

Outlier-insensitive Kalman Smoothing and Marginal Message Passing

Federico Wadehn*, Lukas Bruderer*, Justin Dauwels†, Vijay Sahdeva*, Hang Yu†, and Hans-Andrea Loeliger*

*Dept. of Information Technology and Electrical Engineering, ETH Zurich, Switzerland

†School of Electrical and Electronics Engineering, Nanyang Technological University, Singapore

Abstract—We propose a new approach to outlier-insensitive Kalman smoothing based on normal priors with unknown variance (NUV). In contrast to prior work, the actual computations amount essentially to iterations of a standard Kalman smoother (with few extra computations). The proposed approach is easily extended to nonlinear estimation problems by combining the outlier detection with an extended Kalman smoother. For the Kalman smoothing, we consider both a Modified Bryson-Frazier smoother and the recently proposed Backward Information Filter Forward Marginal smoother, neither of which requires matrix inversions.

I. INTRODUCTION

The performance of standard Kalman filters and smoothers is severely impaired by outliers in the measurements due to the sensitivity of the least-squares cost function to deviations from nominal noise [1], [2]. Real measurements however, are commonly contaminated with outliers, such as motion artifacts in physiological signals or glint noise in radar signals [1], [3].

Prior work on dealing with robust filtering and smoothing includes [1], [3], where Kalman smoothing is formulated as a linear regression problem and outlier detection is performed via a sparsifying ℓ_1 -penalty. Along the same lines in [4], [5], outliers are assumed to be non-Gaussian, and the mean squared error criterion is replaced by an appropriate score function in a robust M-estimation framework.

In this paper we propose a new approach to outlier-insensitive Kalman smoothing (OIKS): using an idea from sparse Bayesian learning [6], we model outliers as Gaussian impulses with unknown variances and estimate these variances by expectation maximization (EM). In contrast to the mentioned prior works, our approach is parameter-free and amounts essentially to iterated standard Kalman smoothing (with additional simple update steps for the unknown variances), i.e., we effectively stay within the linear Gaussian framework. The extension of the proposed framework to nonlinear estimation problems is straightforward and will be discussed in section IV-F.

For the actual Kalman smoothing, we consider two algorithms, both without matrix inversions: The Modified Bryson-Frazier smoother (MBF) [7], and the recently proposed BIFM (backward information filter, forward marginal); the latter incurs the least computational costs [8]. In contrast to the commonly used Rauch-Tung-Striebel (RTS) smoother [9], neither the MBF nor the BIFM require inverting the state covariance matrix at every time step. In consequence, both

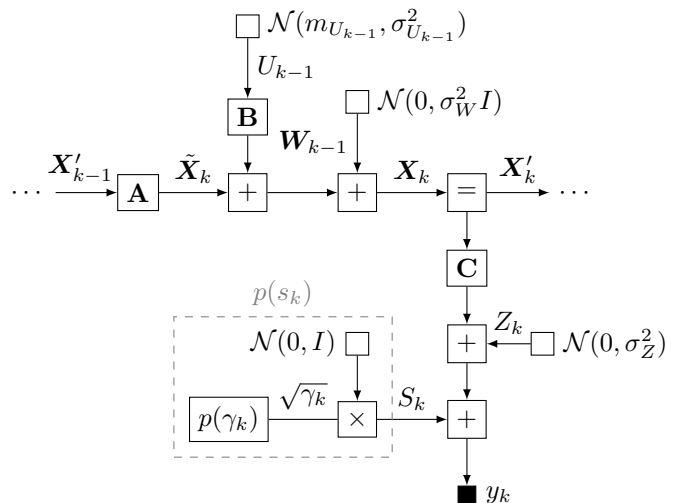


Fig. 1: Factor graph segment of the system model. (The whole factor graph consists of many such sections, one for each time step.)

smoothers have better numerical properties and require fewer matrix multiplications than the RTS smoother.

