
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Tronarp, Filip; Särkkä, Simo

Iterative statistical linear regression for Gaussian smoothing in continuous-time non-linear stochastic dynamic systems

Published in:
Signal Processing

DOI:
[10.1016/j.sigpro.2019.01.013](https://doi.org/10.1016/j.sigpro.2019.01.013)

Published: 01/06/2019

Document Version
Peer reviewed version

Published under the following license:
CC BY-NC-ND

Please cite the original version:
Tronarp, F., & Särkkä, S. (2019). Iterative statistical linear regression for Gaussian smoothing in continuous-time non-linear stochastic dynamic systems. *Signal Processing*, 159, 1-12.
<https://doi.org/10.1016/j.sigpro.2019.01.013>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Iterative Statistical Linear Regression for Gaussian Smoothing in Continuous-Time Non-linear Stochastic Dynamic Systems

Filip Tronarp and Simo Särkkä

Department of Electrical Engineering and Automation, Aalto University, Rakentajanaukio 2, 02150 Espoo, Finland

Abstract

This paper considers approximate smoothing for discretely observed non-linear stochastic differential equations. The problem is tackled by developing methods for linearising stochastic differential equations with respect to an arbitrary Gaussian process. Two methods are developed based on 1) taking the limit of statistical linear regression of the discretised process and 2) minimising an upper bound to a cost functional. Their difference is manifested in the diffusion of the approximate processes. This in turn gives novel derivations of pre-existing Gaussian smoothers when Method 1 is used and a new class of Gaussian smoothers when Method 2 is used. Furthermore, based on the aforementioned development the iterative Gaussian smoothers in discrete-time are generalised to the continuous-time setting by iteratively re-linearising the stochastic differential equation with respect to the current Gaussian process approximation to the smoothed process. The method is verified in two challenging tracking problems, a reentry problem and a radar tracked coordinated turn model with state dependent diffusion. The results show that the method has competitive estimation accuracy with state-of-the-art smoothers.

Keywords: stochastic differential equations, statistical linear regression, iterative methods, continuous-discrete Gaussian smoothing.

1. Introduction

Inference in continuous-time stochastic dynamic systems is a frequently occurring topic in disciplines such as navigation, tracking, and time series modelling [1, 2, 3, 4]. The system is typically described in terms of a latent Markov process $\{X(s)\}_{s \geq 0}$, governed by a stochastic differential equation (SDE) [5]. Furthermore, the process $\{X(s)\}_{s \geq 0}$ is assumed to be measured at a set of time instants $\{t_k\}_{k=1}^K$ by a collection of random variables $\{Y(t_k)\}_{k=1}^K$, each having a conditional distribution with respect to the process outcome at the corresponding time stamp. For the special case of an affine, Gaussian system, the calculation of predictive, filtering, and smoothing distributions amount to manipulating the joint moments of the latent process and the measurement process. In the

case of filtering this procedure is known as Kalman-Bucy filtering [6]. It was subsequently shown that the smoothing moments can be expressed in terms of ordinary differential equations with the filter moment as inputs [7, 8].

While the theory of filtering and smoothing in linear dynamic systems is mature, the case of non-linear systems is still an area of intense research and a common strategy is to find a suitable linearisation of the system which enables the aforementioned methods for affine systems. An early approach was to linearise the system around the mean trajectory using truncated Taylor series [9]. Another, fairly recent, approach was to apply the Rauch-Tung-Striebel smoother [7] to a discretisation of the process and applying the limit $\delta t \rightarrow 0$ [10]. The aforementioned approaches belong

to the class of Gaussian smoothers, which were further studied in [11], where the non-linear smoothing theory [8, 12] was used to derive the Type I smoother. The Type I smoother of [11] was subsequently re-derived in [13] using the projection methods developed in [14, 15].

Another line of research has been into variational Gaussian smoothers, which imposes a Gaussian process posterior by fixed-form variational Bayes [16]. This class of smoothers operate by iteratively refining the approximate smoothing solutions by taking gradient steps associated with the evidence lower bound. However, the development in [16] requires a non-singular, state-independent diffusion. The variational smoother of [16] was extended in [17] by allowing singular diffusions if the drift in the singular subspace is an affine function of the state. Nonetheless, both formulations of variational smoothers require a state independent diffusion [16, 17]. The requirement of state independent diffusion was later circumvented by relaxing the condition of the fixed-form posterior to only have Gaussian marginals in [18]. However, this adds the requirement of solving a stochastic partial differential equation to obtain an initial value to their optimal control formulation [18]. This can be computationally challenging even for moderately large state spaces, which is why the method was only validated on problems of small state space (one dimension). Another approach in a similar vein is the continuous-time expectation propagation algorithm [19], which is also applicable when the diffusion is state dependent. However, accurate approximations of the moments with respect to the so-called *tilted distributions*, which becomes a non-trivial problem as the dimension of the state-space grows.

A recent advance, in discrete time inference, is the iterated posterior linearisation smoother (IPLS) [20] (see also [21]) which generalises the iterated extended Kalman smoother [22] to sigma-point methods. This was done on the basis of statistical linear regression [23], where a given smoothing solution is improved upon by re-linearising the system using the current Gaussian smoother approximation and then running the smoother again [20, 21]. However, an analogue for continuous-time smoothing has yet to appear.

The purpose of this paper is thus to generalise the

discrete-time smoother of [20, 21] to the continuous-time case. This is accomplished by generalising the statistical linear regression method [23] to the setting of stochastic differential equations. Two ways of doing this are discovered, 1) taking the limit of the statistical linear regression solution of a discretised process and 2) minimising an upper bound to a cost functional. This gives novel derivations of the smoothers in [11] when Method 1 is used and a new kind of smoothers when Method 2 is used. Furthermore, using the aforementioned linearisation methods iterative Gaussian smoothers, analogous to the discrete-time iterative Gaussian smoother [22, 20, 21], are developed for stochastic differential equations.

In Section 2 the smoothing problem is formally posed, and linear smoothing theory and previous approaches to smoothing in non-linear systems is reviewed. Lastly, the present contribution is outlined. In Section 3, the statistical linear regression method is generalised to stochastic differential equations. It is derived both as a discrete time limit and as a minimiser to a certain cost functional. The development in Section 3 is subsequently combined with a linear smoothing theory to arrive at novel derivations of the smoothers presented in [10, 11]. The main result is presented in Section 5 where the discrete time iterative Gaussian smoothers [20, 21, 22] are generalised to continuous-time models. In Section 6 the iterative smoothers are demonstrated in two non-linear and high-dimensional target tracking problems. The manuscript ends with the conclusion in Section 7.

2. Problem Formulation

The setting is as follows, there is a latent Markov process $\{X(t)\}_{t \geq 0}$, $X(t) \in \mathbb{R}^{d_x}$ which is assumed to evolve according to the following discretely observed stochastic differential equation (SDE) model

$$dX(t) = \mu(t, X(t)) dt + \sigma(t, X(t)) dW(t), \quad (1a)$$

$$Y(t_k) = h(t_k, X(t_k)) + V(t_k), \quad (1b)$$

$$\mathbb{C}[V(t_k), V(t_l)] = \delta_{k,l} R, \quad V(t_k) \sim \mathcal{N}(0, R),$$

where $f_{X(0)}(x) = \mathcal{N}(x; \bar{x}(0^-), \Sigma(0^-))$, $\mu: \mathbb{R}_+ \times \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$ is a drift function, $\sigma: \mathbb{R}_+ \times \mathbb{R}^{d_x} \rightarrow$

$\mathbb{R}^{d_X \times d_W}$ is a diffusion matrix, $\{W(s)\}_{s \geq 0}$ is a d_W -dimensional standard Brownian motion, $h: \mathbb{R}_+ \times \mathbb{R}^{d_X} \rightarrow \mathbb{R}^{d_Y}$ is a measurement function, and $\delta_{k,l}$ is Kronecker's delta function. Furthermore, given a measurement series $\{y(t_k)\}_{k=1}^K$, $t_{k+1} > t_k$. The set of measurements up to just before time t and the set of measurements up to precisely time t are denoted by $\mathcal{Y}(t^-) = \{y(t_k): t_k < t\}$ and $\mathcal{Y}(t) = \{y(t_k): t_k \leq t\}$, respectively.

The inference problem for $X(\tau)$ is then in the Bayesian sense to find a family of conditional densities

$$f_{X(\tau)|\mathcal{Y}(t_k)}(x), \quad k = 1, \dots, K. \quad (2)$$

When $\tau < t_k$ the probability density function in Equation (2) is said to be a smoothing distribution, if, in particular, $\tau = t_k$ it is a filtering distribution and if $\tau > t_k$ then the density in Equation (2) is said to be a predictive distribution. Moreover, the expectation, cross-covariance, and covariance operators are denoted by $\mathbb{E}[\cdot]$, $\mathbb{C}[\cdot, \cdot]$ and $\mathbb{V}[\cdot]$. We use the following notation:

$$\bar{x}(t) = \mathbb{E}[X(t) | \mathcal{Y}(t)], \quad (3a)$$

$$\bar{x}(t^-) = \mathbb{E}[X(t) | \mathcal{Y}(t^-)], \quad (3b)$$

$$\Sigma(t) = \mathbb{V}[X(t) | \mathcal{Y}(t)], \quad (3c)$$

$$\Sigma(t^-) = \mathbb{V}[X(t) | \mathcal{Y}(t^-)], \quad (3d)$$

and similarly for the smoothing moments based on the entire measurement series:

$$\hat{x}(t) = \mathbb{E}[X(t) | \mathcal{Y}(t_K)], \quad (4a)$$

$$\Omega(t) = \mathbb{V}[X(t) | \mathcal{Y}(t_K)]. \quad (4b)$$

2.1. Prior Work

Smoothing in state space models has endured long and considerable efforts in the past 50 years [7, 8, 12, 24]. First the linear smoothing theory will be reviewed and subsequently the more prominent approaches to approximate smoothers.

