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

Raitoharju, Matti; Svensson, Lennart; Garcia Fernandez, Angel; Piche, Robert
Damped Posterior Linearization Filter

Published in:
IEEE Signal Processing Letters

DOI:
[10.1109/LSP.2018.2806304](https://doi.org/10.1109/LSP.2018.2806304)

Published: 14/02/2018

Document Version
Peer reviewed version

Please cite the original version:
Raitoharju, M., Svensson, L., Garcia Fernandez, A., & Piche, R. (2018). Damped Posterior Linearization Filter. *IEEE Signal Processing Letters*, 25(4), 536-540. <https://doi.org/10.1109/LSP.2018.2806304>

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.

Damped Posterior Linearization Filter

Matti Raitoharju, Lennart Svensson, Ángel F. García-Fernández, and Robert Piché

Abstract—In this letter, we propose an iterative Kalman type algorithm based on posterior linearization. The proposed algorithm uses a nested loop structure to optimize the mean of the estimate in the inner loop and update the covariance, which is a computationally more expensive operation, only in the outer loop. The optimization of the mean update is done using a damped algorithm to avoid divergence. Our simulations show that the proposed algorithm is more accurate than existing iterative Kalman filters.

Index Terms—Bayesian state estimation; nonlinear; estimation; Kalman filters

I. INTRODUCTION

In Bayesian state estimation, a state that evolves stochastically in time is estimated from noisy measurements. In this letter, we concentrate on the measurement update stage of the Bayesian filter. In the measurement update, a prior (the dynamic model's state prediction) is updated using information from a measurement. If the measurement model is linear and Gaussian the posterior density can be computed analytically using the Kalman filter [1], but for a general measurement model, the computation of the posterior density is intractable. One approximate approach is the general Gaussian filter (GGF), which represents the joint state and measurement distribution by a multivariate Gaussian using moment matching, and then uses standard marginalization and conditioning formulas to compute the conditional state distribution given the measurement to obtain the posterior estimate [2]. The GGF moment matching is also intractable, but there exist several Kalman filter extensions (KFEs) that are approximations of the GGF. The standard way to apply GGF is to do statistical linear regression (SLR), or an approximation of the SLR, in the prior; this can be interpreted as the optimal linearization with the given prior. The GGF approach does not work well for severe nonlinearities or low measurement noise [3] so there is interest in developing alternatives.

Iterative algorithms generate approximative solutions based on previous solutions with the goal of improving the approximation. The iterated extended Kalman filter (IEKF) [4, pp. 349-351] produces a sequence of mean estimates by

making first-order Taylor approximations of the measurement function. In [5], it was shown that the IEKF measurement update is equivalent to the Gauss-Newton (GN) algorithm for computing the maximum a posteriori (MAP) estimate. IEKF performs well in situations where the true posterior is close to a Gaussian. Convergence of the GN algorithm is not guaranteed. Better convergence is offered by damped (or descending) GN algorithms [6], which ensure that the cost function is nonincreasing at each iteration. An alternative to damped GN is the Levenberg-Marquadt algorithm [7].

There are many other iterated filters in the literature. The iterated sigma point Kalman filter [8], at each iteration, uses a linearization that mixes SLR with respect to the prior and analytical linearization at the current MAP estimate. The recursive update filter (RUF) uses several updates of the prior with down-weighted Kalman gain [9]. Progressive Gaussian Filtering is a homotopy continuation method that updates the prior starting from an easily computable measurement likelihood that gradually evolves into the true measurement likelihood [10]–[12]. The Kullback-Leibler partitioned update Kalman filter (KLPUKF) can be used when the measurement is multidimensional [13]. It sequentially updates using nonlinearity-minimizing linear combinations of the measurements.

In this paper, we focus on the iterated posterior linearization filter (IPLF), which uses SLR w.r.t. the posterior density. SLR's linearization is based on a larger area determined by a probability density function (PDF) instead of a point as in the Taylor linearization, which improves the accuracy of the filter. In [14], it was shown that it is better for the linearization to be w.r.t. the posterior instead of the prior. The basic idea in the IPLF is to compute a posterior estimate using the prior, then use this estimate to compute a better linearization and posterior estimate, and so on. In this letter, we observe that the IPLF does not always converge, and we propose a damped version of the algorithm with improved convergence properties.

We show in simulations how the proposed algorithm is less prone to filter divergence than the original IPLF. Furthermore, we compare the posterior estimate accuracy with other iterative Kalman filter methods and show that the proposed algorithm is more accurate in the simulations.

