

Sigma-Point Filtering and Smoothing Based Parameter Estimation in Nonlinear Dynamic Systems

JUHO KOKKALA
ARNO SOLIN
SIMO SÄRKKÄ

We consider approximate maximum likelihood parameter estimation in nonlinear state-space models. We discuss both direct optimization of the likelihood and expectation-maximization (EM). For EM, we also give closed-form expressions for the maximization step in a class of models that are linear in parameters and have additive noise. To obtain approximations to the filtering and smoothing distributions needed in the likelihood-maximization methods, we focus on using Gaussian filtering and smoothing algorithms that employ sigma-points to approximate the required integrals. We discuss different sigma-point schemes based on the third, fifth, seventh, and ninth order unscented transforms and the Gauss-Hermite quadrature rule. We compare the performance of the methods in two simulated experiments: a univariate nonlinear growth model as well as tracking of a maneuvering target. In the experiments, we also compare against approximate likelihood estimates obtained by particle filtering and extended Kalman filtering based methods. The experiments suggest that the higher-order unscented transforms may in some cases provide more accurate estimates.

Manuscript received March 31, 2015; revised August 31, 2015; released for publication November 9, 2015.

Refereeing of this contribution was handled by Jindrich Dunik.

Authors' addresses: Aalto University, Espoo, Finland. (e-mail: juho.kokkala@aalto.fi, arno.solin@aalto.fi, simo.sarkka@aalto.fi).

This work was supported by grants from the Academy of Finland (266940, 273475) and by the Emil Aaltonen foundation.

The authors declare no conflict-of-interest.

1557-6418/16/\$17.00 © 2016 JAIF

I. INTRODUCTION

This paper is an extended version of our article [24] where we considered parameter estimation in state-space models using expectation-maximization (EM) algorithms based on sigma-point and particle smoothers. In this paper, we extend our interest from EM algorithms to so called direct maximum likelihood based parameter estimation methods, where instead of using the EM algorithm, the marginal likelihood of the parameters is directly approximated using nonlinear filtering methods. In particular, we focus our interest to sigma-point filters which use high-order unscented Kalman filters and Gauss-Hermite Kalman filters to approximate the likelihood surface.

We consider state-space models of the following form:

$$\begin{aligned}\mathbf{x}_k &= \mathbf{f}(\mathbf{x}_{k-1}, \boldsymbol{\theta}) + \mathbf{q}_{k-1}, \\ \mathbf{y}_k &= \mathbf{h}(\mathbf{x}_k, \boldsymbol{\theta}) + \mathbf{r}_k,\end{aligned}\quad (1)$$

where $\mathbf{x}_k \in \mathbb{R}^n$ is the discrete-time state sequence with an initial distribution $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{x}_0 \mid \mathbf{m}_0(\boldsymbol{\theta}), \mathbf{P}_0(\boldsymbol{\theta}))$, $\mathbf{y}_k \in \mathbb{R}^d$ is the measurement sequence, $\mathbf{q}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}(\boldsymbol{\theta}))$ is the Gaussian process noise sequence, $\mathbf{r}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}(\boldsymbol{\theta}))$ is the Gaussian measurement error sequence, and $\boldsymbol{\theta} \in \mathbb{R}^m$ is a static parameter vector. Typically, one is interested in computing the posterior distribution of the state \mathbf{x}_k given measurements up to time k , $p(\mathbf{x}_k \mid \mathbf{y}_1, \dots, \mathbf{y}_k)$, known as the filtering problem, or computing the posterior distribution of the state \mathbf{x}_k given all measurements, $p(\mathbf{x}_k \mid \mathbf{y}_1, \dots, \mathbf{y}_T)$, where $k \leq T$, known as the smoothing problem. In the general case, analytical expressions do not exist and we have to resort to approximative algorithms such as the sigma-point methods. See, for example, [32] for a general overview of Bayesian filtering and smoothing.

While many filtering and smoothing algorithms are formulated assuming fixed static parameters $\boldsymbol{\theta}$, in practice optimal values for these parameters are generally unknown. Therefore, methods for estimating the parameters from the data are desired. In this paper, we concentrate on maximum-likelihood methods, where the parameters are selected by maximizing the marginal likelihood, or equivalently the logarithm of the marginal likelihood, that is

$$\boldsymbol{\theta}_{\text{ML}} = \arg \max_{\boldsymbol{\theta}} \log p(\mathbf{y}_{1:T} \mid \boldsymbol{\theta}). \quad (2)$$

In linear systems with additive Gaussian noise, the likelihood can be evaluated using the Kalman filter [17, 21]. Many optimization algorithms utilize also the gradient of the log-likelihood. The gradient can be evaluated by so-called sensitivity equations, a recursion that is obtained by differentiating the Kalman filter recursion [15]. Alternatively, due to Fisher's identity, the gradient may be evaluated by differentiating an auxiliary function that can be computed during the smoothing

pass [28, 35]. Instead of directly optimizing the likelihood, the expectation-maximization (EM) algorithm [9] can be used to optimize parameters. The EM algorithm consists of iterating the expectation (E) step where a bound of the log-likelihood is computed using the current parameter estimates, and the maximization (M) step where the bound is maximized with respect to the parameters. The evaluation of the bound in the E-step is obtained by solving the smoothing problem. See [36] for a discussion of applying the EM algorithm in state-space models. Note that in the linear-Gaussian case both gradient evaluation methods as well as the EM algorithm in principle converge to the same solution, namely, the parameter value that maximizes the log-likelihood.

In this paper, our interest lies in estimating the static parameters by maximum-likelihood estimation in the case of nonlinear state-space models with additive Gaussian noise, that is model (1). Formally, the marginal likelihood can be computed by marginalizing out the states from the joint distribution of the measurements and states using nonlinear filtering equations and the prediction error decomposition (see, e.g., [5, 32]), leading to similar methods as in the linear-Gaussian case. However, since the state variables \mathbf{x} cannot in general be marginalized out analytically, one needs to employ approximative methods. In the so called direct likelihood methods, the likelihood is approximated directly using approximative nonlinear filtering methods (see, e.g., [5, 22, 29, 32, 37]) and its maximum is found via nonlinear optimization. Similarly, the expectation-maximization (EM) algorithm can be employed, but the E-step cannot be solved exactly. Instead, the E-step is approximated with nonlinear smoothing algorithms (see, e.g., [12, 24, 31, 34, 41]).

The aim of this paper is to extend the results of our paper [24] by showing how high-order (i.e., third, fifth, seventh, and ninth order) unscented transforms and Gauss-Hermite integration based sigma-point methods can be used for approximate direct likelihood and EM-based parameter estimation in nonlinear state-space models. For EM, we also give closed-form expressions for the maximization step in a class of models that are linear in parameters and have additive noise. We compare the unscented transform and Gauss-Hermite based sigma-point methods to linearization-based extended Kalman filter algorithms and Monte Carlo based particle filtering algorithms. We also provide an algorithm for computing the gradients required by the gradient-based optimization methods. Although we focus on maximum likelihood estimation, the provided algorithms can be easily extended to computation of maximum a posteriori estimates by including a prior distribution to the objective function.

The remainder of the paper is organized as follows. In Section II, we present the sigma-point filters and smoothers based on assumed Gaussian density filtering

and smoothing. In Section III, we discuss parameter estimation based on direct likelihood maximization and EM. Numerical experiments are presented in Section IV. Finally, Section V presents concluding discussion.

II. SIGMA-POINT FILTERING AND SMOOTHING

Under our interpretation, sigma-point filtering and smoothing is derived by assuming Gaussian approximations for the state distributions, which enables the use of a Kalman filter like filtering recursion and a Rauch-Tung-Striebel backward pass for the smoothing distributions. The Gaussian filtering and smoothing equations contain expectations over Gaussian distributions which cannot be generally evaluated in closed form. The sigma-points arise from approximating these Gaussian integrals by weighted sums determined by some cubature (multi-dimensional quadrature) formula. Hence, we interpret the different sigma-point methods as incarnations of different integral approximations.

In the following, we first present the assumed Gaussian density filtering and smoothing framework. Then, we discuss various different cubature rules for approximating the Gaussian integrals. Finally, we show how the cubature rules are applied to the assumed Gaussian density filtering and smoothing framework to obtain the filtering and smoothing equations explicitly in the sigma-point form.

A. General Gaussian Filtering and Smoothing

Assumed density Gaussian filtering (see [16, 32, 33]) is based on assuming that the filtering distributions are approximately Gaussian, that is, assuming means $\mathbf{m}_{k|k}$ and covariances $\mathbf{P}_{k|k}$ such that

$$p(\mathbf{x}_k | \mathbf{y}_{1:k}) \approx \mathbb{N}(\mathbf{x}_k | \mathbf{m}_{k|k}, \mathbf{P}_{k|k}) \quad (3)$$

as well as means $\mathbf{m}_{k|k+1}$ and covariances $\mathbf{P}_{k|k+1}$ such that

$$p(\mathbf{x}_{k+1} | \mathbf{y}_{1:k}) \approx \mathbb{N}(\mathbf{x}_{k+1} | \mathbf{m}_{k|k+1}, \mathbf{P}_{k|k+1}). \quad (4)$$

The filtering equations of the resulting Gaussian filter [16, 43] consist of a *prediction step* and an *update step*. In the prediction step, we compute the state mean and covariance of the distribution $p(\mathbf{x}_k | \mathbf{y}_{1:k-1})$ using the Gaussian approximation for $p(\mathbf{x}_{k-1} | \mathbf{y}_{1:k-1})$. The resulting equations are

$$\begin{aligned} \mathbf{m}_{k|k-1} &= \mathbb{E}[\mathbf{f}(\mathbf{x}_{k-1})], \\ \mathbf{P}_{k|k-1} &= \mathbb{E}[(\mathbf{f}(\mathbf{x}_{k-1}) - \mathbf{m}_{k|k-1}) \\ &\quad \times (\mathbf{f}(\mathbf{x}_{k-1}) - \mathbf{m}_{k|k-1})^T] + \mathbf{Q}, \end{aligned} \quad (5)$$

where the expectations are taken with respect to the distribution $\mathbf{x}_{k-1} \sim \mathbb{N}(\mathbf{m}_{k-1|k-1}, \mathbf{P}_{k-1|k-1})$.

In the corresponding update step, we assume a Gaussian density $p(\mathbf{x}_k | \mathbf{y}_{1:k-1}) = \mathbb{N}(\mathbf{x}_k | \mathbf{m}_{k|k-1}, \mathbf{P}_{k|k-1})$ and

compute the state mean and covariance for the distribution $p(\mathbf{x}_k | \mathbf{y}_{1:k})$. The resulting equations are

$$\begin{aligned}
\boldsymbol{\mu}_k &= \mathbb{E}[\mathbf{h}(\mathbf{x}_k)], \\
\mathbf{S}_k &= \mathbb{E}[(\mathbf{h}(\mathbf{x}_k) - \boldsymbol{\mu}_k)(\mathbf{h}(\mathbf{x}_k) - \boldsymbol{\mu}_k)^\top] + \mathbf{R}, \\
\mathbf{C}_k &= \mathbb{E}[(\mathbf{x}_k - \mathbf{m}_{k|k-1})(\mathbf{h}(\mathbf{x}_k) - \boldsymbol{\mu}_k)^\top], \\
\mathbf{K}_k &= \mathbf{C}_k \mathbf{S}_k^{-1}, \\
\mathbf{m}_{k|k} &= \mathbf{m}_{k|k-1} + \mathbf{K}_k (\mathbf{y}_k - \boldsymbol{\mu}_k), \\
\mathbf{P}_{k|k} &= \mathbf{P}_{k|k-1} - \mathbf{K}_k \mathbf{S}_k \mathbf{K}_k^\top,
\end{aligned} \tag{6}$$

where $\boldsymbol{\mu}_k$ and \mathbf{S}_k are the expectation and variance, respectively, of the measurement \mathbf{y}_k and \mathbf{C}_k is the covariance of the state \mathbf{x}_k and the measurement \mathbf{y}_k . These expectations are taken with respect to the distribution $\mathbf{x}_k \sim \mathbb{N}(\mathbf{m}_{k|k-1}, \mathbf{P}_{k|k-1})$.