II. SYSTEM MODEL AND BASIC IDEA

The given discrete-time signal $\mathbf{y} = (y_1, \dots, y_N)$ is modeled as the output of the linear state space model

$$\begin{aligned} \mathbf{X}_k &= \mathbf{A}\mathbf{X}_{k-1} + \mathbf{B}U_{k-1} + \mathbf{W}_{k-1} \\ Y_k &= \mathbf{C}\mathbf{X}_k + Z_k + S_k \end{aligned} \quad (1)$$

with $\mathbf{A} \in \mathbb{R}^{d \times d}$, $\mathbf{X}_k, \mathbf{W}_k, \mathbf{B}, \mathbf{C}^T \in \mathbb{R}^d$, and $U_k, Y_k \in \mathbb{R}$. The input signal $\mathbf{U} = (U_1, \dots, U_N)$ is a sequence of independent Gaussian random variables (not necessarily i.i.d., but with known means and variances). Modeling uncertainties are accounted for with i.i.d. zero-mean Gaussian process noise $\mathbf{W} = (W_1, \dots, W_N)$.

We assume two sources of observation noise: the normal distributed observation noise $\mathbf{Z} = (Z_1, \dots, Z_N)$ and the outlier-causing impulsive noise $\mathbf{S} = (S_1, \dots, S_N)$. The former is i.i.d. zero-mean Gaussian with constant variance σ_Z^2 . The impulsive noise terms (S_1, \dots, S_N) are modeled as independent zero-mean Gaussians, each with its own variance γ_k , with $k \in \{1, \dots, N\}$. The point of this model is that, both the maximum-likelihood estimate and the MAP estimate with a suitable prior $p(\gamma)$ for $(\gamma_1, \dots, \gamma_N)$, tend to be sparse [8],

i.e. most of the estimated components $\hat{\gamma}_k$ will be exactly zero, which implies $S_k = 0$.

The unknown variances $(\gamma_1, \dots, \gamma_N)$ will be estimated by an EM algorithm, as described in Section IV. In each E-step of the EM algorithm, the variances $(\gamma_1, \dots, \gamma_N)$, are (momentarily) fixed; the resulting model is purely linear Gaussian, and each E-step amounts to one round of Kalman smoothing. We first turn to the Kalman smoothing procedure.

III. INVERSION-FREE KALMAN SMOOTHING

To describe the two Kalman smoother variations, we will use the factor graph block diagram representation of the state space model (1), shown in Fig. 1. Using this representation, tabulated message passing rules as shown in Table I can be devised, which allow the composition of algorithms ranging from state estimation to output smoothing (interpolation) and input estimation [8]. Let $\gamma = (\gamma_1, \dots, \gamma_N)$ in Fig. 1 be fixed to zero. Inference in linear state space models, i.e. Kalman filtering and smoothing, boils down to sum-product message passing on the given factor graph [10]. There are two general approaches to compute the smoothing distribution [11]. Firstly, by a two-filter (2-Filt.) smoother [10], based on the observation that:

$$p(\mathbf{x}_k | \mathbf{y}_{1:N}) \propto p(\mathbf{x}_k | \mathbf{y}_{1:k}) p(\mathbf{y}_{k+1:N} | \mathbf{x}_k), \quad (2)$$

which means that the smoothing distribution can be computed by a forward pass and an independent backward pass (Kalman filter that runs backwards in time). Secondly, via recursive Bayesian estimation as in (3), which requires the computation of the filtering distribution $p(\mathbf{x}_k | \mathbf{y}_{1:k})$ beforehand.

$$p(\mathbf{x}_k | \mathbf{y}_{1:N}) = p(\mathbf{x}_k | \mathbf{y}_{1:k}) \int \frac{p(\mathbf{x}_{k+1} | \mathbf{x}_k) p(\mathbf{x}_{k+1} | \mathbf{y}_{1:N})}{p(\mathbf{x}_{k+1} | \mathbf{y}_{1:k})} d\mathbf{x}_{k+1} \quad (3)$$

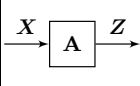
The RTS, MBF and the BIFM smoother are all instances of the latter representation of recursive Bayesian estimation.

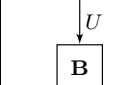
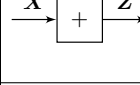
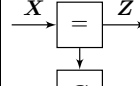
A. Notation

