Stochastic Filtering (SS2016) Exercise Sheet 3

Lecture and Exercises: JProf. Dr. Philipp Harms
Due date: May 11, 2016

General setup

(X,Y) is a HMM with state kernel $P(x,dx')$ and observation kernel $K(x,dy)$. The observation kernel is non-degenerate, i.e., $K(x,dy) = \lambda(x,y)\phi(dy)$ for a positive function λ and a probability measure ϕ. We write $\pi_{k|n}$ for the regular conditional distribution of X_k given $Y_{0:n}$, and we set $\pi_k = \pi_{k|k}$. The unnormalized filters are denoted by $\sigma_{k|n}$ and σ_k.

3.1. Filtering recursion

The filtering recursion expresses π_{k+1} in terms of π_k.

a) A direct approach to deriving the filtering recursion is to apply Bayes' formula to an increment of the process (X,Y) as follows. Assume for fixed $k \in \mathbb{N}$ that $(X_k,Y_k) \sim \pi_k(dx_k)K(x_k,dy_k)$ for some given probability measure π_k. Calculate the regular conditional distribution $\pi_{k+1} = P_{X_{k+1}|Y_{k+1}}$ under the HMM.

b) Contrast this with the derivation of the filtering recursion shown in the lecture, where one applies Bayes' formula to the joint distribution of $(X_{0:k},Y_{0:k})$ and compares the results for k and $k+1$.

3.2. Prediction recursion

The prediction recursion expresses $\pi_{k+1|n}$ in terms of $\pi_{k|n}$ for $k \geq n$.

a) A direct approach to deriving the prediction recursion is to apply Bayes' formula to an increment of the process (X,Y) as follows. Assume for fixed $k \in \mathbb{N}$ that $(X_k, Y_k) \sim \pi_k(dx_k)K(x_k, dy_k)$ for some given probability measure π_k. Calculate the distribution of X_{k+1} conditional on (X_k, Y_k) under the HMM.

b) Show that the prediction recursion coincides with the filtering recursion if the observation kernel $K(x, dy)$ does not depend on the state x.

3.3. Predictor and corrector step

Each step $\pi_k \rightsquigarrow \pi_{k+1}$ of the filtering recursion can be split in a predictor step $\pi_k \rightsquigarrow \pi_{k+1|k}$ followed by a corrector step $\pi_{k+1|k} \rightsquigarrow \pi_{k+1}$.

a) Provide explicit formulas for the predictor and corrector steps.

b) Which steps depend on the state kernel and which ones on the observation kernel?

3.4. Kalman Filter

Let us consider linear Gaussian HMMs: the state and observation processes are assumed to evolve according to

$$X_{k+1} = AX_k + R\xi_k, \qquad Y_k = BX_k + S\eta_k, \qquad k \geq 0$$

with initial state $X_0 \sim \mathcal{N}(m_0, \Sigma_0)$ and parameters $A, R \in \mathbb{R}^{p\times p}$, $B \in \mathbb{R}^{q\times p}$, and $S \in \mathbb{R}^{q\times q}$ with S invertible. The variables $\xi_i \sim \mathcal{N}(0, \mathbb{1}_{p\times p})$, $\eta_j \sim \mathcal{N}(0, \mathbb{1}_{q\times q})$, and X_0 are independent. The filtering distributions of this HMM are Gaussian (see Exercise 1.4.b)), and we write $\pi_k \sim \mathcal{N}(\hat{X}_k, \hat{\Sigma}_k)$ and $\pi_{k+1|k} \sim \mathcal{N}(\hat{X}_{k+1|k}, \hat{\Sigma}_{k+1|k})$.

a) Predictor step: assume that $\pi_k \sim \mathcal{N}(\hat{X}_k, \hat{\Sigma}_k)$ and show that $\pi_{k+1|k} \sim \mathcal{N}(\hat{X}_{k+1|k}, \hat{\Sigma}_{k+1|k})$, where

$$\hat{X}_{k+1|k} = A\hat{X}_k, \qquad \hat{\Sigma}_{k+1|k} = A\hat{\Sigma}_k A^\top + RR^\top.$$

b) Corrector step: assume that $\pi_{k+1|k} \sim \mathcal{N}(\hat{X}_{k+1|k}, \hat{\Sigma}_{k+1|k})$ and show that $\pi_{k+1} \sim \mathcal{N}(\hat{X}_{k+1}, \hat{\Sigma}_{k+1})$, where

$$\hat{X}_{k+1} = \hat{X}_{k+1|k} - \hat{\Sigma}_{k+1|k} B^\top (SS^\top + B\hat{\Sigma}_{k+1|k} B^\top)^{-1} (B\hat{X}_{k+1|k} - Y_{k+1}),$$
$$\hat{\Sigma}_{k+1} = \hat{\Sigma}_{k+1|k} - \hat{\Sigma}_{k+1|k} B^\top (SS^\top + B\hat{\Sigma}_{k+1|k} B^\top)^{-1} B\hat{\Sigma}_{k+1|k}.$$

Hint. You may look this up in the reference of your choice or [1].

3.5. Normalized versus unnormalized recursions

a) Come up with an example of a HMM where the unnormalized filter σ_k quickly exceeds machiner precision, whereas the normalized filter π_k does not.

b) Explain in which sense the normalized filtering recursion is non-linear and the unnormalized one linear.

References

[1] Olivier Cappé, Eric Moulines, and Tobias Ryden. *Inference in Hidden Markov Models*. Springer Series in Statistics. Springer Verlag, New York, 2005.