STATISTICALLY LINEARIZED RECURSIVE LEAST SQUARES

Matthieu Geist and Olivier Pietquin

IMS Research Group
Supélec
Metz, France

ABSTRACT

This article proposes a new interpretation of the sigma-point kalman filter (SPKF) for parameter estimation as being a statistically linearized recursive least-squares algorithm. This gives new insight on the SPKF for parameter estimation and particularly this provides an alternative proof for a result of Van der Merwe. On the other hand, it legitimates the use of statistical linearization and suggests many ways to use it for parameter estimation, not necessarily in a least-squares sens.

Index Terms— Recursive least-squares, statistical linearization, parameter estimation.

1. INTRODUCTION

The Unscented Kalman Filter (UKF) [1] has recently been introduced as an efficient derivative-free alternative to the Extended Kalman Filter (EKF) [2] for the nonlinear filtering problem. The basic idea behind UKF is that it is easier to approximate an arbitrary random variable rather than an arbitrary nonlinear function. It uses an approximation scheme, the so-called unscented transform (UT) [3], to approximate the statistics of interest involved by Kalman equations (the filter being seen as the optimal linear state estimator minimizing the expected mean-square error conditioned on past observations). More generally, a Kalman filter for which statistics of interest are computed by approximating the random variable rather than the nonlinear function is called a Sigma-Point Kalman Filter (SPKF) [4]. UKF, but also Divided Difference Filter (DVF) [5] or Central Difference Filter (CDF) [6] for example, belong to the SPKF family.

A special form of SPKF, which is the case of interest of this paper, is SPKF for parameter estimation [4]. In this setting, the aim is to estimate a set of stationary parameters instead of tracking a hidden state. It is a simpler case than the general SPKF, because the evolution equation of the corresponding state-space model is at most a random walk. However, it is a case of interest, notably for the Machine Learning community, as it is an efficient derivative-free learning method providing an uncertainty information which can be useful (*e.g.*, for active learning or the exploration/exploitation dilemma in reinforcement learning [7]). SPKF for parameter estimation has been used successfully for supervised learning [8] and even for reinforcement learning [9]. In this article, it is shown how the SPKF for parameter estimation can be obtained from a least-squares perspective using statistical linearization. This gives new insights by linking recursive least-squares to SPKF and suggests that statistical linearization can provide useful for other optimization problems than pure L_2 minimization. An interpretation of (general) UKF as performing a statistical linearization has been proposed before [10]: the state-space formulation of the filtering problem is statistically linearized, which is quite different from the proposed least-squares-based approach.

Sec. 2 introduce some necessary preliminaries about statistical linearization and recursive least-squares. Sec. 3 provides the derivation of the proposed statistically linearized recursive least-squares approach. Sec. 4 shows that the proposed method actually allows to obtain any form of SPKF for parameter estimation. It is also shown that the proposed approach can be seen as an alternative proof for a result of [4], stating that SPKF for parameter estimation is a maximum a posteriori (MAP) estimator. Sec. 5 proposes some perspectives of this work.

2. PRELIMINARIES

In this section statistical linearization and linear recursive least-squares are briefly reminded. For ease of notations, a scalar output is assumed through this article, however the presented results extend easily to the vectorial case.

2.1. Statistical linearization

Let $g : \mathbf{x} \in \mathbb{R}^n \rightarrow y = g(\mathbf{x}) \in \mathbb{R}$ be a nonlinear function. Assume that it is evaluated in r points $(\mathbf{x}^{(j)}, y^{(j)} =$