We will express the filtering and smoothing update steps in message passing notation (cf. [10]). The Gaussian forward filtering distribution $p(\mathbf{x}_k | \mathbf{y}_{1:k})$, computed with the covariance Kalman filter, is parametrized by the mean $\vec{\mathbf{m}}_{\mathbf{x}_k}$ and covariance $\vec{\mathbf{V}}_{\mathbf{x}_k}$. The backward filtering distribution $p(\mathbf{x}_k | \mathbf{y}_{k:N})$, computed with the backward information filter, will be parametrized by the precision matrix $\overleftarrow{\mathbf{W}}_{\mathbf{x}_k}$ (inverse of covariance matrix) and the transformed mean $\overleftarrow{\boldsymbol{\xi}}_{\mathbf{x}_k} \triangleq \overleftarrow{\mathbf{W}}_{\mathbf{x}_k} \vec{\mathbf{m}}_{\mathbf{x}_k}$. The right-pointing arrow is used for quantities in the forward filter that depend only on past and present observations [10], the left pointing arrow for quantities that depend only on future and present observations. Message quantities without an arrow (such as $\mathbf{m}_{\mathbf{x}_k}$ and $\mathbf{V}_{\mathbf{x}_k}$) indicate parameters of the marginal distribution $p(\mathbf{x}_k | \mathbf{y}_{1:N})$.

B. MBF and BIFM Smoother

Table I shows the message update rules (for both the MBF smoother and the BIFM smoother) for the critical blocks of the factor graph. Both smoothers consist of three message updates,

Node	Update rule
	$\mathbf{m}_Z = \mathbf{A} \mathbf{m}_X, \quad \mathbf{V}_Z = \mathbf{A} \mathbf{V}_X \mathbf{A}^\top \quad (4)$

	$\overleftarrow{\boldsymbol{\xi}}_X = \mathbf{A}^\top \overleftarrow{\boldsymbol{\xi}}_Z, \quad \overleftarrow{\mathbf{W}}_X = \mathbf{A}^\top \overleftarrow{\mathbf{W}}_Z \mathbf{A} \quad (5)$
	$\mathbf{m}_Z = \overleftarrow{\mathbf{F}}^\top \mathbf{m}_X + \mathbf{B} \mathbf{H} (\mathbf{B}^\top \overleftarrow{\boldsymbol{\xi}}_Z + \overleftarrow{\boldsymbol{\xi}}_U) \quad (6)$
	$\mathbf{V}_Z = \overleftarrow{\mathbf{F}}^\top \mathbf{V}_X \overleftarrow{\mathbf{F}} + \mathbf{B} \mathbf{H} \mathbf{B}^\top \quad (7)$
	$\overleftarrow{\mathbf{F}} = \mathbf{I} - \overleftarrow{\mathbf{W}}_Z \mathbf{B} \mathbf{H} \mathbf{B}^\top \quad (8)$
	$\mathbf{H} = (\overleftarrow{\mathbf{W}}_U + \mathbf{B}^\top \overleftarrow{\mathbf{W}}_Z \mathbf{B})^{-1} \quad (9)$
	-----
	$\overleftarrow{\boldsymbol{\xi}}_X = \overleftarrow{\boldsymbol{\xi}}_Z, \quad \overleftarrow{\mathbf{W}}_X = \overleftarrow{\mathbf{W}}_Z \quad (10)$
	$\mathbf{m}_Z = \mathbf{m}_X, \quad \mathbf{V}_Z = \mathbf{V}_X \quad (11)$

	$\overleftarrow{\boldsymbol{\xi}}_X = \overleftarrow{\mathbf{F}}^\top \overleftarrow{\boldsymbol{\xi}}_Z + \mathbf{C}^\top \mathbf{G} (\mathbf{C} \vec{\mathbf{m}}_X - \vec{\mathbf{m}}_Y) \quad (12)$
	$\overleftarrow{\mathbf{W}}_X = \overleftarrow{\mathbf{F}}^\top \overleftarrow{\mathbf{W}}_Z \overleftarrow{\mathbf{F}} + \mathbf{C}^\top \mathbf{G} \mathbf{C} \quad (13)$
	$\overleftarrow{\mathbf{F}} = \mathbf{I} - \overleftarrow{\mathbf{V}}_X \mathbf{C}^\top \mathbf{G} \mathbf{C} \quad (14)$
	$\mathbf{G} = (\overleftarrow{\mathbf{V}}_Y + \mathbf{C} \overleftarrow{\mathbf{V}}_X \mathbf{C}^\top)^{-1} \quad (15)$

TABLE I: Update rules for BIFM smoother (above the dashed line) and MBF smoother (below the dashed line).