The smoothing distributions $p(\mathbf{x}_k | \mathbf{y}_{1:T})$ are obtained from a backward pass, that is, starting from $k = T$ and iterating backwards in time. On each step, the smoothing density of \mathbf{x}_{k+1} is assumed to be Gaussian: $p(\mathbf{x}_{k+1} | \mathbf{y}_{1:T}) = \mathbb{N}(\mathbf{x}_{k+1} | \mathbf{m}_{k+1|T}, \mathbf{P}_{k+1|T})$. The mean and covariance for $p(\mathbf{x}_k | \mathbf{y}_{1:T})$ are then computed from the previous Gaussian smoothing density and the Gaussian filtering densities using the Rauch-Tung-Striebel backward pass [13, 30] as follows [33]:

$$\begin{aligned}
\mathbf{m}_{k+1|k} &= \mathbb{E}[\mathbf{f}(\mathbf{x}_k)], \\
\mathbf{P}_{k+1|k} &= \mathbb{E}[(\mathbf{f}(\mathbf{x}_k) - \mathbf{m}_{k+1|k}) \\
&\quad \times (\mathbf{f}(\mathbf{x}_k) - \mathbf{m}_{k+1|k})^\top] + \mathbf{Q}, \\
\mathbf{D}_{k+1} &= \mathbb{E}[(\mathbf{x}_k - \mathbf{m}_{k|k})(\mathbf{f}(\mathbf{x}_k) - \mathbf{m}_{k+1|k})^\top], \\
\mathbf{G}_k &= \mathbf{D}_{k+1} [\mathbf{P}_{k+1|k}]^{-1}, \\
\mathbf{m}_{k|T} &= \mathbf{m}_{k|k} + \mathbf{G}_k (\mathbf{m}_{k+1|T} - \mathbf{m}_{k+1|k}), \\
\mathbf{P}_{k|T} &= \mathbf{P}_{k|k} + \mathbf{G}_k (\mathbf{P}_{k+1|T} - \mathbf{P}_{k+1|k}) \mathbf{G}_k^\top,
\end{aligned} \tag{7}$$

where \mathbf{G}_k is known as the smoother gain and the expectations are taken with respect to the distribution $\mathbf{x}_k \sim \mathbb{N}(\mathbf{m}_{k|k}, \mathbf{P}_{k|k})$. The pairwise joint smoothing distributions $p(\mathbf{x}_k, \mathbf{x}_{k-1} | \mathbf{y}_{1:T})$ are also of interest since they are used in the expectation-maximization algorithm (see Section III-B). Gaussian approximations for these distributions are obtained as a by-product of the smoothing backward pass results as follows (see, e.g., [32, p. 189]).

$$\begin{aligned}
p(\mathbf{x}_k, \mathbf{x}_{k-1} | \mathbf{y}_{1:T}) &\approx \\
&\mathbb{N} \left(\begin{pmatrix} \mathbf{x}_k \\ \mathbf{x}_{k-1} \end{pmatrix} \middle| \begin{pmatrix} \mathbf{m}_{k|T} \\ \mathbf{m}_{k-1|T} \end{pmatrix}, \begin{pmatrix} \mathbf{P}_{k|T} & \mathbf{P}_{k|T} \mathbf{G}_{k-1}^\top \\ \mathbf{G}_{k-1} \mathbf{P}_{k|T} & \mathbf{P}_{k-1|T} \end{pmatrix} \right).
\end{aligned} \tag{8}$$

B. Approximating the Gaussian Integrals

As we saw in the previous section, during the evaluation of the prediction and update steps of the Gaussian

filter and smoother, we need to solve a set of Gaussian integrals on each step. These integrals are of the following form:

$$\mathbb{E}[\mathbf{g}(\mathbf{x})] = \int_{\mathbb{R}^n} \mathbf{g}(\mathbf{x}) \mathbb{N}(\mathbf{x} | \mathbf{m}, \mathbf{P}) d\mathbf{x}, \tag{9}$$

where $\mathbf{g}: \mathbb{R}^n \rightarrow \mathbb{R}^d$ is the integrand and the weighting function $\mathbb{N}(\mathbf{x} | \mathbf{m}, \mathbf{P})$ is a multi-dimensional Gaussian density with mean \mathbf{m} and covariance matrix \mathbf{P} . In this paper, these integrals are computed by using multi-dimensional generalizations of Gaussian quadratures—also referred to as Gaussian *cubatures* [6]. They give approximations of the form

$$\mathbb{E}[\mathbf{g}(\mathbf{x})] \approx \sum_i w_i \mathbf{g}(\mathbf{x}_i), \tag{10}$$

where the weights w_i and sigma-points \mathbf{x}_i are functions of the mean \mathbf{m} and covariance \mathbf{P} of the Gaussian weighting function. The sigma-points are positioned as follows:

$$\mathbf{x}_i = \mathbf{m} + \mathbf{L} \boldsymbol{\xi}_i, \tag{11}$$

where $\boldsymbol{\xi}_i$ are method specific unit sigma-points, and \mathbf{L} is a matrix square-root factor such that $\mathbf{P} = \mathbf{L} \mathbf{L}^\top$ (e.g., the Cholesky decomposition of \mathbf{P}). The differences in the methods come from different choices of weights and unit sigma-points.

In the following we briefly introduce a number of schemes for choosing the weights and sigma-points. The difference between these schemes stems from a trade-off between the number of sigma-points (required function evaluations) versus precision in the approximation. The degree of approximation is quantified by the highest polynomial order, p , for which the method is exact.

Unscented transform. The *unscented transform* (UT, [19, 20]) uses a set of $2n + 1$ cubature points located in the center and on the surface of an n -sphere. The radius and the weights can be controlled using a set of parameters. The cubature points are given by:

$$\begin{aligned}
\boldsymbol{\xi}_0 &= \mathbf{0}, \\
\boldsymbol{\xi}_i &= \begin{cases} \sqrt{\lambda + n} \mathbf{e}_i, & i = 1, \dots, n, \\ -\sqrt{\lambda + n} \mathbf{e}_{i-n}, & i = n + 1, \dots, 2n, \end{cases}
\end{aligned}$$

where \mathbf{e}_i denotes a unit vector to the direction of coordinate axis i , and the weights are defined as follows:

$$\begin{aligned}
w^{(0)} &= \begin{cases} \frac{\lambda}{n + \lambda}, & \text{for mean terms,} \\ \frac{\lambda}{n + \lambda} + (1 - \alpha^2 + \beta), & \text{for covariance terms,} \end{cases} \\
w^{(i)} &= \frac{1}{2(n + \lambda)}, \quad i = 1, \dots, 2n,
\end{aligned}$$

where $\lambda = \alpha^2(n + \kappa) - n$ and α , β , and κ are parameters of the method.

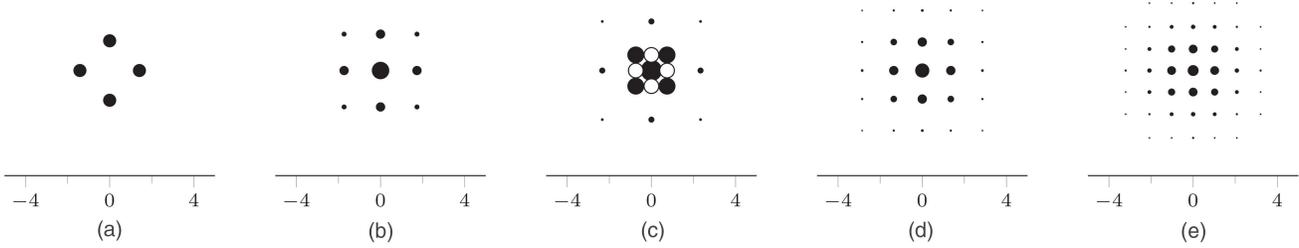


Fig. 1. Unit sigma-points in two dimensions for each of the methods. The absolute value of the weights are indicated by the point size, positive weights being black, negative weights white. (a) The symmetric cubature rule of order $p = 3$ with 4 points. (b) The symmetric cubature rule of order $p = 5$ with 9 points. (c) The symmetric cubature rule of order $p = 7$ with 17 points. (d) The symmetric cubature rule of order $p = 9$ with 25 points. (e) For comparison, the Gauss-Hermite (order $p = 9$) sigma-points (81 points, many of which with very small weights) are also shown.

Symmetric, 3rd order. A widely applicable sigma-point scheme is constructed by setting the unscented transform parameters to $\alpha = \pm 1$, $\beta = 0$, and $\kappa = 0$ [43]. This is also known as the 3rd order symmetric spherical-radial cubature method (CKF, [3]; see [38] for the explicit connection). This method utilizes a scaled and rotated set of $2n$ points, which are selected to be at the intersections of an n -sphere and the coordinate axes:

$$\xi_i = \begin{cases} \sqrt{n}\mathbf{e}_i, & i = 1, 2, \dots, n, \\ -\sqrt{n}\mathbf{e}_{i-n}, & i = n + 1, \dots, 2n. \end{cases}$$

The weights are defined as $w_i = 1/(2n)$ for $i = 1, 2, \dots, 2n$. The number of evaluation points is a linear function of the state dimension. The corresponding sigma-point filter is referred to as UKF 3.

Symmetric, 5th order. Building upon the work of McNamee and Stenger [26], it is possible to find explicit fully symmetric integration formulas of higher order than three. These integration schemes are exact for symmetric polynomials up to a given order p . For order $p = 5$, the number of required sigma-points is $2n^2 + 1$. The corresponding sigma-point filter is referred to as UKF 5.

Symmetric, 7th order. For order $p = 7$, the number of required sigma-points is $\frac{1}{3}(4n^3 + 8n + 3)$, meaning that they scale cubically with the number of state dimensions. The corresponding sigma-point filter is referred to as UKF 7.

Symmetric, 9th order. For order $p = 9$, the number of required sigma-points is $\frac{1}{3}(2n^4 - 4n^3 + 22n^2 - 8n + 3)$. The corresponding sigma-point filter is referred to as UKF 9. If required, even higher order methods can be constructed in the spirit of [26].

Gauss-Hermite. The n -dimensional Gauss-Hermite quadrature method forms the sigma-points as a Cartesian product of the one-dimensional Gauss-Hermite quadratures, and the weights are simply products of the one-dimensional weights [6, 16, 43]. The disadvantage of this method is that with a p th order GH approximation (exact for polynomials up to order p), the required number of evaluation points is p^n , the number growing

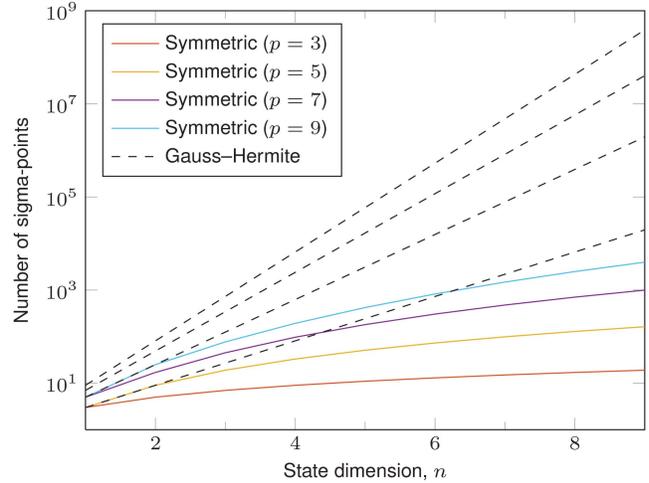


Fig. 2. Scaling of the number of sigma-points for each of the symmetric methods (solid lines). The required number of sigma-points for the Gauss-Hermite cubature of corresponding order p is visualized by the dashed lines.

exponentially with state dimension n . The corresponding filter is referred to as GHKF.

The exact formulas for the higher-order methods become lengthy and have been omitted here for brevity (see, [25, 26, 43], for implementation details and discussion). Figure 1 gives a pictorial example of how the points and weights are placed in two dimensions ($n = 2$) for each of the methods. Note that even though the 5th and 9th order methods do not have negative weights when $n = 2$, they have negative weights in other dimensions.

For higher state dimensions, Figure 2 shows how the number of required points scale in each of the schemes. The exponentially growing number of evaluation points for Gauss-Hermite is apparent in Figure 2. In the UKFs, the number of evaluation points grow polynomially. McNamee and Stenger provide the following bound for the number of evaluation points for the fully symmetric integration formulas of arbitrary degree $p = 2k + 1$ in n -space: $O((2n)^k/k!)$. Note that while in this paper we focus on the higher-order methods based on McNamee and Stenger, alternative cubature rules have also been suggested (see [7, 18]).

C. Sigma-Point Filtering and Smoothing

The following sigma-point filtering and smoothing equations are obtained by selecting a cubature rule, for example, one of the rules discussed in Section II-B and substituting it in place of the expectations in the Gaussian filtering and smoothing equations (Sec. II-A, Eqs. 5–7).

