0}^t \rho(X_i, \hat{X}_i)\}], \quad (5)$$

where γ is a strictly positive (risk-sensitive) parameter.

As discussed in the introduction, the exponential criterion allows for the inclusion of higher order moments of the cost and the approximation $E[\exp(\gamma\Psi)] \cong 1 + \gamma E\Psi$ shows that for small values of γ , the additive cost criterion is recovered. The structure of ρ allows for the following simplification of (4):

$$E\left[\prod_{i=0}^t \rho^*(X_i, \hat{X}_i)\right] \quad (6)$$

$$\rho^*(u, v) = \begin{cases} 1 & \text{if } u=v; \\ e^{-\gamma} & \text{otherwise.} \end{cases} \quad (7)$$

(7) can be considered as the "multiplicative" cost. Define an *information state*:

$$\sigma_{t+1}^\gamma = N_Y \cdot \bar{Q}(Y_{t+1}) \mathcal{D}^T(\hat{X}_t) \cdot \sigma_t^\gamma, \quad (8)$$

where $\bar{Q}(y) := \text{diag}(q_{i,y})$, \mathcal{A}^T denotes the transpose of the matrix \mathcal{A} and the matrix \mathcal{D} is defined by

$$[\mathcal{D}(v)]_{i,j} := a_{i,j} \cdot \exp(\gamma\rho(i, v)). \quad (9)$$

σ_0^γ is set equal to $N_Y \cdot \bar{Q}(Y_0) p_0$, where p_0 is the initial distribution of the state and is assumed to be known.

Theorem 1 *The optimization problem (6) is solved recursively by*

$$\hat{X}_t(\sigma) = i \quad \text{if } \sigma^i \geq \sigma^j, \quad \forall j \neq i,$$

where $\sigma = (\sigma^1, \dots, \sigma^{N_X})$ is the value the information state (8) takes at time t .

The estimator obtained is termed Risk-sensitive Maximum A Posteriori Probability (RMAP) estimator.

Proof See (Ramezani 2001).

We next obtain a simplex preserving operator F^γ by assuming that no observation is available and that the initial probability distribution is given. In the risk-neutral context, this operator is simply \mathcal{A}^T which governs the evolution of probability distribution; as the risk-sensitive cost is a generalization of the risk-neutral one, one might expect that this new operator which governs the evolution of "risk-sensitive probability" to be a generalization of \mathcal{A}^T . Setting $N_Y \cdot \bar{Q}(Y_t)$ equal to the identity matrix I corresponds to the case when no observation is available. It can be shown that the information state is independent of scaling; i.e., if σ is

an information state so is $\alpha\sigma$ for every $\alpha > 0$ and replacing it one with the other does not change the resulting estimate of the state. Associate with each $i \in X$, a unit vector in R_X^N whose i^{th} component is 1. Denote the ‘‘risk-sensitive probability’’ U_t as the normalized information state(6) when no observation is available. We have the following theorem:

Theorem 2 *Let $N_Y \cdot \bar{Q}(Y_t) = I$, then the estimator which minimizes (5) is given by*

$$\hat{X}_t = \operatorname{argmax}_{i \in S_X} \langle U_t, e_i \rangle,$$

where U_t evolves according to

$$U_{t+1} = \mathcal{A}^T H \{ \operatorname{diag} (\exp(\gamma \langle e_{\operatorname{argmax}_i U_t^i}, e_j \rangle)) \cdot U_t \} := F^\gamma(U_t), \quad (10)$$

and $H(X) = \frac{X}{\sum_i X_i}$ and $U_0 = p_0$.

Proof The scale independence property allows the introduction of $H(X)$, and the result follows from straightforward computation afforded by the assumption $N_Y \cdot \bar{Q}(Y_t) = I$.

