Technical report from Automatic Control at Linköpings universitet

ML Estimation of Process Noise Variance in Dynamic Systems

Patrik Axelsson, Umut Orguner, Fredrik Gustafsson, Mikael Norrlöf

Division of Automatic Control

E-mail: axelsson@isy.liu.se, umut@isy.liu.se, fredrik@isy.liu.se, mino@isy.liu.se

8th October 2010

Report no.: LiTH-ISY-R-2969
Submitted to the 18th IFAC World Congress 2011

Address:
Department of Electrical Engineering
Linköpings universitet
SE-581 83 Linköping, Sweden

WWW: http://www.control.isy.liu.se

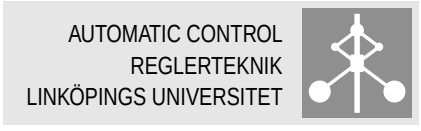

Abstract

The performance of a non-linear filter hinges in the end on the accuracy of the assumed non-linear model of the process. In particular, the process noise covariance Q is hard to get by physical modeling and dedicated system identification experiments. We propose a variant of the expectation maximization (EM) algorithm which iteratively estimates the unobserved state sequence and Q based on the observations of the process. The extended Kalman smoother (EKS) is the instrument to find the unobserved state sequence. Our contribution fills a gap in literature, where previously only the linear Kalman smoother and particle smoother have been applied. The algorithm will be important for future industrial robots with more flexible structures, where the particle smoother cannot be applied due to the high state dimension. The proposed method is compared to two alternative methods on a simulated robot.

Keywords: Robotic manipulators, Extended Kalman filters, Smoothing filters, Identification, Maximum likelihood, Covariance matrices

ML Estimation of Process Noise Variance in Dynamic Systems

Patrik Axelsson, Umut Orguner, Fredrik Gustafsson, Mikael Norrlöf

Department of Electrical Engineering, Linköping University, SE-581 83 Linköping, Sweden (e-mail: {axelsson, umut, fredrik, mino}@isy.liu.se).

Abstract: The performance of a non-linear filter hinges in the end on the accuracy of the assumed non-linear model of the process. In particular, the process noise covariance Q is hard to get by physical modeling and dedicated system identification experiments. We propose a variant of the expectation maximization (EM) algorithm which iteratively estimates the unobserved state sequence and Q based on the observations of the process. The extended Kalman smoother (EKS) is the instrument to find the unobserved state sequence. Our contribution fills a gap in literature, where previously only the linear Kalman smoother and particle smoother have been applied. The algorithm will be important for future industrial robots with more flexible structures, where the particle smoother cannot be applied due to the high state dimension. The proposed method is compared to two alternative methods on a simulated robot.

Keywords: Robotic manipulators, Extended Kalman filters, Smoothing filters, Identification, Maximum likelihood, Covariance matrices

1. INTRODUCTION

Joint parameter identification and state estimation in state space model is an important branch of system identification, Ljung [1999]. During the last decade, subspace based approaches for estimating fully parameterized linear state space models (so called black box models) have been well explored, Ljung [1999]. At the same time, the theory of grey-box identification of uncertain parameters in physical models has been developed, Bohlin [2006]. The model is here a non-linear state space model without process noise. The basic idea is that the system can be simulated for each value of the parameter vector, and the simulated output can be compared to the observed measurements, where for instance the maximum likelihood estimate (MLE) is computed. The situation with process noise is considerably harder, since the simulated has to be replaced with a smoothing filter for the state sequence. A further problem is that the likelihood function becomes rather complicated. The EM algorithm in Dempster et al. [1977] provides a method to compute the MLE by separating the smoothing and parameter estimation problems. It has been explored for linear Gaussian models, where the system matrices (A, C, Q, R) are estimated using the Kalman smoother as the state estimator, Cappé et al. [2005]. For non-linear models, there is on-going research, Schön et al. [2010], on using the particle smoother to estimate the parameters in a non-linear dynamic model. However, the particle smoother is not applicable for models with high state dimension.

Here we propose to use a linearised model for state estimation, leading to an extended Kalman smoother (EKS). The EM algorithm will thus be approximate in the same way as the EKS. We focus on the process noise covariance matrix, which is the hardest one to assess in the modeling phase. Our application in mind is industrial robots, where inertia, flexibilities and friction parameters in each joint are all rather straightforwardly identified by dedicated experiments, see Wernholt [2007] and Carvalho Bittencourt et al. [2010]. The sensor noise covariance is also quite easy to get. Process noise, on the other hand, models quite complex phenomena as well as model uncertainties.

The motivation to do this for industrial robots is the development of new robots with increased elasticity and larger individual variations, such as variation of gearbox stiffness or in the parameters describing the mechanical arm. To maintain or improve the robot performance, the motion control must be improved for this new generation of robots. For robots with traditional measurement systems, where only the motor angular position is measured, this can be obtained by improving the model-based control as described in Björkman et al. [2008]. Another option is to use inertial sensors to improve the estimation of the robot tool position and velocity, which requires good knowledge about the noise. The estimated state trajectories can be used for on-line feedback control as a mean of increasing both the robust and the nominal performance of the robot. Another possible use of tool estimation is iterative learning control, (ILC) Wallén et al. [2009].

Section 2 gives a short introduction to the problem and the EM algorithm. The calculations for the EM algorithm are then given in Section 3. Two alternative methods are presented in Section 4. Section 5 describes a robot model which are used in Section 6 to compare the three methods. Finally, Section 7 concludes the paper.