II. BACKGROUND

We consider measurements of the form

$$y = h(x) + \varepsilon, \quad (1)$$

where y is the n_y dimensional real measurement value, $h(\cdot)$ is the measurement function, x is the n_x dimensional real random state vector and ε is zero mean Gaussian noise with covariance R .

This article is carried out with support from a mobility grant from COST Action TU1302 (SAPPART). Financial support by the Academy of Finland under grant no. #287792 (OpenKin) is hereby gratefully acknowledged.

M. Raitoharju is with the laboratory of Automation and Hydraulic Engineering, Tampere University of Technology and Department of Electrical Engineering and Automation, Aalto University, Finland. E-mail: matti.raitojarju@tut.fi. L. Svensson is with the department of Signals and Systems, Chalmers University of Technology, Sweden. Á. F. García-Fernández is with the Department of Electrical Engineering and Electronics, University of Liverpool, UK, Robert Piché is with the laboratory of Automation and Hydraulic Engineering, Tampere University of Technology, Finland

This is an author accepted version of a letter published in IEEE Signal Processing Letters. DOI:10.1109/LSP.2018.2806304 ©2018 IEEE.

A. GGF

The GGF [2] for measurement model (1) uses the expected value of predicted measurement $\hat{y}$, cross covariance of measurement and state $P_{h(x)x}$, and measurement covariance $P_{h(x)h(x)}$. These moments are

$$\hat{y} = \int h(x) p_N(x; \mu^-, P^-) dx \quad (2)$$

$$P_{h(x)x} = \int (h(x) - \hat{y}) (x - \mu^-)^T p_N(x; \mu^-, P^-) dx \quad (3)$$

$$P_{h(x)h(x)} = \int (h(x) - \hat{y}) (h(x) - \hat{y})^T p_N(x; \mu^-, P^-) dx, \quad (4)$$

where $p_N(x|\mu^-, P^-)$ is the multivariate normal PDF of the prior that has mean μ^- and covariance P^- . The update equations for the mean and covariance of the posterior are then

$$\mu^+ = \mu^- + K(y - \hat{y}) \quad (5)$$

$$P^+ = P^- - KSK^T, \quad (6)$$

where

$$S = P_{h(x)h(x)} + R \quad (7)$$

$$K = P_{xh(x)} S^{-1}. \quad (8)$$

The integrals (2)-(4) do not have in general closed form solutions. Many KFEs, such as unscented Kalman filter (UKF) [15] and cubature Kalman filter (CKF) [16], can be interpreted as numerical approximations of GGF.

B. IPLF

In [14], it is shown that a lower bound of the Kullback-Leibler divergence (KLD) of a GGF with respect to the true posterior is minimized if the moments (2)-(4) are computed using the posterior instead of the prior. This update cannot be computed directly with the GGF equations (2)-(8), but instead a linear model that has correct moments is generated and is then applied to the prior.

The IPLF algorithm iteratively approximates the posterior using SLR, creating a sequence of updated means μ_i and covariance matrices P_i as follows. First, set $\mu_1 = \mu^-$ and $P_1 = P^-$. Then, at the i th iteration, compute the SLR of $h(\cdot)$ with respect to (μ_i, P_i) by first computing moments (2)-(4) and then defining the linear measurement that corresponds to those moments [14]:

$$\hat{h}(x) = J_i x + b_i + \varepsilon_{\Omega_i} + \varepsilon, \quad (9)$$

where

$$J_i = P_{xh(x)_i}^T P_i^{-1} \quad (10)$$

$$b_i = \hat{y}_i - J_i \mu_i \quad (11)$$

$$\Omega_i = P_{h(x)h(x)_i} - J_i P_i J_i^T \quad (12)$$

$$\varepsilon_{\Omega_i} \sim N(0, \Omega_i), \quad (13)$$

ε_{Ω_i} is an independent noise whose covariance Ω_i is the covariance of the linearization error. For linear systems Ω_i is 0.

TABLE I
MEAN ESTIMATES OF FIRST 6 ITERATIONS OF IPLF AND IEKF

Iteration	1	2	3	4	5	6
IPLF	-2.51	5.28	-31.75	17.52	-40.56	11.96
IEKF	-7.64	58.28	-1.77	2.60	-6.66	48.47

Finally, compute the posterior mean and covariance that correspond to the linearized measurement function (9):

$$S = J_i P^- J_i^T + R + \Omega_i \quad (14)$$

$$K_i = P^- J_i^T S^{-1} \quad (15)$$

$$\mu_{i+1} = \mu^- + K_i(y - J_i \mu^- - b_i) \quad (16)$$

$$P_{i+1} = P^- - K_i S_i K_i^T. \quad (17)$$

The obtained posterior is used for the next linearization. The process is repeated for a predetermined number of steps or until a convergence criterion is met, e.g. until the KLD of two consecutive estimates is below a threshold.