2.1.1. Linear smoothing theory

The linear smoothing theory applies to systems of the following form:

$$dX(t) = (A(t)X(t) + b(t))dt + \sigma(t)dW(t), \quad (5a)$$

$$Y(t_k) = C(t_k)X(t_k) + d(t_k) + V(t_k), \quad (5b)$$

$$\mathbb{C}[V(t_k), V(t_l)] = \delta_{t_k, t_l} R, \quad V(t_k) \sim \mathcal{N}(0, R),$$

where $W(t)$ is a standard Wiener process. Since the collection $(X(t_{1:K}), Y(t_{1:K}))$ is jointly Gaussian the conditioning reduces to projections in a finite dimensional space. Furthermore, this can be implemented in a sequential manner where alternations between update and predictions are carried out [6].

Starting with the predictive distribution, $X(t_k^-) \sim \mathcal{N}(\bar{x}(t_k^-), \Sigma(t_k^-))$, the parameters of the filtering distribution are then computed according to [4]

$$S(t_k) = C(t_k)\Sigma(t_k^-)C(t_k)^\top + R, \quad (6a)$$

$$K(t_k) = \Sigma(t_k^-)C(t_k)^\top S^{-1}(t_k), \quad (6b)$$

$$\hat{v}(t_k) = y(t_k) - C(t_k)\bar{x}(t_k^-) - d(t_k), \quad (6c)$$

$$\bar{x}(t_k) = \bar{x}(t_k^-) + K(t_k)\hat{v}(t_k), \quad (6d)$$

$$\Sigma(t_k) = \Sigma(t_k^-) - K(t_k)S(t_k)K(t_k)^\top. \quad (6e)$$

The predictive distribution at t_{k+1}^- is then given by solving [5, 6]

$$\frac{d\bar{x}(t)}{dt} = A(t)\bar{x}(t) + b(t), \quad (7a)$$

$$\frac{d\Sigma(t)}{dt} = A(t)\Sigma(t) + \Sigma(t)A(t)^\top + Q(t), \quad (7b)$$

on the interval $\tau \in [t_k, t_{k+1}]$ with initial conditions $(\bar{x}(t_k), \Sigma(t_k))$, where $Q(t) = \sigma(t)\sigma(t)^\top$. The differential equations for the smoothing moments can then be expressed in terms of the filtering moments according to [7, 8, 12]

$$\begin{aligned} \frac{d\hat{x}(t)}{dt} &= A(t)\hat{x}(t) + b(t) \\ &\quad + Q(t)\Sigma^{-1}(t)(\hat{x}(t) - \bar{x}(t)), \end{aligned} \quad (8a)$$

$$\begin{aligned} \frac{d\Omega(t)}{dt} &= [A(t) + Q(t)\Sigma^{-1}(t)]\Omega(t) \\ &\quad + \Omega(t)[A(t) + Q(t)\Sigma^{-1}(t)]^\top - Q(t). \end{aligned} \quad (8b)$$

2.1.2. The Non-linear smoothing theory approach

Now consider the smoothing problem for the non-linear model in Equation (1). The expectations with respect to the smoothing distribution of some test

function $\psi(X(t))$ was studied in [8, 12], and the backwards differential equation

$$\frac{d}{dt}\mathbb{E}[\psi(X(t)) \mid \mathcal{Y}(t_K)] = \mathbb{E}[\mathcal{K}_1\psi(X(t)) \mid \mathcal{Y}(t_K)]$$

is obtained, where the operator \mathcal{K}_1 is given by

$$\begin{aligned} \mathcal{K}_1\psi(x(t)) &= \sum_p \mu_p(t, x(t)) \partial_p \psi(x(t)) \\ &- \sum_{p,r} \partial_p \psi(x(t)) \partial_r [Q(t, x(t))]_{p,r} \\ &- \frac{1}{2} \sum_{p,r} [Q(t, x(t))]_{p,r} \partial_{p,r}^2 \psi(x(t)) \\ &- \frac{\sum_{p,r} [Q(t, x(t))]_{p,r} \partial_p \psi(x(t)) \partial_r \psi(x(t))}{f_{X(t)|\mathcal{Y}(t)}(x(t))}, \end{aligned}$$

where ∂_r is the partial derivative operator with respect to the r :th coordinate in X and $\partial_{p,r}^2$ is the composition of ∂_r and ∂_p . Approaches to implement this was not covered in [8] and in [12] a Taylor expansion was used. However, by plugging in a Gaussian approximation to $f_{X(t)|\mathcal{Y}(t)}(x(t))$, the Type I smoother of [11] is derived. The Type II and III smoothers are also discussed by [11], which originate from the smoother formulation developed in [10].

2.1.3. The Kullback-Leibler approach

Another approach to smoothing in non-linear systems is based on minimising the Kullback-Leibler divergence between the posterior measure, $\mathbb{P}_{X|Y}$, and a fixed form Gaussian measure, $\mathbb{Q}_{X|Y}$ [16]. This requires (i) $\sigma(t, X(t)) = \sigma(t)$ and (ii) $Q(t)^{-1} = (\sigma(t)\sigma^\top(t))^{-1}$ exists. The Kullback-Leibler divergence is given by

$$\begin{aligned} \text{KL}[\mathbb{Q}_{X|Y} \parallel \mathbb{P}_{X|Y}] &= \frac{1}{2} \int_0^T E(t) dt + \frac{Kd_X}{2} \\ &+ \frac{K}{2} \log \det R + \log Z, \end{aligned} \quad (9)$$

where Z is a normalisation constant, $E(t) = E_X(t) + E_{Y|X}(t)$, and

$$\begin{aligned} E_X(t) &= \mathbb{E}^{\mathbb{Q}} \left[\|\mu(t, X(t)) - A(t)X(t) - b(t)\|_{Q^{-1}(t)}^2 \right], \end{aligned}$$

$$\begin{aligned} E_{Y|X}(t) &= \sum_{k=1}^K \mathbb{E}^{\mathbb{Q}} \left[\|y(t_k) - h(t_k, X(t_k))\|_{R^{-1}}^2 \right] \delta(t - t_k), \end{aligned}$$

where $\|\cdot\|_W$ is a weighted Euclidean norm with weighting matrix W . Now, it is clear that the above functional is not well defined if $Q(t)$ is singular. This was extended in [17] to the case of singular $Q(t)$ under the assumption that the drift function of the singular sub-space is an affine function of the state. This was further extended to the case of state dependent Q in [18], where they minimise Equation (9) subject to $\mathbb{Q}_{X|Y}$ being a diffusion process with a prescribed marginal law (i.e Gaussian). However, this requires computing $w(x, 0)$, with $w(x, t)$ being the solution to the Kolmogorov backward equation.

Furthermore, the minimisation objective of continuous-time expectation propagation [19] is retrieved by exponentiating the argument to the expectations in the definition of $E_{Y|X}(t)$, along with some other minor modifications.

2.2. The Contribution

The aim of this paper is to develop iterative techniques for obtaining smoothing estimates of the system in Equations (1a) and (1b). More specifically the following contributions are put forth:

1. The statistical linear regression [23, 25] method is generalised to the setting of stochastic differential equations. Two alternative versions of this procedure are provided, 1) is based on applying standard statistical linear regression to a discretised version of an auxiliary SDE and passing to the continuous limit, while 2) sets up a least squares problem for the difference between the auxiliary SDE and an affine approximation, which results in the minimisation of a quadratic functional, which is solved by methods in variational calculus [26].
2. It is shown that the Type II and III smoothers of [11] can be derived by using Method 1 for linearising the stochastic differential equation together with the linear smoothing results [7, 8]. Using Method 2 gives a new class of smoothers. These

shall be separated by the suffixes *of the first kind* and *of the second kind* for Method 1 and Method 2 of linearising the stochastic differential equation, respectively.

- Iterated smoothers are developed on the basis of using Contribution 1 above to re-linearise the SDE with respect to the current best Gaussian process approximation of the smoothed process, which gives a continuous-time generalisation of the iterative Gaussian smoothers [20, 21].

3. Statistical Linear Regression For Stochastic Differential Equations

In this section, the statistical linear regression method [23] (see also [25]) is generalised to the case of affine approximations of stochastic differential equations. Let $X(t)$ be driven by the SDE in Equation (1a). Then a Gaussian process, $\{\widehat{X}(s)\}_{s \geq 0}$ can be used to approximate the evolution of $X(t)$ according to

$$dX(t) \approx [A(t)X(t) + b(t)] dt + \bar{\sigma}(t) d\widehat{W}(t), \quad (11)$$

where $\widehat{W}(t)$ is a standard Wiener process. The procedures for doing this is given in Algorithms 1 and 2. The role of the Gaussian process, $\{\widehat{X}(s)\}_{s \geq 0}$, in Algorithms 1 and 2 is essentially, to approximate the drift function and the diffusion matrix according to Equation (11). This is done by considering an auxiliary process defined by

$$d\widetilde{X}(t) = \mu(t, \widehat{X}(t)) dt + \sigma(t, \widehat{X}(t)) d\widetilde{W}(t), \quad (12)$$

where $\widetilde{W}(t)$ is a standard Brownian motion with that is independent of $\widehat{X}(t)$. The parameters $A(t)$, $b(t)$, $\bar{\sigma}(t)$ are then found by making an affine approximation of Equation (12) in terms of $\widehat{X}(t)$. This can be done in two ways, (i) performing SLR [23] on an Euler-Maruyama discretisation of Equation (12) and passing to the limit, which gives Algorithm 1, and (ii) using Equation (12) to set up a cost functional for $A(t)$, $b(t)$, and $\bar{\sigma}(t)$, which results in Algorithm 2. Detailed derivations of these approaches are developed in Sections 3.1 and 3.2.

