An odd behaviour of the particle filter

Jakob Åslund
Department of Electrical Engineering
Linköping University, Sweden
jakob.aslund@liu.se

Fredrik Gustafsson
Department of Electrical Engineering
Linköping University, Sweden
fredrik.gustafsson@liu.se

Gustaf Hendeby
Department of Electrical Engineering
Linköping University, Sweden
gustaf.hendeby@liu.se

I. Abstract

This work is about an odd behavior which can occur in the particle filter (PF). Typically, the more particles are used in the filter, the better the results are, and as the number of particles approaches infinity, the results become optimal. We are instead interested in a case where the error at first increases, meaning that more computation is used to get a worse result.

II. Introduction

The PF is a useful filter for estimating the state in many nonlinear systems of the type

$$x_{k+1} = f(x_k) + Gw_k \tag{1a}$$
$$y_k = h(x_k) + v_k, \tag{1b}$$

where $x_k \in \mathbb{R}^{n_x}$, $y_k \in \mathbb{R}^{n_y}$, f and h are functions $\mathbb{R}^{n_x} \to \mathbb{R}^{n_x}$ and $\mathbb{R}^{n_x} \to \mathbb{R}^{n_y}$, respectively, and w_k and v_k are noises of dimensions n_x and n_y, respectively, and with known probability distributions [1], [2].

The filter works by using N samples, referred to as particles and sequentially updating them based on a proposal, and a weighting function. For this work, the bootstrap particle filter is analyzed. In order to keep the particles in the region of interest, the particles are also resampled at every iteration (with probability of selection proportional to their weights).

There have been plenty of surveys and tutorials published on the PF, and they all state that it converges to the true posterior distribution as N goes to infinity [3], [4]. Less is known about the results for lower numbers of particles. Based on the convergence at infinite particles and on the fact that the method is based around approximating the probability distribution with a number of samples, the expectation would be that the more particles one uses, the better results one can expect. This is also common knowledge among practitioners.

This work is instead interested in a category of systems for which this does not hold and the error of the estimate instead initially becomes worse as more particles are added.

III. Illustrative example

To highlight this, we focus on two systems

System 1:

$$x_{k+1} = \underbrace{\begin{pmatrix} 0.6 & -1 \\ 0 & 0.5 \end{pmatrix}}_{F_1} x_k + \underbrace{\begin{pmatrix} 1 \\ 1 \end{pmatrix}}_{G} w_k \tag{2a}$$
$$y_k = \underbrace{\begin{pmatrix} 1 & 0 \end{pmatrix}}_{H} x_k + v_k, \tag{2b}$$

System 2:

$$x_{k+1} = \underbrace{\begin{pmatrix} 0.6 & 1 \\ 0 & 0.5 \end{pmatrix}}_{F_2} x_k + \underbrace{\begin{pmatrix} 1 \\ 1 \end{pmatrix}}_{G} w_k \tag{3a}$$
$$y_k = \underbrace{\begin{pmatrix} 1 & 0 \end{pmatrix}}_{H} x_k + v_k, \tag{3b}$$

where w_k and v_k are assumed Gaussian with zero mean and variance 1 and 0.2 respectively. Both systems are stable, and even have the same poles 0.6 and 0.5. A practitioner might notice that the measurement noise is low which is generally not good for the standard PF. Further, the process noise only affects part of the state. While this might cause issues, there are so far no reasons to believe that the above understanding of the PF would break.

Simulating these systems for different numbers of particles results in Fig. 1. While the results for system 2 show the expected behavior, with the error quickly decreasing and approaching the optimal estimate, here given by the Kalman filter, the results for system 1 show an entirely different behavior. There, we can see that the error markedly increases for low numbers of particles before finally decreasing and approaching the optimal estimate at about 10000 particles. Note that this system also requires a significantly larger amount of particles to get a good estimate. These differences can be described by what is here referred to as a projected instability.

IV. Projected instability

Assume a system of the form (1). Next, assume there is state noise $\breve{w}_k^{(i)}$ that propagates the current particle $\breve{x}_k^{(i)}$ to a new particle

$$\breve{x}_{k+1}^{(i)} = f(\breve{x}_k^{(i)}) + G\breve{w}_k^{(i)}, \tag{4}$$

that gives a perfect fit to the measurement

$$y_{k+1} = h(\breve{x}_{k+1}^{(i)}), \tag{5}$$

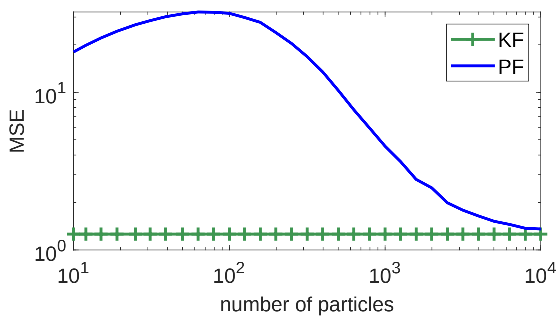

(a) MSE for System 1 in (2).