$g(\mathbf{x}^{(j)})$). The following statistics are defined :

$$\bar{\mathbf{x}} = \frac{1}{r}\sum_{j=1}^{r}\mathbf{x}^{(j)}, \quad \bar{y} = \frac{1}{r}\sum_{j=1}^{r}y^{(j)} \tag{1}$$

$$\mathbf{P}_{\mathbf{xx}} = \frac{1}{r}\sum_{j=1}^{r}(\mathbf{x}^{(j)} - \bar{\mathbf{x}})(\mathbf{x}^{(j)} - \bar{\mathbf{x}})^T \tag{2}$$

$$\mathbf{P}_{\mathbf{x}y} = \frac{1}{r}\sum_{j=1}^{r}(\mathbf{x}^{(j)} - \bar{\mathbf{x}})(y^{(j)} - \bar{y}) \tag{3}$$

$$P_{yy} = \frac{1}{r}\sum_{j=1}^{r}(y^{(j)} - \bar{y})^2 \tag{4}$$

Statistical linearization consists in linearizing $y = g(\mathbf{x})$ around $\bar{\mathbf{x}}$ by adopting a statistical point of view. It finds a linear model $y = \mathbf{A}\mathbf{x} + b$ by minimizing the sum of squared errors between values of nonlinear an linearized functions in the regression points:

$$\min_{\mathbf{A},b}\sum_{j=1}^{r}e_j^T e_j \text{ with } e_j = y^{(j)} - (\mathbf{A}\mathbf{x}^{(j)} + b) \tag{5}$$

The solution of Eq. (5) is given by [11]:

$$\mathbf{A} = \mathbf{P}_{y\mathbf{x}}\mathbf{P}_{\mathbf{xx}}^{-1} \quad b = \bar{y} - \mathbf{A}\bar{\mathbf{x}} \tag{6}$$

Moreover, it is easy to check that the covariance matrix of the error is given by:

$$P_{ee} = \frac{1}{r}\sum_{j=1}^{r}e_j e_j^T \tag{7}$$

$$= P_{yy} - \mathbf{A}\mathbf{P}_{\mathbf{xx}}\mathbf{A}^T \tag{8}$$

For now, how to choose regression points has not been discussed. This is left for Sec. 3.

2.2. Linear recursive least-squares

Assume the following linear observation model:

$$y_i = \theta^T\mathbf{x}_i + v_i \tag{9}$$

where $\mathbf{x}_i \in \mathbb{R}^n$, $y_i \in \mathbb{R}$ and v_i is a white observation noise of variance P_{vv_i}. Least-squares approach seeks at estimating the $n \times 1$ parameter vector θ by minimizing the squared error among i observed samples $((\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_i, y_i))$:

$$\theta_i^{\text{LS}} = \underset{\theta}{\operatorname{argmin}}\, J_i(\theta) \tag{10}$$

$$\text{with } J_i(\theta) = \sum_{j=1}^{i}\frac{1}{P_{vv_j}}\left(y_j - \theta^T\mathbf{x}_j\right)^2 \tag{11}$$

The least-squares solution is classically obtained by zeroing the gradient of the cost function $J_i(\theta)$, which gives the least-squares (LS) estimate:

$$\theta_i^{\text{LS}} = \left(\sum_{j=1}^{i}\frac{1}{P_{vv_j}}\mathbf{x}_j\mathbf{x}_j^T\right)^{-1}\sum_{j=1}^{i}\frac{1}{P_{vv_j}}\mathbf{x}_j y_j \tag{12}$$

The parameter vector θ_i can be estimated online by being updated for each new observation. For this, the matrix $\mathbf{P}_i = (\sum_{j=1}^{i}\frac{1}{P_{vv_j}}\mathbf{x}_j\mathbf{x}_j^T)^{-1} = (P_{i-1}^{-1} + \mathbf{x}_i\mathbf{x}_i^T)^{-1}$ is computed recursively by using the Sherman-Morrison formula:

$$\mathbf{P}_i = \mathbf{P}_{i-1} - \frac{\mathbf{P}_{i-1}\mathbf{x}_i\mathbf{x}_i^T\mathbf{P}_{i-1}}{P_{vv_i} + \mathbf{x}_i^T\mathbf{P}_{i-1}\mathbf{x}_i} \tag{13}$$