However, the IPLF, like the GN, can diverge, as illustrated in the next example.

Example 1: The measurement model is

$$y = \arctan(x) + \varepsilon. \quad (18)$$

State x has prior mean 2.75 and variance 1, the measurement value is $y = 0$ and measurement noise variance is $R = 10^{-4}$. The integrals (2)-(4) are computed using Monte Carlo (MC) integration with 10^5 samples and also with IEKF. Table I gives the means in the first 6 iterations. Neither of the algorithms converge to the true posterior, which is close to 0, even after 50 iterations. In the supplementary material, we provide another example, whose moments have closed form solutions, where the IPLF does not converge.

C. IEKF as a GN algorithm

IEKF is similar to IPLF, but it uses a first order Taylor series approximation of the measurement function instead of SLR. The IEKF iteration formulas are as in Section II-B, but with (10)-(12) replaced by

$$J_i = \left. \frac{dh(x)}{dx} \right|_{\mu_{i-1}} \quad (19)$$

$$b_i = h(\mu_{i-1}) - J_i \mu_i \quad (20)$$

$$\Omega_i = 0. \quad (21)$$

In [5], it is shown that the IEKF for measurements of form (1) is equal to the GN algorithm [6] that minimizes the cost function

$$q(\mu) = \frac{1}{2} (h(\mu) - y)^T R^{-1} (h(\mu) - y) + \frac{1}{2} (\mu - \mu^-)^T P^{-1} (\mu - \mu^-). \quad (22)$$

To decrease the possibility of the divergence of GN, the damped (or descending) version of the GN algorithm can be used [6]. A damped IEKF that uses line search [6] in minimization of (22) is given in [17]. The mean update is

$$\mu_{i+1} = (1 - \alpha) \mu_i + \alpha (\mu_0 + K_i(y - h(\mu_i))) \quad (23)$$

with step length parameter α selected so that $0 \leq \alpha \leq 1$ and $q(\mu_{i+1}) \leq q(\mu_i)$. Finally, in the IEKF, the posterior covariance matrix is computed based on the linearization at the MAP point after the last iteration [5].

III. DAMPED ITERATED POSTERIOR LINEARIZATION FILTER (DIPLF)

The main weakness of the IPLF is that it sometimes diverges. In this section, we present a damped version of the IPLF, which has substantially better convergence properties. The key insight is that the update of the mean can be seen as an optimization problem and solved with a damped GN algorithm that decreases the cost function in every iteration. A detailed motivation of the algorithm is presented in Section III-A, followed by a description of line search required to implement it. Further implementation details are discussed in the Supplementary material.

A. Derivation of DIPLF

Ideally, we want the posterior mean and covariance to be used in the SLR to linearize the measurement function [14]. Our approach to iteratively finding the mean and covariance for the linearization is to use a separate loop for the optimization of the mean. Thus, the proposed filter computes the estimate using two nested loops. The inner loop uses a damped GN algorithm to optimize the mean μ_i while keeping P_j and Ω_j fixed. The outer loop updates the covariances P_j and Ω_j . Index i increases in every iteration of the inner loop and j in the outer loop.

1) *Inner loop*: The purpose of the inner loop is to find optimal μ_i when P_j and Ω_j are fixed. At an optimum μ_i coincides with the mean used in the SLR and we can write:

$$\mu_i = \mu^- + K(\mu_i)(y - J(\mu_i)\mu^- - b(\mu_i)), \quad (24)$$

where we used (μ_i) to identify vectors and matrices that depend on the mean. In the supplementary material, we show that the above equation is fulfilled when

$$q(\mu_i) = \frac{1}{2}(\hat{y}(\mu_i) - y)^T(R + \Omega_j)^{-1}(\hat{y}(\mu_i) - y) + \frac{1}{2}(\mu_i - \mu^-)^T(P^-)^{-1}(\mu_i - \mu^-) \quad (25)$$

achieves its minimum. Therefore, we use it as the cost function for the inner loop. The optimization problem is analogous to the optimization of the cost function of the IEKF (22), but using $\hat{y}$ instead of $h(\mu)$. Because the step is taken in a descent direction of the target function we can use a scaling factor α ($0 < \alpha \leq 1$) in the update as in (23):