through the “equality”-, “multiplier”- and “plus”-node. The smoothing distribution computed via the MBF is parametrized by the dual precision $\overleftarrow{\mathbf{W}}_X \triangleq (\overleftarrow{\mathbf{V}}_X + \overleftarrow{\mathbf{V}}_X)^{-1}$ and the dual mean $\overleftarrow{\boldsymbol{\xi}}_X \triangleq \overleftarrow{\mathbf{W}}_X (\vec{\mathbf{m}}_X - \vec{\mathbf{m}}_X)$ as described in [8]. The (forward) filtering distribution $p(\mathbf{x}_k | \mathbf{y}_{k:N})$, parametrized by $\vec{\mathbf{m}}_{\mathbf{x}_k}$ and $\vec{\mathbf{V}}_{\mathbf{x}_k}$ is computed via the standard covariance Kalman filter. Note that the invariance of the message parameters $\overleftarrow{\mathbf{W}}_X$ and $\overleftarrow{\boldsymbol{\xi}}_X$ at the “plus”-node, renders the MBF an ideal candidate for systems with multiple inputs and state noise as well as for input estimation as in [12].

For the BIFM in contrast, the smoothing distribution is parametrized by the mean \mathbf{m}_X and the covariance matrix \mathbf{V}_X . To determine the filtering distribution $p(\mathbf{x}_k | \mathbf{y}_{k:N})$, parametrized by $\overleftarrow{\boldsymbol{\xi}}_{\mathbf{x}_k}$ and $\overleftarrow{\mathbf{W}}_{\mathbf{x}_k}$, the BIFM uses the backward information filter [10]. Note that due to the invariance of the marginal mean and the covariance at the “equality”-node, the BIFM is well suited for systems with multiple outputs.

The duality between BIFM and the MBF can be appreciated by observing that the MBF smoother is inversion-free when working with scalar observations, whereas BIFM is inversion-free when working with scalar inputs. However, in our simulations we did not find any significant difference in numerical precision between these two smoothers. Table II shows the computational complexity, the number of matrix inversion and the storage requirements of the RTS, MBF, BIFM and the 2-Filter smoother. The MBF and BIFM compare favourably with other Kalman smoothers and we advocate their usage.

IV. OUTLIER-INSENSITIVE KALMAN SMOOTHER

The proposed Outlier-Insensitive Kalman Smoother (OIKS) combines MBF/BIFM smoothing with the estimation of the time-varying noise variance of the impulsive noise. The multiplier node (cf. [10]) in Fig. 1, multiplies the prior $p(\gamma_k)$ of the variance γ_k with the standard normal distribution,

	matrix mult.	matrix inv.	storage
RTS	4 (2)	1 (1)	$\vec{V}_{\mathbf{x}_k}, \vec{V}_{\mathbf{x}_k}$
2-Filter	10 (10)	0	$\vec{V}_{\mathbf{x}_k}$
MBF	6 (0)	0	$\vec{V}_{\mathbf{x}_k}$
BIFM	4 (0)	0	$\vec{W}_{\mathbf{x}_k}$

TABLE II: Matrix multiplications and inversions at every time-step, as well as storage requirements of Kalman filter quantities, for the smoothing part of different Kalman smoothers. If the variance of the estimate is not required, unlike the RTS and the 2-Filter smoother [Alg. E [10]], both the MBF and BIFM require no matrix multiplications (shown in parenthesis).




Fig. 2: Top: Ground truth (reference) and observation \mathbf{y} with outliers. Bottom: Kalman smoother (KS), Kalman smoother with median filter (MedKS) and outlier-insensitive Kalman smoother (OIKS).

yielding the impulsive noise distribution $p(s_k)$. Given a signal $\mathbf{y} = (y_1, \dots, y_N)$, the MAP/ML estimates

$$\hat{\gamma}_k \triangleq \underset{\gamma_k}{\operatorname{argmax}} p(\mathbf{y} | \gamma_1, \dots, \gamma_k, \dots, \gamma_N) \prod_{\ell=1}^N p(\gamma_\ell) \quad (16)$$

will result in $\gamma_k > 0$ when an outlier is detected at time step k and 0 otherwise. To compute the ML estimate in a local manner, akin to message passing, we devise EM [13], [14] updates that iteratively compute all $\hat{\gamma}_k$ in parallel.