By injecting Eq. (13) into Eq. (12) and by assuming that some priors θ_0 and $\mathbf{P}_0$ are chosen, the recursive least-squares (RLS) algorithm is obtained:

$$\mathbf{K}_i = \frac{\mathbf{P}_{i-1}\mathbf{x}_i}{P_{vv_i} + \mathbf{x}_i^T\mathbf{P}_{i-1}\mathbf{x}_i} \tag{14}$$

$$\theta_i^{\text{RLS}} = \theta_{i-1}^{\text{RLS}} + \mathbf{K}_i\left(y_i - (\theta_{i-1}^{\text{RLS}})^T\mathbf{x}_i\right) \tag{15}$$

$$\mathbf{P}_i = \mathbf{P}_{i-1} - \mathbf{K}_i\left(P_{vv_i} + \mathbf{x}_i^T\mathbf{P}_{i-1}\mathbf{x}_i\right)\mathbf{K}_i^T \tag{16}$$

This RLS formulation provides useful for the statistical linearization of nonlinear least-squares presented in the next section. Notice that this estimator does not minimize $J_i(\theta)$, but a regularized version of this cost function:

$$\theta_i^{\text{RLS}} = \underset{\theta}{\operatorname{argmin}}\left(J_i(\theta) + (\theta - \theta_0)^T\mathbf{P}_0^{-1}(\theta - \theta_0)\right) \tag{17}$$

From a Bayesian point of view, the LS estimate can be seen as a maximum likelihood (ML) estimate whereas the RLS estimate can be seen as a maximum a posteriori (MAP) estimate. From now, this difference is no longer specified, as it is clear from the context (batch or recursive estimation).

3. STATISTICALLY LINEARIZED RECURSIVE LEAST-SQUARES

Assume the following nonlinear observation model:

$$y_i = f_\theta(\mathbf{x}_i) + v_i \tag{18}$$

where $\mathbf{x}_i \in \mathbb{R}^n$, $y_i \in \mathbb{R}$, v_i is a white observation noise of variance P_{vv_i} and f_θ is a parametric function approximator of interest (for example an artificial neural network for which θ specifies synaptic weights [12]). For a set of i observed samples, the least-squares solution is given by:

$$\theta_i = \underset{\theta}{\operatorname{argmin}}\sum_{j=1}^{i}\frac{1}{P_{vv_j}}\left(y_j - f_\theta(\mathbf{x}_j)\right)^2 \tag{19}$$

To address this nonlinear least-squares problem, the nonlinear observation model is statistically linearized (see Sec. 2.1):

$$y_i = \mathbf{A}_i\theta + b_i + e_i + v_i \tag{20}$$

At this point, it should be noted that a set of points has to be sampled so as to perform statistical linearization, that is to compute $\mathbf{A}_i$, b_i and e_i. For now, this is left as an open question, this problem being addressed later.

Let $u_i = e_i + v_i$ be the noise associated to observation model (20). Noises v_i and e_i being independent, the variance of u_i is given by $P_{uu_i} = P_{vv_i} + P_{ee_i}$. Observation models (18) and (20) being equivalent, the least-squares solution can be rewritten as:

$$\theta_i = \underset{\theta}{\operatorname{argmin}} \sum_{j=1}^{i} \frac{1}{P_{uu_j}} \left(y_j - (\mathbf{A}_j\theta + b_j)\right) \tag{21}$$

$$= \left(\sum_{j=1}^{i} \frac{1}{P_{uu_j}} \mathbf{A}_j^T \mathbf{A}_j\right)^{-1} \sum_{j=1}^{i} \frac{1}{P_{uu_j}} \mathbf{A}_j^T \left(y_j - b_j\right) \tag{22}$$

Using the Sherman-Morrison formula, a recursive formulation of this estimation can be obtained (see Sec. 2.2):

$$\mathbf{K}_i = \frac{\mathbf{P}_{i-1}\mathbf{A}_i^T}{P_{uu_i} + \mathbf{A}_i\mathbf{P}_{i-1}\mathbf{A}_i^T} \tag{23}$$

$$\theta_i = \theta_{i-1} + \mathbf{K}_i \left(y_i - b_i - \mathbf{A}_i\theta_{i-1}\right) \tag{24}$$

$$\mathbf{P}_i = \mathbf{P}_{i-1} - \mathbf{K}_i \left(P_{uu_i} + \mathbf{A}_i\mathbf{P}_{i-1}\mathbf{A}_i^T\right)\mathbf{K}_i^T \tag{25}$$

