

OPTIMAL MEASUREMENT BUDGET ALLOCATION FOR PARTICLE FILTERING

Antoine ASPEEL, Amaury GOUVERNEUR, Raphaël M. JUNGERS and Benoît MACQ

ICTEAM Institute, UCLouvain, Avenue Georges Lemaître 4-6, Louvain-la-Neuve, Belgium

ABSTRACT

Particle filtering is a powerful tool for target tracking. When the budget for observations is restricted, it is necessary to reduce the measurements to a limited amount of samples carefully selected. A discrete stochastic nonlinear dynamical system is studied over a finite time horizon. The problem of selecting the optimal measurement times for particle filtering is formalized as a combinatorial optimization problem. We propose an approximated solution based on the nesting of a genetic algorithm, a Monte Carlo algorithm and a particle filter. Firstly, an example demonstrates that the genetic algorithm outperforms a random trial optimization. Then, the interest of non-regular measurements versus measurements performed at regular time intervals is illustrated and the efficiency of our proposed solution is quantified: better filtering performances are obtained in 87.5% of the cases and on average, the relative improvement is 27.7%.

Index Terms— Optimal measurement times, Particle filtering, Sequential Monte Carlo methods, Sparse measurements, Genetic algorithm.

1. INTRODUCTION

Stochastic nonlinear dynamical systems have shown their ability to model a number of real-world problems [1]. Particle filtering is an efficient approach to estimate the state of such systems from a set of noisy measurements [2]. This tool has been largely used, among others, in computer vision [3, 4, 5]. In practice, performing measurements may be difficult due to energy consumption, economical constraints or health hazards. For instance, in tumor tracking based on X-ray images, the number of X-ray acquisitions should be minimized in order to limit patients' exposure to harmful radiations [6].

Under such constraints, the problem is then to select the best moments to measure the system *a priori*, i.e. before any measurement acquisition. In other words, one has a measurement budget and has to choose when to acquire measurements. The optimality criterion is to minimize the expected

filtering *mean squared error* (MSE) over the complete time horizon.

In the particular case of linear systems subject to Gaussian noise processes, the selection of optimal measurement times over a finite time horizon has been studied using the Kalman filtering framework, in both discrete [7] and continuous-time [8, 9] settings. However, more general formulations have received little attention in the literature. This paper addresses the problem of providing optimal measurement times in the discrete-time nonlinear case with perturbation and measurement noise processes following arbitrary distributions. Our approach relies on particle filtering, the efficiency of which has been widely demonstrated in nonlinear dynamical systems [10, 11, 12].

The two main contributions of this paper are (i) to propose an efficient algorithm to solve the problem of optimal measurement times selection and (ii) to show the interest of non-regular measurements in particle filtering.

This paper is organised as follows, section 2 presents how to implement particle filtering with intermittent measurements (subsection 2.1); the criterion to select a good set of measurement times (subsection 2.2) and how to compute them (subsections 2.3 and 2.4). An example is presented and discussed in section 3. Finally, section 4 concludes and discusses possible improvements and perspectives.

A MATLAB (MathWorks, Natick, Massachusetts, USA) implementation of all the presented algorithms and the code that generate all figures are available on GITHUB at [GITHUB.COM/AMAURYGOUVERNEUR/OPTIMAL_MEASUREMENT_BUDGET_ALLOCATION_FOR_PARTICLE_FILTERING](https://github.com/AMAURYGOUVERNEUR/OPTIMAL_MEASUREMENT_BUDGET_ALLOCATION_FOR_PARTICLE_FILTERING).

2. MATERIALS AND METHODS

2.1. Intermittent particle filter

A discrete stochastic nonlinear dynamical system describes the evolution of a state $x(t)$ over the finite time horizon $t = 0, \dots, T$. One wants to estimate a quantity $z(t)$ related to $x(t)$ and has access to previously acquired noisy measurements $y(t)$ of $x(t)$. Measurements are not available at each time step. More formally, a measurement $y(t)$ is only available for $t \in \mathcal{M}$, where $\mathcal{M} \subseteq \{0, \dots, T\}$ is of size N , i.e.

A.A and A.G contributed equally. A.A. is supported by the Walloon Region, its grant is RW-DGO6-Biowin-Bidmed. R.J. is a FNRS Research Associate. He is supported by the Walloon Region and the Innoviris Foundation.