2. PROBLEM FORMULATION

This paper is focused on a model structure given by

$$x_{k+1} = F_1(x_k, u_k) + F_2(x_k)v_k \tag{1a}$$
$$y_k = h(x_k, u_k) + e_k \tag{1b}$$

where $x_k \in \mathbb{R}^n$, $y_k \in \mathbb{R}^m$, $v_k \sim \mathcal{N}(0, Q)$ and $e_k \sim \mathcal{N}(0, R)$. All model parameters are assumed to be known except for $Q \in S_+^p$ and $R \in S_{++}^m$ [1]. Assume also that $F_2(x_k)$ has the following structure

$$F_2(x_k) = \begin{pmatrix} 0 \\ \tilde{F}_2(x_k) \end{pmatrix}. \tag{2}$$

This type of model structure is common for mechanical systems derived by Newton's law or Lagrange's equation.

One approach to find the covariance matrices R and Q is to use the well known Maximum Likelihood (ML) method. The idea with the ML method is to find the unknown parameters θ such that the measurements $y_{1:N} = \{y_1 \ \dots \ y_N\}$ become as likely as possible. In other words,

$$\hat{\theta}^{ML} = \arg\max_{\theta \in \Theta} p_\theta(y_{1:N}), \tag{3}$$

where $p_\theta(y_{1:N})$ is the probability density function (pdf) of the observations parametrised with the parameter θ. It is common to take the logarithm of the pdf,

$$L_\theta(y_{1:N}) = \log p_\theta(y_{1:N}), \tag{4}$$

and find the parameter θ that maximises (4), i.e.,

$$\hat{\theta}^{ML} = \arg\max_{\theta \in \Theta} L_\theta(y_{1:N}). \tag{5}$$

These two problems are equivalent since the logarithm is a monotonic function. The ML problem can be solved using a standard optimisation method, see e.g. Boyd and Vandenberghe [2009]. The solution can however be hard to find which has lead to the development of the Expectation Maximisation (EM) algorithm.

The EM algorithm was originally invented by Dempster et al. [1977]. Theory and examples for the EM algorithm can be found in McLachlan and Krishnan [2008], see also Gibson and Ninness [2005] for robust estimation of LTI state space models. In Schön et al. [2010], a solution of the more complicated problem to estimate non-linear state space models is given, using a particle smoother. As mentioned before, the particle smoother cannot be applied if the state dimension is too high.

2.1 The Expectation Maximisation Algorithm

The principal idea with the EM algorithm is to introduce variables $x_{1:N} = \{x_1 \ \dots \ x_N\}$ which are not observed directly. The variables $x_{1:N}$ can instead be observed indirectly from $y_{1:N}$. Take now the joint log-likelihood function

$$L_\theta(y_{1:N}, x_{1:N}) = \log p_\theta(y_{1:N}, x_{1:N}) \tag{6}$$

of the observed variables $y_{1:N}$ and the variables $x_{1:N}$. Equation (6) cannot be used directly since $x_{1:N}$ are unknown. Instead, calculate

$$\Gamma(\theta; \theta_l) = E_{\theta_l}\left[\log p_\theta(y_{1:N}, x_{1:N}) | y_{1:N}\right], \tag{7}$$

where $E_{\theta_l}[\cdot|\cdot]$ is the conditional mean with respect to a pdf defined by the parameter θ_l. It can now be shown, see e.g. Schön et al. [2010], that any θ, such that

$$\Gamma(\theta; \theta_l) > \Gamma(\theta_l; \theta_l), \tag{8}$$

implies that

$$L_\theta(y_{1:N}) > L_{\theta_l}(y_{1:N}). \tag{9}$$

Hence, the EM algorithm provides a sequence θ_l, $l = 1, 2, \dots$, which approximates $\hat{\theta}^{ML}$ better and better for every iteration. The EM algorithm can now be summarised in the following steps.

Algorithm 1. The EM Algorithm

(1) Select an initial value θ_0 and set $l = 0$.
(2) Expectation Step (E-step): Calculate
$$\Gamma(\theta; \theta_l) = E_{\theta_l}\left[\log p_\theta(y_{1:N}, x_{1:N}) | y_{1:N}\right].$$
(3) Maximisation Step (M-step): Compute
$$\theta_{l+1} = \arg\max_{\theta \in \Theta} \Gamma(\theta; \theta_l).$$
(4) If converged, stop. If not, set $l = l + 1$ and go to step 2.

A possible stop criterion for the EM algorithm could be when the change in θ, between two iterations, is small enough.

3. EM FOR COVARIANCE ESTIMATION

This section describes how the covariance matrices for the process and measurement noise in (1) can be estimated using the EM algorithm. It is not possible to simultaneously estimate both the covariance matrix Q for the process noise and the covariance matrix R for the measurement noise, since it is the scaling between them that affects the estimate when an extended Kalman filter (EKF), Anderson and Moore [1979], is used. Therefore, estimate first R with dedicated experiments and then Q with the EM algorithm. The matrix R can be estimated according to the following steps.

(1) Perform simple experiments with a known trajectory.
(2) Calculate $e_k = y_k - \bar{y}_k$, where $\bar{y}_k$ is the true trajectory.
(3) Finally,
$$R = \frac{1}{N}\sum_{k=1}^{N} e_k^T e_k. \tag{10}$$

The matrix R can now be used in the EM algorithm to estimate Q.