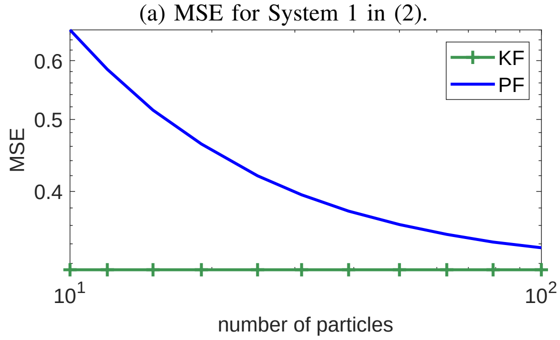

(b) MSE for System 2 in (3).

Fig. 1: A comparison of the MSE for the two studied systems, averaged over x_1 and x_2. The green line is the KF MSE.

or, if no such solution exists, one for which the norm of the difference is as small as possible.

The system contains a projected instability if during this process some part of the state will increase over time. For the linear case

$$x_{k+1} = Fx_k + Gw_k$$
$$y_k = Hx_k + v_k,$$

this can be relatively easily computed as

$$\bar{F} = F - G(HG)^{\dagger}HF,$$

containing eigenvalues outside the unit circle. Looking again at the example systems, we can see that for system 1 $\bar{F}$ has an eigenvalue in 1.5, meaning that it has a projected instability, while for system 2, the eigenvalues of $\bar{F}$ are -0.5 and 0, meaning that it does not contain a projected instability.

V. Simulation study

To make the connection more clear between the projected instability and the MSE increasing when adding more particles, 1000 systems with maximum eigenvalues of $\bar{F}$ between 0.5 and 2 were randomized and categorized based on if the error increased or not. The result can be seen in Fig. 2. Note that the red dots represent systems where the particle filter failed completely, meaning that no particles were even close to the measurements (all weights became 0). The y-axis displays the magnitude of the process noise (the covariance of w_k).

First off, it is clear that, independently of the size of Q, no PF displays the behavior or diverges for a system not containing a projected instability. Secondly, while this

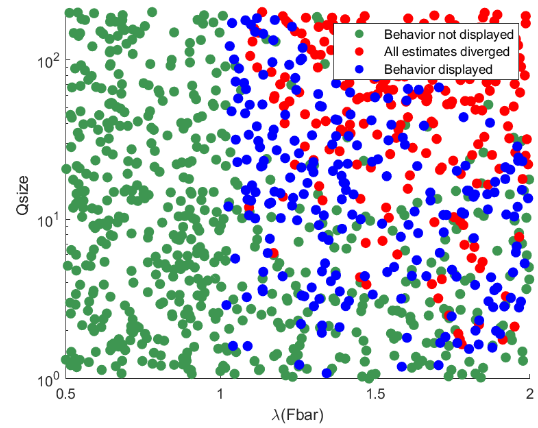

Fig. 2: The PF behaviors for various simulated systems

behavior is most commonly found in systems with a large process noise, it is also possible for it to occur in systems with a more normal process noise. Finally, for systems where $\lambda(\bar{F}) > 1$, increasing the size of Q or further increasing the eigenvalue of $\bar{F}$ results in first the studied behavior to increase in likelihood, and then in the likelihood of the filter diverging completely to increase.

VI. Connection to zeroes

As a final point, it can be shown that the eigenvalues of $\bar{F}$ are the same as the zeros of the transfer function from w_{k-1} to y_k, when the difference of the degree of the numerator and the denominator of the transfer function is 1. With some minor changes, similar connections can be done between a projected instability and the zeroes of the transfer function when the difference in degree is larger than 1. This means that the projected instability can also be interpreted as a non-minimum-phase problem.

VII. Conclusion

We have here shown that a system containing a projected instability (or equivalently, a non-minimum phase) can cause issues for the PF, and can even result in the error increasing as the number of particles increases. A second way of viewing this is that if a system contains a projected instability, it is likely that significantly more particles will be needed to get a good estimate, as that is the remedy to both the error increasing (for system 1, about 10000 particles are needed), or for the filter diverging entirely.

References

[1] J. S. Liu and R. Chen, "Sequential Monte Carlo methods for dynamic systems," *Journal of the American Statistical Association*, vol. 93, no. 443, pp. 1032–1044, 1998.

[2] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, 2002.

[3] A. Doucet and A. Johansen, "A tutorial on particle filtering and smoothing: Fifteen years later," *Handbook of Nonlinear Filtering*, vol. 12, 01 2009.

[4] F. Gustafsson, "Particle filter theory and practice with positioning applications," *IEEE Aerospace and Electronic Systems Magazine*, vol. 25, no. 7, pp. 53–82, 2010.