In the following equations, we denote the lower triangular matrix square-root (Cholesky) factor of a covariance matrix \mathbf{P} by \mathbf{L} so that for example $\mathbf{P}_{k|k-1} = \mathbf{L}_{k|k-1} \mathbf{L}_{k|k-1}^T$. The prediction step is

$$\begin{aligned} \mathbf{m}_{k|k-1} &= \sum_i w_i \mathbf{f}(\mathbf{m}_{k-1|k-1} + \mathbf{L}_{k-1|k-1} \boldsymbol{\xi}_i), \\ \mathbf{P}_{k|k-1} &= \sum_i \{w_i (\mathbf{f}(\mathbf{m}_{k-1|k-1} + \mathbf{L}_{k-1|k-1} \boldsymbol{\xi}_i) - \mathbf{m}_{k|k-1}) \\ &\quad \times (\mathbf{f}(\mathbf{m}_{k-1|k-1} + \mathbf{L}_{k-1|k-1} \boldsymbol{\xi}_i) - \mathbf{m}_{k|k-1})^T\} + \mathbf{Q} \end{aligned} \quad (12)$$

and the update step is

$$\begin{aligned} \boldsymbol{\mu}_k &= \sum_i w_i \mathbf{h}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_i), \\ \mathbf{S}_k &= \sum_i \{w_i (\mathbf{h}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_i) - \boldsymbol{\mu}_k) \\ &\quad \times (\mathbf{h}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_i) - \boldsymbol{\mu}_k)^T\} + \mathbf{R}, \\ \mathbf{C}_k &= \sum_i \{w_i \mathbf{L}_{k|k-1} \boldsymbol{\xi}_i \\ &\quad \times (\mathbf{h}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_i) - \boldsymbol{\mu}_k)^T\}, \\ \mathbf{K}_k &= \mathbf{C}_k \mathbf{S}_k^{-1}, \\ \mathbf{m}_{k|k} &= \mathbf{m}_{k|k-1} + \mathbf{K}_k (\mathbf{y}_k - \boldsymbol{\mu}_k), \\ \mathbf{P}_{k|k} &= \mathbf{P}_{k|k-1} - \mathbf{K}_k \mathbf{S}_k \mathbf{K}_k^T. \end{aligned} \quad (13)$$

The Rauch-Tung-Striebel smoother equations are

$$\begin{aligned} \mathbf{m}_{k+1|k} &= \sum_i \{w_i \mathbf{f}(\mathbf{m}_{k|k} + \mathbf{L}_{k|k} \boldsymbol{\xi}_i)\}, \\ \mathbf{P}_{k+1|k} &= \sum_i \{w_i (\mathbf{f}(\mathbf{m}_{k|k} + \mathbf{L}_{k|k} \boldsymbol{\xi}_i) - \mathbf{m}_{k+1|k}) \\ &\quad \times (\mathbf{f}(\mathbf{m}_{k|k} + \mathbf{L}_{k|k} \boldsymbol{\xi}_i) - \mathbf{m}_{k+1|k})^T\} + \mathbf{Q}, \\ \mathbf{D}_{k+1} &= \sum_i \{w_i \mathbf{L}_{k|k} \boldsymbol{\xi}_j (\mathbf{f}(\mathbf{m}_{k|k} + \mathbf{L}_{k|k} \boldsymbol{\xi}_i) - \mathbf{m}_{k+1|k})\}, \\ \mathbf{G}_k &= \mathbf{D}_{k+1} [\mathbf{P}_{k+1|k}]^{-1}, \\ \mathbf{m}_{k|T} &= \mathbf{m}_{k|k} + \mathbf{G}_k (\mathbf{m}_{k+1|T} - \mathbf{m}_{k+1|k}), \\ \mathbf{P}_{k|T} &= \mathbf{P}_{k|k} + \mathbf{G}_k (\mathbf{P}_{k+1|T} - \mathbf{P}_{k+1|k}) \mathbf{G}_k^T. \end{aligned} \quad (14)$$

To evaluate expectations with respect to the pairwise smoothing distributions (Eq. 8), the required $2n$ -dimensional sigma-points need to be generated separately as they are not used in the smoother pass. The

sigma-points used for the pairwise smoothing distributions are of the form

$$\begin{pmatrix} \mathbf{x}_k^{(i)} \\ \mathbf{x}_{k-1}^{(i)} \end{pmatrix} = \begin{pmatrix} \mathbf{m}_{k|T}^{(i)} \\ \mathbf{m}_{k-1|T}^{(i)} \end{pmatrix} + \sqrt{\begin{pmatrix} \mathbf{P}_{k|T} & \mathbf{P}_{k|T} \mathbf{G}_{k-1}^T \\ \mathbf{G}_{k-1} \mathbf{P}_{k|T} & \mathbf{P}_{k-1|T} \end{pmatrix}} \boldsymbol{\xi}_i^{(2n)}, \quad (15)$$

where $\boldsymbol{\xi}_i^{(2n)}$ are the $2n$ -dimensional unit sigma-points. Then, expectation a function $\mathbf{f}(\mathbf{x}_k, \mathbf{x}_{k-1})$ is approximated as

$$\mathbb{E}(\mathbf{f}(\mathbf{x}_k, \mathbf{x}_{k-1}) | \mathbf{y}_{1:T}) = \sum_i w_i^{(2n)} \mathbf{f}(\mathbf{x}_k^{(i)}, \mathbf{x}_{k-1}^{(i)}), \quad (16)$$

where $w_i^{(2n)}$ are the corresponding weights of the $2n$ -dimensional sigma-point scheme.

III. PARAMETER ESTIMATION

In this section, we consider methods for estimating the static parameters $\boldsymbol{\theta}$ of the state-space model (1). All methods discussed target the maximum likelihood solution, that is, aim to maximize $p(\mathbf{y}_{1:T} | \boldsymbol{\theta})$, or equivalently the log-likelihood:

$$\boldsymbol{\theta}_{\text{ML}} = \arg \max_{\boldsymbol{\theta}} \log p(\mathbf{y}_{1:T} | \boldsymbol{\theta}). \quad (17)$$

Since the state variables $\mathbf{x}_{0:T}$ cannot in general be marginalized in closed-form, approximative numeric methods are needed.

Here, we focus on three approaches where sigma-point filtering and smoothing is used to approximate the likelihood. First, we consider a so-called direct-likelihood approach, where the sigma-point algorithm is used to directly approximate the log-likelihood and its gradient, which are then used in numeric optimization algorithms such as conjugate-gradient optimization. Second, the expectation-maximization (EM) algorithm which is based on a lower bound for the log-likelihood and iterating optimization of parameters with respect to the lower bound and updating the lower bound with new parameters. The third approach is a modification of the direct-likelihood optimization where Fisher's identity is used to express the gradient of the log-likelihood using the same lower bound function that appears in the EM algorithm. Note that the third approach is otherwise similar to the first, but since it is based on the EM lower bound, we present the methods in this order. Each of these three approaches may be used in combination with any of the sigma-point rules discussed in Section II-B.

Note that all the algorithms presented in this section are easily extended to maximum a posteriori estimation since maximizing the posterior density is equivalent to maximizing the (unnormalized) log-posterior. That is, the sum of log-likelihood and log-prior:

$$\boldsymbol{\theta}_{\text{MAP}} = \arg \max_{\boldsymbol{\theta}} [\log p(\mathbf{y}_{1:T} | \boldsymbol{\theta}) + \log p(\boldsymbol{\theta})]. \quad (18)$$

Since the log-prior is known, approximations of the unnormalized log-posterior as well as its gradient and

lower bounds are immediately obtained from the corresponding approximations for the log-likelihood.

A. Direct Likelihood Based Parameter Estimation

The marginal log-likelihood function can be formulated as the following sum

$$\mathcal{L}_T(\boldsymbol{\theta}) = \sum_{k=1}^T \log p(\mathbf{y}_k | \mathbf{y}_{1:k-1}, \boldsymbol{\theta}). \quad (19)$$

Furthermore, the terms of the sum on the right hand side may in principle be evaluated by

$$p(\mathbf{y}_k | \mathbf{y}_{1:k-1}, \boldsymbol{\theta}) = \int p(\mathbf{y}_k | \mathbf{x}_k, \boldsymbol{\theta}) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}, \boldsymbol{\theta}) d\mathbf{x}_k, \quad (20)$$

that is, integrating the measurement model $p(\mathbf{y}_k | \mathbf{x}_k, \boldsymbol{\theta})$ over the predicted state distribution $p(\mathbf{x}_k | \mathbf{y}_{1:k-1}, \boldsymbol{\theta})$ which is computed during the Bayesian filtering recursion. In assumed density Gaussian filtering is used, we get the approximation

$$p(\mathbf{y}_k | \mathbf{y}_{1:k-1}, \boldsymbol{\theta}) \approx \mathbb{N}(\mathbf{y}_k | \boldsymbol{\mu}_k, \mathbf{S}_k), \quad (21)$$

whence the marginal log-likelihood expression in Equation (19) evaluates to

$$\begin{aligned} \mathcal{L}_T(\boldsymbol{\theta}) &= \log p(\mathbf{y}_{1:T} | \boldsymbol{\theta}) \approx -\frac{1}{2} \sum_{k=1}^T \log |2\pi \mathbf{S}_k| \\ &\quad - \frac{1}{2} \sum_{k=1}^T (\mathbf{y}_k - \boldsymbol{\mu}_k)^T \mathbf{S}_k^{-1} (\mathbf{y}_k - \boldsymbol{\mu}_k), \end{aligned} \quad (22)$$

where the quantities $\boldsymbol{\mu}_k$ and \mathbf{S}_k are evaluated during the filtering recursion, in the case of sigma-point methods by Equation (13).

To enable use of gradient-based optimization algorithms, we also need a method for evaluating the gradients of the marginal log-likelihood. This is based on the so-called sensitivity equations [14, 32] that are obtained by differentiating the filtering equations. Namely, the gradient of the log-likelihood is obtained by the recursion

$$\begin{aligned} \frac{\partial \mathcal{L}_k(\boldsymbol{\theta})}{\partial \theta_i} &= \frac{\partial \mathcal{L}_{k-1}(\boldsymbol{\theta})}{\partial \theta_i} - \frac{1}{2} \text{tr} \left(\mathbf{S}_k^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{S}_k(\boldsymbol{\theta})}{\partial \theta_i} \right) \\ &\quad - \mathbf{v}_k^T(\boldsymbol{\theta}) \mathbf{S}_k^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{v}_k(\boldsymbol{\theta})}{\partial \theta_i} \\ &\quad + \frac{1}{2} \mathbf{v}_k^T(\boldsymbol{\theta}) \mathbf{S}_k^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{S}_k(\boldsymbol{\theta})}{\partial \theta_i} \mathbf{S}_k^{-1}(\boldsymbol{\theta}) \mathbf{v}_k(\boldsymbol{\theta}), \end{aligned} \quad (23)$$

where $\mathbf{v}_k = \mathbf{y}_k - \boldsymbol{\mu}_k$. The derivatives $\partial \mathbf{S}_k(\boldsymbol{\theta}) / \partial \theta_i$ and $\partial \mathbf{v}_k(\boldsymbol{\theta}) / \partial \theta_i$ are computed along the filtering pass by the equations shown in Figure 3.

B. Expectation-Maximization Based Parameter Estimation

Expectation-maximization (EM), proposed by Dempster *et al.* [9] is an iterative algorithm for find-

ing maximum likelihood parameter estimates in settings with some unobserved variables, such as the state variables \mathbf{x} in the state-space context. The motivation is that the so-called full-data likelihood of the observed and unobserved variables is easier to compute, and a lower bound for the marginal likelihood of the observed variables may be obtained based on expected full-data log-likelihood. In the following, we present the EM algorithm following the formulation by Neal and Hinton [27] and the notation of Schön *et al.* [34].

The EM algorithm is based on the following lower bound of the log-likelihood:

$$\log p(\mathbf{y}_{1:T} | \boldsymbol{\theta}) \geq \int q(\mathbf{x}_{0:T}) \log \frac{p(\mathbf{x}_{0:T}, \mathbf{y}_{1:T} | \boldsymbol{\theta})}{q(\mathbf{x}_{0:T})} d\mathbf{x}_{0:T}, \quad (24)$$

where q is an arbitrary probability density over the states $\mathbf{x}_{0:T}$. The idea is to iteratively maximize this lower bound with respect to q (holding $\boldsymbol{\theta}$ fixed) and with respect to $\boldsymbol{\theta}$ (holding q fixed). Furthermore, when $\boldsymbol{\theta} = \boldsymbol{\theta}^{(n)}$ is fixed, the maximum with respect to q is obtained by

$$q(\mathbf{x}_{0:T}) := p(\mathbf{x}_{0:T} | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)}). \quad (25)$$

By substituting this into Equation (24), the bound becomes

$$\begin{aligned} &\int p(\mathbf{x}_{0:T} | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)}) \log \frac{p(\mathbf{x}_{0:T}, \mathbf{y}_{1:T} | \boldsymbol{\theta})}{p(\mathbf{x}_{0:T} | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)})} d\mathbf{x}_{0:T} \\ &= \int p(\mathbf{x}_{0:T} | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)}) \log p(\mathbf{x}_{0:T}, \mathbf{y}_{1:T} | \boldsymbol{\theta}) d\mathbf{x}_{0:T} \\ &\quad - \int p(\mathbf{x}_{0:T} | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)}) \log p(\mathbf{x}_{0:T} | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)}) d\mathbf{x}_{0:T}. \end{aligned}$$

The latter term is independent of $\boldsymbol{\theta}$ and may thus be omitted when maximizing the lower bound with respect to $\boldsymbol{\theta}$. The first term is the conditional expectation of $\log p(\mathbf{y}_{1:T}, \mathbf{x}_{0:T} | \boldsymbol{\theta})$ conditional on $\boldsymbol{\theta}^{(n)}$ and $\mathbf{y}_{1:T}$. Thus, the step of maximizing the lower bound (Eq. 24) may be replaced by computing the following function:

$$\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) = \mathbb{E}[\log p(\mathbf{x}_{0:T}, \mathbf{y}_{1:T} | \boldsymbol{\theta}) | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)}]. \quad (26)$$

The EM algorithm in its general form thus consists of initializing the parameters to $\boldsymbol{\theta}^{(0)}$ and for $n = 0, 1, \dots$ iterating the following two steps:

- E-step: compute $\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)})$.
- M-step: $\boldsymbol{\theta}^{(n+1)} \leftarrow \arg \max_{\boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)})$.