Equation (1) can also be expressed in the more general conditional densities as

$$x_{k+1} \sim p(x_{k+1}|x_k) = \mathcal{N}\left(x_{k+1}; F_{1,k}, F_{2,k}QF_{2,k}^T\right) \tag{11a}$$
$$y_k \sim p(y_k|x_k) = \mathcal{N}\left(y_k; h_k, R\right) \tag{11b}$$

where $\mathcal{N}(\cdot)$ is the multivariate normal distribution function. The multivariate normal distribution for the n-dimensional variable μ with mean $\bar{\mu}$ and covariance Σ is defined according to

$$\mathcal{N}(\mu; \bar{\mu}, \Sigma) \triangleq \frac{1}{(2\pi)^{n/2}|\Sigma|^{1/2}} e^{-\frac{1}{2}(\mu-\bar{\mu})^T\Sigma^{-1}(\mu-\bar{\mu})}. \tag{12}$$

The short notation

$$F_{1,k} = F_1(x_k, u_k), \quad F_{2,k} = F_2(x_k), \quad h_k = h(x_k, u_k)$$

is sometimes used for simplicity in the sequel.

Proceed now with the derivation of the expectation and maximisation steps in Algorithm 1 where $\theta = Q$ is the sought parameter.

[1] S_{++}^p $\left(S_+^p\right)$ is the set of all symmetric positive definite (semidefinite) $p \times p$ matrices

3.1 Expectation step

The joint likelihood can easily be written as

$$p_Q(y_{1:N}, x_{1:N}) = p_Q(x_1, y_1) \prod_{i=2}^{N} p_Q(y_i|x_i) p_Q(x_i|x_{i-1}), \tag{13}$$

where

$$\begin{aligned} p(x_k, y_k|x_{1:k-1}, y_{1,k-1}) &= p(x_k, y_k|x_{k-1}) \\ &= p(y_k|x_k) p(x_k|x_{k-1}) \end{aligned} \tag{14}$$

has been used repeatedly. Taking the logarithm of (13) and using (11) together with (12) give

$$\begin{aligned} L_Q(y_{1:N}, x_{1:N}) &= \log p_Q(y_{1:N}, x_{1:N}) = \tilde{L} \\ &+ \frac{1}{2} \sum_{i=2}^{N} \log \left(\prod_{\lambda_j \neq 0} \lambda_j \left(\left(F_{2,i-1} Q F_{2,i-1}^T \right)^\dagger \right) \right) \\ &- \frac{1}{2} \sum_{i=2}^{N} \tilde{F}_{1,i}^T \left(F_{2,i-1} Q F_{2,i-1}^T \right)^\dagger \tilde{F}_{1,i}, \end{aligned} \tag{15}$$

where $\tilde{L}$ is an expression independent of Q, $\dagger$ is the Moore-Penrose inverse, Mitra and Rao [1971],

$$\tilde{F}_{1,i} = x_i - F_{1,i-1} \tag{16}$$

and $\prod_{\lambda_j \neq 0} \lambda_j(\cdot)$ means to take the product of all nonzero eigenvalues. The structure of F_2 in (2) gives

$$\left(F_{2,i-1} Q F_{2,i-1}^T \right)^\dagger = \begin{pmatrix} 0 & 0 \\ 0 & \tilde{F}_{2,i-1} Q \tilde{F}_{2,i-1}^T \end{pmatrix}^\dagger \tag{17}$$

which leads to

$$\prod_{\lambda_j \neq 0} \lambda_j \left(\left(F_{2,i-1} Q F_{2,i-1}^T \right)^\dagger \right) = \left| \left(\tilde{F}_{2,i-1} Q \tilde{F}_{2,i-1}^T \right)^\dagger \right|. \tag{18}$$

The joint log-likelihood function (15) can now be written as

$$\begin{aligned} L_Q(y_{1:N}, x_{1:N}) &= \tilde{L} + \frac{1}{2} \sum_{i=2}^{N} \log \left(\left| \left(\tilde{F}_{2,i-1} Q \tilde{F}_{2,i-1}^T \right)^\dagger \right| \right) \\ &- \frac{1}{2} \sum_{i=2}^{N} \tilde{F}_{1,i}^T \left(F_{2,i-1} Q F_{2,i-1}^T \right)^\dagger \tilde{F}_{1,i}. \end{aligned} \tag{19}$$

Next step is to calculate the expectation of $L_Q(y_{1:N}, x_{1:N})$ to obtain $\Gamma(Q; Q_l)$.

$$\begin{aligned} \Gamma(Q; Q_l) &= E_{Q_l} \left[L_Q(y_{1:N}, x_{1:N}) | y_{1:N} \right] = \bar{L} \\ &+ \frac{1}{2} \sum_{i=2}^{N} E_{Q_l} \left[\log \left(\left| \left(\tilde{F}_{2,i-1} Q \tilde{F}_{2,i-1}^T \right)^\dagger \right| \right) \middle| y_{1:N} \right] \\ &- \frac{1}{2} \mathrm{Tr} \sum_{i=2}^{N} E_{Q_l} \left[\left(F_{2,i-1} Q F_{2,i-1}^T \right)^\dagger \tilde{F}_{1,i} \tilde{F}_{1,i}^T \middle| y_{1:N} \right] \end{aligned} \tag{20}$$

where $\bar{L}$ is independent of Q. The trace operator comes from the fact that

$$\begin{aligned} &\tilde{F}_{1,i}^T \left(F_{2,i-1} Q F_{2,i-1}^T \right)^\dagger \tilde{F}_{1,i} \\ &= \mathrm{Tr} \left(\tilde{F}_{1,i}^T \left(F_{2,i-1} Q F_{2,i-1}^T \right)^\dagger \tilde{F}_{1,i} \right). \end{aligned} \tag{21}$$

Start with the calculations of the first expectation in (20).