A. Expectation Step

By the Markov property of the state space model, the expectation in the i -th iteration is:

$$\begin{aligned} Q(\gamma) &= \sum_{\ell=1}^N Q(\gamma_\ell) = E_{\mathbf{x} | \gamma^{i-1}, \mathbf{y}} [\log(p(\mathbf{y}, \mathbf{x}, \gamma))] \\ &= \sum_{\ell=1}^N E_{\mathbf{x}_\ell | \gamma^{i-1}, \mathbf{y}} [\log(p(y_\ell | \mathbf{x}_\ell, \gamma_\ell))] + \log(p(\gamma_\ell)) + \text{const} \quad (17) \end{aligned}$$

Introducing $\nu_k \triangleq \sigma_Z^2 + \gamma_k$ and the second moment $\boldsymbol{\mu}_{\mathbf{x}_k}^{\text{II}} = (\mathbf{V}_{\mathbf{x}_k} + \mathbf{m}_{\mathbf{x}_k} \mathbf{m}_{\mathbf{x}_k}^T)$, the conditional expectation at time step k is:

$$\begin{aligned} E_{\mathbf{x}_k | \gamma^{i-1}, \mathbf{y}} [\log(p(y_k | \mathbf{x}_k, \gamma_k))] &\propto \\ \log(\nu_k) + (y_k^2 - 2y_k \mathbf{C} \mathbf{m}_{\mathbf{x}_k} + \mathbf{C} \boldsymbol{\mu}_{\mathbf{x}_k}^{\text{II}} \mathbf{C}^T) / \nu_k &\quad (18) \end{aligned}$$

The marginal mean and covariance $\mathbf{m}_{\mathbf{x}_k}$ and $\mathbf{V}_{\mathbf{x}_k}$ are computed via a forward-backward sweep of the MBF or the BIFM. For the first iteration γ^0 is initialized as the all-zero vector.

B. Maximization Step

The maximization of (17) w.r.t. to γ_k yields:

$$\frac{1}{\nu_k} - \frac{(y_k^2 - 2y_k \mathbf{C} \mathbf{m}_{\mathbf{x}_k} + \mathbf{C} \boldsymbol{\mu}_{\mathbf{x}_k}^{\text{II}} \mathbf{C}^T)}{\nu_k^2} = -\frac{p'(\gamma_k)}{p(\gamma_k)} \quad (19)$$

When using an improper prior, i.e. $p(\gamma_k) \propto 1$, which imposes an improper outlier distribution $p(s_k) \propto 1/s_k$ that is heavily sparsifying, we obtain an analytic expression for the variance updates:

$$\gamma_k^i = \max \left(y_k^2 - \sigma_Z^2 - 2\mathbf{C} y_k \mathbf{m}_{\mathbf{x}_k} + \mathbf{C} \boldsymbol{\mu}_{\mathbf{x}_k}^{\text{II}} \mathbf{C}^T, 0 \right) \quad (20)$$

C. Termination Criterion

We iterate the EM steps until convergence, defined as:

$$\max_{k \in \{1, \dots, N\}} |\gamma_k^i - \gamma_k^{i-1}| \leq 10^{-4} \gamma_k^{i-1}, \quad (21)$$

or up to a maximum number of iterations ($i < 10$ in our case).

D. Noise Floor Estimation:

To estimate the noise floor $V_Z = \sigma_Z^2$, we resort again to EM and alternate between estimating γ and V_Z . For estimating V_Z , we use only observations that were not detected as outliers in the previous iteration. We use an initial noise floor estimate $\sigma_{Z_0}^2$ (which can for instance be inferred from the output signal's energy) to mark an observation y_k as outlier if $\gamma_k \geq 10\sigma_{Z_0}^2$. The resulting update equation for V_Z , with N_o being the number of outliers and N the total number of observations, is therefore:

$$V_Z^i = \frac{1}{N - N_o} \sum_{k: \gamma_k^i < 10\sigma_{Z_0}^2} y_k^2 - 2y_k \mathbf{C} \mathbf{m}_{\mathbf{x}_k} + \mathbf{C} \boldsymbol{\mu}_{\mathbf{x}_k}^{\text{II}} \mathbf{C}^T \quad (22)$$

E. Extension to Multi-Channel Observations

Given a multi-dimensional observation vector $\mathbf{y}_k = (y_k^{(1)}, \dots, y_k^{(L)})^T \in \mathbb{R}^L$, which is affected by outliers (cf. Fig. 3), we can apply the same procedure as in the scalar observation case if we assume that both the vector-valued noise floor $\mathbf{Z}[\mathbf{k}]$ and impulsive noise source $\mathbf{S}[\mathbf{k}]$ on each channel are independent. This allows us to split the \mathbf{C} -matrix into row vectors (cf. Fig. 3) and use the same computation rules described for the scalar output case.