In state-space models, the \mathcal{Q} -function can be decomposed by employing the Markov property of the state sequence and the conditional independence of the measurements:

$$\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) = \mathcal{I}_1(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) + \mathcal{I}_2(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) + \mathcal{I}_3(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}), \quad (27)$$

$$\begin{aligned}
\frac{\partial \mathbf{m}_{k|k-1}}{\partial \theta_i} &= \sum_j \left\{ w_j \left[\mathbf{F}_x(\mathbf{m}_{k-1|k-1} + \mathbf{L}_{k-1|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) \right. \right. \\
&\quad \times \left(\frac{\partial \mathbf{m}_{k-1|k-1}}{\partial \theta_i} + \frac{\partial \mathbf{L}_{k-1|k-1}}{\partial \theta_i} \boldsymbol{\xi}_j \right) \\
&\quad \left. \left. + \frac{\partial \mathbf{f}}{\partial \theta_i}(\mathbf{m}_{k-1|k-1} + \mathbf{L}_{k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) \right] \right\}, \\
\frac{\partial \mathbf{P}_{k|k-1}}{\partial \theta_i} &= \sum_j \left\{ w_j \left[\mathbf{F}_x(\mathbf{m}_{k-1} + \mathbf{L}_{k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) \right. \right. \\
&\quad \times \left(\frac{\partial \mathbf{m}_{k-1|k-1}}{\partial \theta_i} + \frac{\partial \mathbf{L}_{k-1|k-1}}{\partial \theta_i} \boldsymbol{\xi}_j \right) \\
&\quad + \frac{\partial \mathbf{f}}{\partial \theta_i}(\mathbf{m}_{k-1|k-1} + \mathbf{L}_{k-1|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) - \frac{\partial \mathbf{m}_{k|k-1}}{\partial \theta_i} \Big] \\
&\quad \times [\mathbf{f}(\mathbf{m}_{k-1|k-1} + \mathbf{L}_{k-1|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) - \mathbf{m}_{k|k-1}]^\top \\
&\quad + [\mathbf{f}(\mathbf{m}_{k-1|k-1} + \mathbf{L}_{k-1|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) - \mathbf{m}_{k|k-1}] \\
&\quad \times [\mathbf{F}_x(\mathbf{m}_{k-1} + \mathbf{L}_{k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) \\
&\quad \times \left(\frac{\partial \mathbf{m}_{k-1|k-1}}{\partial \theta_i} + \frac{\partial \mathbf{L}_{k-1|k-1}}{\partial \theta_i} \boldsymbol{\xi}_j \right) \\
&\quad \left. \left. + \frac{\partial \mathbf{f}}{\partial \theta_i}(\mathbf{m}_{k-1|k-1} + \mathbf{L}_{k-1|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) - \frac{\partial \mathbf{m}_{k|k-1}}{\partial \theta_i} \right]^\top + \frac{\partial \mathbf{Q}}{\partial \theta_i} \right\}, \\
\frac{\partial \boldsymbol{\mu}_k}{\partial \theta_i} &= \sum_j \left\{ w_j \left[\mathbf{H}_x(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) \right. \right. \\
&\quad \left. \left. + \frac{\partial \mathbf{h}}{\partial \theta_i}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) \right] \right\}, \\
\frac{\partial \mathbf{v}_k}{\partial \theta_i} &= -\frac{\partial \boldsymbol{\mu}_k}{\partial \theta_i}, \\
\frac{\partial \mathbf{S}_k}{\partial \theta_i} &= \sum_j \left\{ w_j \left[\mathbf{H}_x(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) \right. \right. \\
&\quad \times \left(\frac{\partial \mathbf{m}_{k|k-1}}{\partial \theta_i} + \frac{\partial \mathbf{L}_{k|k-1}}{\partial \theta_i} \boldsymbol{\xi}_j \right) \\
&\quad + \frac{\partial \mathbf{h}}{\partial \theta_i}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) - \frac{\partial \boldsymbol{\mu}_k}{\partial \theta_i} \Big] \\
&\quad \times [\mathbf{h}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) - \boldsymbol{\mu}_k]^\top \\
&\quad + [\mathbf{h}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) - \boldsymbol{\mu}_k] \\
&\quad \times [\mathbf{H}_x(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) \\
&\quad \left. \left. + \frac{\partial \mathbf{h}}{\partial \theta_i}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) - \frac{\partial \boldsymbol{\mu}_k}{\partial \theta_i} \right]^\top \right\} + \frac{\partial \mathbf{R}}{\partial \theta_i}, \\
\frac{\partial \mathbf{C}_k}{\partial \theta_i} &= \sum_j \left\{ w_j \left[\frac{\partial \mathbf{L}_{k|k-1}}{\partial \theta_i} \boldsymbol{\xi}_j \left(\mathbf{h}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) - \boldsymbol{\mu}_k \right)^\top \right. \right. \\
&\quad + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j \left[\mathbf{H}_x(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) \right. \\
&\quad \times \left(\frac{\partial \mathbf{m}_{k|k-1}}{\partial \theta_i} + \frac{\partial \mathbf{L}_{k|k-1}}{\partial \theta_i} \boldsymbol{\xi}_j \right) \\
&\quad \left. \left. + \frac{\partial \mathbf{h}}{\partial \theta_i}(\mathbf{m}_{k|k-1} + \mathbf{L}_{k|k-1} \boldsymbol{\xi}_j, \boldsymbol{\theta}) - \frac{\partial \boldsymbol{\mu}_k}{\partial \theta_i} \right]^\top \right] \Big\}, \\
\frac{\partial \mathbf{K}_k}{\partial \theta_i} &= \frac{\partial \mathbf{C}_k}{\partial \theta_i} \mathbf{S}_k^{-1} - \mathbf{C}_k \mathbf{S}_k^{-1} \frac{\partial \mathbf{S}_k}{\partial \theta_i} \mathbf{S}_k^{-1}, \\
\frac{\partial \mathbf{m}_{k|k}}{\partial \theta_i} &= \frac{\partial \mathbf{m}_{k|k-1}}{\partial \theta_i} + \frac{\partial \mathbf{K}_k}{\partial \theta_i} \mathbf{v}_k + \mathbf{K}_k \frac{\partial \mathbf{v}_k}{\partial \theta_i}, \\
\frac{\partial \mathbf{P}_{k|k}}{\partial \theta_i} &= \frac{\partial \mathbf{P}_{k|k-1}}{\partial \theta_i} - \frac{\partial \mathbf{K}_k}{\partial \theta_i} \mathbf{S}_k \mathbf{K}_k^\top - \mathbf{K}_k \frac{\partial \mathbf{S}_k}{\partial \theta_i} \mathbf{K}_k^\top - \mathbf{K}_k \mathbf{S}_k \frac{\partial \mathbf{K}_k^\top}{\partial \theta_i}.
\end{aligned}$$

Fig. 3. Recursion for computing the derivatives of the prediction and update steps. \mathbf{F}_x is the Jacobian of $\mathbf{f}(\mathbf{x}, \boldsymbol{\theta})$ as a function of \mathbf{x} and \mathbf{H}_x is the Jacobian of $\mathbf{h}(\mathbf{x}, \boldsymbol{\theta})$ as a function of \mathbf{x} . Algorithms for computing the derivatives of the Cholesky factors \mathbf{L} such that $\mathbf{P} = \mathbf{L}\mathbf{L}^\top$ are omitted here (see [32]).

where the terms are

$$\mathbf{I}_1(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) = \mathbb{E}[\log p(\mathbf{x}_0 | \boldsymbol{\theta}) | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)}], \quad (28)$$

$$\mathbf{I}_2(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) = \sum_{k=1}^T \mathbb{E}[\log p(\mathbf{x}_k | \mathbf{x}_{k-1}, \boldsymbol{\theta}) | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)}], \quad (29)$$

$$\mathbf{I}_3(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) = \sum_{k=1}^T \mathbb{E}[\log p(\mathbf{y}_k | \mathbf{x}_k, \boldsymbol{\theta}) | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)}]. \quad (30)$$

To evaluate this expression, one needs the smoothing distributions $p(\mathbf{x}_t | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)})$ and the joint smoothing distributions of consecutive states $p(\mathbf{x}_k, \mathbf{x}_{k+1} | \mathbf{y}_{1:T}, \boldsymbol{\theta}^{(n)})$. Sigma-point approximations to the EM algorithm are then obtained by replacing the expectations over the smoothing distributions by their sigma-point smoother approximations. The Gaussian smoother approximation for \mathcal{Q} is

$$\begin{aligned}
\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) &\approx \\
&-\frac{1}{2} \log |2\pi \mathbf{P}_0| - \frac{T}{2} \log |2\pi \mathbf{Q}| - \frac{T}{2} \log |2\pi \mathbf{R}| \\
&-\frac{1}{2} \text{tr} \{ \mathbf{P}_0^{-1} [\mathbf{P}_{0|T} + (\mathbf{m}_{0|T} - \mathbf{m}_0)(\mathbf{m}_{0|T} - \mathbf{m}_0)^\top] \} \\
&-\frac{1}{2} \sum_{k=1}^T \text{tr} \{ \mathbf{Q}^{-1} \mathbb{E}[(\mathbf{x}_k - \mathbf{f}(\mathbf{x}_{k-1}))(\mathbf{x}_k - \mathbf{f}(\mathbf{x}_{k-1}))^\top | \mathbf{y}_{1:T}] \} \\
&-\frac{1}{2} \sum_{k=1}^T \text{tr} \{ \mathbf{R}^{-1} \mathbb{E}[(\mathbf{y}_k - \mathbf{h}(\mathbf{x}_k))(\mathbf{y}_k - \mathbf{h}(\mathbf{x}_k))^\top | \mathbf{y}_{1:T}] \}, \quad (31)
\end{aligned}$$

where $\mathbf{P}_0, \mathbf{Q}, \mathbf{R}$ and the model functions $\mathbf{f}(\cdot), \mathbf{h}(\cdot)$ depend on the parameters $\boldsymbol{\theta}$. The smoothing distribution means and covariances $\mathbf{m}_{k|T}, \mathbf{P}_{k|T}$ are obtained during the smoothing backward pass. The expectations over the smoothing distribution in the latter two terms are evaluated by using the sigma-point approximations for Gaussian integrals as follows. The second expectation depends only on the smoothing distribution $\mathbb{N}(\mathbf{x}_k | \mathbf{m}_{k|T}, \mathbf{P}_{k|T})$ and is computed as follows:

$$\begin{aligned}
&\mathbb{E}[(\mathbf{y}_k - \mathbf{h}(\mathbf{x}_k))(\mathbf{y}_k - \mathbf{h}(\mathbf{x}_k))^\top | \mathbf{y}_{1:T}] \\
&\approx \sum_i w_i (\mathbf{y}_k - \mathbf{h}(\mathbf{m}_{k|T}) + \mathbf{L}_{k|T} \boldsymbol{\xi}_i) \\
&\quad \times (\mathbf{y}_k - \mathbf{h}(\mathbf{m}_{k|T}) + \mathbf{L}_{k|T} \boldsymbol{\xi}_i)^\top. \quad (32)
\end{aligned}$$

The first expectation depends on the pairwise joint smoothing distribution $p(\mathbf{x}_k, \mathbf{x}_{k-1} | \mathbf{y}_{1:T})$ (cf. Sec. II-A, Eq. 8). Thus, to evaluate it we need to use $2n$ -dimensional sigma-points as discussed in Section II-C, Equation 15.

In general, maximizing \mathcal{Q} in the M-step requires the use numerical optimization, for example, using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [11]. However, using numerical optimization inside EM is quite cumbersome, because with the same effort we

could numerically optimize the approximate likelihood directly. Hence the benefit of EM is in the situation when the optimization can be performed in closed form. This kind of special case is the class of models where the parameters appear linearly although the model itself might be nonlinear.

In the following, we present closed-form solutions for the special case where the model functions are linear combinations of the parameters where the parameters appear as coefficients of the linear combinations and/or the covariances. That is, we consider models that can be represented as follows:

$$\mathbf{x}_k = \tilde{\mathbf{A}}\tilde{\mathbf{f}}(\mathbf{x}_{k-1}) + \mathbf{q}_k, \quad (33)$$

$$\mathbf{y}_k = \tilde{\mathbf{H}}\tilde{\mathbf{h}}(\mathbf{x}_k) + \mathbf{r}_k, \quad (34)$$

where $\tilde{\mathbf{f}}(\cdot)$ and $\tilde{\mathbf{h}}(\cdot)$ are functions containing the nonlinearities and the parameters are a subset of $\{\mathbf{A}, \mathbf{H}, \mathbf{Q}, \mathbf{R}, \mathbf{m}_0, \mathbf{P}_0\}$.