4 SIMULATION RESULTS

The operator F^γ can be viewed as a non-linear generalization of the linear operator \mathcal{A}^T . It is apparent that this operator plays the same role in the context of risk-sensitive estimation as the operator \mathcal{A}^T does in the risk-neutral case. Thus, one might expect that the risk-sensitive properties of the exponential criterion be reflected in the action of F^γ .

First, observe that both operators are simplex preserving and $F^\gamma \rightarrow \mathcal{A}^T$ as $\gamma \rightarrow 0$. It is well known that under primitivity of the matrix \mathcal{A} , the dynamical system defined by

$$p_{n+1} = \mathcal{A}^T p_n, \quad (11)$$

for every choice of the initial probability distribution p_0 , converges to p^* which satisfies $\mathcal{A}^T p^* = p^*$ (Shriyayev 1984).

Compare (10) and (11); how does the ‘‘product count’’ of errors manifest itself in the behavior of the RMAP estimator of an HMM? Let us begin with a basic example.

Example 1: Let $N_Y = N_X = 2$, $q_{11} = q_{22} = q$. (See Figure 1 and 2.) Also, assume that the chain has an initial distribution $p = (p_1, p_2)$ with $p_1 + p_2 = 1$. q controls the amount of the available information, in particular $q = 1/2$ implies that no information is available, (the case of pure ‘‘prediction’’) and $q=1$ corresponds to the case of ‘‘perfect observation’’.

Let $q = 1/2$, $a_{12} = a_{21} = 0$ and allow γ to vary: $1.02 \leq e^\gamma \leq 1.22$. Figure 2 shows the behavior of the estimator. We have stacked up the sample paths starting at e_1 for a




Figure 1: Example 1




Figure 2: $a_{12} = a_{21} = 0$, $q = 1/2$

range of values of e^γ ; on each path the lower value indicates e_1 and the higher value corresponds to e_2 .

When $p_1 = p_2$, an oscillatory behavior appears only after one step. As we increase p_1 , the onset of this behavior is delayed, but it is inevitable, unless $p_2 = 0$.

Now let $a_{12} = 0.1$ and $a_{21} = 0.15$ with everything else kept the same (Figure 3). The oscillatory behavior is delayed and relatively suppressed. This appears counter intuitive: the initial setting ($a_{12} = a_{21} = 0$) does not allow any transition between the states, but the RMAP estimator is oscillatory. The second set of parameters ($a_{12} = 0.1$ and $a_{21} = 0.15$) allows transitions between the states, yet the estimator’s behavior is less oscillatory and closer to that of MAP.

Now let $p_1=0.6$, $a_{12} = a_{21} = 0$, but $q = 0.6$. Figure 4 shows a series of sample paths with the state initialized at e_1 and $1.02 \leq e^\gamma \leq 1.22$ as before.

Once again, the oscillations are suppressed. It is evident that the transition probabilities representing the underlying dynamics, the risk-sensitivity parameter and the availability of information are coupled in the behavior of the sample




Figure 3: $a_{21} = 0.15, a_{12} = 0.1, q = 1/2$




Figure 5: RS-Probability




Figure 4: $a_{12} = a_{21} = 0, q = 0.6$

paths. The dependence of the period on transition probabilities is shown next. Let

$$\mathcal{A} = \begin{bmatrix} 0.9 - \epsilon & 0.1 & \epsilon \\ 0.4 & 0.6 & 0.0 \\ 0.0 & \epsilon & 1.0 - \epsilon \end{bmatrix} \quad (12)$$

and $e^{\mathcal{A}} = 101$. The periodic orbit appear to be independent of the initial conditions but the period can depend strongly on ϵ and as $\epsilon \rightarrow 0$ the period appears to be unbounded.

Figure 5 shows the values of the first component of the RS-probability vs. time for $\epsilon = 0.001$. (There are 2000 data points and hence some apparent overlaps) Our simulations show that as $\epsilon \rightarrow \infty$, the period increases rapidly. One possible explanation is that ϵ controls the mixing properties of (12); the matrix \mathcal{A} is primitive only for strictly positive values of ϵ and as ϵ approaches zero, (12) "approaches"




Figure 6: The Multi-Scale Representation

a non-mixing dynamical system and hence its stationary behavior becomes less "stable".