The problem of choosing a specific statistical linearization is now addressed. With the recursive formulation, θ_{i-1} and $\mathbf{P}_{i-1}$ are known, and the issue is to compute $\mathbf{A}_i$ and b_i. A first thing is to choose around what point to linearize and with which magnitude. Recall that the previous estimate θ_{i-1} is known. Moreover, the matrix $\mathbf{P}_{i-1}$ can be interpreted as the variance matrix associated to θ_{i-1}. It is thus legitimate to sample r points $(\theta_i^{(j)}, y_i^{(j)} = f_{\theta_i^{(j)}}(\mathbf{x}_i))$ such that $\bar{\theta}_i = \theta_{i-1}$ and $\mathbf{P}_{\theta\theta_i} = \mathbf{P}_{i-1}$. The following statistics are thus available (how to sample these r points is discussed in Sec. 4):

$$\bar{\theta}_i = \theta_{i-1} = \frac{1}{r}\sum_{j=1}^{r} \theta_i^{(j)}, \quad \bar{y}_i = \frac{1}{r}\sum_{j=1}^{r} y_i^{(j)} \tag{26}$$

$$\mathbf{P}_{\theta\theta_i} = \mathbf{P}_{i-1} = \frac{1}{r}\sum_{j=1}^{r} (\theta^{(j)} - \bar{\theta})(\theta^{(j)} - \bar{\theta})^T \tag{27}$$

$$\mathbf{P}_{\theta y_i} = \frac{1}{r}\sum_{j=1}^{r} (\theta_i^{(j)} - \bar{\theta}_i)(y_i^{(j)} - \bar{y}_i) \tag{28}$$

$$P_{yy_i} = \frac{1}{r}\sum_{j=1}^{r} (y_i^{(j)} - \bar{y}_i)^2 \tag{29}$$

The solution to the statistical linearization problem is thus (see Sec. 2.1):

$$\mathbf{A}_i = \mathbf{P}_{y\theta_i}\mathbf{P}_{\theta\theta_i}^{-1} = \mathbf{P}_{y\theta_i}\mathbf{P}_{i-1}^{-1} \tag{30}$$

$$b_i = \bar{y}_i - \mathbf{A}_i\bar{\theta}_i = \bar{y}_i - \mathbf{A}_i\theta_{i-1} \tag{31}$$

The noise variance induced by the statistical linearization is given by (see again Sec. 2.1):

$$P_{ee_i} = P_{yy_i} - \mathbf{A}_i\mathbf{P}_{i-1}\mathbf{A}_i^T \tag{32}$$

Injecting Eq. (30) and (32) into Eq. (23) gives (recall also that $P_{uu_i} = P_{vv_i} + P_{ee_i}$):

$$\mathbf{K}_i = \frac{\mathbf{P}_{i-1}\mathbf{A}_i^T}{P_{uu_i} + \mathbf{A}_i\mathbf{P}_{i-1}\mathbf{A}_i^T} \tag{33}$$

$$= \frac{\mathbf{P}_{i-1}\left(\mathbf{P}_{y\theta_i}\mathbf{P}_{i-1}^{-1}\right)^T}{P_{vv_i} + P_{yy_i} - \mathbf{A}_i\mathbf{P}_{i-1}\mathbf{A}_i^T + \mathbf{A}_i\mathbf{P}_{i-1}\mathbf{A}_i^T} \tag{34}$$

$$= \frac{P_{\theta y_i}}{P_{vv_i} + P_{yy_i}} \tag{35}$$

Injecting Eq. (30-31) into Eq. (24) gives:

$$\theta_i = \theta_{i-1} + \mathbf{K}_i \left(y_i - b_i - \mathbf{A}_i\theta_{i-1}\right) \tag{36}$$

