Automatica 80 (2017) 54–61

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Complete offline tuning of the unscented Kalman filter☆

Leonardo Azevedo Scardua [a,b], José Jaime da Cruz [b]

[a] Federal Institute of Espírito Santo, Serra, ES, Brazil

[b] Automation and Control Laboratory, University of São Paulo, São Paulo, SP, Brazil

ARTICLE INFO

Article history:
Received 1 June 2015
Received in revised form
20 November 2016
Accepted 16 January 2017

Keywords:
State estimation
Nonlinear filtering
Optimization
Unscented Kalman filtering

ABSTRACT

The unscented Kalman filter (UKF) is a widely used nonlinear Gaussian filter. It has the potential to deal with highly nonlinear dynamic systems, while displaying computational cost of the same order of magnitude as that of the extended Kalman filter (EKF). The quality of the estimates produced by the UKF is dependent on the tuning of both the parameters that govern the unscented transform (UT) and the two noise covariance matrices of the system model. In this paper, the tuning of the UKF is framed as an optimization problem. The tuning problem is solved by a new stochastic search algorithm and by a standard model-based optimizer. The filters tuned with the proposed algorithm and with the standard model-based optimizer are numerically tested against other nonlinear Gaussian filters, including two UKF tuned with state-of-the-art tuning strategies. One of these strategies relies on online tuning and the other on offline tuning.

1. Introduction

The UKF (Julier & Uhlmann, 1997) and the extended Kalman filter (EKF) (Maybeck, 1979) are two widely used approximate solutions to the problem of filtering nonlinear dynamical systems. The UKF has the potential to achieve better estimation performance than the EKF, while displaying the same order of computational complexity (Wan & van der Merwe, 2000). The heart of the UKF is the (scaled) unscented transform (UT) (Julier & Industries, 2002), which is aimed at propagating the mean and covariance of a random variable that undergoes a nonlinear transformation. For the UKF to work properly, it is thus necessary to tune the three scalar parameters of the UT. The state estimation performance of the UKF is also influenced by the tuning of the two noise covariance matrices of the system model used by the filter, since these noises can be used to model uncertainties about the true dynamic system (see Chapter 1 of Sarkka, 2013). The fact that the tuning of this important filter is such a difficult problem has motivated researchers to devise methods aimed at either tuning solely the UT, solely the noise covariance matrices, or both the UT and the noise covariance matrices. Let us now briefly review some of these methods.

In Julier, Uhlmann, and Durrant-Whyte (2000), for the formulation of the UT that features only the scaling parameter κ,[1] the user should set $\kappa = 3 - n_x$, where n_x is the dimension of the state vector. In this paper, the UKF tuned with such rule is called UkfD.

In Sakai and Kuroda (2010), all three UT parameters plus the noise covariance matrices are obtained by gradient-free optimization. The optimization is carried out offline, implying no extra computational cost for the filter at runtime. The optimization process, though, demands the availability of highly accurate measurements of the state vector.

In Dunik, Simandl, and Straka (2010), Dunik, Simandl, and Straka (2012) and Straka, Dunik, and Simandl (2012), for the formulation of the UT that features only parameter κ, a value for this parameter is selected in real time from a user-defined discrete set of possible values. The UKF that results from this approach works according to the following basic steps. Each time a measurement is received by the filter, the measurement predictive moments are calculated for each one of the possible values of κ. The value of κ that yields the measurement predictive moments that maximize a given performance criterion, such as the likelihood of the measurement received by the filter, is chosen. The filter then uses the

☆ José Jaime da Cruz was partially supported by a CNPq grant. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Antonio Vicino under the direction of Editor Torsten Söderström.

E-mail addresses: lascardua@ifes.edu.br, lascardua@gmail.com (L.A. Scardua), jaime@lac.usp.br (J.J. da Cruz).