$$\begin{aligned} &E_{Q_l} \left[\log \left(\left| \left(\tilde{F}_{2,i-1} Q \tilde{F}_{2,i-1}^T \right)^\dagger \right| \right) \middle| y_{1:N} \right] \\ &= \int \log \left(\left| \left(\tilde{F}_2(x_{i-1}) Q \tilde{F}_2^T(x_{i-1}) \right)^\dagger \right| \right) \\ &\quad \times p_{Q_l}(x_{i-1}|y_{1:N}) dx_{i-1}. \end{aligned} \tag{22}$$

The integral cannot be solved analytically. Instead, an approximation has to be made. The smoothed density of x_{i-1} has a high peak around the smoothed estimate if the sampling frequency and the SNR are high. Here, we use the extended Kalman smoother (EKS), see Yu et al. [2004]. We can therefore approximate x_{i-1} with the smoothed states $\hat{x}_{i-1|N}^s$, in other words,

$$\begin{aligned} &E_{Q_l} \left[\log \left(\left| \left(\tilde{F}_{2,i-1} Q \tilde{F}_{2,i-1}^T \right)^\dagger \right| \right) \middle| y_{1:N} \right] \\ &\approx \log \left(\left| \left(\tilde{F}_2(\hat{x}_{i-1|N}^s) Q \tilde{F}_2^T(\hat{x}_{i-1|N}^s) \right)^\dagger \right| \right). \end{aligned} \tag{23}$$

The second expectation in (20) can be written as

$$\begin{aligned} &E_{Q_l} \left[\left(F_{2,i-1} Q F_{2,i-1}^T \right)^\dagger \tilde{F}_{1,i} \tilde{F}_{1,i}^T \middle| y_{1:N} \right] \\ &= \int \left(F_2(x_{i-1}) Q F_2^T(x_{i-1}) \right)^\dagger \tilde{F}_{1,i} \tilde{F}_{1,i}^T \\ &\quad \times p_{Q_l}(x_i, x_{i-1}|y_{1:N}) dx_i dx_{i-1}. \end{aligned} \tag{24}$$

Now use the smoothed density again and let

$$\left(F_2(x_{i-1}) Q F_2^T(x_{i-1}) \right)^\dagger \approx \left(F_2(\hat{x}_{i-1|N}^s) Q F_2^T(\hat{x}_{i-1|N}^s) \right)^\dagger. \tag{25}$$

We have now

$$\begin{aligned} &E_{Q_l} \left[\left(F_{2,i-1} Q F_{2,i-1}^T \right)^\dagger \tilde{F}_{1,i} \tilde{F}_{1,i}^T \middle| y_{1:N} \right] \\ &\approx \left(F_2(\hat{x}_{i-1|N}^s) Q F_2^T(\hat{x}_{i-1|N}^s) \right)^\dagger \\ &\quad \times \int \tilde{F}_{1,i} \tilde{F}_{1,i}^T p(x_i, x_{i-1}|y_{1:N}) dx_i dx_{i-1}, \end{aligned} \tag{26}$$

where $p_{Q_l}(x_i, x_{i-1}|y_{1:N})$ can be seen as the smoothed density of the augmented state vector $\xi_i = \begin{pmatrix} x_{i-1}^T & x_i^T \end{pmatrix}^T$, i.e.,

$$p_{Q_l}(x_i, x_{i-1}|y_{1:N}) = p_{Q_l}(\xi_i|y_{1:N}) = \mathcal{N}\left(\xi_i; \hat{\xi}_{i|N}^s, P_{i|N}^{\xi,s} \right). \tag{27}$$

The first and second order moments of the smoothed ξ_i can be expressed as

$$\hat{\xi}_{i|N}^s = \begin{pmatrix} \hat{x}_{i-1|N}^s \\ \hat{x}_{i|N}^s \end{pmatrix}, \tag{28a}$$

$$P_{i|N}^{\xi,s} = \begin{pmatrix} P_{i-1|N}^s & P_{i-1,i|N}^s \\ \left(P_{i-1,i|N}^s \right)^T & P_{i|N}^s \end{pmatrix}, \tag{28b}$$

where $\hat{x}_{i-1|N}^s$, $\hat{x}_{i|N}^s$, $P_{i-1|N}^s$ and $P_{i|N}^s$ are the first and second order moments of the smoothed $\hat{x}_{i-1}$ and $\hat{x}_i$ respectively. These are obtained if the augmented model

$$\xi_{k+1} = \begin{pmatrix} x_k \\ F_1(x_k, u_k) \end{pmatrix} \tag{29}$$

is used in the EKS.

The integral in (26) cannot be solved analytically. Instead, a first order Taylor expansion of $F_1(x_{i-1})$ around $\hat{x}_{i-1|N}^s$ is used. The expression $\tilde{F}_{1,i} \tilde{F}_{1,i}^T$ in (26) can now be written as

$$
\begin{aligned}
\tilde{F}_{1,i}\tilde{F}_{1,i}^T &= (x_i - F_1(x_{i-1}))(x_i - F_1(x_{i-1}))^T \\
&= \left(x_i - F_1(\hat{x}^s_{i-1|N}) - J_1(x_{i-1} - \hat{x}^s_{i-1|N})\right) \\
&\quad \times \left(x_i - F_1(\hat{x}^s_{i-1|N}) - J_1(x_{i-1} - \hat{x}^s_{i-1|N})\right)^T \\
&= (-J_1 \ I)\left(\xi_i - \hat{\xi}^s_{i|N}\right)\left(\xi_i - \hat{\xi}^s_{i|N}\right)^T (-J_1 \ I)^T \\
&\quad + (-J_1 \ I)\left(\xi_i - \hat{\xi}^s_{i|N}\right)\left(\hat{x}^s_{i|N} - F_1(\hat{x}^s_{i-1|N})\right)^T \\
&\quad + \left(\hat{x}^s_{i|N} - F_1(\hat{x}^s_{i-1|N})\right)\left(\xi_i - \hat{\xi}^s_{i|N}\right)^T (-J_1 \ I)^T \\
&\quad + \left(\hat{x}^s_{i|N} - F_1(\hat{x}^s_{i-1|N})\right)\left(\hat{x}^s_{i|N} - F_1(\hat{x}^s_{i-1|N})\right)^T, \quad (30)
\end{aligned}
$$