Furthermore, it should be noted that Algorithms 1 and 2 suggests an iterative scheme for smoothing,

analogous to the discrete-time case [20, 21]. That is, the linearising process $\widehat{X}(t)$ is taken to be the current best approximation of the smoothing solution of Equation (1). The parameters that are retrieved can then be used in conjunction with Equations (6) to (8) to retrieve a new approximate smoothing solution. This issue that shall be revisited in Section 5.

Algorithm 1 Statistical Linear Regression I (Discrete-time limit)

Input: The marginal moment functions of a Gaussian process $\widehat{X}(s)$, $\{(\mathbb{E}[\widehat{X}(s)], \mathbb{V}[\widehat{X}(s)])\}_{s=0}^T$, drift function $\mu(t, X(t))$, and diffusion matrix $\sigma(t, X(t))$.

Output: Approximate drift function, $A(t)X(t)+b(t)$, and a diffusion matrix $\bar{\sigma}_1(t)$.

$$A(t) \leftarrow \mathbb{C}[\mu(t, \widehat{X}(t)), \widehat{X}(t)] \mathbb{V}[\widehat{X}(t)]^{-1}$$

$$b(t) \leftarrow \mathbb{E}[\mu(t, \widehat{X}(t))] - A(t) \mathbb{E}[\widehat{X}(t)]$$

$$\bar{\sigma}_1(t) \leftarrow \mathbb{E}[\sigma(t, \widehat{X}(t)) \sigma^T(t, \widehat{X}(t))]^{1/2}$$

Algorithm 2 Statistical Linear Regression II (functional minimisation)

Input: The marginal moment functions of a Gaussian process $\widehat{X}(s)$, $\{(\mathbb{E}[\widehat{X}(s)], \mathbb{V}[\widehat{X}(s)])\}_{s=0}^T$, drift function $\mu(t, X(t))$, and diffusion matrix $\sigma(t, X(t))$.

Output: Approximate drift function, $A(t)X(t)+b(t)$, and a diffusion matrix $\bar{\sigma}_2(t)$.

$$A(t) \leftarrow \mathbb{C}[\mu(t, \widehat{X}(t)), \widehat{X}(t)] \mathbb{V}[\widehat{X}(t)]^{-1}$$

$$b(t) \leftarrow \mathbb{E}[\mu(t, \widehat{X}(t))] - A(t) \mathbb{E}[\widehat{X}(t)]$$

$$\bar{\sigma}_2(t) \leftarrow \mathbb{E}[\sigma(t, \widehat{X}(t))]$$

Remark 1. If $X(\tau) \sim \mathcal{N}(\mathbb{E}[X(\tau)], \mathbb{V}[X(\tau)])$ then Algorithms 1 and 2 can be used to obtain a Gaussian approximation of $X(\tau + \delta)$ by approximating the drift function and diffusion matrix at time τ using the moments of $X(\tau)$. This is the usual procedure in continuous-time Gaussian filtering (cf. [11]), where Algorithm 1 has previously been used implicitly.

3.1. Discrete Time Limit

Here, a short-term variant of the procedure in Algorithm 1 is derived (see Remark 1). Using an Euler-

Maruyama discretisation [27] of Equation (12) gives

$$\tilde{X}(t + \delta) = \tilde{X}(t) + \mu(t, \hat{X}(t))\delta + \sigma(t, \hat{X}(t))\delta\tilde{W}(t), \quad (13)$$

where $\delta\tilde{W}(t) = \tilde{W}(t + \delta) - \tilde{W}(t)$ is a standard Wiener increment of size δ . The statistical linear regression approach then involves forming an affine approximation to $\tilde{X}(t + \delta)$ according to

$$\tilde{X}_a(t + \delta) = \tilde{X}_a(t) + [A(t)\hat{X}(t) + b(t)]\delta + \Xi(t, \delta), \quad (14)$$

where $\Xi(t, \delta)$ is a zero mean random variable with covariance matrix $\Gamma(t, \delta)$ accounting for the error, assumed to be Gaussian (see e.g [25]), and $\tilde{X}_a(t) \triangleq \tilde{X}(t)$. The parameters $A(t)$ and $b(t)$ are then found by minimising the mean squared error of the residual:

$$\mathbb{E} \left[\left\| \tilde{X}(t + \delta) - \tilde{X}_a(t) - [A(t)\hat{X}(t) + b(t)]\delta \right\|^2 \right].$$

This is simply a quadratic optimisation problem and the parameters $A(t)$ and $b(t)$ are thus (c.f [23, 25])

$$\begin{aligned} A(t) &= \mathbb{C}[\mu(t, \hat{X}(t)), \hat{X}(t)]\mathbb{V}[\hat{X}(t)]^{-1}, \\ b(t) &= \mathbb{E}[\mu(t, \hat{X}(t))] - A(t)\mathbb{E}[\hat{X}(t)]. \end{aligned}$$

Furthermore, the residual is given by

$$\tilde{R}(t, \delta) = \tilde{X}(t + \delta) - \tilde{X}_a(t) - [A(t)\hat{X}(t) + b(t)]\delta.$$

Straight-forward calculations gives the moments of $\tilde{R}(t, \delta)$ as

$$\begin{aligned} \mathbb{E}[\tilde{R}(t, \delta)] &= 0, \\ \mathbb{V}[\tilde{R}(t, \delta)] &= \mathbb{E}[\sigma(t, \hat{X}(t))\sigma^\top(t, \hat{X}(t))]\delta + o(\delta), \end{aligned}$$

which are taken to be the moments of $\Xi(t, \delta)$. That is,

$$\Gamma(t, \delta) = \mathbb{E}[\sigma(t, \hat{X}(t))\sigma^\top(t, \hat{X}(t))]\delta + o(\delta). \quad (18)$$

Now define $\bar{\sigma}_1(t) = \mathbb{E}[\sigma(t, \hat{X}(t))\sigma^\top(t, \hat{X}(t))]^{1/2}$. Then the increment of $\tilde{X}_a(t)$ is approximately given by

$$\begin{aligned} \tilde{X}_a(t + \delta) - \tilde{X}_a(t) &= \left(A(t)\tilde{X}_a(t) + b(t) \right)\delta \\ &\quad + \bar{\sigma}_1(t)\delta\widehat{W}(t) + o(\delta), \end{aligned}$$

where $\delta\widehat{W}(t) = \widehat{W}(t + \delta) - \widehat{W}(t)$ is a standard Wiener increment of size δ , independent of $\tilde{X}_a(t)$ and $W(t)$, that matches the variance of the part of $\Xi(t, \delta)$ that does not vanish faster than δ as $\delta \rightarrow 0$. Now, passing to the limit, $\delta \rightarrow 0$, gives the following differential

$$d\tilde{X}_a(t) = \left(A(t)\tilde{X}_a(t) + b(t) \right) dt + \bar{\sigma}_1(t) d\widehat{W}(t),$$

which makes the procedure in Algorithm 1 apparent.

3.2. A Variational Formulation

Another approach for arriving at Equation (11) is by employing variational calculus [26] as follows. Define an approximating process to Equation (12), $\tilde{X}_a(t)$, with $\tilde{X}_a(t_b) = \tilde{X}(t_b)$ and $t_b > 0$, given by

$$d\tilde{X}_a(t) = (A(t)\hat{X}(t) + b(t)) dt + \bar{\sigma}_2(t) d\tilde{W}(t). \quad (19)$$

For $t_e > t_b$, the mean square error between $\tilde{X}_a(t_e)$ and $\tilde{X}(t_e)$ is given by

$$\begin{aligned} &\mathbb{E}[\|\tilde{X}_a(t_e) - \tilde{X}(t_e)\|^2] \\ &= \mathbb{E} \left[\left\| \int_{t_b}^{t_e} A(\tau)\hat{X}(\tau) + b(\tau) - \mu(\tau, \hat{X}(\tau)) d\tau \right\|^2 \right] \\ &\quad + \mathbb{E} \left[\left\| \int_{t_b}^{t_e} \bar{\sigma}_2(\tau) - \sigma(\tau, \hat{X}(\tau)) d\tilde{W}(\tau) \right\|^2 \right] \end{aligned}$$

where the independence between $\hat{X}(t)$ and $\tilde{W}(t)$ was used to eliminate the cross-term. Furthermore, employing Jensen's inequality and Itô isometry gives

$$\begin{aligned} &\mathbb{E}[\|\tilde{X}_a(t_e) - \tilde{X}(t_e)\|^2] \\ &\leq \mathbb{E} \left[\int_{t_b}^{t_e} \|A(\tau)\hat{X}(\tau) + b(\tau) - \mu(\tau, \hat{X}(\tau))\|^2 d\tau \right] \\ &\quad + \mathbb{E} \left[\int_{t_b}^{t_e} \|\bar{\sigma}_2(\tau) - \sigma(\tau, \hat{X}(\tau))\|_F^2 d\tau \right], \end{aligned}$$

where $\|\cdot\|_F$ is the Frobenius norm. Therefore, an appropriate cost functional for fitting A , b , and $\bar{\sigma}_2$ may be defined as

$$\begin{aligned} \mathcal{J}(A, b, \bar{\sigma}_2) &= \frac{1}{2} \int_{t_b}^{t_e} \mathbb{E} \left[\|\bar{\sigma}_2(\tau) - \sigma(\tau, \hat{X}(\tau))\|_F^2 \right] d\tau \\ &\quad + \frac{1}{2} \int_{t_b}^{t_e} \mathbb{E} \left[\|A(\tau)\hat{X}(\tau) + b(\tau) - \mu(\tau, \hat{X}(\tau))\|^2 \right] d\tau \end{aligned} \quad (20)$$