$$\mu_{i+1} = (1 - \alpha)\mu_i + \alpha(\mu_0 + K_i(y - J_i\mu_0 - b_i)) \quad (26)$$

and select α using line search so that the value of our target function (25) decreases. Thus, the next mean is a weighted sum of previous mean and the mean provided by IPLF. This makes the algorithm locally convergent on almost all nonlinear least squares problems, provided that the line search is carried out appropriately. In fact, it is usually globally convergent [6, Chapter 9.2.1]. To facilitate the stopping of the inner loop,

Algorithm 1: DIPLF

```

1  $i \leftarrow 0, j \leftarrow 0, \mu_0 \leftarrow \mu^-, P_0 \leftarrow P^-$ 
2 Compute  $\hat{y}_i, J_i$ , and  $\Omega_j$  at prior using (2), (10), and (12)
3 while  $p_N(\hat{y}_i|y, R + \Omega_j)p_N(\mu_i|\mu^-, P^-)$  increases significantly
  do
4   while  $q(\mu_i, P_j) < \beta q(\mu_{i-1}, P_j)$  do
5      $i \leftarrow i + 1$ 
6     Find  $0 < \alpha_i \leq 1$  for (26) so that (25) becomes smaller
       using line search (Section III-B)
7     Compute  $\mu_i$  using (26)
8     Compute  $J_i$  using (3) and (10)
9     Compute  $\hat{y}_i$  using (2)
10  end
11   $j \leftarrow j + 1$ 
12  Compute  $P_j$  using (17)
13  Compute  $\Omega_j$  using (12)
14 end
15 Use the mean and covariance from the second last round of the
    outer loop, i.e. those that had highest
     $p_N(\hat{y}_i|y, R + \Omega_j)p_N(\mu_i|\mu^-, P^-)$ , as the posterior estimate
```

we propose that the algorithm should repeat the inner loop as long as (25) reduces significantly at each iteration i.e. while

$$q(\mu_{i+1}, P_j) < \beta q(\mu_i, P_j) \quad (27)$$

with, e.g., $\beta = 0.9$.

2) *Outer loop*: The outer loop updates the covariances P_j and Ω_j for the next iteration of the inner loop. To define a stopping condition for the outer loop, we first exponentiate function (25) and normalize to get a product of normal distributions:

$$\frac{1}{\sqrt{(2\pi)^{n_y}|R + \Omega_j|(2\pi)^{n_x}|P^-|}} e^{-q(x)} \quad (28)$$

$$= p_N(\hat{y}_i|y, R + \Omega_j)p_N(\mu_i|\mu^-, P^-).$$

If we treat P_i and Ω_i as constants, (25) achieves its minimum when (28) achieves its maximum. We thus continue the outer loop as long as (28) is increasing significantly. Compared to using (25), the factor containing Ω_i in (28) causes the stopping condition based on (28) to favor estimates with smaller nonlinearity.

3) *Full algorithm*: The algorithm is summarized in Algorithm 1. The inner loop does the optimization of the target function (25) while keeping the state covariance and nonlinearity covariance constant. The direction of change in a GN algorithm is a descent direction of the cost function [6, Chapter 9.2.1]. Thus, with small enough α the cost function value should decrease, unless the estimate is in a local extremum or a saddle point. By using $\alpha = 1$ and exiting the inner loop after a single iteration, Algorithm 1 becomes equivalent to an IPLF algorithm.

The outer loop changes the optimization problem of the inner loop as Ω changes and the covariance, for which the expected measurement value $\hat{y}$ is computed, changes.

B. Line search

The value of α can be computed with different line search techniques [6, Chapter 9.2.1]. Backtracking line search, also known as the Armijo-Goldstein step length principle, is one

TABLE II
KLDs OF ESTIMATES PRODUCED BY DIFFERENT FILTERS IN arctan TEST

	GGF	RUF	IPLF	DIPLF
MC	15.9	0.04	88.69	$3 \cdot 10^{-6}$
EKF	4009.10	0.01	65.12	10^{-6}
CKF	3370.78	0.01	64.39	10^{-6}
UKF	92.55	0.01	10^{-6}	10^{-6}

of the simplest and most widely used options, and is done as follows.