where the Taylor expansion

$$
F_1(x_{i-1}) \approx F_1(\hat{x}^s_{i-1|N}) + J_1(x_{i-1} - \hat{x}^s_{i-1|N}), \quad (31a)
$$

$$
J_1 = \left.\frac{\partial F_1(x)}{\partial x}\right|_{x=\hat{x}^s_{i-1|N}}, \quad (31b)
$$

has been used.

The integral in (26) now becomes

$$
\begin{aligned}
M &\triangleq \int \tilde{F}_{1,i}\tilde{F}_{1,i}^T p_{Q_l}(x_i, x_{i-1}|y_{1:N})dx_i dx_{i-1} \\
&= (-J_1 \ I)\, P^{\xi,s}_{i|N}\, (-J_1 \ I)^T \\
&\quad + \left(\hat{x}^s_{i|N} - F_1(\hat{x}^s_{i-1|N})\right)\left(\hat{x}^s_{i|N} - F_1(\hat{x}^s_{i-1|N})\right)^T. \quad (32)
\end{aligned}
$$

It is thus possible to calculate $\Gamma(Q; Q_l)$ according to

$$
\begin{aligned}
\Gamma(Q; Q_l) &= \bar{L} + \frac{1}{2}\sum_{i=2}^{N} \log\left(\left|\left(\tilde{F}_2^\dagger(\hat{x}^s_{i-1|N})\right)^T\right|\right) \\
&\quad + \frac{1}{2}\sum_{i=2}^{N}\left[\log\left|Q^\dagger\right| + \log\left(\left|\tilde{F}_2^\dagger(\hat{x}^s_{i-1|N})\right|\right)\right] \\
&\quad - \frac{1}{2}\mathrm{Tr}Q^\dagger \sum_{i=2}^{N} F_2^\dagger(\hat{x}^s_{i-1|N}) M \left(F_2^\dagger(\hat{x}^s_{i-1|N})\right)^T \quad (33)
\end{aligned}
$$

where we have used the fact that

$$
\left(\tilde{F}_{2,i-1}Q\tilde{F}_{2,i-1}^T\right)^\dagger = \left(\tilde{F}_{2,i-1}^\dagger\right)^T Q^\dagger \tilde{F}_{2,i-1}^\dagger. \quad (34)
$$

In (34) it is used that $\tilde{F}_{2,i-1}^T$ and Q have full row rank, and $\tilde{F}_{2,i-1}$ and $\tilde{F}_{2,i-1}Q$ have full column rank, see Mitra and Rao [1971].

3.2 Maximisation step

Maximisation with respect to Q is the same as maximisation with respect to $Q^\dagger = Q^{-1}$. Take the derivative of (33) with respect to Q^{-1} and let the result be equal to zero gives

$$
\begin{aligned}
\frac{\partial \Gamma(Q; Q_l)}{\partial Q^{-1}} &= \frac{N-1}{2}Q \\
&\quad - \frac{1}{2}\sum_{i=2}^{N} F_2^\dagger(\hat{x}^s_{i-1|N}) M \left(F_2^\dagger(\hat{x}^s_{i-1|N})\right)^T = 0. \quad (35)
\end{aligned}
$$

The solution of the maximisation step is now obtained as

$$
Q_{l+1} = \frac{1}{N-1}\sum_{i=2}^{N} F_2^\dagger(\hat{x}^s_{i-1|N}) M \left(F_2^\dagger(\hat{x}^s_{i-1|N})\right)^T. \quad (36)
$$

3.3 Stop Criterion

The stop criterion can be chosen in different ways. Section 2.1 suggests that the EM algorithm stops when the difference in the new and previous estimate is less than a threshold. Another way is to use $L_Q(y_{1:N})$ in (4). The main problem is to maximise $L_Q(y_{1:N})$, therefore stop the algorithm when no increase in $L_Q(y_{1:N})$ can be observed. Equation (4) can be written as

$$
\begin{aligned}
L_Q(y_{1:N}) &= \log p_Q(y_{1:N}) = \log\left(p(y_1)\prod_{i=1}^{N-1} p_Q(y_{i+1}|y_{1:i})\right) \\
&= \log p(y_1) + \sum_{i=1}^{N-1} \log p_Q(y_{i+1}|y_{1:i}), \quad (37)
\end{aligned}
$$

where Bayes' rule has been used repeatedly. Here, $\log p(y_1)$ is a constant and can be omitted in the sequel for simplicity. The pdf $p_Q(y_{i+1}|y_{1:i})$ is identified as the pdf for the innovations which can be calculated as

$$
p_Q(y_{i+1}|y_{1:i}) = \mathcal{N}\left(y_{i+1}; h(\hat{x}_{i+1|i}), HP_{i+1|i}H^T + R\right), \quad (38)
$$

$$
H = \left.\frac{\partial h(x)}{\partial x}\right|_{x=\hat{x}_{i+1|i}}, \quad (39)
$$

where $\hat{x}_{i+1|i}$ and $P_{i+1|i}$ are calculated in the EKF during the measurement update. The algorithm can now be stopped when

$$
\left|L_{Q_l}(y_{1:N}) - L_{Q_{l-m}}(y_{1:N})\right| \leq \gamma, \quad (40)
$$

where m and γ are parameters to choose.