F. Extension to Nonlinear State Space Models

Assuming additive input, state and measurement noise, the outlier insensitive state estimation can be extended to nonlinear dynamical systems. Given the discrete-time nonlinear dynamics, together with an additive impulsive noise source, we use an extended Kalman smoother combined with the sparsifying outlier noise variance estimation. Here we will limit ourselves to autonomous systems with nonlinear state dynamics $\mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}) + \mathbf{w}_{k-1}$, where $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a nonlinear function and \mathbf{w} is the process noise. The

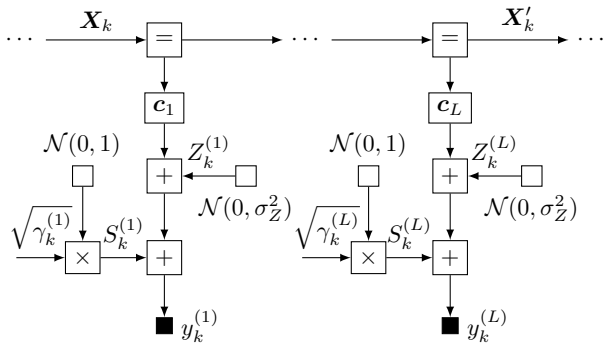


Fig. 3: The outlier insensitive smoother can be extended to multi-channel outputs by splitting the C matrix and assuming i.i.d. outliers.

measurement process is assumed to be linear, as considered before. The main step in the linearization of the nonlinear state dynamics is:

$$\mathbf{f}(\mathbf{X}_k) \approx \mathbf{f}(\vec{\mathbf{m}}_{\mathbf{X}'_{k-1}}) + \mathbf{A}_{k-1}(\mathbf{X}_k - \vec{\mathbf{m}}_{\mathbf{X}'_{k-1}}), \quad (23)$$

where \mathbf{A}_{k-1} is the Jacobi matrix of $\mathbf{f}(\cdot)$ evaluated at $\vec{\mathbf{m}}_{\mathbf{X}'_{k-1}}$.

The extended Kalman smoother (EKS) can be implemented either via an extended version of the MBF or BIFM. In the MBF-EKS version, the filtering problem is solved using the standard extended Kalman filter [11], and the backward step is performed on the linearized system (23), using the MBF smoother. Conversely, the BIFM-EKS uses the backward extended information filter [15] and the BIFM update rules on the linearized system as described in Table I.

V. NUMERICAL SIMULATIONS

A. Linear Time-Invariant System

To evaluate the OIKS we will first use a linear single-input single-output system with system matrices $\mathbf{A} = [1.12, -0.49, 0.11, -0.35; 1, 0, 0, 0; 0, 1, 0, 0; 0, 0, 1, 0]$, $\mathbf{B} = [-0.38, 0.59, 0.51, 0.3]$, $\mathbf{C} = [1, 0, 0, 0]$. Note that this system is both stable and observable.

Fig. 2 shows the noisy output \mathbf{y} and the corresponding noise-free output (reference) of this model subject to sinusoidal inputs. Uncertainty in the dynamics and measurements are accounted for by Gaussian i.i.d. input, state and measurement noise, whereas the outliers are modeled with Laplace distributed noise. The output's SNR when considering only the Gaussian noise floor is 7.4 dB, whereas if also impulsive noise is accounted for (here, 20% Laplace-distributed outliers) the signal-to-interference-plus-noise ratio SINR drops to -6.1 dB. The estimation performance of the OIKS is significantly less deteriorated by impulsive noise (cf. Fig. 2) compared to a regular Kalman smoother (KS), implemented as MBF smoother, and even compared to an adhoc robustified Kalman smoother denoted MedKS. To remove outliers the MedKS filters the signal \mathbf{y} with a 5-th order median filter and then runs a regular Kalman smoother.