First, α is set to 1, and a value of μ is found using (23). If this value does not decrease (25), w.r.t. the previous value of μ , then α is multiplied by a constant factor $0 < \tau < 1$ and the procedure is repeated. To reduce the computational load one can terminate the inner loop when α is smaller than a predetermined value (we use 2^{-4}) as then the change is the estimate is negligible.

IV. SIMULATIONS

In our first test, we evaluate the situation in Example 1. We compare posterior estimates obtained with estimation algorithms that are based on approximating integrals (2)-(4). We use MC integration with 10^5 samples, extended Kalman filter (EKF), UKF with parameters from [15], and third-degree CKF [16] to compute the moment approximations. The algorithms used are GGF, RUF [9], IPLF [14], and DIPLF. The GGF with different moment computation algorithms is equivalent to the algorithms from which the moment computation was taken. RUF applies the measurement using down-weighted Kalman gains and uses 10 iterations in our tests. When EKF is used to compute the moments GGF is EKF, RUF is equivalent to the algorithm in [18], DIPLF is IEKF, and DIPLF is damped IEKF. We use $\tau = 0.5$ for reducing the step length. For the threshold of inner loop, we use $\beta = 0.9$ and we exit the outer loop if $0.999p_N(\hat{y}_i|y, R + \Omega_j)$ is not larger than the corresponding likelihood of the previous iteration. A more thorough discussion on the selection of parameters is given in the supplementary material.

Results are given in Table II. DIPLF gives good posterior estimates with all moment computation methods. IPLF converges towards an estimate within 50 iterations only when using UKF for moment computation. Though the UKF and CKF usually have similar accuracy, in this case, the IPLF did not converge when using the CKF; changing the prior mean a little may make it to converge with different moment computation algorithms. RUF has slightly worse estimates than DIPLF and GGF does not provide good estimates with any of the tested moment computation methods.

In our second test, we computed estimates for two dimensional positioning using three range measurements in a single measurement update. The prior has zero mean with covariance I . Ranges were computed to beacons located at $[-1 \ 0]^T$, $[0 \ 1]^T$ and $[1 \ -2]^T$. Measurement noise covariance was I and measurements were generated by sampling a true location from the prior and then generating corresponding measurements. In this test, we used the same filters as in the previous test and KLUKF [13]. KLUKF uses a

TABLE III
MEAN KLDs OF ESTIMATES IN RANGE TEST

	GGF	KLUKF	RUF	IPLF	DIPLF
MC	0.25	0.17	0.21	0.26	0.17
EKF	0.48	0.60	0.48	0.55	0.55
CKF	0.28	0.22	0.29	0.38	0.23
UKF	0.35	0.34	0.33	0.37	0.26

linear transformation to decorrelate measurement elements in such a way that the nonlinearity of a measurement element is minimized and applies measurement elements sequentially. For the first test with only a scalar measurement, KLUKF would have been identical to GGF.

The test is repeated 1000 times using different true location and measurement values. We analyze only state after one update step, but a more accurate update step can be expected to lead to improved stability in the filtering recursion. The mean KLDs, which are computed numerically using a dense grid, are presented in Table III. The iterative algorithms that use EKF moment computation, except RUF, produce worse estimates than EKF itself. This is probably because the posterior may be multimodal in this test. With other moment computation methods the IPLF is worse than GGF, but DIPLF is better than either of those. In this test, KLUKF and DIPLF were the most accurate algorithms, except when tested with EKF. The residuals of the mean estimates compared to true locations were similar for all algorithms, thus the largest improvement of using DIPLF comes from better posterior covariance estimation. In this test, the computational time of DIPLF was 8 times higher than the computational cost of KLUKF. Our implementations of other algorithms had runtimes between these, IPLF taking similar time as DIPLF. The exact comparison of computational complexity of algorithms evaluated in this letter cannot be done, because algorithms have different conditions for loops. These conditions depend on the estimation problem and computational complexity cannot be stated as a function of dimensionalities of the state and the measurement vector.

V. CONCLUSIONS AND FUTURE WORK

We have shown that by fixing the covariance matrices of the IPLF algorithm, the optimal posterior mean in the posterior linearization approach is the solution of an optimization problem. We solved this optimization problem using a damped GN algorithm. Then we updated the covariance matrices P_j and Ω_j and repeated the mean optimization. Compared to IPLF, this algorithm is less prone to diverge. The proposed algorithm was more accurate than other tested algorithms in our simulations.