4. ALTERNATIVE WAYS TO FIND THE COVARIANCE MATRIX OF THE PROCESS NOISE

There are many alternative ways of estimating the covariance matrix for the process noise and here two examples will be described. These two alternatives, which are less complicated than the EM algorithm, will be compared to the result of the EM algorithm in Section 6.

The first method minimises the path error

$$
\mathsf{e}_k = \sqrt{|\mathsf{x}_k - \hat{\mathsf{x}}_k|^2 + |\mathsf{y}_k - \hat{\mathsf{y}}_k|^2}, \quad (41)
$$

where x_k and y_k are the true coordinates for the tool, and $\hat{\mathsf{x}}_k$ and $\hat{\mathsf{y}}_k$ are the estimated coordinates for the tool. To simplify the problem, the covariance matrix is parametrised as a diagonal matrix.

Algorithm 2. Minimisation of the path error

(1) Select an initial diagonal matrix $Q_0 \in \mathbb{R}^{4\times 4}$.
(2) Minimise $\sqrt{\sum_{k=1}^{N} |\mathsf{e}_k|^2}$ subject to $\lambda_j > 0, j = 1, \ldots, 4$ $Q = \mathrm{diag}(\lambda_1, \lambda_2, \lambda_3, \lambda_4)Q_0$ and $(\hat{\mathsf{x}}, \hat{\mathsf{y}}) = \mathrm{EKF}(Q)$.
(3) The sought covariance matrix Q are obtained as $Q = \mathrm{diag}(\lambda_1^*, \lambda_2^*, \lambda_3^*, \lambda_4^*)Q_0$, where λ_j^* is the optimal value from step 2.

The notation $(\hat{\mathsf{x}}, \hat{\mathsf{y}}) = \mathrm{EKF}(Q)$ in Algorithm 2 denotes that the estimated position is a function of Q. A standard optimisation method can be used to solve step 2 in Algorithm 2, see e.g. Boyd and Vandenberghe [2009]. The problem is not convex, i.e., the solution is not guaranteed to give a global optimum. However, the method is straightforward and has been used before, see Henriksson

et al. [2009] and Axelsson [2009]. One disadvantage with the method is that the true tool position is required.

The second method starts with an initial guess Q_0. The smoothed states are then calculated using Q_0. After that, equation (1a) and the smoothed states are used in order to derive the noise v_k, $k = 1, \ldots, N$. The covariance matrix is finally obtained from the vector $v_{1:N} = \{v_1 \ \ldots \ v_N\}$. The method is repeated with the new Q-matrix until convergence is obtained. The method is summarised in the following steps.

Algorithm 3. Iterative covariance estimation with EKS

(1) Select an initial value Q_0 and set $l = 0$.
(2) Use the EKS with Q_l.
(3) Calculate the noise according to

$$v_k = F_2^{\dagger}\left(\hat{x}^s_{k|N}\right)\left(\hat{x}^s_{k+1|N} - F_1\left(\hat{x}^s_{k|N}, u_k\right)\right).$$

(4) Let Q_{l+1} be the covariance matrix for v_k according to

$$Q_{l+1} = \frac{1}{N}\sum_{k=1}^{N} v_k^T v_k.$$

(5) If converged, stop, If not, set $l = l+1$ and go to step 2.

The first method, based on the minimisation of the path error, is called Alg. 2 and the second method is called Alg. 3 in the rest of the paper.

5. APPLICATION TO INDUSTRIAL ROBOTS

The robot model is a joint flexible two axes model from Moberg et al. [2008]. The model assumes rigid links and flexible joints. Each joint is a two mass system consisting of two angles, the arm angle q_{ai}, and the motor angle q_{mi}. Let the state vector be given by

$$x = \begin{pmatrix} x_1^T & x_2^T & x_3^T & x_4^T \end{pmatrix}^T = \begin{pmatrix} q_{\mathrm{a}}^T & q_{\mathrm{m}}^T & \dot{q}_{\mathrm{a}}^T & \dot{q}_{\mathrm{m}}^T \end{pmatrix}^T, \qquad (42)$$

where $q_a = \begin{pmatrix} q_{\mathrm{a}1} & q_{\mathrm{a}2} \end{pmatrix}^T$, $q_m = \begin{pmatrix} q_{\mathrm{m}1} & q_{\mathrm{m}2} \end{pmatrix}^T$, contain the arm angles q_{a} and the motor angles q_{m} of both joints. The model accounts for flexibilities in the joints via non-linear stiffness and linear viscous damping. The model also includes non-linear friction. A continuous-time state space model of the system is given by,

$$\dot{x} = \begin{pmatrix} x_3 \\ x_4 \\ M_a^{-1}(x_1)\big(-C(x_1, x_3) - G(x_1) - A(x) + v_a\big) \\ M_m^{-1}\big(A(x) + \kappa(x_4) + u + v_m\big) \end{pmatrix}, \qquad (43)$$

where $A(x) = D(x_3 - x_4) + \tau_s(x_1, x_2)$. $A(x)$ accounts for the flexibilities in the joints, via the linear viscous damping $D(x_3 - x_4)$ and the non-linear stiffness $\tau_s(x_1, x_2)$. In other words, if we dispense with $A(x)$, we are back at a standard rigid robot model. Furthermore, $M_a(x_1)$ and M_m are the mass matrices for the arm and motor, $C(x_1, x_3)$ accounts for the centrifugal and centripetal torques, and $G(x_1)$ accounts for the effect of gravity on the links. The non-linear friction is described by $\kappa(x_3)$, u represents the motor torque applied to the robot and $v = \begin{pmatrix} v_a & v_m \end{pmatrix}^T$ is the process noise. An Euler forward approximation of (43) gives a discretised model according to (1) and (2). The rank conditions in order to use (34) are also satisfied.