Perturbing A , b , and $\bar{\sigma}_2$ by arbitrary functions ε_b , ε_A , and $\varepsilon_{\bar{\sigma}_2}$, respectively gives

$$\begin{aligned}
& \mathcal{J}(A + \varepsilon_A, b, \bar{\sigma}_2) - \mathcal{J}(A, b, \bar{\sigma}_2) = r_A(\varepsilon_A) \\
& + \int_{t_b}^{t_e} \text{tr}\{\mathbb{E}[\widehat{X}(\tau)(\widehat{X}^\top(\tau)A^\top(\tau) + b^\top(\tau))\varepsilon_A(\tau)]\} d\tau \\
& - \int_{t_b}^{t_e} \text{tr}\{\mathbb{E}[\widehat{X}(\tau)\mu^\top(\tau, \widehat{X}(\tau))\varepsilon_A(\tau)]\} d\tau \\
& \mathcal{J}(A, b + \varepsilon_b, \bar{\sigma}_2) - \mathcal{J}(A, b, \bar{\sigma}_2) = r_b(\varepsilon_b) \\
& + \int_{t_b}^{t_e} \mathbb{E}[A(\tau)\widehat{X}(\tau) + b(\tau) - \mu(\tau, \widehat{X}(\tau))]^\top \varepsilon_b(\tau) d\tau \\
& \mathcal{J}(A, b, \bar{\sigma}_2 + \varepsilon_{\bar{\sigma}_2}) - \mathcal{J}(A, b, \bar{\sigma}_2) = r_{\bar{\sigma}_2}(\varepsilon_{\bar{\sigma}_2}) \\
& + \int_{t_b}^{t_e} \text{tr}\{(\bar{\sigma}_2(\tau) - \mathbb{E}[\sigma(\tau, \widehat{X}(\tau))])\} d\tau,
\end{aligned}$$

where r_A , r_b , and $r_{\bar{\sigma}_2}$ contain the higher order terms of ε_A , ε_b , and $\varepsilon_{\bar{\sigma}_2}$, respectively. The sufficient conditions for minima are given by [26] :

$$\begin{aligned}
0 &= \mathbb{E}[A(\tau)\widehat{X}(\tau) + b(\tau) - \mu(\tau, \widehat{X}(\tau))], \\
0 &= \mathbb{E}[\widehat{X}(\tau)(\widehat{X}^\top(\tau)A^\top(\tau) + b^\top(\tau) - \mu^\top(\tau, \widehat{X}(\tau)))] , \\
0 &= \bar{\sigma}_2(\tau) - \mathbb{E}[\sigma(\tau, \widehat{X}(\tau))].
\end{aligned}$$

Therefore, the minimisers are given by

$$A(\tau) = \mathbb{C}[\mu(\tau, \widehat{X}(\tau)), \widehat{X}(\tau)]\mathbb{V}[\widehat{X}(\tau)]^{-1}, \quad (23a)$$

$$b(\tau) = \mathbb{E}[\mu(\tau, \widehat{X}(\tau))] - A(\tau)\mathbb{E}[\widehat{X}(\tau)], \quad (23b)$$

$$\bar{\sigma}_2(\tau) = \mathbb{E}[\sigma(\tau, \widehat{X}(\tau))]. \quad (23c)$$

Thus the procedure in Algorithm 2 is obtained.

Remark 2. Note that the cost functional in Equation (20) is defined on any interval $[t_b, t_e]$ with $t_e > t_b$. Thus one can take $t_b = 0$ and $t_e = t_K$ to make Algorithm 2 globally defined. This is in contrast to Algorithm 1 which is defined by stitching together local approximations.

3.3. The difference between the discrete-time limit and the variational formulation

While the difference between the discrete-time limit approach (Algorithm 1) and the variational approach (Algorithm 2) is small it is nonetheless interesting

to highlight. The only difference being the diffusion matrices, $\bar{\sigma}_1(t)$ and $\bar{\sigma}_2(t)$ for Algorithm 1 and Algorithm 2, respectively. The following holds.

Proposition 1. Assume the same Gaussian process, $\widehat{X}(t)$, is used to obtain $\bar{\sigma}_1(t)$ and $\bar{\sigma}_2(t)$ then the following inequality holds

$$\text{tr}\{\bar{\sigma}_1(t)\bar{\sigma}_1^\top(t)\} \geq \text{tr}\{\bar{\sigma}_2(t)\bar{\sigma}_2^\top(t)\}. \quad (24)$$

Proof. Plugging in the expressions from Algorithms 1 and 2 gives

$$\begin{aligned}
& \text{tr}\{\mathbb{E}[\sigma(t, \widehat{X}(t))\sigma^\top(t, \widehat{X}(t))]\} \\
& \geq \text{tr}\{\mathbb{E}[\sigma(t, \widehat{X}(t))]\mathbb{E}[\sigma(t, \widehat{X}(t))]^\top\},
\end{aligned} \quad (25)$$

which is Jensen's inequality. \square

The result in Proposition 1 essentially says that $\bar{\sigma}_2$ will be smaller than $\bar{\sigma}_1$, in Frobenius sense. Equality can be retrieved when $\sigma(t, X(t)) = \sigma(t)$ or when $\mathbb{V}[\widehat{X}(t)] \rightarrow 0$. Furthermore, it is clear that approximate implementations of Algorithms 1 and 2, employing Taylor series expansions of the integrand up to first order around $\mathbb{E}[\widehat{X}(t)]$ also give $\bar{\sigma}_1 = \bar{\sigma}_2$.

4. Continuous-Discrete Gaussian Smoothers

The linearisation technique presented in Section 3 allows for the formulation of approximate smoothers to the system Equation (1) by simply plugging in $A(t)$, $b(t)$, and $\bar{Q}_i(t) = \bar{\sigma}_i(t)\bar{\sigma}_i^\top(t)$, $i \in \{1, 2\}$ into the linear smoothing equations in Equation (8). Furthermore, it has not been specified what Gaussian process, $\widehat{X}(t)$ is used to compute $A(t)$, $b(t)$ and $\bar{\sigma}_i(t)$.

In any case, the smoothers developed in this paper rely on a Gaussian approximation to the filtering distribution which can be constructed on the fly during filtering by using Algorithm 1 or Algorithm 2 to compute the linearisation parameters (see Remark 1), which are then given by

$$A(t) = \mathbb{C}[\mu(t, X(t)), X(t) | \mathcal{Y}(t)]\Sigma^{-1}(t), \quad (26a)$$

$$b(t) = \mathbb{E}[\mu(t, X(t)) | \mathcal{Y}(t)] - A(t)\bar{x}(t), \quad (26b)$$

$$\bar{\sigma}_i(t) = \begin{cases} \mathbb{E}[\sigma(t, X(t))\sigma^\top(t, X(t)) | \mathcal{Y}(t)]^{1/2}, & i = 1, \\ \mathbb{E}[\sigma(t, X(t)) | \mathcal{Y}(t)], & i = 2. \end{cases} \quad (26c)$$

Plugging Equation (26) into Equation (7) gives

$$\frac{d\bar{x}(t)}{dt} = \mathbb{E}[\mu(t, X(t)) | \mathcal{Y}(t)], \quad (27a)$$

$$\begin{aligned} \frac{d\Sigma(t)}{dt} &= \bar{Q}_i(t) + \mathbb{C}[\mu(t, X(t)), X(t) | \mathcal{Y}(t)] \\ &\quad + \mathbb{C}[\mu(t, X(t)), X(t) | \mathcal{Y}(t)]^\top, \end{aligned} \quad (27b)$$

which is the prediction equations for the Gaussian filter, the filter update is handled by conventional discrete-time methods, see [4, 28, 25, 21]. Note that Equation (27) gives rise to two kinds of Gaussian filters depending on whether Algorithm 1 or Algorithm 2 are used. These shall be referred to as the Gaussian filter of the first kind and the Gaussian filter of the second kind for $i = 1$ and $i = 2$, respectively, with $i = 1$ giving the classical Gaussian smoother (c.f [11]). This terminology will be used throughout the paper to distinguish between methods derived from Algorithms 1 and 2, respectively.