One topic for future research is to see how posterior linearization could be done with other optimization algorithms than GN, such as Levenberg-Marquadt. Another interesting research topic would be extending the algorithm to work with non-additive noise. We have done some preliminary simulations in time series, where we found that the proposed algorithm is more accurate than other iterative algorithms, but all iterative algorithms have a risk of converging to a wrong local solution. Investigating this is one future topic.

REFERENCES

- [1] Y. Ho and R. Lee, "A Bayesian approach to problems in stochastic estimation and control," *IEEE Transactions on Automatic Control*, vol. 9, no. 4, pp. 333 – 339, October 1964, doi:10.1109/TAC.1964.1105763.
- [2] S. Särkkä, *Bayesian filtering and smoothing*. Cambridge University Press, 2013, vol. 3. [Online]. Available: http://users.aalto.fi/~ssarkka/pub/cup_book_online_20131111.pdf
- [3] M. Morelande and Á.F. García-Fernández, "Analysis of Kalman filter approximations for nonlinear measurements," *IEEE Transactions on Signal Processing*, vol. 61, no. 22, pp. 5477–5484, Nov 2013, doi:10.1109/TSP.2013.2279367.
- [4] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*, ser. Mathematics in Science and Engineering. Academic Press, 1970, vol. 64.
- [5] B. Bell and F. Cathey, "The iterated Kalman filter update as a Gauss-Newton method," *Automatic Control, IEEE Transactions on*, vol. 38, no. 2, pp. 294–297, 1993, doi:10.1109/9.250476.
- [6] Å. Björck, *Numerical Methods for Least Squares Problems*. Society for Industrial and Applied Mathematics, 1996, doi:10.1137/1.9781611971484.
- [7] R. L. Bellare, E. W. Kamen, and S. M. Zabin, "New nonlinear iterated filter with applications to target tracking," vol. 2561, 1995, pp. 240–251, doi:10.1117/12.217701.
- [8] G. Sibley, G. Sukhatme, and L. Matthies, "The iterated sigma point Kalman filter with applications to long range stereo," in *Proceedings of Robotics: Science and Systems*, Philadelphia, USA, August 2006.
- [9] Y. Huang, Y. Zhang, N. Li, and L. Zhao, "Design of sigma-point Kalman filter with recursive updated measurement," *Circuits, Systems, and Signal Process.*, pp. 1–16, August 2015, doi:10.1007/s00034-015-0137-y.
- [10] U. D. Hanebeck and J. Steinbring, "Progressive Gaussian filtering based on Dirac mixture approximations," in *2012 15th International Conference on Information Fusion*, July 2012, pp. 1697–1704.
- [11] U. D. Hanebeck, K. Briechle, and A. Rauh, "Progressive Bayes: a new framework for nonlinear state estimation," in *AeroSense 2003*. International Society for Optics and Photonics, 2003, pp. 256–267.
- [12] Y. Huang, Y. Zhang, N. Li, and L. Zhao, "Gaussian approximate filter with progressive measurement update," in *2015 54th IEEE Conference on Decision and Control (CDC)*, Dec 2015, pp. 4344–4349.
- [13] M. Raitoharju, Á. F. García-Fernández, and R. Piché, "Kullback–Leibler divergence approach to partitioned update Kalman filter," *Signal Processing*, vol. 130, pp. 289 – 298, 2017, doi:10.1016/j.sigpro.2016.07.007.
- [14] Á.F. García-Fernández, L. Svensson, M. Morelande, and S. Särkkä, "Posterior linearization filter: Principles and implementation using sigma points," *IEEE Transactions on Signal Processing*, vol. 63, no. 20, pp. 5561–5573, Oct 2015, doi:10.1109/TSP.2015.2454485.
- [15] E. Wan and R. Van der Merwe, "The unscented Kalman filter for nonlinear estimation," in *Proceedings of the Adaptive Systems for Signal Processing, Communications, and Control Symposium. AS-SPCC.*, 2000, pp. 153–158, doi:10.1109/ASSPCC.2000.882463.
- [16] I. Arasaratnam and S. Haykin, "Cubature Kalman filters," *IEEE Transactions on Automatic Control*, vol. 54, no. 6, pp. 1254–1269, June 2009, doi:10.1109/TAC.2009.2019800.
- [17] M. A. Skoglund, G. Hendeb, and D. Axehill, "Extended Kalman filter modifications based on an optimization view point," in *Information Fusion (Fusion), 2015 18th International Conference on*, July 2015, pp. 1856–1861.
- [18] R. Zanetti, "Recursive update filtering for nonlinear estimation," *IEEE Transactions on Automatic Control*, vol. 57, no. 6, pp. 1481–1490, June 2012, doi:10.1109/TAC.2011.2178334.