$$= \theta_{i-1} + \mathbf{K}_i \left(y_i - (\bar{y}_i - \mathbf{A}_i\theta_{i-1}) - \mathbf{A}_i\theta_{i-1}\right) \tag{37}$$

$$= \theta_{i-1} + \mathbf{K}_i \left(y_i - \bar{y}_i\right) \tag{38}$$

Injecting Eq. (32) into Eq. (25) gives (recall again that $P_{uu_i} = P_{vv_i} + P_{ee_i}$):

$$\mathbf{P}_i = \mathbf{P}_{i-1} - \mathbf{K}_i \left(P_{uu_i} + \mathbf{A}_i\mathbf{P}_{i-1}\mathbf{A}_i^T\right)\mathbf{K}_i^T \tag{39}$$

$$= \mathbf{P}_{i-1} - \mathbf{K}_i \left(P_{vv_i} + P_{yy_i}\right)\mathbf{K}_i^T \tag{40}$$

Eq. (35), (38) and (40) define the statistically linearized recursive least squares (SL-RLS) algorithm:

$$\mathbf{K}_i = \frac{P_{\theta y_i}}{P_{vv_i} + P_{yy_i}} \tag{41}$$

$$\theta_i = \theta_{i-1} + \mathbf{K}_i \left(y_i - \bar{y}_i\right) \tag{42}$$

$$\mathbf{P}_i = \mathbf{P}_{i-1} - \mathbf{K}_i \left(P_{vv_i} + P_{yy_i}\right)\mathbf{K}_i^T \tag{43}$$

The last question to answer is how to sample the r points $\theta_i^{(j)}$ such that $\bar{\theta}_i = \theta_{i-1}$ and $\mathbf{P}_{\theta\theta_i} = \mathbf{P}_{i-1}$.

4. LINKS TO SPKF FOR PARAMETER ESTIMATION

A first natural idea to sample these r points is to assume a Gaussian distribution of mean θ_{i-1} and of variance matrix $\mathbf{P}_{i-1}$ and to compute statistics of interest using a Monte Carlo approach. However, more efficient methods exist, notably the unscented transform [3]. It consists

in deterministically sampling a set of $2n+1$ so-called sigma-points as follows:

$$\theta_i^{(j)} = \theta_{i-1} \qquad j = 0 \quad (44)$$

$$\theta_i^{(j)} = \theta_{i-1} + \left(\sqrt{(n+\kappa)\mathbf{P}_{i-1}}\right)_j \qquad 1 \leq j \leq n \quad (45)$$

$$\theta_i^{(j)} = \theta_{i-1} - \left(\sqrt{(n+\kappa)\mathbf{P}_{i-1}}\right)_j \qquad j+1 \leq j \leq 2n \quad (46)$$

as well as associated weights:

$$w_0 = \frac{\kappa}{n+\kappa} \text{ and } w_j = \frac{1}{2(n+\kappa)} \forall j > 0 \quad (47)$$

where κ is a scaling factor which controls the accuracy of the unscented transform [3] and $\left(\sqrt{(n+\kappa)\mathbf{P}_{i-1}}\right)_j$ is the j^{th} column of the Cholesky decomposition of the matrix $(n+\kappa)\mathbf{P}_{i-1}$. The image of each of these sigma-points is computed:

$$y_i^{(j)} = f_{\theta_i^{(j)}}(\mathbf{x}_i), \forall j \quad (48)$$

and statistics of interest are computed as follows:

$$\bar{y}_i = \sum_{j=0}^{2n} w_j y_i^{(j)} \quad (49)$$

$$\mathbf{P}_{\theta y_i} = \sum_{j=0}^{2n} w_j \left(\theta_i^{(j)} - \theta_{i-1}\right)\left(y_i^{(j)} - \bar{y}_i\right) \quad (50)$$

$$P_{yy_i} = \sum_{j=0}^{2n} w_j \left(y_i^{(j)} - \bar{y}_i\right)^2 \quad (51)$$

As a non-equiweighted sum can be rewritten as an equiweighted sum by considering some of the terms more than one time (by assuming that weights are rational numbers, which is not a too strong hypothesis), the unscented transform can be interpreted as a form of statistical linearization.