Since the filtering distribution and the smoothing distribution are equal at $t = t_K$, there are two notable options for deriving approximations to the smoothing distribution, namely (i) linearising with respect to the filtering distribution or (ii) linearising with respect to the smoothing distribution on the fly. In the latter case, the linearisation parameters are given by

$$A(t) = \mathbb{C}[\mu(t, X(t)), X(t) | \mathcal{Y}(t_K)]\Omega^{-1}(t), \quad (28a)$$

$$b(t) = \mathbb{E}[\mu(t, X(t)) | \mathcal{Y}(t_K)] - A(t)\hat{x}(t), \quad (28b)$$

$$\bar{\sigma}_i(t) = \begin{cases} \mathbb{E}[\sigma(t, X(t))\sigma^\top(t, X(t)) | \mathcal{Y}(t_K)]^{1/2}, & i = 1, \\ \mathbb{E}[\sigma(t, X(t)) | \mathcal{Y}(t_K)], & i = 2. \end{cases} \quad (28c)$$

Plugging Equation (28) into the linear smoothing equations Equation (8) gives

$$\begin{aligned} \frac{d\hat{x}(t)}{dt} &= \mathbb{E}[\mu(t, X(t)) | \mathcal{Y}(t_K)] \\ &\quad + \bar{Q}_i(t)\Sigma^{-1}(t)[\hat{x}(t) - \bar{x}(t)], \end{aligned} \quad (29a)$$

$$\begin{aligned} \frac{d\Omega(t)}{dt} &= \bar{Q}_i(t)\Sigma^{-1}(t)\Omega(t) + \Omega(t)\Sigma^{-1}(t)\bar{Q}_i(t) - \bar{Q}_i(t) \\ &\quad + \mathbb{C}[\mu(t, X(t)), X(t) | \mathcal{Y}(t_K)] \\ &\quad + \mathbb{C}[\mu(t, X(t)), X(t) | \mathcal{Y}(t_K)]^\top. \end{aligned} \quad (29b)$$

The smoother in Equation (29) are referred to as Type I* of the first kind and Type I* of the second kind for $i = 1$ and $i = 2$, respectively. This, because their similarity to the Type I smoother of [11], their connection is elaborated on in Proposition 2.

Proposition 2. *Let $X(t)$ and $\{Y(t_k)\}_{k=1}^K$ be governed by the system in Equation (1) and assume $\sigma(t, X(t)) = \sigma(t)$. Then the smoothing equations for the Type I [11] and Type I* smoothers of the first and second kind agree.*

Proof. First note that $\bar{Q}_1 = \bar{Q}_2$ since the diffusion is state independent. The statement then follows by direct comparison to [11, Equation (27)]. \square

When the linearisation is done with respect to the filtering distribution the smoothing moments are retrieved by plugging Equation (26) into Equation (8)

$$\begin{aligned} \frac{d\hat{x}(t)}{dt} &= \mathbb{C}[\mu(t, X(t)), X(t) | \mathcal{Y}(t)] \\ &\quad \times \Sigma^{-1}(t)(\hat{x}(t) - \bar{x}(t)) \\ &\quad + \bar{Q}_i(t)\Sigma^{-1}(t)(\hat{x}(t) - \bar{x}(t)) \\ &\quad + \mathbb{E}[\mu(t, X(t)) | \mathcal{Y}(t)], \end{aligned} \quad (30a)$$

$$\begin{aligned} \frac{d\Omega(t)}{dt} &= \left(\mathbb{C}[\mu(t, X(t)), X(t) | \mathcal{Y}(t)] + \bar{Q}_i(t) \right) \\ &\quad \times \Sigma^{-1}(t)\Omega(t) \\ &\quad + \Omega(t)\Sigma^{-1}(t) \\ &\quad \times \left(\mathbb{C}[\mu(t, X(t)), X(t) | \mathcal{Y}(t)]^\top + \bar{Q}_i(t) \right) \\ &\quad - \bar{Q}_i(t), \end{aligned} \quad (30b)$$

which corresponds to the Type II smoother of [11] when $i = 1$ and when $i = 2$ another smoother is obtained. These shall, again, be referred to as Type II of the first kind and Type II of the second kind for $i = 1$ and $i = 2$, respectively. Furthermore, by the same argument as in [11], the Type II formulation may be converted to a Type III formulation, where only forward-time ODEs need to be solved. This argument is not repeated here, but the result is simply (see [11])

$$\frac{d\bar{x}(t)}{dt} = A(t)\bar{x}(t) + b(t), \quad (31a)$$

$$\frac{d\Sigma(t)}{dt} = A(t)\Sigma(t) + \Sigma(t)A^\top(t) + \bar{Q}_i(t), \quad (31b)$$

$$\frac{dH_l(t)}{dt} = H_l(t)A^\top(t), \quad (31c)$$

$$G_{l+1} = H_l(t_{l+1}^-)\Sigma^{-1}(t_{l+1}^-), \quad (31d)$$

$$\hat{x}(t_l) = \bar{x}(t_l) + G_{l+1}\left(\hat{x}(t_{l+1}) - \bar{x}(t_{l+1}^-)\right), \quad (31e)$$

$$\Omega(t_l) = G_{l+1}\left(\Omega(t_{l+1}) - \Sigma(t_{l+1}^-)\right)G_{l+1}^\top + \Sigma(t_l). \quad (31f)$$

Remark 3. *The Type II and III smoothers of the first kind are precisely the Type II and Type III smoothers as described by [11].*

To conclude this section we note that Proposition 1 indicates that the diffusion term in the smoothers of the first kind will be larger than the diffusion term for the smoothers of the second kind. This leads to the expectation that the smoothers of the first kind will report a larger uncertainty than those of the second kind.

5. Continuous-Discrete Iterative Gaussian Smoothers

In this section, the linearisation techniques of Section 3 are combined with the Type I*/II/III smoothers (of the first and second kind) of Section 4 to develop iterative Gaussian smoothers. This is done in an analogous manner to the discrete-time iterative smoothers [20, 21, 22]. The basic idea is that given a Gaussian process, $\hat{X}(t)$, Algorithm 1 or Algorithm 2 can readily be applied to the system in Equations (1a) and (1b) (using standard statistical linear regression for the measurement equation [25]), which yields an approximate affine system for which inference is straight-forward. An iterative scheme is then obtained by alternating between linearisation and Gaussian smoothing, where $\hat{X}(t)$ is always chosen as the current best approximation to the smoothing process. This defines an iterative scheme reminiscent of the Gauss-Newton method [22].

5.1. Iterative Smoothers

Let $\{\hat{X}^{(j)}(s)\}_{s \geq 0}$ be a Gaussian process approximating the smoothed process at iteration j , with moment

functions, $\hat{x}^{(j)}(t)$ and $\Omega^{(j)}(t)$. Moreover, for a test function, $\psi(X)$, denote the expectation of $\psi(\hat{X}^{(j)}(t))$ by $\mathbb{E}^{(j)}[\psi(\hat{X}^{(j)}(t))]$. The linearisation parameters, $A^{(j)}$, $b^{(j)}$, and $\bar{\sigma}_i^{(j)}(t)$ can then be obtained by either using Algorithm 1 or Algorithm 2 and the linearisation of the measurement model is given by [20, 21, 25]

$$\begin{aligned} C^{(j)}(t_k) &= \mathbb{C}^{(j)}[h(t_k, \hat{X}(t_k)), \hat{X}(t_k)](\Omega^{(j)}(t_k))^{-1}, \\ d^{(j)}(t_k) &= \mathbb{E}^{(j)}[h(t_k, \hat{X}(t_k))] - C^{(j)}(t_k)\hat{x}^{(j)}(t_k), \\ \Delta^{(j)}(t_k) &= \mathbb{V}^{(j)}[h(t_k, X(t))] + R, \\ &\quad - C^{(j)}(t_k)\Omega^{(j)}(t_k)(C^{(j)}(t_k))^\top, \end{aligned}$$

where $\Delta^{(j)}(t_k)$ is the variance of the residual at iteration j . The approximate smoothed process at iteration $j+1$ is then obtained by considering the system:

$$\begin{aligned} dX(t) &= \left(A^{(j)}(t)X(t) + b^{(j)}(t)\right) dt \\ &\quad + \bar{\sigma}^{(j)}(t) d\widehat{W}^{(j)}(t), \end{aligned} \quad (33a)$$

$$Y(t_k) = C^{(j)}(t_k)X(t_k) + d^{(j)}(t_k) + \widehat{V}^{(j)}(t_k), \quad (33b)$$

$$\begin{aligned} \mathbb{C}[\widehat{V}^{(j)}(t_k), \widehat{V}^{(j)}(t_l)] &= \delta_{k,l}\Delta^{(j)}(t_k), \\ \widehat{V}^{(j)}(t_k) &\sim \mathcal{N}(0, \Delta^{(j)}(t_k)), \end{aligned} \quad (33c)$$

where $\widehat{W}^{(j)}(t)$ is a standard Wiener process. An approximation to the filtered process at iteration $j+1$, $\{\bar{X}^{(j+1)}(s)\}_{s \geq 0}$, is then obtained by using the linear filter defined by Equations (6) and (7), after which any of the smoother formulations Equations (29) to (31) may be used to obtain $\{\hat{X}^{(j+1)}(s)\}_{s \geq 0}$.

Remark 4. *In practice, Algorithms 1 and 2 can not be implemented in closed form. Standard approaches to approximate expectations with respect to a Gaussian density is by first order Taylor series or sigma-points [4]. If the first order Taylor series method is used together with Algorithm 1 or Algorithm 2 then a continuous-time iterated extended Kalman smoother is obtained.*

5.2. Fixed Point Characterisation

A convergence analysis of the proposed iteration scheme is beyond the scope of this paper. However, for the smoothers of the first kind, one can discretise

the system, apply the analysis of the discrete time case [20, 21], and assume the limits $J \rightarrow \infty$ and $\delta t \rightarrow 0$ can be interchanged, in which case convergence is guaranteed if the iterations are initialised sufficiently close to a fix point.