SUPPLEMENTARY MATERIAL TO DAMPED POSTERIOR LINEARIZATION FILTER

A. Analytical example of a situation where the IPLF does not converge

Consider a one dimensional state with prior $x_0 \sim \mathcal{N}(\mu^-, P^-)$ with $\mu^- = 1$ and $P^- = 1$ and a measurement model

$$h(x) = y^2 + \varepsilon, \quad (\text{S-1})$$

where $\varepsilon \sim \mathcal{N}(0, R)$ with $R = 2^2$. Because the measurement function is a second order polynomial, the second order EKF [1] produces closed form SLR. For (S-1) the moments (2)-(4) in paper are

$$\hat{y} = \mu_i + P_j \quad (\text{S-2})$$

$$P_{h(x)x} = 2\mu_i P_j \quad (\text{S-3})$$

$$P_{h(x)h(x)} = 2P_j^2. \quad (\text{S-4})$$

Consider a measurement with value $y = -4$. After the first iteration, the mean of IPLF is -0.2 and, after the third iteration, the mean is 1.34 , thus the mean has moved to other side of the initial mean after the first iteration. When the iterations are continued, the estimates' means oscillate between -0.20 and 1.35 . On the other hand, the damped iterated posterior linearization filter (DIPLF) has the same estimate after the first iteration, but after that it converges to a single estimate with mean 0.36 . Figure S-1 shows how the mean estimates converge during the iterations in IPLF and DIPLF. Figure S-2 shows the PDFs of the true posterior, the oscillating estimates of the IPLF and the estimate made by DIPLF. Figure shows how the mean and mode of the estimate provided by the DIPLF are close to the mean and mode of the true posterior. However, due to the nonlinearity of the measurement function, DIPLF overestimates the covariance.

Fig. S-1. Evolution of estimates of IPLF and DIPLF

Fig. S-2. True posterior compared to the oscillating estimates of IPLF and the estimate provided by DIPLF

B. Derivation of the cost function

Here we derive the cost function for updating the inner loop. First, we modify the Kalman update similarly as in [2] to get a GN update from IEKF update, but taking the Ω_j term of IPLF into account:

$$\begin{aligned}
\mu_i &= \mu^- + K(\mu_i)(y - J(\mu_i)\mu^- - b(\mu_i)) \\
&= \mu^- + (J(\mu_i)^T(R + \Omega_j)^{-1}J(\mu_i) + P^{-1})^{-1}J(\mu_i)^T(R + \Omega_j)^{-1}(y - J(\mu_i)\mu^- - \hat{y}(\mu_i) + J(\mu_i)\mu_i) \\
&= (J(\mu_i)^T(R + \Omega_j)^{-1}J(\mu_i) + P^{-1})^{-1}(J(\mu_i)^T(R + \Omega_j)^{-1}(y - \hat{y}(\mu_i) + J(\mu_i)\mu_i) + P^{-1}\mu^-) \\
&= (J(\mu_i)^T(R + \Omega_j)^{-1}J(\mu_i) + P^{-1})^{-1}(J(\mu_i)^T(R + \Omega_j)^{-1}(y - \hat{y}(\mu_i) + J(\mu_i)\mu_i) + P^{-1}\mu_i - P^{-1}\mu_i + P^{-1}\mu^-) \\
&= \mu_i + (J(\mu_i)^T(R + \Omega_j)^{-1}J(\mu_i) + (P^-)^{-1})^{-1}(J(\mu_i)^T(R + \Omega_j)^{-1}(y - \hat{y}(\mu_i)) + P^{-1}(\mu^- - \mu_i)).
\end{aligned} \tag{S-5}$$

From the last row we can see that the IPLF update for the mean with fixed Ω_j has a stationary point when

$$J(\mu_i)^T(R + \Omega_j)^{-1}(\hat{y}(\mu_i) - y) + P^{-1}(\mu_i - \mu^-) = 0. \tag{S-6}$$

Differentiating the expected value of a measurement

$$\hat{y}(\mu_i) = \int h(x)p_N(x; \mu_i, P_j)dx \tag{S-7}$$

with respect to the mean μ_i gives

$$\begin{aligned}
&\frac{d \int h(x)p(x; \mu_i, P_j)dx}{d\mu_i} \\
&= \int h(x) \frac{dp(x; \mu_i, P_j)}{d\mu_i} dx \\
&= \int h(x)(x - \mu_i)^T p(x; \mu_i, P_j) dx P_j^{-1}.
\end{aligned} \tag{S-8}$$