Table 1. Max and min of the 2-norm of the path error in mm for the three different methods.

	Max	Min
EM	0.2999	0.2996
Alg. 2	3.3769	1.5867
Alg. 3	2.6814	2.6814

6. SIMULATION RESULTS

The method in Section 3 is evaluated and compared to the two alternative methods described in Section 4. The model given in Section 2 is first simulated, according to Axelsson [2009], to get all the required quantities, i.e., u_k, y_k, x_k and y_k. In system identification, it is common to estimate a certain parameter or parameters starting at different initial values and see if the true one is obtained. Here, on the other hand, there is no information about the true covariance matrix, even for simulated data. Instead, the estimated covariance matrices, for different initial values, are used to calculate the path error according to (41). When the path error differs a lot with different initial values it means that the method converges to different local optima. There is however no guarantee that a solution is in a global optimum although the path errors do not differ. Here, the maximum and minimum of the 2-norm of the path error is used to see how much the solutions differ with different initial values. It is preferred to have a method that results in small and similar path errors for different initial values.

Table 1 shows that the maximal and minimal path errors for the EM algorithm are more or less the same. The same concerns Alg. 3. The EM algorithm gives however a lower path error. Alg. 2 gives, on the other hand, path errors that differs considerably. This can also be seen in Fig. 1. This means that Alg. 2 gets stuck in different local optima. A comparison between the path errors for the EM algorithm, Alg. 3 and the best solution of Alg. 2 is shown in Fig. 2. The EM algorithm is clearly much better than the two alternatives.

It is also interesting to see how (37) looks like for Q_l, $l = 0, \ldots$, both for the EM algorithm and Alg. 3. The EM algorithm and Alg. 3 were therefore forced to take more iterations than necessary. The log-likelihood function (37) can be seen in Fig. 3 for 100 iterations. We see that the curve for the EM algorithm flattens out somewhere around 50 iterations and stays constant after that. It means that it is unnecessary to continue to more than about 60 iterations. One thing to comment is the peak around 10 iterations in the curve. This contradicts the property of the EM algorithm that the sequence Q_l, $l = 0, \ldots$, approximates $\hat{Q}^{ML}$ better and better. This can be explained by the approximations that have been made during the Expectation step and that the calculation of (37) in the EKF is approximately. The curve for Alg. 3 flattens out after 10 iterations and stays constant after that. Alg. 3 is also without any peak and the stationary value is lower than for the EM algorithm.

7. CONCLUSIONS AND FUTURE WORK

Three different methods to estimate the covariance matrices have been compared. The EM algorithm derived in

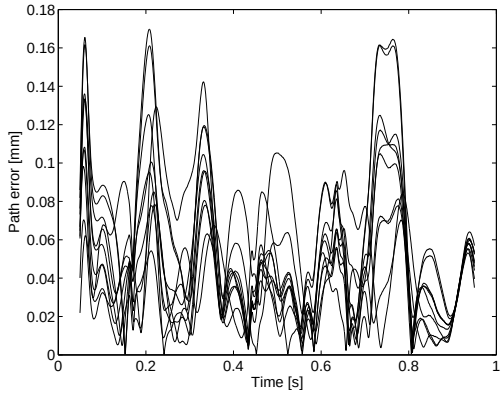

Fig. 1. The path error for 10 Monte Carlo simulations of Alg. 2.

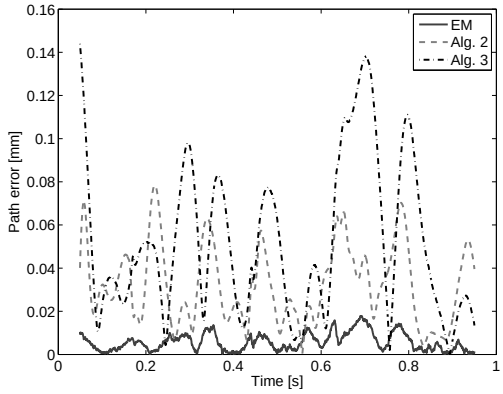

Fig. 2. The best path error for the EM algorithm and the two alternative methods.

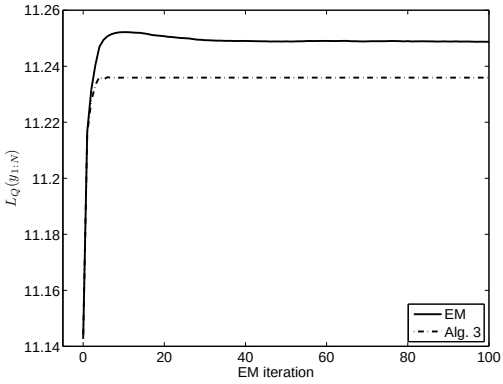

Fig. 3. The log-likelihood function for the first 100 iterations in the EM algorithm and Alg. 3.

Section 3 gives a lower path error, considering the true path and the estimated path from an EKF. The EM algorithm is also robust to changes in the initial value. One advantage with the EM algorithm is that no true tool position is needed, which is the case for Alg. 2. A next step is to use the EM algorithm on experimental data to estimate the covariance matrices.