If the unscented transform is considered as the statistical linearization process, than the SL-RLS algorithm, that is Eq. (41-43), is exactly the UKF when no evolution model is considered in the state-space model. In other words, SL-RLS is the UKF for parameter estimation. In a similar way, if the scaled unscented transform [13] is used to perform statistical linearization, SL-RLS is the scaled UKF for parameter estimation. If the statistical linearization is performed using a Sterling's interpolation, SL-RLS is DDF or CDF for parameter estimation. So, generally speaking, depending on the scheme chosen to perform statistical linearization, SL-RLS is SPKF for parameter estimation.

This interpretation of the SPKF for parameter estimation as being a statistically linearized recursive least-squares algorithm allows to provide an alternative, simpler and requiring less assumptions proof of a result of Van der Merwe. This result [4, Ch. 4.5.1] states that the SPKF for parameter estimation algorithm is equivalent to a MAP estimate of the underlying parameters under a Gaussian posterior (and noise distribution) assumption.

Theorem 1 (SL-RLS estimate is a MAP estimate). *Assume that prior and noise distributions are Gaussian. Then the statically linearized recursive least-squares estimate is equivalent to the maximum a posteriori estimate.*

Proof. By construction, with priors defined by θ_0 and $\mathbf{P}_0$, the SL-RLS estimate minimizes the following regularized cost function:

$$\theta_i = \underset{\theta}{\text{argmin}} \Big(\sum_{j=0}^{i} \frac{1}{P_{vv_i} + P_{ee_i}} \left(y_i - (\mathbf{A}_i\theta + b_i)\right)^2 + (\theta - \theta_0)^T \mathbf{P}_0^{-1} (\theta - \theta_0) \Big) \quad (52)$$

$$= \underset{\theta}{\text{argmin}} \Big(\sum_{j=0}^{i} \frac{1}{P_{vv_i}} \left(y_i - f_\theta(\mathbf{x}_i)\right)^2 + (\theta - \theta_0)^T \mathbf{P}_0^{-1} (\theta - \theta_0) \Big) \quad (53)$$

On the other hand, the MAP estimator is defined as (and using the Bayes rule):

$$\theta_i^{\text{MAP}} = \underset{\theta}{\text{argmax}}\, p(\theta | y_{1:i}) \quad (54)$$

$$= \underset{\theta}{\text{argmax}} \frac{p(y_{1:i}|\theta)p(\theta)}{p(y_{1:i})} \quad (55)$$

As the observation noise is white, the joint likelihood is the product of local likelihoods, and the probability $p(y_{1:i})$ does not depend on θ, so:

$$\theta_i^{\text{MAP}} = \underset{\theta}{\text{argmax}} \left(p(\theta) \prod_{j=1}^{i} p(r_j|\theta) \right) \quad (56)$$

Prior and noise distributions are assumed to be Gaussian, thus:

$$p(\theta) \propto \exp\left(-\frac{1}{2}(\theta - \theta_0)^T \mathbf{P}_0^{-1} (\theta - \theta_0)\right) \quad (57)$$

$$p(y_j|\theta) \propto \exp\left(-\frac{(y_j - f_\theta(\mathbf{x}_j))^2}{2P_{vv_j}}\right) \quad (58)$$

Finally, maximizing a product of probability distributions is equivalent to minimizing the sum of the negatives of their logarithms, which gives the result:

$$\theta_i^{\text{MAP}} = \underset{\theta}{\text{argmin}} \Big(\sum_{j=0}^{i} \frac{1}{P_{vv_i}} \left(y_i - f_\theta(\mathbf{x}_i)\right)^2 + (\theta - \theta_0)^T \mathbf{P}_0^{-1} (\theta - \theta_0) \Big) = \theta_i \quad (59)$$

□

This alternative proof is shorter than the original one. But above all it does not assume that the posterior distribution is Gaussian, which is a very strong assumption for a nonlinear observation model.