[1] This simplified UT is obtained when setting the other parameters of the scaled UT (Wan & van der Merwe, 2000) as $\alpha = 1$ and $\beta = 0$.

http://dx.doi.org/10.1016/j.automatica.2017.01.008
0005-1098/

selected value of κ to perform its regular update and predict steps. The proposed approach thus imposes extra computational cost to the UKF at runtime. The amount of extra computation depends on the number of possible values of κ. This approach was further explored and refined in Straka, Dunik, and Simandl (2014), but the numerical results presented showed that the UKF enhanced with the proposed approach was still significantly slower than the regular UKF. In this paper, the UKF tuned with the approach described in Dunik et al. (2010), Dunik et al. (2012) and Straka et al. (2012) is called UkfK.

In Turner and Rasmussen (2010) and Turner and Rasmussen (2012), model-based optimization is used to pick values for the three UT parameters. The optimization process is guided by the maximization of the upper confidence bound (UCB) (Cox & John, 1997), which is a confidence-bound acquisition criterion (see Section 3). The only data needed from the original dynamic system are the measurements. The optimization is executed offline, so the approach does not increase the runtime computational cost of the UKF. The optimization process is performed by a gradient-based optimizer that relies on the derivative of the UCB. However, reliance on gradient information of the acquisition criteria may prevent the adoption of other relevant acquisition criteria. This is important because the choice of the acquisition criteria influences the optimization results. In this paper, the UKF tuned with this approach is called UkfO.

We approach the tuning of the UKF parameters as an optimization problem. The parameter space is composed by the UT scalar parameters and the main diagonal elements of the two noise covariance matrices of the system model used by the filter. A point in this space is represented by the vector $\boldsymbol{\theta} \in \Re^{n_\theta}$. The first dimension of $\boldsymbol{\theta}$ corresponds to the UT parameter α, the second to parameter β and the third to parameter κ. The other dimensions correspond to the diagonal elements of the noise covariance matrices. Two algorithms are used to perform the optimization.

The first algorithm is a standard model-based optimizer (Forrester, Sobester, & Keane, 2008) that uses a genetic algorithm (Goldberg, 1989) to maximize the UCB criterion. The UKF that results from this approach is here called UkfM. The second is new a tuning method based on stochastic search. This algorithm was used in Scardua and Cruz (2015) to tune solely the UT parameters. The UKF tuned with the second algorithm is called UkfP. For both algorithms, the tuning process is guided solely by the measurements obtained from the nonlinear dynamical system that is to be filtered.

To assess the impact of the tuning on the state estimation performance of the UKF, the two proposed tuning strategies (UkfM and UkfP) are numerically tested against the EKF, the cubature Kalman filter (CKF) (Arasaratnam & Haykin, 2009), the UkfD, the UkfK, and the UkfO. The EKF was chosen because it is a benchmark against which nonlinear filters are usually tested (Psiaki, 2013). UkfD and CKF both result from different tunings of the three UT scalar parameters, while UkfO and UkfK result from state-of-the-art tuning strategies. The numerical tests are performed on three well-known nonlinear filtering problems. To improve readability, we list some of the abbreviations used in the rest of the paper:

- UkfD—UKF tuned according to Julier et al. (2000);
- UkfP—UKF tuned with the proposed optimizer, described in Algorithm 3;
- UkfK—UKF tuned with the online approach of Dunik et al. (2010), Dunik et al. (2012) and Straka et al. (2012);
- UkfO—UKF tuned with the model-based approach of Turner and Rasmussen (2010) and Turner and Rasmussen (2012);
- UkfM—UKF tuned with the classical model-based optimizer, described in Algorithm 2.

The rest of this paper is organized as follows. Section 2 frames the tuning of the UKF parameters as an optimization problem. Section 3 provides background on model-based optimization and describes the UkfM algorithm. Section 4 describes algorithm UkfP. Section 5 analyzes the computational burden of Algorithm 3. Section 6 describes and analyzes the numerical experiments, and Section 7 presents the final comments.