For these models, the expression for \mathcal{Q} can be written as

$$\begin{aligned} \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) = & \\ & -\frac{1}{2} \log |2\pi \mathbf{P}_0| - \frac{T}{2} \log |2\pi \mathbf{Q}| - \frac{T}{2} \log |2\pi \mathbf{R}| \\ & -\frac{1}{2} \text{tr}\{\mathbf{P}_0^{-1}[\mathbf{P}_{0|T} + (\mathbf{m}_{0|T} - \mathbf{m}_0)(\mathbf{m}_{0|T} - \mathbf{m}_0)^T]\} \\ & -\frac{T}{2} \text{tr}\{\mathbf{Q}^{-1}[\boldsymbol{\Sigma} - \mathbf{C}\mathbf{A}^T - \mathbf{A}\mathbf{C}^T + \mathbf{A}\boldsymbol{\Phi}\mathbf{A}^T]\} \\ & -\frac{T}{2} \text{tr}\{\mathbf{R}^{-1}[\mathbf{D} - \mathbf{B}\mathbf{H}^T - \mathbf{H}\mathbf{B}^T + \mathbf{H}\boldsymbol{\Theta}\mathbf{H}^T]\}, \end{aligned}$$

where the model parameters to be optimized are some subset of $\{\mathbf{A}, \mathbf{H}, \mathbf{Q}, \mathbf{R}, \mathbf{m}_0, \mathbf{P}_0\}$ and $\boldsymbol{\Sigma}, \boldsymbol{\Phi}, \boldsymbol{\Theta}, \mathbf{B}, \mathbf{C}, \mathbf{D}$ can be evaluated based on the latest E-step sigma-point smoother results as follows:

$$\boldsymbol{\Sigma} = \frac{1}{T} \sum_{k=1}^T \mathbf{P}_{k|T} + \mathbf{m}_{k|T}[\mathbf{m}_{k|T}]^T, \quad (35)$$

$$\boldsymbol{\Phi} = \frac{1}{T} \sum_{k=1}^T \mathbb{E}[\tilde{\mathbf{f}}(\mathbf{x}_{k-1})\tilde{\mathbf{f}}^T(\mathbf{x}_{k-1}) | \mathbf{y}_{1:T}], \quad (36)$$

$$\boldsymbol{\Theta} = \frac{1}{T} \sum_{k=1}^T \mathbb{E}[\tilde{\mathbf{h}}(\mathbf{x}_k)\tilde{\mathbf{h}}^T(\mathbf{x}_k) | \mathbf{y}_{1:T}], \quad (37)$$

$$\mathbf{B} = \frac{1}{T} \sum_{k=1}^T \mathbf{y}_k \mathbb{E}[\tilde{\mathbf{h}}^T(\mathbf{x}_k) | \mathbf{y}_{1:T}], \quad (38)$$

$$\mathbf{C} = \frac{1}{T} \sum_{k=1}^T \mathbb{E}[\mathbf{x}_k \tilde{\mathbf{f}}^T(\mathbf{x}_{k-1}) | \mathbf{y}_{1:T}], \quad (39)$$

$$\mathbf{D} = \frac{1}{T} \sum_{k=1}^T \mathbf{y}_k \mathbf{y}_k^T. \quad (40)$$

Using these values, the optimal parameters in the M-step, that is, the maximum points of the $\mathcal{Q}(\cdot, \boldsymbol{\theta}^{(n)})$ -function are

- When $\boldsymbol{\theta} = \mathbf{A}$, we get

$$\mathbf{A}^* = \mathbf{C}\boldsymbol{\Phi}^{-1}.$$

- When $\boldsymbol{\theta} = \mathbf{H}$, we get

$$\mathbf{H}^* = \mathbf{B}\boldsymbol{\Theta}^{-1}.$$

- When $\boldsymbol{\theta} = \mathbf{Q}$, we get

$$\mathbf{Q}^* = \boldsymbol{\Sigma} - \mathbf{C}\mathbf{A}^T - \mathbf{A}\mathbf{C}^T + \mathbf{A}\boldsymbol{\Phi}\mathbf{A}^T.$$

- When $\boldsymbol{\theta} = \mathbf{R}$, we get

$$\mathbf{R}^* = \mathbf{D} - \mathbf{H}\mathbf{B}^T - \mathbf{B}\mathbf{H}^T + \mathbf{H}\boldsymbol{\Theta}\mathbf{H}^T.$$

- When $\boldsymbol{\theta} = \mathbf{m}_0$, we get

$$\mathbf{m}_0^* = \mathbf{m}_{0|T}. \quad (41)$$

- Finally, the maximum with respect to the initial covariance $\boldsymbol{\theta} = \mathbf{P}_0$ is

$$\mathbf{P}_0^* = \mathbf{P}_{0|T} + (\mathbf{m}_{0|T} - \mathbf{m}_0)(\mathbf{m}_{0|T} - \mathbf{m}_0)^T.$$

C. Evaluating the Gradient Based on Fisher's Identity

The expected log-likelihood that appears in the EM algorithm may also be used as a basis of an alternative approach for evaluating the gradient in direct optimization. Based on Fisher's identity, the gradient of the marginal log-likelihood may be expressed as

$$\frac{\partial \mathcal{L}_T(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \frac{\partial \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)})}{\partial \boldsymbol{\theta}} \Big|_{\boldsymbol{\theta}^{(n)} = \boldsymbol{\theta}}, \quad (42)$$

where \mathcal{Q} is the function defined in the EM algorithm (Eq. 27). When the \mathcal{Q} -function is approximated with sigma-point smoothers, we obtain an alternative approximation of the gradient of the marginal log-likelihood that may be used in place of the approximation derived in Section III-A. For linear state-space models, this approach was suggested by Segal and Weinstein [35] and later by Olsson *et al.* [28] who called the approach the 'easy gradient recipe.' See [5, 32] for discussions of the nonlinear case.

IV. EXPERIMENTS

In this section, we demonstrate the different sigma-point schemes and different parameter estimation algorithms with two example models. First, we use a one-dimensional model (the univariate nonstationary growth model, UNGM, [1, 23]) to illustrate the approximate likelihood curves obtained by different methods. Second, we compare the performance of different algorithms with simulated data in a problem of tracking a maneuvering target with bearings-only measurements. In this example, we focus on estimating the sensor variances and compare the variance estimates as well as the actual tracking error.

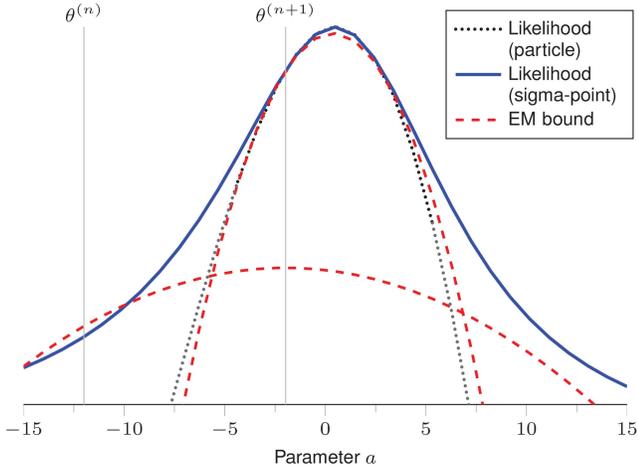


Fig. 4. Visualization of the one-step evolution of the EM algorithm for the univariate estimation of parameter a . The dotted line represents the particle filter log-likelihood estimate (extrapolated in the tails), while the solid line is the sigma-point filter log-likelihood approximation. The dashed lines correspond to the sigma-point EM bounds for iterations n and $n + 1$.

Codes for replicating the experiments are available on the author homepages.

A. Simple Nonlinear Growth Model

We simulated a realization with $T = 100$ data from the following model:

$$x_{k+1} = ax_k + b \frac{x_k}{1 + x_k^2} + c \cos(1.2k) + q_k, \quad (43)$$

$$y_k = dx_k + r_k. \quad (44)$$

with $a = 0.5$, $b = 25$, $c = 8$, $d = \sqrt{0.05}$, $q_k \sim \mathcal{N}(0, 10)$, $r_k \sim \mathcal{N}(0, 0.01)$, $x_0 \sim \mathcal{N}(0, 0.01)$. This is the univariate nonstationary growth model [1, 23] except that we changed the measurement model to linear as the model with the typically used quadratic measurement model is known to be challenging for sigma-point algorithms [42].

First, we estimated the likelihood of parameter a , holding other parameters fixed at their ground-truth values. Likelihood curves obtained by direct likelihood estimation with various sigma-point rules as well as the EM lower bounds for two iterations are shown in Figure 4. For comparison, a likelihood estimate obtained by particle filtering (1000 particles and the optimal importance distribution) is also shown. The EM iterations seem to converge toward the maximum of the likelihood curve and the second EM bound is rather close to the particle filter likelihood estimate. The EM lower bounds are mostly below the sigma-point likelihood curve as expected, except that the first EM lower bound slightly exceeds the sigma-point likelihood approximation in the vicinity of the initial parameter. However, both the evaluation of \mathcal{Q} and the sigma-point estimate of the likelihood are approximations.

Second, to compare the different sigma-point rules, we considered estimation of the parameter b with other

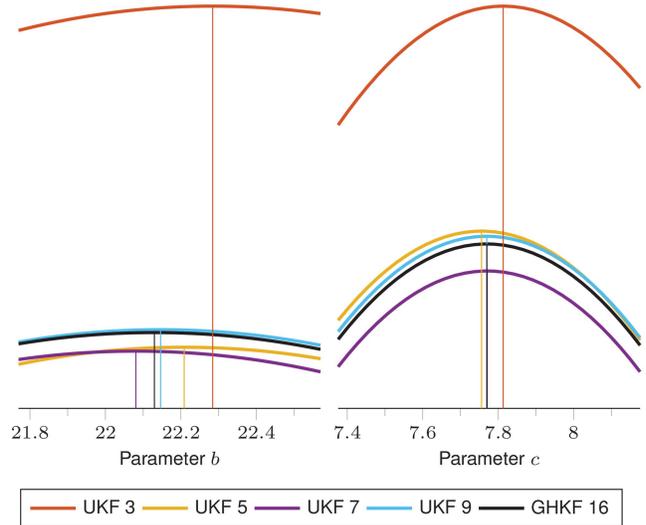


Fig. 5. Log-likelihood curves for parameters b and c evaluated by five different sigma-point methods. The vertical line indicates the location of the maximum.

parameters fixed and parameter c with other parameters fixed, using a grid of parameter values with close proximity to the maximum likelihood values. Namely, for b we used 32 evenly spaced points between 21.7698 and 22.5698 and for c we used 32 points between 7.376 and 8.176. These are shown in Figure 5. The estimate obtained by the Gauss-Hermite rule and the estimates obtained by the higher-order UKFs are rather close to each other while the estimate by the 3rd order UKF is farther in both parameters.

B. Coordinated-Turn Model

In this section, we compare the performance of the parameter estimation methods discussed in this article using a more practical example. The problem is tracking a target maneuvering according to the coordinated turn model [2, 4, 33, 38] with bearings-only sensor measurements. The state is 5-dimensional:

$$\mathbf{x} = (x_1 \quad x_2 \quad \dot{x}_1 \quad \dot{x}_2 \quad \omega)^T, \quad (45)$$

where (x_1, x_2) is the location of the target in 2-dimensional Cartesian coordinates, (\dot{x}_1, \dot{x}_2) is the corresponding speed, and ω is the turn rate. The dynamic model is

$$\mathbf{x}_{k+1} = \begin{pmatrix} 1 & 0 & \frac{\sin(\omega_k \Delta t)}{\omega_k} & \frac{\cos(\omega_k \Delta t) - 1}{\omega_k} & 0 \\ 0 & 1 & -\frac{\cos(\omega_k \Delta t) - 1}{\omega_k} & \frac{\sin(\omega_k \Delta t)}{\omega_k} & 0 \\ 0 & 0 & \cos(\omega_k \Delta t) & -\sin(\omega_k \Delta t) & 0 \\ 0 & 0 & \sin(\omega_k \Delta t) & \cos(\omega_k \Delta t) & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \mathbf{x}_k + \mathbf{q}_k. \quad (46)$$

The process noise is $\mathbf{q}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$, where

$$\mathbf{Q} = \begin{pmatrix} q_c \Delta t^3 / 3 & 0 & q_c \Delta t^2 / 2 & 0 & 0 \\ 0 & q_c \Delta t^3 / 3 & 0 & q_c \Delta t^2 / 2 & 0 \\ q_c \Delta t^2 / 2 & 0 & q_c \Delta t & 0 & 0 \\ 0 & q_c \Delta t^2 / 2 & 0 & q_c \Delta t & 0 \\ 0 & 0 & 0 & 0 & q_w \Delta t \end{pmatrix}. \quad (47)$$

The measurements are angles from the sensors with additive Gaussian noise:

$$\mathbf{y}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{r}_k, \quad (48)$$

where the measurement noise is $\mathbf{r}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$. The covariance matrix \mathbf{R} is naturally assumed diagonal, as the measurement errors of separate sensors should be independent. For each sensor i at location \mathbf{s}_i the measurement is given by

$$h_i(\mathbf{x}_k) = \text{atan2}(x_{2,k} - s_{2,i}, x_{1,k} - s_{1,i}), \quad (49)$$

where atan2 is the four-quadrant inverse tangent. We focus on estimating the measurement noise variances while keeping other parameters fixed. That is, the sensor locations and dynamic model covariance are assumed to be known and the initial state distribution fixed.