ACKNOWLEDGEMENTS

This work was supported by the Swedish Research Council under the Linnaeus Center CADICS and by Vinnova Excellence Center LINK-SIC.

REFERENCES

Brian D. O. Anderson and John B. Moore. *Optimal Filtering*. Information and System Sciences Series. Prentice Hall Inc., Englewood Cliffs, New Jersey, USA, 1979.

Patrik Axelsson. A simulation study on the arm estimation of a joint flexible 2 dof robot arm. Technical Report LiTH-ISY-R-2926, Department of Electrical Enginering, Linköping University, SE-581 83 Linköping, Sweden, December 2009.

Mattias Björkman, Torgny Brogårdh, Sven Hanssen, Sven-Erik Lindström, Stig Moberg, and Mikael Norrlöf. A new concept for motion control of industrial robots. In *Proceedings of 17th IFAC World Congress*, pages 15714–15715, Seoul, Korea, July 2008.

Torsten Bohlin. *Practical Grey-box Process Identification, Theory and Applications*. Advances in Industrial Control. Springer, London, UK, 2006.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, United Kingdom, first edition, 2009.

Olivier Cappé, Eric Moulines, and Tobias Rydén. *Inference in Hidden Markov Models*. Springer Series in Statistics. Springer, New York, USA, 2005.

André Carvalho Bittencourt, Erik Wernholt, Shiva Sander-Tavallaey, and Torgny Brogårdh. An extended friction model to capture load and temperature effects in robot joints. Technical Report LiTH-ISY-R-2948, Department of Electrical Enginering, Linköping University, SE-581 83 Linköping, Sweden, April 2010.

A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.

Stuart Gibson and Brett Ninness. Robust maximum-likelihood estimation of multivariable dynamic systems. *Automatica*, 41(10): 1667–1682, October 2005.

Robert Henriksson, Mikael Norrlöf, Stig Moberg, Erik Wernholt, and Thomas B. Schön. Experimental comparison of observers for tool position estimation of industrial robots. In *Proceedings of 48th IEEE Conference on Decision and Control*, pages 8065–8070, Shanghai, China, December 2009.

Lennart Ljung. *System Identification, Theory for the User*. Information and System Sciences Series. Prentice Hall Inc., Upper Saddle River, New Jersey, USA, second edition, 1999.

Geoffrey J. McLachlan and Thriyambakam Krishnan. *The* EM *Algorithm and Extensions*. Wiley Series in Probability and Statistics. John Wiley & Sons, Hoboken, New Jersey, USA, second edition, 2008.

Sujit Kumar Mitra and C. Radhakrishna Rao. *Generalized Inverse of Matrices and its Applications*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, 1971.

Stig Moberg, Jonas Öhr, and Svante Gunnarsson. A benchmark problem for robust control of a multivariable nonlinear flexible manipulator. In *Proceedings of 17th IFAC World Congress*, pages 1206–1211, Seoul, Korea, July 2008. URL: `http://www.robustcontrol.org`.

Thomas B. Schön, Adrian Wills, and Brett Ninness. System identification of nonlinear state-space models. *Automatica*, 2010. (Accepted for publication).

Johanna Wallén, Svante Gunnarsson, Robert Henriksson, Stig Moberg, and Mikael Norrlöf. ILC applied to a flexible two-link robot model using sensor-fusion-based estimates. In *Proceedings of the 48th IEEE Conference on Decision and Control*, pages 458–463, Shanghai, China, December 2009.

Erik Wernholt. *Multivariable Frequency-Domain Identification of Industrial Robots*. Linköping studies in science and technology. dissertations. no. 1138, SE-581 83 Linköping, Sweden, November 2007.

Byron M. Yu, Krishna V. Shenoy, and Maneesh Sahani. Derivation of extended kalman filtering and smoothing equations, October 2004. URL: `www-npl.stanford.edu/~byronyu/papers/derive_eks.pdf`.

Avdelning, Institution
Division, Department

Division of Automatic Control
Department of Electrical Engineering

Datum
Date

2010-10-08

Språk
Language

☐ Svenska/Swedish
☒ Engelska/English

☐ _____

Rapporttyp
Report category

☐ Licentiatavhandling
☐ Examensarbete
☐ C-uppsats
☐ D-uppsats
☒ Övrig rapport
☐ _____

ISBN
—

ISRN
—

Serietitel och serienummer
Title of series, numbering

ISSN
1400-3902

LiTH-ISY-R-2969

URL för elektronisk version

http://www.control.isy.liu.se

Titel
Title

ML Estimation of Process Noise Variance in Dynamic Systems

Författare
Author

Patrik Axelsson, Umut Orguner, Fredrik Gustafsson, Mikael Norrlöf

Sammanfattning
Abstract

The performance of a non-linear filter hinges in the end on the accuracy of the assumed non-linear model of the process. In particular, the process noise covariance Q is hard to get by physical modeling and dedicated system identification experiments. We propose a variant of the expectation maximization (EM) algorithm which iteratively estimates the unobserved state sequence and Q based on the observations of the process. The extended Kalman smoother (EKS) is the instrument to find the unobserved state sequence. Our contribution fills a gap in literature, where previously only the linear Kalman smoother and particle smoother have been applied. The algorithm will be important for future industrial robots with more flexible structures, where the particle smoother cannot be applied due to the high state dimension. The proposed method is compared to two alternative methods on a simulated robot.

Nyckelord
Keywords

Robotic manipulators, Extended Kalman filters, Smoothing filters, Identification, Maximum likelihood, Covariance matrices