4.1 MULTI-SCALE REPRESENTATION AND RISK-SENSITIVITY

By a *Multi-scale representation* of a Markov chain, we mean a partition of the states of that chain. Each element of the partition will be called a "scale" or a "cluster".

Example 2: Consider the Markov chain given by Figure 6. Consider the observation matrices C_u and C_l :

$$C_l \times 10^{-3} =$$

$$\begin{bmatrix} 991 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 991 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 496 & 496 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 496 & 496 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 991 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 991 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 991 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 991 & 1 & 1 \end{bmatrix}$$

$$C_u \times 10^{-3} =$$

$$\begin{bmatrix} 991 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 991 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 701 & 291 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 291 & 701 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 991 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 496 & 496 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 496 & 496 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 496 & 496 & 1 & 1 & 1 \end{bmatrix}$$

We say that a “cluster” of states is resolved if it is possible from the observations to determine whether or not a sample path has assumed any of the states within that cluster. Examining C_u and C_l shows that both provide nearly perfect observation at the scales shown in the Figure 6 (M1 to M5). Within the clusters, C_u provides partial observation for the components of M1, e_3 and e_4 while keeping M5 “unresolved” so that it remains impossible to distinguish between the states e_6 , e_7 and e_8 . C_l leaves M1 unresolved while nearly resolving M5 into its components. Which one of these matrices should we choose to “better” observe our Markov chain and is this choice invariant under risk-sensitivity?

C_l provides nearly perfect observation for the less likely path while C_u provides only partial observation but for the more likely path. Our simulations presented in Graph 7 (D is the cost under C_u subtracted from the cost under C_l via Monte Carlo simulation for the cost function averaged over 10,000 sample paths) show that the choice depends on the amount of risk considered and is *not* invariant under risk-sensitivity. (On the x-axis, the parameter u is $e^\gamma - 1$.) Let us look at the example and the simulations closely to understand why. For MAP and RMAP with small values of γ , C_u , as the graph shows, is a better option but as γ increases choosing C_l quickly becomes the better option. The switch happens at a value of u^* between $u = 0.14$ and $u = 0.15$ (See Graph 7); changing the transition probabilities among the states of M1 and among those of M2 may change the




Figure 7: The Cost Difference

value of u at which we switch from C_u to C_l . For example, when we set $a_{34} = a_{43} = 0.1$ and $a_{33} = a_{44} = .9$, the value of u^* is less than half the previous value.

Do we want to see the “details” of M1 or M5? A more “conservative” estimator prefers to resolve M5 because M5 contains more branches (e_6 to e_8). The estimator can choose one branch at a time and will return to it only after it has visited the other branches. If we provide no information and leave M5 completely unresolved then, as γ increases, the estimator is forced to visit each state of M5 successively and thus each branch of M5 is ignored for two successive steps. This turns out to be costly as the value of γ is increased. If we increase the number of the branches in M5, the value of u^* at which the switch happens quickly decreases.

5 CONCLUSIONS

In this paper, we introduced and studied risk-sensitive estimators as a generalization of Maximum A Posteriori Probability (MAP) estimator for Hidden Markov Models. We introduced the notion of risk-sensitive probability. We considered the relationship between risk-sensitivity, transition probabilities and information. We showed via simulation that the asymptotic behavior of risk-sensitive probability could be periodic. We showed that the value and interpretation of information depends on the degree to which higher order moments are taken into account. These simulations were carried out for ergodic Markov chains. The non-ergodic case is the subject of further research.

6 ACKNOWLEDGMENTS

Research supported in part by the Air Force Office of Scientific Research under grants F496200110161 and FA95500410210 and by the National Science Foundation under grant DMI0323220.