Noting that $\int \hat{y}(\mu_i)(x - \mu_i)^T p(x; \mu_i, P_j)dx = 0$ we can write

$$\int (h(x) - \hat{y}(\mu_i))(x - \mu_i)^T p(x; \mu_i, P_j)dx P_j^{-1} = P_{h(x)x} P_j^{-1}. \tag{S-9}$$

From this we see that the Jacobian computed from SLR

$$J(\mu_i) = P_{xh(x)}^T P_j^{-1} \tag{S-10}$$

is the Jacobian of the expected value of the measurement i.e. $J(\mu_i) = \frac{d\hat{y}(\mu_i)}{d\mu_i}$. Using this we see that the left hand side of (S-6) is the derivative of the function:

$$q(\mu_i) = \frac{1}{2}(\hat{y}(\mu_i) - y)^T(R + \Omega_j)^{-1}(\hat{y}(\mu_i) - y) + \frac{1}{2}(\mu_i - \mu^-)^T(P^-)^{-1}(\mu_i - \mu^-), \tag{S-11}$$

which has a local extremum when (S-6) holds. Thus, (S-11) can be used as the cost function of the inner loop.

C. Implementation considerations

1) *Implementation with sigma-point filters:* Integrals (2)-(4) do not have closed form solutions in general. Sigma-point filtering is a common framework for approximating these integrals [1]. The formulas are

$$\hat{y} \approx \sum_{k=1}^m w_k h(\mu + \Delta_k) \tag{S-12}$$

$$P_{h(x)x} \approx \sum_{k=1}^m w_k (h(\mu + \Delta_k) - \hat{y}) \Delta_k^T \tag{S-13}$$

$$P_{h(x)h(x)} \approx \sum_{k=1}^m w_k (h(\mu + \Delta_k) - \hat{y}) (h(\mu + \Delta_k) - \hat{y})^T, \tag{S-14}$$

where w_k is the weight of k th sigma-point, m is the number of sigma points, and Δ_k is the translation of the k th sigma-point from the state mean. The translation vectors Δ_k depend on the state dimension and the state covariance matrix P and their computation involves taking a matrix square root (e.g. Cholesky decomposition) of P . In Algorithm 1 the state covariance matrix is constant in the inner loop and so the Δ_k vectors need to be recomputed only in the outer loop.

Fig. S-3. Effect of parameters τ and β on the mean KLD and computing time

2) *Efficient computation of the Jacobian*: The updated state covariance (17) can be written in information form using the matrix inversion lemma:

$$P_{j+1} = \left(J_i^T (R + \Omega_j)^{-1} J_i + (P^-)^{-1} \right)^{-1}. \quad (\text{S-15})$$

Using this in (10) we get

$$J_{i+1} = P_{h(x)x} \left(J_i^T (R + \Omega_j)^{-1} J_i + (P^-)^{-1} \right). \quad (\text{S-16})$$

Since P^- does not change between iterations, $(P^-)^{-1}$ needs to be computed only once for the inner loop and matrices that have to be inverted at every iteration of the inner loop have dimension $n_y \times n_y$. Thus, when $n_y \ll n_x$ (S-16) is faster than (10).

3) *Selection of parameters τ and β* : The parameters τ and β govern the behavior of inner loop of the DIPLF. Setting a low β makes the algorithm exit faster from the inner loop. Using a large τ changes the value of α less and makes the algorithm slower, while setting a low τ reduces the step length fast. Figure S-3 shows the effect of the parameters τ and β on the KLD and on the computing time in our second test in Section IV. Figure S-3 shows that using $\tau < 0.12$ reduces the accuracy, unless β is close to 1. When τ is larger than 0.12 there is not much difference in the accuracy, but the runtime increases fast. The parameter β has a smaller effect on the runtime than τ and for the accuracy the optimal value is near 0.9. Naturally, the given values are valid only for the example in this paper.

REFERENCES

- [1] S. Särkkä, *Bayesian filtering and smoothing*. Cambridge University Press, 2013, vol. 3. [Online]. Available: http://users.aalto.fi/~ssarkka/pub/cup_book_online_20131111.pdf
- [2] B. Bell and F. Cathey, "The iterated Kalman filter update as a Gauss-Newton method," *Automatic Control, IEEE Transactions on*, vol. 38, no. 2, pp. 294–297, 1993, doi:10.1109/9.250476.