5. PERSPECTIVES

In this article a statistically linearized recursive least-squares algorithm has been introduced. It has been shown to be actually the SPKF for parameter estimation algorithm. This gives new insights to sigma-point Kalman filters by showing that they are generalizations of a statstically linearized least-squares approach. This new point of view allowed to provide an alternative proof of a result stating that without evolution model, SPKF estimate is the maximum a posteriori estimate. The proof proposed in Sec. 4 is shorter and above all it does not assume a Gaussian posterior, which is a very strong hypothesis in the case of a nonlinear evolution model.

The technique of statistical linearization can be applied in much more general problems than the L_2 minimization addressed in this paper. The fact that statistically linearized recursive least-squares is indeed a special form of sigma-point Kalman filtering tends to justify this approach. Interesting perspectives can be (but are not limited to) the application of this general statistical linearization to L_1 minimization (*e.g.*, [14]), L_1 regularization (*e.g.*, [15]) or fixed-point approximation (*e.g.*, [16, 17]).

Acknowledgments

The authors thank the European Community (FP7/2007-2013, grant agreement 216594, CLASSiC project : www.classic-project.org) and the Région Lorraine for financial support.

6. REFERENCES

[1] S. J. Julier and J. K. Uhlmann, "A new extension of the Kalman filter to nonlinear systems," in *Int. Symp. Aerospace/Defense Sensing, Simul. and Controls 3*, 1997.

[2] D. Simon, *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches*, Wiley & Sons, August 2006.

[3] S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, 2004.

[4] R. van der Merwe, *Sigma-Point Kalman Filters for Probabilistic Inference in Dynamic State-Space Models*, Ph.D. thesis, OGI School of Science & Engineering, Oregon Health & Science University, Portland, OR, USA, 2004.

[5] M. Noorgaard, N. Poulsen, and O. Ravn, "New Developments in State Estimation for Nonlinear Systems," *Automatica*, 2000.

[6] K. Ito and K. Xiong, "Gaussian Filters for Nonlinear Filtering Problems," *IEEE Transactions on Automatic Control*, vol. 25, no. 5, pp. 910–927, 2000.

[7] M. Geist and O. Pietquin, "Managing uncertainty within value function approximation in reinforcement learning," in *Active Learning and Experimental Design workshop (collocated with AISTATS 2010)*, Sardinia (Italy), 2010.

[8] S. Haykin, *Kalman Filtering and Neural Networks*, Wiley, 2001.

[9] M. Geist, O. Pietquin, and G. Fricout, "Kalman Temporal Differences: the deterministic case," in *IEEE International Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL 2009)*, Nashville, TN, USA, April 2009.

[10] T. Lefebvre, H. Bruyninckx, and J. De Shutter, "Comments on "A New Method for the Nonlinear Transformation of Means and Covariances in Filters and Estimators"," *IEEE Transactions on Automatic Control*, vol. 47, no. 8, pp. 1406–1409, 2002.

[11] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis*, Wiley, 1984.

[12] C. M. Bishop, *Neural Networks for Pattern Recognition*, Oxford University Press, New York, USA, 1995.

[13] S. J. Julier, "The scaled unscented transformation," in *American Control Conference*, 2002, vol. 6, pp. 4555–4559.

[14] G. O. Wesolowsky, "New Descent Algorithm for the Least Absolute Value Regression Problem," *Communications in Statistics - Simulation and Computation*, , no. 5, pp. 479–491, 1981.

[15] R. Tibshirani, "Regression Shrinkage and Selection via the LASSO," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.

[16] S. J. Bradtke and A. G. Barto, "Linear Least-Squares algorithms for temporal difference learning," *Machine Learning*, vol. 22, no. 1-3, pp. 33–57, 1996.

[17] A. Nedić and D. P. Bertsekas, "Least Squares Policy Evaluation Algorithms with Linear Function Approximation," *Discrete Event Dynamic Systems: Theory and Applications*, vol. 13, pp. 79–110, 2003.