2. Approaching the tuning of the UKF as an optimization problem

Consider the discrete-time nonlinear dynamic system

$$\begin{aligned} \mathbf{x}_k &= f(\mathbf{x}_{k-1}) + \mathbf{w}_{k-1}, \\ \mathbf{y}_k &= h(\mathbf{x}_k) + \mathbf{v}_k, \end{aligned} \tag{1}$$

where $\mathbf{x}_k \in \Re^{n_x}$ and $\mathbf{y}_k \in \Re^{n_y}$ are respectively the state of the dynamic system and its corresponding noisy measurement, $f(\cdot)$ and $h(\cdot)$ are known functions, $\mathbf{w}_{k-1}$ is the process noise at discrete time step $k-1$, and $\mathbf{v}_k$ is the measurement noise at discrete time step k. The distributions of the process and measurement noises are supposed to be Gaussian with zero means and unknown covariance matrices $\mathbf{Q}_{k-1}$ and $\mathbf{R}_k$. These assumptions are respectively written as $\mathbf{w}_{k-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_{k-1})$ and $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_k)$. The initial state $\mathbf{x}_0$ and the two noises are considered to be independent random variables. Both noises are also assumed to be white.

We are interested in improving the quality of the UKF state estimates for (1), without increasing the computational burden of the filter during runtime.

The UKF is a Gaussian filter (Ito & Xiong, 2000). The general form of the Gaussian filter uses a Kalman-like structure to address the estimation task for system (1) (Ito & Xiong, 2000). Assuming that $\mathbf{m}_{k-1}$ and $\mathbf{P}_{k-1}$ are respectively the mean and covariance of the state estimation at time-step $k-1$, the equations of the general form are (Algorithm 6.3 of Sarkka, 2013):

2.1. Prediction

$$\begin{aligned} \mathbf{m}_k^- &= \int f(\mathbf{x}_{k-1})\mathcal{N}(\mathbf{x}_{k-1}|\mathbf{m}_{k-1}, \mathbf{P}_{k-1})d\mathbf{x}_{k-1}, \\ \mathbf{P}_k^- &= \int (f(\mathbf{x}_{k-1}) - \mathbf{m}_k^-)(f(\mathbf{x}_{k-1}) - \mathbf{m}_k^-)^T \\ &\quad \times \mathcal{N}(\mathbf{x}_{k-1}|\mathbf{m}_{k-1}, \mathbf{P}_{k-1})d\mathbf{x}_{k-1} + \mathbf{Q}_{k-1}. \end{aligned} \tag{2}$$

2.2. Update

$$\begin{aligned} \boldsymbol{\mu}_k &= \int h(\mathbf{x}_k)\mathcal{N}(\mathbf{x}_k|\mathbf{m}_k^-, \mathbf{P}_k^-)d\mathbf{x}_k, \\ \mathbf{S}_k &= \int (h(\mathbf{x}_k) - \boldsymbol{\mu}_k)(h(\mathbf{x}_k) - \boldsymbol{\mu}_k)^T \mathcal{N}(\mathbf{x}_k|\mathbf{m}_k^-, \mathbf{P}_k^-)d\mathbf{x}_k + \mathbf{R}_k, \\ \mathbf{C}_k &= \int (\mathbf{x}_k - \mathbf{m}_k^-)(h(\mathbf{x}_k) - \boldsymbol{\mu}_k)^T \mathcal{N}(\mathbf{x}_k|\mathbf{m}_k^-, \mathbf{P}_k^-)d\mathbf{x}_k, \\ \mathbf{K}_k &= \mathbf{C}_k\mathbf{S}_k^{-1}, \\ \mathbf{m}_k &= \mathbf{m}_k^- + \mathbf{K}_k(\mathbf{y}_k - \boldsymbol{\mu}_k), \\ \mathbf{P}_k &= \mathbf{P}_k^- - \mathbf{K}_k\mathbf{S}_k\mathbf{K}_k^T. \end{aligned} \tag{3}$$

Different approaches to solving the moment integrals in (2) and (3) give rise to different Gaussian filters (Wu, Hu, Wu, & Hu, 2006). The UKF originates from using the UT to approximately calculate