We compared the performance of our OIKS with a regular Kalman smoother, the MedKS and the so called *doubly robust*

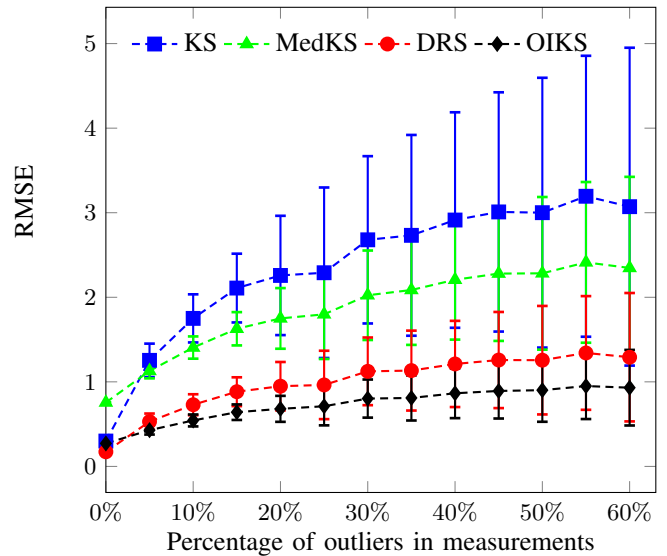


Fig. 4: Comparison of RMSE of state estimates of a regular Kalman smoother (KS), the MedKS, the DRS from [1] and the OIKS against percentage of outliers in the measurements.

smoother from [1], which is based on an ℓ_1 -penalty on the outliers. We performed a sweep of the output contamination levels with the state space model and noise floor level as described before. In contrast to [1], where a parameter λ , setting the sparsity level, has to be determined via a grid-search, the OIKS is free of tuning parameters. In Fig. 4, we can see that both the robust smoother from [1] and the OIKS clearly outperform both the regular Kalman smoother and the ad hoc robustified version. For low contamination levels the OIKS' performance is comparable to [1], but for high contamination levels the OIKS outperforms the algorithm from [1]. Unlike [1] which estimates the mean of the outlier value (which becomes increasingly hard when contamination levels are very high), the OIKS does not follow a generative approach, but instead assigns to outliers a large variance and thus discards them from the smoothing procedure.

For low contamination levels the performance of the MedKS is comparable to the OIKS, but with an increasing number of outliers (or contiguous noise bursts) the strength of the model-based approach becomes evident. Note that when a signal is (almost) free of outliers, most γ 's are set to zero and the OIKS reverts to a regular KS.

B. Nonlinear State Space Model: Van der Pol Oscillator

The Van der Pol oscillator describes a stable oscillation subject to nonlinear damping, which can be used to model a variety of relaxation oscillations such as for instance ECG waveforms [16]. The discrete-time model can be obtained via the Euler method resulting in:

$$\begin{aligned} X_1[k+1] &= X_1[k] + X_2[k]\Delta + W_1[k] \\ X_2[k+1] &= X_2[k] + [\mu(1 - X_1^2[k])X_2[k] - X_1[k]]\Delta + W_2[k] \\ \mathbf{Y}[k] &= \mathbf{C}\mathbf{X}[k] + \mathbf{Z}[k] + \mathbf{S}[k], \end{aligned} \quad (24)$$




Fig. 5: The OIEKS detects most of the outliers on both channels and assigns them a large variance γ . Correctly detected outliers are marked with circles.

where $\mathbf{Y}[k]$ is the multichannel output signal and $\mathbf{C} = [1, 0; 0, 1]$, μ the damping coefficient and Δ the step size. Assuming diagonal covariance matrices for both the noise floor $\mathbf{Z}[k]$ and impulsive noise source $\mathbf{S}[k]$, we can treat each channel separately as described in section IV-E.

Fig. 5 shows the simulated two-channel output of a Van der Pol oscillator with $\mu = 0.5$ and $\Delta = 0.01$ subject to impulsive noise with a contamination level of 10%. The outlier insensitive extended Kalman smoother, implemented as an MBF-EKS, assigns large variances to observations affected by outliers. The remaining impulsive noise variances are set to zero, indicating absence of outliers.

VI. CONCLUSION

We have presented an outlier-insensitive Kalman smoother denoted OIKS which can be used both for linear systems as well as nonlinear systems when combined with an extended Kalman smoother. Based on ideas from sparse Bayesian learning, outliers are modeled as Gaussians with unknown variances, which are estimated by expectation maximization; the resulting prior is a sparsifying improper heavy-tailed distribution. The actual computations boil down to iterations of standard Kalman smoothers.

We have considered two Kalman smoothers: the *Modified Bryson-Frasier* (MBF) and the *Backward Information Filter Forward Marginal* (BIFM), neither of which requires inversions of the state covariance matrix. In the simulated examples the numerical stability of both approaches is comparable, but BIFM has a lower complexity.