The parameters of the process noise covariance were set to $q_c = 0.1$, $q_w = 0.1$ and the time step to $\Delta t = 0.01$. The ground-truth measurement noise covariance was $\mathbf{R} = \text{diag}(0.05^2, 0.1^2)$. The two sensors were located at $\mathbf{s}_1 = (-1, 0.5)$ and $\mathbf{s}_2 = (1, 1)$. The parameters of the initial distribution were $\mathbf{m}_0 = (2, 0, 0, 0, 0)^T$ and $\mathbf{P}_0 = \text{diag}(0.5^2, 0.5^2, 0.5^2, 0.5^2, 1^2)$. We simulated 100 different trajectories with $T = 50$ timesteps from this model.

To compare performance of the different sigma-point schemes, we performed direct maximum likelihood estimation of the sensor noise standard deviation of the first sensor, keeping the noise of the second sensor as well as other parameters fixed at their ground truth values. The sigma-point schemes used were UKF 3, UKF 5, and UKF 7, as well as GHKF 3, GHKF 5, and GHKF 7. The 9th order schemes were omitted since the number of sigma-points is already quite high as the state is 5-dimensional. In addition to the sigma-point methods, we also compared against maximum likelihood estimation based on the extended Kalman filter (EKF, see, e.g. [17]).

The optimization was performed with gradient-based optimization using the Matlab optimization toolbox.¹ Furthermore, we investigated how the estimation performance varies as a function of uncertainty of the target's initial location. This was done by using an additional parameter for the per-coordinate standard deviation ($\sigma \in (0, 0.5)$). The first two diagonal components

¹MATLAB version R2014b, the `fminunc` function, quasi-Newton algorithm, initialized with $\sqrt{\mathbf{R}_{1,1}} = 0.1$.

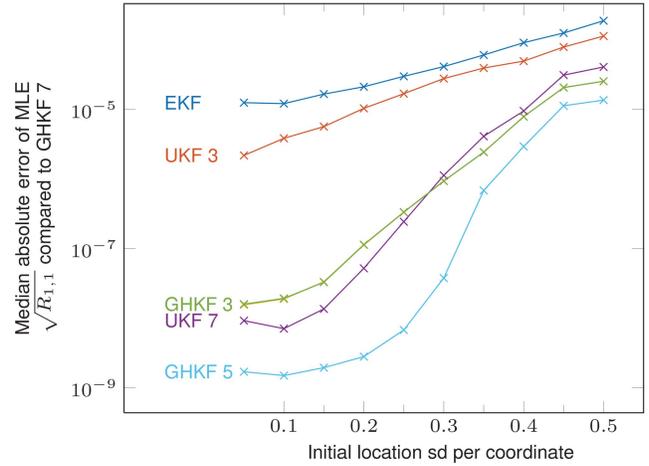


Fig. 6. Median absolute error of the parameter estimates compared to GHKF 7 (median taken over the 100 simulated trajectories) as a function of the initial location prior standard deviation. UKF 5 is essentially indistinguishable from GHKF 3.

of \mathbf{P}_0 were set to σ^2 . Furthermore, the first two components of \mathbf{m}_0 were interpolated between the original \mathbf{m}_0 and the simulated \mathbf{x}_0 to keep the uncertainty of the initial location consistent with the prior.

Since GHKF 7 is the highest-order sigma-point scheme amongst those used in this experiment, we assume it is the most accurate and compare against it. Figure 6 shows comparison of median (over the 100 trajectories) absolute deviation of the MLE estimates obtained by the various filtering schemes compared to the ones obtained by GHKF 7. GHKF 5 is closest, while EKF and UKF 3 are farthest from the baseline. UKF 7, GHKF 3 and UKF 5 have similar performance. UKF 5 and GHKF 3 are essentially identical. This is explained by the observation that in 5 dimensions, all UKF 5 sigma-points are present in the GHKF 3 sigma-point set and the sum of the GHKF 3 weights of these points is 0.79. The contribution of the remaining sigma-points that have total 0.21 weight apparently has a negligible contribution at least with this model. In addition, we also look at track estimation errors using the final parameter estimates by each sigma-point scheme. Figure 7 shows the mean RMSE over the 100 trajectories, that is, for each simulated trajectory, we computed the smoother RMSE and then took the average.

To compare the two different gradient evaluation approaches, sensitivity equations (Section III-A) and the Fisher identity approach (Section III-C), we evaluated the derivative of the log-likelihood with respect to the standard deviation of the error of the first sensor, using UKF 3 and UKF 5 and both gradient evaluation approaches. The results are shown in Figure 8.

To measure the performance as a function of computational cost, we recorded the parameter values as well as the times used at each iteration of the optimization routines. In this experiment, the initial location standard deviation per coordinate was set to 0.5. Median

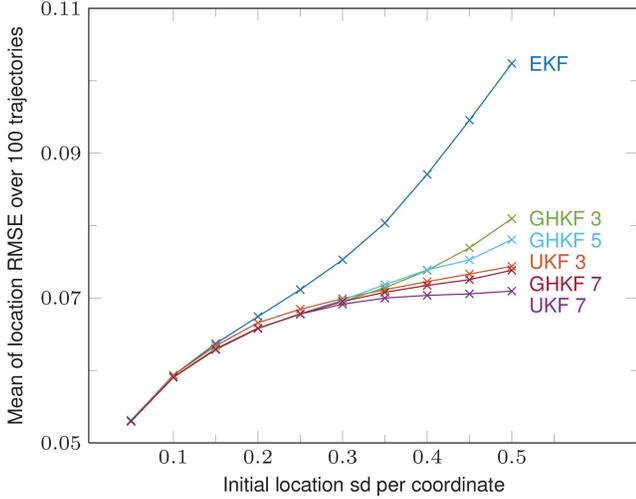


Fig. 7. Mean RMSE of smoothed location (mean taken over the 100 trajectories) using the MLE estimated noise variance of the first sensor. UKF 5 is essentially indistinguishable from GHKF 3.

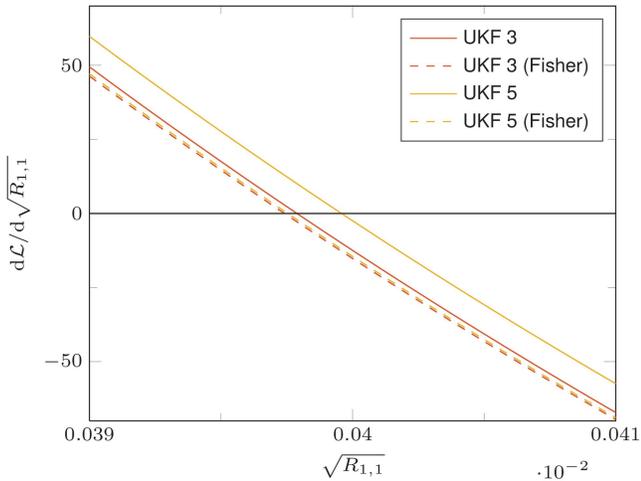


Fig. 8. Derivative of the log-likelihood with respect to the standard deviation of the error of the first sensor. Evaluated using UKF 3 and UKF 5 both with the sensitivity equation approach (Section III-A) and the Fisher identity based approach (Section III-C).

absolute error (compared to final GHKF 7 estimate, as a function of time) is shown in Figures 9 and 10. As one would expect, the higher-order schemes are more computationally demanding and GHKF is more computationally demanding than UKF, but eventually the higher-order GHKF schemes find better parameter estimates.

Figure 11 shows the evolution of EM parameter estimation for one simulated trajectory with $\sigma = 0.5$. The EM algorithm practically converges in a couple of steps with all three sigma-point schemes, and the final parameter estimates are rather close to each other and to the direct MLE estimates. Theoretically, the EM algorithm has linear convergence [9] although it is hard to say whether these convergence results extend to the case where the E-step is approximated using sigma-point smoothers.

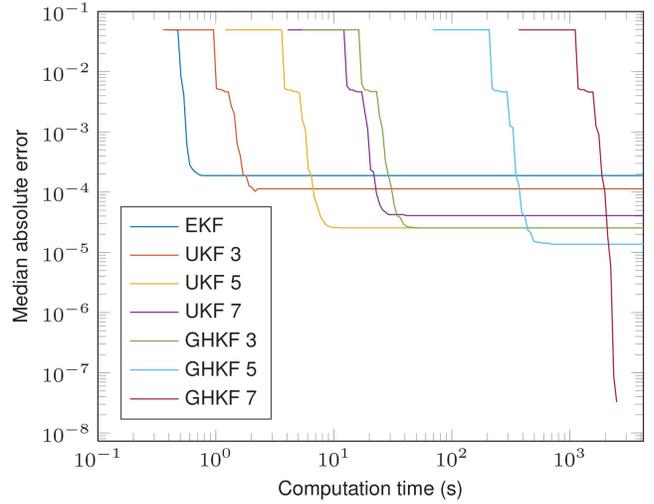


Fig. 9. Median absolute error of the parameter as a function of computation time during the optimization. Median taken over 100 datasets.

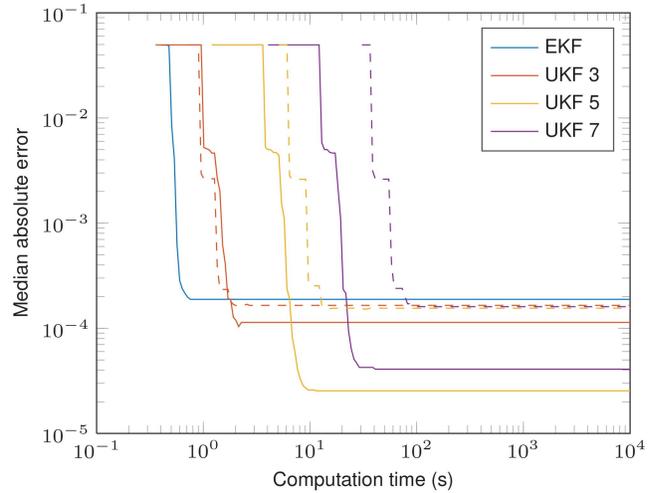


Fig. 10. Median absolute error of the parameter as a function of computation time during the optimization. Median taken over 100 datasets. The solid lines are gradient-based direct optimization, while the dashed lines show EM (run for 32 iterations) with the corresponding sigma-point schemes.

V. CONCLUSION AND DISCUSSION

In this paper together with the complementing conference article [24], we have considered various probabilistic point estimation approaches for parameter estimation in nonlinear system identification. We discussed direct likelihood maximization as well as the expectation-maximization (EM) algorithm coupled with various filtering and smoothing algorithms, namely, sigma-point filters, particle filters, and extended Kalman filters as well as the corresponding smoothers. In this paper, we focused on the differences between different sigma-point filters based on unscented transforms of third, fifth, seventh and ninth orders, and the Gauss-Hermite cubature rules.

In diminishing order of computational complexity and theoretical exactness, the filtering methods would

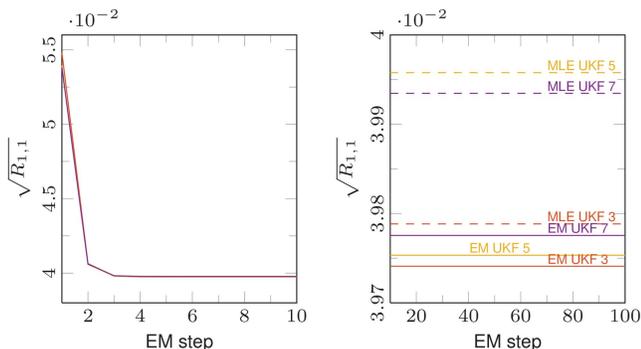


Fig. 11. The plot on the left shows the evolution of the parameter estimate during the first 10 EM steps using UKF 3, UKF 5, and UKF 7. All methods converge essentially at same speed to the same value. The plot on the right shows the evolution of the parameter estimates after the first 10 EM steps as well as the corresponding direct MLE estimates (the dashed lines).

rank as follows: particle filter, sigma-point filter, extended Kalman filter based direct likelihood approximation. In theory, particle filters converge to the exact filtering solution as the number of particles increases, while the other methods considered are based on assuming a Gaussian density and using the Kalman filter equations. However, especially in high-dimensional cases, the computational cost may prohibit the use of particle filters. In practice, the assumed density Gaussian filtering approach may have satisfactory performance if the nonlinearity is not too high. In principle, all sigma-point filters are based on assuming Gaussian density, and a higher-order cubature rule should lead to more accurate approximation of the Gaussian integrals at the cost of higher computational burden. However, typically in the literature (e.g. [2]) it has been claimed that when the Gaussian density approximation is already inaccurate, more accurate computation of the Gaussian integrals is not beneficial.