Another topic of investigation is the relationship between the different types of smoothers at the fixed point. More specifically, the relationship between the Type II and Type I* smoother is illuminated. The smoothing moments for the Type II smoother at iteration $j + 1$ are given by

$$\begin{aligned} \frac{d\hat{x}^{(j+1)}(t)}{dt} &= \mathbb{E}^{(j)}[\mu(t, X(t)) \mid \mathcal{Y}(t)] \\ &+ \bar{Q}_i^{(j)}(t)[\Sigma^{(j+1)}(t)]^{-1}(\hat{x}^{(j+1)}(t) - \bar{x}^{(j+1)}(t)) \\ &+ \mathbb{C}^{(j)}[\mu(t, X(t)), X(t)][\Sigma^{(j+1)}(t)]^{-1} \\ &\times (\hat{x}^{(j+1)}(t) - \bar{x}^{(j+1)}(t)) \\ \frac{d\Omega^{(j+1)}(t)}{dt} &= -\bar{Q}_i^{(j)}(t) \\ &+ \left(\mathbb{C}^{(j)}[\mu(t, X(t)), X(t)] + \bar{Q}_i^{(j)}(t) \right) \\ &\times [\Sigma^{(j+1)}(t)]^{-1}\Omega^{(j+1)}(t) \\ &+ \Omega^{(j+1)}(t)[\Sigma^{(j+1)}(t)]^{-1} \\ &\times \left(\mathbb{C}^{(j)}[\mu(t, X(t)), X(t)]^\top + \bar{Q}_i^{(j)}(t) \right). \end{aligned}$$

Proposition 3. *The Type I*, Type II, and Type III smoothers of first and second kinds are equivalent at the fixed point, respectively. That is, they converge to the same point.*

Proof. Assume $(\hat{x}^{(j)}, \Omega^{(j)})$ is a fixed point of the iteration and iterate once again. That is, insert $A^{(j)}$, $b^{(j)}$ and $\bar{Q}_i^{(j)}$ into Equation (30) to obtain

$$\begin{aligned} \frac{d\hat{x}^{(j+1)}(t)}{dt} &= \mathbb{E}^{(j)}[\mu(t, X(t))] \\ &+ \bar{Q}_i^{(j)}(t)[\Sigma^{(j+1)}(t)]^{-1}(\hat{x}^{(j+1)}(t) - \bar{x}^{(j+1)}(t)) \\ &+ A^{(j)}(t)(\hat{x}^{(j+1)}(t) - \hat{x}^{(j)}(t) - \bar{x}^{(j+1)}(t) + \bar{x}^{(j)}(t)), \\ \frac{d\Omega^{(j+1)}(t)}{dt} &= A^{(j)}(t)\Omega^{(j+1)}(t) + \Omega^{(j+1)}(t)(A^{(j)}(t))^\top \\ &+ \bar{Q}_i^{(j)}(t)[\Sigma^{(j+1)}(t)]^{-1}\Omega^{(j+1)}(t) \\ &+ \Omega^{(j+1)}(t)[\Sigma^{(j+1)}(t)]^{-1}\bar{Q}_i^{(j)}(t) - \bar{Q}_i^{(j)}(t). \end{aligned}$$

Now, plugging in the definition of $A^{(j)}$ and using the fact that $\bar{x}^{(j+1)} = \bar{x}^{(j)}$, $\Sigma^{(j+1)} = \Sigma^{(j)}$, $\hat{x}^{(j+1)} = \hat{x}^{(j)}$, and $\Omega^{(j+1)} = \Omega^{(j)}$, since $(\hat{x}^{(j)}, \Omega^{(j)})$ is a fixed point, gives the following

$$\begin{aligned} \frac{d\hat{x}^{(j)}(t)}{dt} &= \mathbb{E}^{(j)}[\mu(t, X(t))] \\ &+ \bar{Q}_i^{(j)}(t)[\Sigma^{(j)}(t)]^{-1}(\hat{x}^{(j)}(t) - \bar{x}^{(j)}(t)), \\ \frac{d\Omega^{(j)}(t)}{dt} &= -\bar{Q}_i^{(j)}(t) + \mathbb{C}^{(j)}[\mu(t, X(t)), X(t)] \\ &+ \mathbb{C}^{(j)}[\mu(t, X(t)), X(t)]^\top \\ &+ \bar{Q}_i^{(j)}(t)[\Sigma^{(j)}(t)]^{-1}\Omega^{(j)}(t) \\ &+ \Omega^{(j)}(t)[\Sigma^{(j)}(t)]^{-1}\bar{Q}_i^{(j)}(t), \end{aligned}$$

which is the differential equations satisfied by a Type I* smoother (see Equation (29)). Since Type III is equivalent to Type II, all the presented smoothers (of the same kind) satisfy the same differential equation at the fixed point. \square

5.3. Computational Complexity and Storage Requirement

It is important to consider the computational complexity and storage requirement of the different types of iterative smoothers. If the time interval, for purposes of numerical solving the ODEs, is sub-divided into N time stamps and K measurements are processed, then for the non-iterative smoothers it was found that Type III is superior to Type I and II in terms of storage requirement, while being comparable in the number of Gaussian integrals needed [11].

However, for the iterative schemes the storage requirements for Type I* and Type II smoothers are doubled due to having to store the smoothing solution of the previous iteration. The change for Type III smoother is more dramatic since the linearisation requires the storage of the smoothing solution of the previous iteration at all of the N time stamp. The computational requirements for the smoothers using J iterations are given in Table 1. Therefore there is no significant difference in computational requirements once iterations are introduced.

Table 1: Computational requirements for the iterative smoothers (of any kind).

Smoothers	Integrals	Storage
Type I*	10NKJ	$2NK(d_X + d_X^2)$
Type II	3NKJ	$2NK(2d_X + 3d_X^2)$
Type III	3NKJ	$2NK(2d_X + 3d_X^2)$

6. Experimental Results

6.1. Reentry

The proposed iterative Gaussian smoother is compared to the variational smoother of [17] in a reentry tracking problem. The state, $U = [X, Y, \dot{X}, \dot{Y}, \Psi]^T$, represents the position (X, Y) , velocity (\dot{X}, \dot{Y}) , and an aerodynamic parameter, Ψ of a vehicle. The dynamic equation is given by

$$dU(t) = \begin{bmatrix} 0 & \mathbf{I}_2 & 0 \\ G(t, U(t))\mathbf{I}_2 & D(t, U(t))\mathbf{I}_2 & 0 \\ 0 & 0 & 0 \end{bmatrix} U(t) dt + [0 \quad \mathbf{I}_3]^T \sigma dW(t),$$

where \mathbf{I}_p is a $p \times p$ identity matrix and the zero entries are zero matrices of appropriate sizes. The functions $G(t, u)$ and $D(t, u)$ are given by

$$G(t, u) = -\frac{Gm_0}{(x^2 + y^2)^{3/2}},$$

$$D(t, u) = -\beta_0 \exp\left(+ \frac{R_0 - (x^2 + y^2)^{1/2}}{H_0} \right) \times (\dot{x}^2 + \dot{y}^2)^{1/2}.$$

The parameters were set to

$$\sigma = \text{diag}[\sqrt{2.4064} \cdot 10^{-5/2}, \sqrt{2.4064} \cdot 10^{-5/2}, 1 \cdot 10^{-3}],$$

$\beta_0 = -0.59783$, $H_0 = 13.406$, $Gm_0 = 3.9860 \cdot 10^5$, and $R_0 = 6374$. The vehicle is measured once per second by a radar at position $[s_x, s_y]^T$ according to

$$Z(t_k) = \begin{bmatrix} [(X(t_k) - s_x)^2 + (Y(t_k) - s_y)^2]^{1/2} \\ \tan^{-1} \left(\frac{Y(t_k) - s_y}{X(t_k) - s_x} \right) \end{bmatrix} + V(t_k),$$

where $V(t_k)$ is a Gaussian white noise sequence with covariance matrix

$$R = \text{diag}[1 \cdot 10^{-3}, 1.7 \cdot 10^{-3}].$$

The initial state, $X(0)$ is Gaussian distributed with moments

$$\bar{x}(0) = [6500.4, 349.14, -1.8093, -6.7967, 0.6932],$$

$$\Sigma(0) = \begin{bmatrix} \mathbf{I}_4 \cdot 10^{-6} & 0 \\ 0 & 1 \end{bmatrix}.$$

The system was simulated 100 times on the interval $t \in [0, 200]$ using the Euler-Maruyama method with a step-size of $1/1000$. The Type III template (Equation (31)) is used for the implementation of the proposed iterative Gaussian smoother,¹ the initial linearisation being with respect to the filtering distributions, with up to 4 subsequent iterations. The ODEs are approximated by constant input between discretisation instants, that is zeroth order hold whereby the equivalent discrete time system is computed using the matrix fraction decomposition (see, e.g, [29]). The performance is compared to the variational smoother of [17]², which uses the standard fourth order Runge-Kutta method for integration, the same expectation approximator, and iterates until the change in Kullback-Leibler divergence is less than 10^{-3} , the adaptive step-size goes below the threshold 10^{-3} , or 20 iterations have been performed. Both smoothers use a step-size of $\delta t = 1/100$ for time integration and the spherical-radial cubature rule [30] to approximate expectations.

Boxplots for the RMSE in position, velocity, and the aerodynamic parameter are shown in Figure 1 for iterations 0 through 4 of the proposed smoother and the variational smoother at convergence. The proposed smoother converges after a few iterations have similar performance to the variational smoother.

The χ^2 -statistic (also known as NEES [31]), averaged over Monte Carlo trials, for the various smoothers is shown in Figure 2. Clearly the initialisation is

¹Note that since the diffusion is state independent the iterative smoothers of the first and second kind are equivalent, see Proposition 2.

²Since the diffusion is constant, the smoothers of [17] and [18] are equivalent.

Table 2: The RMSE in position (POS) (m), velocity (VEL) (m/s), aerodynamic parameter (Ψ) averaged over the Monte Carlo trials, and the average χ^2 -statistic, for the variational smoother at convergence (VB) and iterations 0 through 4 of the proposed iterative Gaussian smoother.