In this paper, we use a per-sample penalty for the outliers. This is tailored at impulsive noise (salt and pepper noise), but not at burst noise, which is characterized by contiguous blocks of outliers. Detection of contiguous noise blocks as they are often found in physiological signals, e.g. due to motion artifacts [17], will be addressed (in a future work) by

enforcing sparsity on changes of the time-dependent output noise variance.

REFERENCES

- [1] S. Farahmand, G.B. Giannakis, and D. Angelosante, "Doubly robust smoothing of dynamical processes via outlier sparsity constraints," *IEEE Transactions on Signal Processing*, vol. 59, no. 10, pp. 4529–4543, Oct. 2011.
- [2] J. Ting, E. Theodorou, and S. Schaal, "A Kalman filter for robust outlier detection," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007, pp. 1514–1519.
- [3] A.Y. Aravkin, B.M. Bell, J.V. Burke, and G. Pillonetto, "An ℓ_1 Laplace robust Kalman smoother," *IEEE Transactions on Automatic Control*, vol. 56, no. 12, pp. 2898–2911, Dec. 2011.
- [4] Z.M. Durovic and B.D. Kovacevic, "Robust estimation with unknown noise statistics," *IEEE Transactions on Automatic Control*, vol. 44, no. 6, pp. 1292–1296, June 1999.
- [5] S.C. Chan, Z.G. Zhang, and K.W. Tse, "A new robust Kalman filter algorithm under outliers and system uncertainties," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2005, pp. 4317–4320.
- [6] D. P. Wipf, B. D. Rao, and S. Nagarajan, "Latent variable Bayesian models for promoting sparsity," *IEEE Transactions on Information Theory*, vol. 57, no. 9, pp. 6236–6255, 2011.
- [7] T. Kailath, A. H. Sayed, and B. Hassibi, *Linear Estimation*, vol. 1, Prentice Hall Upper Saddle River, NJ, 2000.
- [8] H.-A. Loeliger, L. Bruderer, H. Malmberg, F. Wadehn, and N. Zalmi, "On sparsity by NUV-EM, Gaussian message passing, and Kalman smoothing," in *Information Theory and Applications Workshop (ITA)*, San Diego, CA, USA, Feb. 2016.
- [9] H. E Rauch, C.T. Striebel, and F. Tung, "Maximum likelihood estimates of linear dynamic systems," *AIAA Journal*, vol. 3, no. 8, pp. 1445–1450, 1965.
- [10] H.-A. Loeliger, J. Dauwels, J. Hu, S. Korl, L. Ping, and F. R. Kschischang, "The factor graph approach to model-based signal processing," *Proceedings of the IEEE*, vol. 95, no. 6, pp. 1295–1322, 2007.
- [11] S. Särkkä, *Bayesian Filtering and Smoothing*, Number 3. Cambridge University Press, 2013.
- [12] L. Bruderer, H. Malmberg, and H.-A. Loeliger, "Deconvolution of weakly-sparse signals and dynamical-system identification by Gaussian message passing," in *IEEE International Symposium on Information Theory (ISIT)*, July 2015.
- [13] P. Stoica and Y. Selén, "Cyclic minimizers, majorization techniques, and the expectation-maximization algorithm: a refresher," *IEEE Signal Processing Magazine*, vol. 21, no. 1, pp. 112–114, 2004.
- [14] Z. Ghahramani and G. Hinton, "Parameter estimation for linear dynamical systems," Tech. Rep., Technical Report CRG-TR-96-2, University of Toronto, Dept. of Computer Science, 1996.
- [15] S. Thrun, Y. Liu, D. Koller, A. Ng, Z. Ghahramani, and H. Durrant-Whyte, "Simultaneous localization and mapping with sparse extended information filters," *The International Journal of Robotics Research*, vol. 23, no. 7-8, pp. 693–716, 2004.
- [16] BZ Kaplan, I Gabay, G Sarafian, and D Sarafian, "Biological applications of the filtered Van der Pol oscillator," *Journal of the Franklin Institute*, vol. 345, no. 3, pp. 226–232, 2008.
- [17] F. Wadehn, L. Bruderer, D. Waltisberg, T. Keresztfalvi, and H.-A. Loeliger, "Sparse-input detection algorithm with applications in electrocardiography and ballistocardiography," in *International Conference on Bio-inspired Systems and Signal Processing, Lisbon, Portugal*, Jan. 2015.