REFERENCES

- Baras J. S. and M. R. James. 1994. Robust and Risk-sensitive Output Feedback Control for Finite State Machines and Hidden Markov Models. Technical Report No. 1994-63, Institute for Systems Research, University of Maryland at College Park, Technical Research Reports.
- Boel, R. K., M. R. James, and I. R. Petersen. 1997. Robustness and risk-sensitive filtering. In *Proceedings of the 36th IEEE Conference on Decision and Control*, Cat. No. 97CH36124.
- Dey, S. and J. Moore. 1995. Risk sensitive filtering and smoothing for hidden Markov models. *Systems and Control Letters*, 25 (1): 361-366.
- James, M. R. J. S. Baras, and R. J. Elliot. 1994. Risk sensitive control and dynamic games for partially observed discrete-time nonlinear systems. *IEEE Transactions on Automatic Control*, 39 (4): 780-792.
- Poor H. V. 1994. *An Introduction to Signal Detection and Estimation*. Springer Verlag.
- Ramezani, V. R. and Steven I. Marcus. 2002. Estimation of hidden Markov models: Risk-sensitive filter banks and qualitative analysis of their sample paths. *IEEE Transactions on Automatic Control*, 47 (12): 1999-2009.
- Ramezani, V. R. PhD dissertation (2001). Dept. of Electrical and Computer Engineering, University of Maryland, College Park. Institute for Systems Research Technical Reports PhD 2001-7.
- Shiryayev, A. N. 1984. *Probability*. New York: Springer-Verlag.
- Speyer, J. L., C. Fan and R. N. Banava. 1992. Optimal stochastic estimation with exponential cost criteria. *Proceedings of 31st Conference on Decision and Control*: 2293-2298.
- Whittle P. 1990. *Risk-Sensitive Optimal Control*. New York: John Wiley & Sons.
- MICHAEL C. FU** is a Professor in the Robert H. Smith School of Business, with a joint appointment in the Institute for Systems Research and an affiliate appointment in the Department of Electrical and Computer Engineering, all at the University of Maryland. He received degrees in mathematics and EE/CS from MIT, and an M.S. and Ph.D. in applied mathematics from Harvard University. His research interests include simulation and applied probability modelling, particularly with applications towards manufacturing systems, inventory control, and financial engineering. He currently serves as Simulation Area Editor of *Operations Research* and as an Associate Editor for *Management Science*. His book (co-authored with J.Q. Hu), *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*, received the INFORMS College on Simulation Outstanding Publication Award in 1998. His e-mail address is <mfu@rsmith.umd.edu>.

STEVEN I. MARCUS (Fellow, IEEE) received the B.A. degree in electrical engineering and mathematics from Rice University in 1971 and the S.M. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology in 1972 and 1975, respectively. From 1975 to 1991, he was with the Department of Electrical and Computer Engineering at the University of Texas at Austin, where he was the L.B. (Preach) Meaders Professor in Engineering. He was Associate Chairman of the Department during the period 1984-89. In 1991, he joined the University of Maryland, College Park, where he was Director of the Institute for Systems Research until 1996. He is currently a Professor and Chair of the Electrical and Computer Engineering Department at the University of Maryland. Currently, his research is focused on stochastic control and estimation, with applications in semiconductor manufacturing, telecommunication networks, and preventive maintenance. He is Editor-in-Chief of the SIAM Journal on Control and Optimization. His email address is <marcus@umd.edu>.

AUTHOR BIOGRAPHIES

VAHID RAMEZANI received the B.S., M.S. and PhD degrees in electrical engineering from the University of Maryland, College Park, MD. During 1995-1997, he was a design engineer for LCC. Inc. primarily concerned with the design, implementation, testing, and optimization of cellular telephony systems. His research interests include optical communication, risk-sensitive estimation and control, wireless communications and intrusion detection. He is currently with the Institute for Systems Research, University of Maryland, College Park. His email address is <rvahid@isr.umd.edu> and his web page is <www.isr.umd.edu/~rvahid>.