Method	POS	VEL	Ψ	χ^2
0	0.3651	0.0132	0.0208	6.5556
1	0.2968	0.0123	0.0138	4.5173
2	0.2967	0.0123	0.0138	4.4565
3	0.2967	0.0123	0.0138	4.4565
4	0.2967	0.0123	0.0138	4.4565
VB	0.2988	0.0124	0.0142	4.9332

inconsistent. Interestingly, the iterations appear to fall below the lower confidence band, which may be indicative of an overestimated covariance. In contrast, the variational smoother tends to have a slightly larger χ^2 -statistic on average, while still mostly keeping itself within the confidence band.

Lastly, averaged performance metrics are reported in Table 2. Here it can be seen that the proposed smoother again converges rapidly, after 1-2 iterations. The variational smoother took an average of 6.5 iterations to converge, though the convergence criterion was rather strict so it can not be excluded that it can use fewer iterations without making a significant sacrifice in performance.

6.2. Radar Tracked Coordinated Turn

The proposed iterative smoothers are assessed in the radar tracked three dimensional coordinated turn model with state dependent diffusion (see [11]). The latent process, $U = (X, Y, Z, \dot{X}, \dot{Y}, \dot{Z}, \Psi)$, is given by

$$\begin{aligned} \mu^\top(u) &= [\dot{x}, \dot{y}, \dot{z}, -\psi\dot{y}, \psi\dot{x}, 0, 0], \\ dU(t) &= \mu(U(t))dt + \sigma(U(t))dW(t), \end{aligned} \quad (39)$$

where (X, Y, Z) are the position coordinates, $(\dot{X}, \dot{Y}, \dot{Z})$ the corresponding velocities, Ψ is the turn rate, $W(t)$

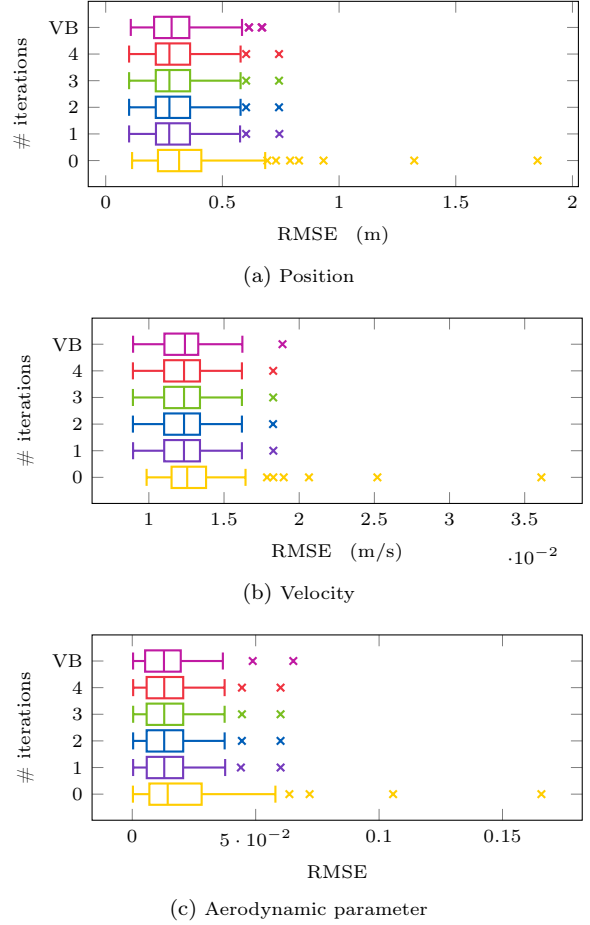


Figure 1: Boxplots of the RMSE distributions over the 100 Monte Carlo trajectories for position (Top), velocity (middle) and, aerodynamic parameter (bottom), for the variational smoother at convergence (VB) and iterations 0 through 4 of the proposed iterative Gaussian smoother (0-4).

is a 4-dimensional Brownian motion, and

$$\begin{aligned} \sigma(u) &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \frac{\dot{x}}{\xi(u)} & \frac{\dot{y}}{\eta(u)} & \frac{\dot{x}\dot{z}}{\xi(u)\eta(u)} & 0 \\ \frac{\dot{y}}{\xi(u)} & -\frac{\dot{x}}{\eta(u)} & \frac{\dot{y}\dot{z}}{\xi(u)\eta(u)} & 0 \\ \frac{\dot{z}}{\xi(u)} & 0 & -\frac{\eta(u)}{\xi(u)} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\ &\times \begin{bmatrix} \sigma_{\parallel} & 0 & 0 & 0 \\ 0 & \sigma_h & 0 & 0 \\ 0 & 0 & \sigma_v & 0 \\ 0 & 0 & 0 & \sigma_{\Psi} \end{bmatrix}, \end{aligned}$$

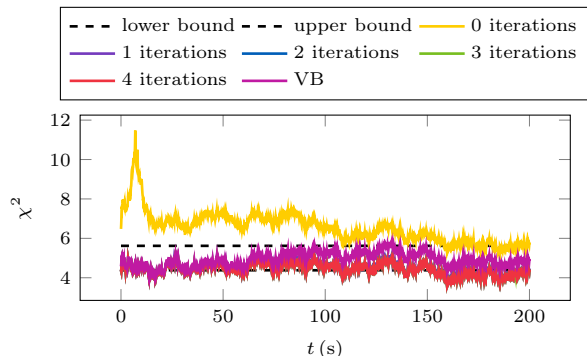


Figure 2: The χ^2 -statistic at each time point, averaged over the Monte Carlo trials with 95% confidence bands for the variational smoother (VB) and iterations 0 through 4 of the proposed iterative Gaussian smoother (0-4).

where

$$\xi(u) = \sqrt{x^2 + y^2 + z^2}, \quad \eta(u) = \sqrt{\dot{x}^2 + \dot{y}^2}.$$

The system is measured according to

$$\begin{bmatrix} \Lambda(t_k) \\ \Theta(t_k) \\ \Phi(t_k) \end{bmatrix} = \begin{bmatrix} \sqrt{X^2(t_k) + Y^2(t_k) + Z^2(t_k)} \\ \tan^{-1}[Y(t_k)/X(t_k)] \\ \tan^{-1} \frac{Z(t_k)}{\sqrt{X^2(t_k) + Y^2(t_k)}} \end{bmatrix} + V(t_k),$$

$$V(t_k) \sim \mathcal{N}\left(0, \text{diag} \left[\sigma_\Lambda^2 \quad \sigma_\Theta^2 \quad \sigma_\Phi^2 \right] \right). \quad (40)$$

The parameters were set as follows, $\sigma_\parallel = \sqrt{100}$, $\sigma_h = \sqrt{0.2}$, $\sigma_v = \sqrt{0.2}$, $\sigma_\Psi = 7 \cdot 10^{-3} \text{rad/s}$, $\sigma_\Lambda = 50 \text{m}$, $\sigma_\Theta = \sigma_\Phi = 0.1\pi/180 \text{rad}$. The statistics of the initial state was set to $U(0) \sim \mathcal{N}(\bar{x}(0^-), \Sigma(0^-))$, where

$$\bar{x}(0^-) = \begin{bmatrix} 1000 & 0 & 2650 & 200 & 0 & 150 & \frac{6\pi}{180} \end{bmatrix}^\top,$$

$$\Sigma(0^-) = 100^2 \text{diag} \left[1, 1, 1, 1, 1, 1, \frac{\pi}{180 \cdot 100^2} \right].$$

It should be noted that the diffusion term in Equation (39) is both singular and state-dependent, hence there exists no Gaussian process with respect to which the probability law of $U(t)$ is absolutely continuous and consequently none of variational smoothers by [16, 17] are applicable. Moreover, as mentioned, the method of [18] requires solving a 7-dimensional

stochastic partial differential equation, which is computationally unattractive. Similarly, for expectation propagation [19] it is not clear how to form good approximations of the expectation with respect to the tilted distribution with a state dimension of 7 and likelihoods in Equation (40). However, this provides a good opportunity to compare the iterative smoothers of the first and second kind (K1 and K2, respectively).

The Euler-Maruyama method was used to generate 100 independent realisations of the system using a step-size of $5/1000$, with the time between measurements set to $\Delta T = 6$, with 26 measurement instants in total, starting from $t = 0$. Both smoothers were implemented in the same manner as the previous experiment, using a step-size of $\delta t = 5/100$.

Boxplots of the root-mean-square error (RMSE) distribution over the 100 Monte-Carlo trajectories is provided in Figure 3 for position, velocity, and turn-rate, respectively. It is clear that iterations can offer a substantial improvement in accuracy. The consistency of the iterative smoothers is assessed using the χ^2 -statistic, which is averaged over the Monte Carlo trajectories and the resulting time series for iterations 0 through 4 is provided in Figure 4. It can be seen that the initialisation of the smoothers is grossly inconsistent and the first iteration provides a massive improvement, while the subsequent iterations provide smaller improvements. Furthermore, the average RMSE and the average χ^2 statistics for the different iterations is shown in Table 3.

The impression is that the smoother converges rather quickly, after two to three iterations in this scenario. Also the iterative smoothers K1 and K2 appear to perform similarly on this problem, with no discernible difference on average for up to 4 significant digits at convergence. However, the iterative smoother of the second kind perform notably worse at initialisation, particularly in terms of consistency, see Table 3. Drawing from the discussion in Section 3.3, we know that both smoothers will be approximately equivalent when the variance in the linearising process is small. As the system noises are fairly small in this experiment this eventually becomes the case as the iterations proceed.