We also tested the methods in two simulated case studies. In the univariate nonstationary growth model, maximum likelihood estimates produced by different sigma-point schemes were similar. However, the estimates obtained by higher-order unscented schemes were closer to the Gauss-Hermite (order 16) baseline than the conventional 3rd order unscented transform. This suggests that the higher-order methods may indeed have some utility.

In the target tracking experiment, we compared the estimates of the noise standard deviation of one of the sensors as a function of prior uncertainty of the target's location using each of the sigma-point schemes. With higher prior uncertainty, there were more differences amongst the methods. This is reasonable because when there is less uncertainty in the model, all the methods obtain more accurate parameter estimates. Furthermore, the nonlinearity of the model has a stronger effect when the state variance is larger. Since no exact maximum likelihood estimate was available, we compared to the highest-order sigma-point scheme, namely,

GHKF 7. Compared to that, the Gauss-Hermite schemes were closer than the unscented transform based schemes and higher-order schemes were closer than lower-order schemes. Thus, the results are consistent with an assumption that the higher-order sigma-point methods produce better approximations to the Gaussian filtering result. However, the magnitude of the differences observed in this experiment is not of practical relevance and the sample size is small. Thus, the experiment may be viewed only as suggestive.

We also measured the performance of the discussed optimization routines as a function of computational time. In direct gradient-based optimization, higher-degree algorithms are more time-consuming but eventually seem to obtain better parameter estimates. The EM algorithm with low-degree sigma-point schemes (UKF 3, UKF 5) was initially faster than the gradient-based optimization, but eventually the gradient-based optimization seems to obtain better values. The EM algorithm with UKF 7 sigma-points was more computationally demanding than the direct optimization with UKF 7. This is due to the fact that EM requires the $2n$ -dimensional sigma-points for the smoothing distribution. This suggests that EM is not applicable in high-dimensional problems combined with high-degree sigma-point schemes. However, when interpreting these findings, it should be noted that we compared to the GHKF 7 estimate since the true maximum-likelihood estimate was not available.

The sigma-point integration schemes are derived by assuming exact integration results for polynomials of certain degrees. Thus, it should be noted that it is not guaranteed that a higher-order integration rule produces a more accurate results, even though it is accurate for higher-order polynomials. Furthermore, it is not guaranteed that a better approximation to the Gaussian filtering result produces a better approximation to the exact maximum likelihood result. On the other hand, there is no reason why in general a lower-order approximation to the Gaussian filtering integrals would produce more accurate approximations to the exact filtering results.

We also compared the actual tracking performance in terms of the smoother root mean square errors of the target locations using the smoother results obtained with the maximum likelihood parameter estimate of each sigma-point filter. There was no clear differences between the different sigma-point schemes in terms of the tracking error in this experiment. The tracking error of EKF increased more rapidly as a function of the initial location uncertainty, which demonstrates the local linearization nature of EKF.

In five dimensions, the UKF 5 sigma-point scheme approximates the GHKF 3 scheme in the sense that all UKF 5 sigma-points are GHKF 3 sigma-points as well, and more than half of total weight is contributed by these points. The target-tracking experiment demonstrated that these two schemes indeed produce almost equal results. However, most GHKF 3 sigma-points are

not used in the UKF 5 scheme, and thus evaluating an integral by UKF 5 requires considerably fewer function evaluations. These results suggest that there is no reason to use GHKF 3 in five dimensions as UKF 5 produces essentially same results with fewer computations. It is possible that similar computationally lighter close approximations exist for other Gauss-Hermite based sigma-point schemes as well.

In the target-tracking experiment, we also investigated how the performance of the EM algorithm varies with the sigma-point scheme used. As a function of EM iterations, the evolution of the parameter estimate was not affected by the choice of sigma-point scheme. However, the number of sigma-points in the UKF 3 rule is a small fraction of the number of sigma-points with the higher-order rules. Thus, measured by model function evaluations, EM with the UKF 3 rule converged faster. This suggests that even when the interest lies in obtaining as accurate parameter estimates as possible, a reasonable computational approach would be to first use EM with a low-order sigma-point scheme, such as UKF 3, to obtain a ballpark estimate. Then, if accuracy is desired, the initial estimate could be refined using a more accurate sigma-point scheme combined with a direct optimization algorithm.

In this paper, we considered only discrete-time state-space models. Different sigma-point schemes may also be used for continuous-discrete state-space models (see, e.g., [8] and references therein). We considered only fixed deterministic sigma-point schemes. An interesting future research topic could be to combine the recently proposed filters based on adapting or randomizing the sigma-points [10, 39, 40] with parameter estimation.

Finally, we attempt to conclude which of the methods considered here one should use in practice. Regarding the choice of sigma-point methods the higher-order unscented transform methods turned out to be quite good in the examples that we considered—but if the best possible accuracy is desired, then Gauss-Hermite methods need to be used. The EM algorithm is indeed useful in situations when the M-step optimization can be done in closed form—of which important special cases are the linear-in-parameters models considered here. However, for models that are not linear-in-parameters, EM might not be a good choice. For these models, it is thus beneficial to directly optimize the log-likelihood, and in that case we have the choice to evaluate the gradients either using the sensitivity equations or using the Fisher's identity. It turned out that the Fisher's identity is often computationally more demanding than the sensitivity equations, due to the requirement of a smoothing pass, which favors the use of the sensitivity equations for this purpose. Furthermore, the sensitivity equations give the exact gradients of the approximate likelihood whereas the Fisher's identity only gives approximate gradients of it. However, the Fisher's identity has the advantage of easy black-box implementation which can sometimes be seen as an advantage.

APPENDIX. M-STEP IN THE LINEAR-IN-PARAMETERS CASE

By substituting the linear-in-parameters model ($\mathbf{f}(\mathbf{x}) := \mathbf{A}\tilde{\mathbf{f}}(\mathbf{x})$, $\mathbf{h}(\mathbf{x}) := \mathbf{H}\tilde{\mathbf{h}}(\mathbf{x})$) into the general expression for $\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)})$, we obtain

$$\begin{aligned} \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) &\approx -\frac{1}{2} \log |2\pi \mathbf{P}_0| - \frac{T}{2} \log |2\pi \mathbf{Q}| - \frac{T}{2} \log |2\pi \mathbf{R}| \\ &\quad - \frac{1}{2} \text{tr} \{ \mathbf{P}_0^{-1} [\mathbf{P}_{0|T} + (\mathbf{m}_{0|T} - \mathbf{m}_0)(\mathbf{m}_{0|T} - \mathbf{m}_0)^T] \} \\ &\quad - \frac{1}{2} \sum_{k=1}^T \text{tr} \{ \mathbf{Q}^{-1} \mathbb{E} [(\mathbf{x}_k - \mathbf{A}\tilde{\mathbf{f}}(\mathbf{x}_{k-1})) \\ &\quad \times (\mathbf{x}_k - \mathbf{A}\tilde{\mathbf{f}}(\mathbf{x}_{k-1}))^T | \mathbf{y}_{1:T}] \} \\ &\quad - \frac{1}{2} \sum_{k=1}^T \text{tr} \{ \mathbf{R}^{-1} \mathbb{E} [(\mathbf{y}_k - \mathbf{H}\tilde{\mathbf{h}}(\mathbf{x}_k)) \\ &\quad \times (\mathbf{y}_k - \mathbf{H}\tilde{\mathbf{h}}(\mathbf{x}_k))^T | \mathbf{y}_{1:T}] \}. \end{aligned} \quad (50)$$

Since trace is linear, the penultimate term can be written as

$$\begin{aligned} &= -\frac{T}{2} \text{tr} \{ \mathbf{Q}^{-1} \frac{1}{T} \sum_{k=1}^T \mathbb{E} [(\mathbf{x}_k - \mathbf{A}\tilde{\mathbf{f}}(\mathbf{x}_{k-1})) \\ &\quad \times (\mathbf{x}_k - \mathbf{A}\tilde{\mathbf{f}}(\mathbf{x}_{k-1}))^T | \mathbf{y}_{1:T}] \} \\ &= -\frac{T}{2} \text{tr} \{ \mathbf{Q}^{-1} \frac{1}{T} \sum_{k=1}^T \mathbb{E} [\mathbf{x}_k \mathbf{x}_k^T - \mathbf{x}_k \tilde{\mathbf{f}}(\mathbf{x}_{k-1})^T \mathbf{A}^T \\ &\quad - \mathbf{A}\tilde{\mathbf{f}}(\mathbf{x}_{k-1}) \mathbf{x}_k^T + \mathbf{A}\tilde{\mathbf{f}}(\mathbf{x}_{k-1}) \tilde{\mathbf{f}}(\mathbf{x}_{k-1})^T \mathbf{A}^T | \mathbf{y}_{1:T}] \}, \end{aligned} \quad (51)$$

which due to linearity of expectation equals

$$\begin{aligned} &= -\frac{T}{2} \text{tr} \left\{ \mathbf{Q}^{-1} \left[\frac{1}{T} \sum_{k=1}^T \mathbb{E} [\mathbf{x}_k \mathbf{x}_k^T | \mathbf{y}_{1:T}] \right. \right. \\ &\quad - \frac{1}{T} \left(\sum_{k=1}^T \mathbb{E} [\mathbf{x}_k \tilde{\mathbf{f}}(\mathbf{x}_{k-1})^T | \mathbf{y}_{1:T}] \right) \mathbf{A}^T \\ &\quad - \mathbf{A} \frac{1}{T} \sum_{k=1}^T \mathbb{E} [\tilde{\mathbf{f}}(\mathbf{x}_{k-1}) \mathbf{x}_k^T | \mathbf{y}_{1:T}] \\ &\quad \left. \left. + \mathbf{A} \frac{1}{T} \sum_{k=1}^T \left(\mathbb{E} [\tilde{\mathbf{f}}(\mathbf{x}_{k-1}) \tilde{\mathbf{f}}(\mathbf{x}_{k-1})^T | \mathbf{y}_{1:T}] \right) \mathbf{A}^T \right] \right\}. \end{aligned} \quad (52)$$

Noting that $\mathbb{E} [\mathbf{x}_k \mathbf{x}_k^T | \mathbf{y}_{1:T}] = \mathbf{m}_{k|T} \mathbf{m}_{k|T}^T + \mathbf{P}_{k|T}$ and substituting in the notation introduced in Equations (35–40), we obtain

$$= -\frac{T}{2} \text{tr} \{ \mathbf{Q}^{-1} [\boldsymbol{\Sigma} - \mathbf{C} \mathbf{A}^T - \mathbf{A} \mathbf{C}^T + \mathbf{A} \boldsymbol{\Phi} \mathbf{A}^T] \}. \quad (53)$$

Similar calculation for the last term in Equation (50), noting that $\mathbb{E}[\mathbf{y}_k \mathbf{y}_k^T | \mathbf{y}_{1:T}] = \mathbf{y}_k \mathbf{y}_k^T$, gives

$$-\frac{T}{2} \text{tr}\{\mathbf{R}^{-1}[\mathbf{D} - \mathbf{B}\mathbf{H}^T - \mathbf{H}\mathbf{B}^T + \mathbf{H}\mathbf{\Theta}\mathbf{H}^T]\}. \quad (54)$$

Substituting Equations (53) and (54) into Equation (50), we get

$$\begin{aligned} \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{(n)}) = & \\ & -\frac{1}{2} \log |2\pi \mathbf{P}_0| - \frac{T}{2} \log |2\pi \mathbf{Q}| - \frac{T}{2} \log |2\pi \mathbf{R}| \\ & -\frac{1}{2} \text{tr}\{\mathbf{P}_0^{-1}[\mathbf{P}_{0|T} + (\mathbf{m}_{0|T} - \mathbf{m}_0)(\mathbf{m}_{0|T} - \mathbf{m}_0)^T]\} \\ & -\frac{T}{2} \text{tr}\{\mathbf{Q}^{-1}[\boldsymbol{\Sigma} - \mathbf{C}\mathbf{A}^T - \mathbf{A}\mathbf{C}^T + \mathbf{A}\boldsymbol{\Phi}\mathbf{A}^T]\} \\ & -\frac{T}{2} \text{tr}\{\mathbf{R}^{-1}[\mathbf{D} - \mathbf{B}\mathbf{H}^T - \mathbf{H}\mathbf{B}^T + \mathbf{H}\mathbf{\Theta}\mathbf{H}^T]\}. \end{aligned}$$

To maximize this with respect to the parameters $(\mathbf{m}_0, \mathbf{P}_0, \mathbf{A}, \mathbf{H}, \mathbf{Q}, \mathbf{R})$ we differentiate with respect to parameter in question and set the derivative to 0. For \mathbf{Q} :

$$\begin{aligned} \frac{d\mathcal{Q}}{d\mathbf{Q}} = & -\frac{T}{2} \frac{d}{d\mathbf{Q}} \log |2\pi \mathbf{Q}| \\ & -\frac{T}{2} \frac{d}{d\mathbf{Q}} \text{tr}\{\mathbf{Q}^{-1}[\boldsymbol{\Sigma} - \mathbf{C}\mathbf{A}^T - \mathbf{A}\mathbf{C}^T + \mathbf{A}\boldsymbol{\Phi}\mathbf{A}^T]\} \\ = & -\frac{T}{2} \mathbf{Q}^{-1} \\ & + \frac{T}{2} \mathbf{Q}^{-1}[\boldsymbol{\Sigma} - \mathbf{C}\mathbf{A}^T - \mathbf{A}\mathbf{C}^T + \mathbf{A}\boldsymbol{\Phi}\mathbf{A}^T] \mathbf{Q}^{-1}. \quad (55) \end{aligned}$$