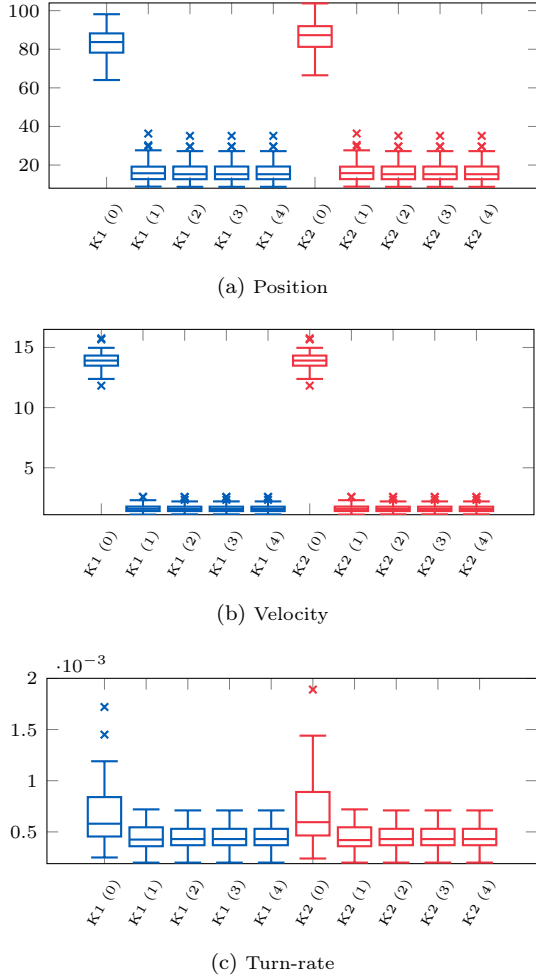




Figure 3: Boxplots of the RMSE distributions over the 100 Monte Carlo trajectories for position (top), velocity (middle), and turn-rate (bottom), for the iterative Gaussian smoothers K1 and K2. The number of iterations after initialisation is shown in parenthesis.

7. Conclusion

The statistical linear regression method was generalised to obtain linear approximations to non-linear SDEs. This allowed for alternate derivation of the Type II and III smoothers [11] for systems with state independent diffusion. It also led to the derivation of the novel Type I* smoother that coincides with the Type I smoother of [11] for state independent diffusions. Furthermore, this linearisation technique



(a) χ^2 -statistics for iterations 1 through 4 for both kinds.



(b) χ^2 -statistics for initialisation for both kinds.

Figure 4: The χ^2 -statistic at each time point, averaged over the Monte Carlo trials with 95% confidence bands for the iterative Gaussian smoothers K1 and K2. The number of iterations after initialisation is shown in parenthesis.

was used to develop a continuous-discrete analogue to the iterated Gaussian smoothers [20, 21, 22]. The method was found to offer considerable improvements in two challenging and high-dimensional target tracking scenarios, being competitive to the variational smoother [17].

Acknowledgment

Financial support by the Academy of Finland through grant #313708 and Aalto ELEC Doctoral School is acknowledged. The authors would also like to thank Juha Ala-Luhtala for sharing his variational smoother codes.

[1] D. H. Titterton, J. L. Weston, Strapdown In-

Table 3: The RMSE in position (POS) (m), velocity (VEL) (m/s), turn-rate (Ψ) (10^{-3} rad/s) averaged over Monte Carlo trials, and the average χ^2 -statistic, for the iterative Gaussian smoothers K1 and K2. The number of iterations after initialising is shown in parenthesis.

Iterations	POS	VEL	Ψ	χ^2
K1 (0)	82.88	13.91	0.658	299.0
K1 (1)	16.52	1.616	0.444	6.906
K1 (2)	16.44	1.611	0.445	6.750
K1 (3)	16.44	1.611	0.445	6.750
K1 (4)	16.44	1.611	0.445	6.750
K2 (0)	86.39	14.31	0.688	458.6
K2 (1)	16.53	1.618	0.445	6.943
K2 (2)	16.44	1.611	0.445	6.750
K2 (3)	16.44	1.611	0.445	6.750
K2 (4)	16.44	1.611	0.445	6.750

ertial Navigation Technology, The Institute of Electrical Engineers, 2004.

- [2] J. L. Crassidis, J. L. Junkins, *Optimal Estimation of Dynamic Systems*, Chapman & Hall/CRC, 2004.
- [3] E. Lindström, H. Madsen, J. N. Nielsen, *Statistics for Finance*, Chapman and Hall/CRC, 2015.
- [4] S. Särkkä, *Bayesian Filtering and Smoothing*, Institute of Mathematical Statistics Textbooks, Cambridge University Press, 2013.
- [5] B. Øksendal, *Stochastic Differential Equations - An Introduction with Applications*, Springer, 2003.
- [6] R. Kalman, R. Bucy, New results in linear filtering and prediction theory, *Transactions of the ASME, Journal of Basic Engineering* 83 (1961) 95–108.
- [7] H. Rauch, F. Tung, C. Striebel, Maximum likelihood estimates of linear dynamic systems, *AIAA Journal* 3 (8) (1965) 1445–1450.
- [8] C. T. Striebel, Partial Differential Equations for the Conditional Distribution of a Markov Process Given Noisy Observations, *Journal of mathematical analysis and applications* 11 (1965) 151–159.
- [9] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*, Academic Press, 1970.
- [10] S. Särkkä, Continuous-time and continuous-discrete-time unscented Rauch-Tung-Striebel smoothers, *Signal Processing* 90 (2010) 225–235.
- [11] S. Särkkä, J. Sarmavuori, Gaussian filtering and smoothing for continuous-discrete dynamic systems, *Signal Processing* 93 (2013) 500–510.
- [12] C. T. Leondes, J. B. Peller, E. B. Stear, Nonlinear smoothing theory, *IEEE Transactions on system science and cybernetics* 6 (1) (1970) 63–71.
- [13] S. Koyama, Projection smoothing for continuous and continuous-discrete stochastic dynamic systems, *Signal Processing* 144 (2018) 333–340.
- [14] D. Brigo, B. Hanzon, F. LeGland, A differential geometric approach to nonlinear filtering: the projection filter, *IEEE Transactions on Automatic Control* 43 (2) (1998) 247–252.
- [15] D. Brigo, B. Hanzon, F. Le Gland, et al., Approximate nonlinear filtering by projection on exponential manifolds of densities, *Bernoulli* 5 (3) (1999) 495–534.
- [16] C. Archambeau, M. Opper, Y. Shen, D. Cornford, J. S. Shawe-taylor, Variational Inference for Diffusion Processes, in: J. C. Platt, D. Koller, Y. Singer, S. T. Roweis (Eds.), *Advances in Neural Information Processing Systems 20*, Curran Associates, Inc., 17–24, 2008.
- [17] J. Ala-Luhtala, S. Särkkä, R. Piché, Gaussian filtering and variational approximations for Bayesian smoothing in continuous-discrete stochastic dynamic systems, *Signal Processing* 111 (Supplement C) (2015) 124 – 136.
- [18] T. Sutter, A. Ganguly, H. Koepl, A variational approach to path estimation and parameter inference of hidden diffusion processes, *The Journal of Machine Learning Research* 17 (1) (2016) 6544–6580.

- [19] B. Cseke, D. Schnoerr, M. Opper, G. Sanguinetti, Expectation propagation for continuous time stochastic processes, *Journal of Physics A: Mathematical and Theoretical* 49 (49) (2016) 494002.
- [20] Á. F. García-Fernández, L. Svensson, S. Särkkä, Iterated posterior linearisation smoother, *IEEE Transactions on Automatic Control* 62 (4) (2017) 2056–2063.
- [21] F. Tronarp, A. F. Garcia-Fernandez, S. Särkkä, Iterative Filtering and Smoothing In Non-Linear and Non-Gaussian Systems Using Conditional Moments, *IEEE Signal Processing Letters* 25 (3) (2018) 408–412.
- [22] B. M. Bell, The iterated Kalman smoother as a Gauss–Newton method, *SIAM Journal on Optimization* 4 (3) (1994) 626–636.
- [23] T. Lefebvre, H. Bruyninckx, J. De Schuller, Comment on “A new method for the nonlinear transformation of means and covariances in filters and estimators” [with authors’ reply], *IEEE Transactions on Automatic Control* 47 (8) (2002) 1406–1409.
- [24] G. Kallianpur, C. Striebel, Estimation of stochastic systems: arbitrary system process with additive white noise observation errors, *Annals of Mathematical Statistics* 39 (3) (1968) 785–801.
- [25] Á. F. García-Fernández, L. Svensson, M. R. Morelande, S. Särkkä, Posterior Linearization Filter: Principles and Implementation Using Sigma Points, *IEEE Transactions on Signal Processing* 63 (20) (2015) 5561–5573.
- [26] R. Weinstock, *Calculus of Variations*, Dover Publications, 1974.
- [27] P. E. Kloeden, E. Platen, *Numerical Solutions of Stochastic Differential Equations*, Springer-Verlag Berlin Heidelberg, 1999.
- [28] B. M. Bell, F. W. Cathey, The iterated Kalman filter update as a Gauss–Newton method, *IEEE Transaction on Automatic Control* 38 (2) (1993) 294–297.
- [29] P. Axelsson, F. Gustafsson, Discrete-time solutions to the continuous-time differential Lyapunov equation with applications to Kalman filtering, *IEEE Transactions on Automatic Control* 60 (3) (2015) 632–643.
- [30] I. Arasaratnam, S. Haykin, Cubature Kalman filters, *IEEE Transactions on Automatic Control* 54 (6) (2009) 1254–1269.
- [31] Y. Bar-Shalom, X. R. Li, T. Kirubarajan, *Estimation with applications to tracking and navigation: theory algorithms and software*, John Wiley & Sons, 2001.