Setting the derivative equal to $\mathbf{0}$, we obtain the equation

$$\frac{T}{2} \mathbf{Q}^{-1} = \frac{T}{2} \mathbf{Q}^{-1}[\boldsymbol{\Sigma} - \mathbf{C}\mathbf{A}^T - \mathbf{A}\mathbf{C}^T + \mathbf{A}\boldsymbol{\Phi}\mathbf{A}^T] \mathbf{Q}^{-1}. \quad (56)$$

Multiplying from right by $(2/T)\mathbf{Q}$ and from left by \mathbf{Q} gives

$$\mathbf{Q} = \boldsymbol{\Sigma} - \mathbf{C}\mathbf{A}^T - \mathbf{A}\mathbf{C}^T + \mathbf{A}\boldsymbol{\Phi}\mathbf{A}^T. \quad (57)$$

The derivations for the optimal solutions of \mathbf{R} and \mathbf{P}_0 are similar. For \mathbf{A} :

$$\begin{aligned} \frac{d\mathcal{Q}}{d\mathbf{A}} = & -\frac{T}{2} \left[-\frac{d}{d\mathbf{A}} \text{tr}(\mathbf{Q}^{-1} \mathbf{C}\mathbf{A}^T) \right. \\ & \left. -\frac{d}{d\mathbf{A}} \text{tr}(\mathbf{Q}^{-1} \mathbf{A}\mathbf{C}^T) + \frac{d}{d\mathbf{A}} \text{tr}(\mathbf{A}\boldsymbol{\Phi}\mathbf{A}^T) \right] \\ = & -\frac{T}{2} \mathbf{Q}^{-1} [2\mathbf{A}\boldsymbol{\Phi} - 2\mathbf{C}] \quad (58) \end{aligned}$$

Since \mathbf{Q}^{-1} is nonsingular, the derivative is zero only if the last factor is zero. If $\boldsymbol{\Phi}$ is invertible, this in turn implies

$$\mathbf{A} = \mathbf{C}\boldsymbol{\Phi}^{-1}. \quad (59)$$

The derivations for the optimal solutions of \mathbf{H} and \mathbf{m}_0 are similar.

If the parameter $\boldsymbol{\theta}$ is any subset of $\{\mathbf{A}, \mathbf{H}, \mathbf{Q}, \mathbf{R}, \mathbf{m}_0, \mathbf{P}_0\}$, it can be optimized by these closed-form expressions.

First, note that (\mathbf{A}, \mathbf{Q}) , (\mathbf{H}, \mathbf{R}) and $(\mathbf{m}_0, \mathbf{P}_0)$ are independent in the sense that, for example, the optimal \mathbf{A} and \mathbf{Q} do not depend on the other four parameters. Furthermore, the optimal \mathbf{A} does not depend on \mathbf{Q} . Thus, \mathbf{A} and \mathbf{Q} can be jointly optimized by first solving the optimal \mathbf{A} and then substituting that into the expression of optimal \mathbf{Q} . Similar reasoning works for (\mathbf{H}, \mathbf{R}) and $(\mathbf{m}_0, \mathbf{P}_0)$.

ACKNOWLEDGMENTS

This work was supported by grants from the Academy of Finland (266940, 273475) and by the Emil Aaltonen foundation. We acknowledge the computational resources provided by the Aalto Science-IT project.

REFERENCES

- [1] M. L. Andrade Netto, L. Gimeno, and M. J. Mendes
A new spline algorithm for non-linear filtering of discrete time systems.
In A Link Between Science and Application of Automatic Control, pages 2123–2130, Helsinki, Finland, 1978.
- [2] I. Arasaratnam and S. Haykin
Cubature Kalman filters.
IEEE Transactions on Automatic Control, 54(6):1254–1269, 2009.
- [3] I. Arasaratnam, S. Haykin, and T. R. Hurd
Cubature Kalman filtering for continuous-discrete systems: Theory and simulations.
IEEE Transactions on Signal Processing, 58(10):4977–4993, 2010.
- [4] Y. Bar-Shalom, X.-R. Li, and T. Kirubarajan
Estimation with Applications to Tracking and Navigation. Wiley Interscience, 2001.
- [5] O. Cappé, E. Moulines, and T. Rydén
Inference in Hidden Markov Models. Springer, 2005.
- [6] R. Cools
Constructing cubature formulae: The science behind the art.
In Acta Numerica, volume 6, pages 1–54. 1997.
- [7] D. Crouse
Basic tracking using nonlinear 3D monostatic and bistatic measurements.
IEEE Aerospace and Electronic Systems Magazine, 29(8):4–53, 2014.
- [8] D. Crouse
Basic tracking using nonlinear continuous-time dynamic models [tutorial].
IEEE Aerospace and Electronic Systems Magazine, 30(2):4–41, 2015.
- [9] A. P. Dempster, N. M. Laird, and D. B. Rubin
Maximum likelihood from incomplete data via the EM algorithm.
Journal of the Royal Statistical Society: Series B (Statistical Methodology), 39(1):1–38, 1977.
- [10] J. Duník, O. Straka, and M. Simandl
Stochastic integration filter.
IEEE Transactions on Automatic Control, 58(6):1561–1566, 2013.
- [11] R. Fletcher
Practical Methods of Optimization. John Wiley & Sons, 2nd edition, 1987.
- [12] M. Gašperin and Đ. Juričić
Application of unscented transformation in nonlinear system identification.
In Proceedings of the 18th IFAC World Congress, volume 18, pages 4428–4433, 2011.

- [13] A. Gelb, editor
Applied Optimal Estimation.
The MIT Press, 1974.
- [14] R. Goodrich and P. E. Caines
Linear system identification from nonstationary cross-sectional data.
IEEE Transactions on Automatic Control, 24(3):403–411, 1979.
- [15] N. K. Gupta and R. K. Mehra
Computational aspects of maximum likelihood estimation and reduction in sensitivity function calculations.
IEEE Transactions on Automatic Control, 19(6):774–783, 1974.
- [16] K. Ito and K. Xiong
Gaussian filters for nonlinear filtering problems.
IEEE Transactions on Automatic Control, 45(5):910–927, 2000.
- [17] A. H. Jazwinski
Stochastic Processes and Filtering Theory.
Academic Press, 1970.
- [18] B. Jia, M. Xin, and Y. Cheng
High-degree cubature Kalman filter.
Automatica, 49(2):510–518, 2013.
- [19] S. J. Julier, J. K. Uhlmann, and H. F. Durrant-Whyte
A new approach for filtering nonlinear systems.
In *Proceedings of the American Control Conference*, volume 3, pages 1628–1632, Seattle, WA, 1995.
- [20] S. J. Julier, J. K. Uhlmann, and H. F. Durrant-Whyte
A new method for the nonlinear transformation of means and covariances in filters and estimators.
IEEE Transactions on Automatic Control, 45(3):477–482, 2000.
- [21] R. E. Kalman
A new approach to linear filtering and prediction problems.
Transactions of the ASME, Journal of Basic Engineering, 82:35–45, March 1960.
- [22] G. Kitagawa
Non-Gaussian state-space modeling of nonstationary time series.
Journal of the American Statistical Association, 82(400):1032–1041, 1987.
- [23] G. Kitagawa
Non-Gaussian state-space modeling of nonstationary time series: Rejoinder.
Journal of the American Statistical Association, 82(400):1060–1063, 1987.
- [24] J. Kokkala, A. Solin, and S. Särkkä
Expectation maximization based parameter estimation by sigma-point and particle smoothing.
In *Proceedings of the 17th International Conference on Information Fusion (FUSION)*, Salamanca, Spain, 2014.
- [25] U. N. Lerner
Hybrid Bayesian networks for reasoning about complex systems.
PhD thesis, Stanford University, Stanford, CA, 2002.
- [26] J. McNamee and F. Stenger
Construction of fully symmetric numerical integration formulas.
Numerische Mathematik, 10(4):327–344, 1967.
- [27] R. M. Neal and G. E. Hinton
A view of the EM algorithm that justifies incremental, sparse, and other variants.
In *Learning in Graphical Models*, pages 355–368. Springer, 1998.
- [28] R. K. Olsson, K. B. Petersen, and T. Lehn-Schiøler
State-space models: From the EM algorithm to a gradient approach.
Neural Computation, 19(4):1097–1111, 2007.
- [29] G. Poyiadjis, A. Doucet, and S. S. Singh
Particle methods for optimal filter derivative: Application to parameter estimation.
In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 5, pages 925–928, Philadelphia, PA, 2005.
- [30] H. E. Rauch, F. Tung, and C. T. Striebel
Maximum likelihood estimates of linear dynamic systems.
AIAA Journal, 3(8):1445–1450, 1965.
- [31] S. Roweis and Z. Ghahramani
Learning nonlinear dynamical systems using the expectation-maximization algorithm.
In S. Haykin, editor, *Kalman Filtering and Neural Networks*, chapter 6, pages 175–220. Wiley-Interscience, 2001.
- [32] S. Särkkä
Bayesian Filtering and Smoothing, volume 3 of *Institute of Mathematical Statistics Textbooks*.
Cambridge University Press, 2013.
- [33] S. Särkkä and J. Hartikainen
On Gaussian optimal smoothing of non-linear state space models.
IEEE Transactions on Automatic Control, 55(8):1938–1941, 2010.
- [34] T. B. Schön, A. Wills, and B. Ninness
System identification of nonlinear state-space models.
Automatica, 47(1):39–49, 2011.
- [35] M. Segal and E. Weinstein
A new method for evaluating the log-likelihood gradient (score) of linear dynamic systems.
IEEE Transactions on Automatic Control, 33(8):763–766, 1988.
- [36] R. H. Shumway and D. S. Stoffer
An approach to time series smoothing and forecasting using the EM algorithm.
Journal of Time Series Analysis, 3(4):253–264, 1982.
- [37] H. Singer
Parameter estimation of nonlinear stochastic differential equations: Simulated maximum likelihood versus extended Kalman filter and Itô-Taylor expansion.
Journal of Computational and Graphical Statistics, 11(4):972–995, 2002.
- [38] A. Solin
Cubature integration methods in non-linear Kalman filtering and smoothing.
Bachelor’s thesis, Faculty of Information and Natural Sciences, Aalto University, Finland, 2010.
- [39] O. Straka, J. Duník, M. Šimandl, and E. Blasch
Randomized unscented transform in state estimation of non-Gaussian systems: Algorithms and performance.
In *Proceedings of the 15th International Conference on Information Fusion (FUSION)*, Singapore, 2012.
- [40] O. Straka, J. Duník, M. Šimandl, and E. Blasch
Comparison of adaptive and randomized unscented Kalman filter algorithms.
In *Proceedings of the 17th International Conference on Information Fusion (FUSION)*, Salamanca, Spain, 2014.
- [41] V. Väinänen
Gaussian filtering and smoothing based parameter estimation in nonlinear models for sequential data.
Master’s thesis, School of Electrical Engineering, Aalto University, Finland, 2012.

- [42] Y. Wu, D. Hu, M. Wu, and X. Hu
Unscented Kalman filtering for additive noise case: Augmented vs. non-augmented.
In *Proceedings of the American Control Conference*, volume 6, pages 4051–4055, Portland, OR, 2005.

- [43] Y. Wu, D. Hu, M. Wu, and X. Hu
A numerical-integration perspective on Gaussian filters.
IEEE Transactions on Signal Processing, 54(8):2910–2921, 2006.



Juho Kokkala received his M.Sc. (Tech.) degree from Aalto University, Finland, in 2010, majoring in Systems and Operations Research. He is currently a doctoral student at Aalto University and will defend his thesis in 2016. His research interests are Bayesian statistics, especially Bayesian methods for nonlinear state-space models.



Arno Solin received his M.Sc. (Tech.) degree (with distinction) in engineering physics and mathematics from Aalto University, Espoo, Finland, in 2012. He is scheduled to defend his Doctoral of Science (Tech.) degree in computational science at Aalto University in early 2016. He is currently a researcher at the Department of Neuroscience and Biomedical Engineering and the Department of Computer Science at Aalto University. His research interests are in probabilistic methods for spatio-temporal data with applications to machine learning, signal processing, and spatial statistics with applications in tracking and brain imaging.



Simo Särkkä received his M.Sc. degree in engineering physics and mathematics, and D.Sc. degree in electrical and communications engineering from Helsinki University of Technology, Espoo, Finland, in 2000 and 2006, respectively. Currently he is an Associate Professor and Academy Research Fellow with Aalto University, and Technical Advisor and Director of IndoorAtlas Ltd. His research interests are in multi-sensor data processing systems with applications in location sensing, health technology, machine learning, inverse problems, and brain imaging.