$|\mathcal{M}| = N$. This is modelled as

$$x(t+1) = f_t(x(t), w(t)) \text{ for } t = 0, \dots, T-1, \quad (1)$$

$$y(t) = g_t(x(t), v(t)) \text{ for } t \in \mathcal{M}, \quad (2)$$

$$z(t) = h_t(x(t)) \text{ for } t = 0, \dots, T, \quad (3)$$

$$x(0) \sim \mathcal{F}, \quad (4)$$

where $x(t) \in \mathbb{R}^n$, $y(t) \in \mathbb{R}^m$ and $z(t) \in \mathbb{R}^p$. In addition, $w(t)$ and $v(t)$ are random processes with known probability density functions. Functions $f_t(\cdot, \cdot)$, $g_t(\cdot, \cdot)$ and $h_t(\cdot)$ are known and have compatible dimensions. The initial state $x(0)$ follows a known distribution \mathcal{F} .

For instance, in a tumor tracking problem based on X-ray images, $x(t) \in \mathbb{R}^6$ can be a state vector containing the tumor's position and velocity in the 3-dimensional space, $y(t) \in \mathbb{R}^2$ can be the 2-dimensional projection of the target and $z(t) \in \mathbb{R}^3$ the position of the mass center in the 3-dimensional space.

An estimate $\hat{z}_{\mathcal{M}}(t) = PF[\{y(\tau)\}_{\tau \leq t, \tau \in \mathcal{M}}]$ of $z(t)$ based on previous measurements $\{y(\tau)\}_{\tau \leq t, \tau \in \mathcal{M}}$, can be computed by a particle filter $PF[\cdot]$. It is the expectation of the estimated probability density function.

Essentially, a particle filter algorithm alternates between (i) a *prediction step* (also called mutation), used to estimate the state at the next time step from the estimate at the current step; and (ii) a *correction step* (also called selection) that updates the state estimation to incorporate the information acquired in the last measurement. To deal with intermittent measurements, the correction step (ii) is skipped when no measurement is available, i.e. when $t \notin \mathcal{M}$.

In this paper, we use the sampling importance resampling particle filter (see Algorithm 4 in [2]).

2.2. Optimal intermittent particle filter

The optimal intermittent particle filter is the particle filter for which the set of N measurement times \mathcal{M} minimizes the expected filtering MSE. This is formalized as

$$\min_{\mathcal{M} \subseteq \{0, \dots, T\}} \mathbb{E}_{\text{MSE}}[\mathcal{M}] := \mathbb{E} \left[\frac{1}{T+1} \sum_{t=0}^T \|z(t) - \hat{z}_{\mathcal{M}}(t)\|^2 \right] \quad (5)$$

subject to $|\mathcal{M}| = N$ and equations (1) to (4),

where $\hat{z}_{\mathcal{M}}(t)$ is obtained from the particle filter and $\|\cdot\|$ is the Euclidean norm (note that it could be any other norm). The expectation is on the random variables $x(t)$ and $v(t)$. The dependency of the cost function $\mathbb{E}_{\text{MSE}}[\mathcal{M}]$ on these two quantities can be expressed explicitly using equations (2)-(3), it gives

$$\mathbb{E} \left[\frac{1}{T+1} \sum_{t=0}^T \|h_t(x(t)) - PF[\{g_\tau(x(\tau), v(\tau))\}_{\tau \in \mathcal{M}, \tau \leq t}]\|^2 \right].$$

Solving problem (5) yields the best measurement times *a priori*, i.e. before any measurement acquisition.

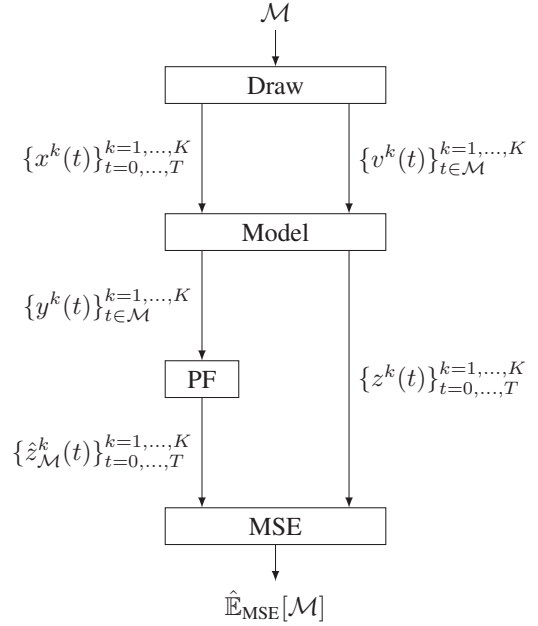


Fig. 1. Representation of the Monte Carlo algorithm which computes $\mathcal{M} \mapsto \hat{\mathbb{E}}_{\text{MSE}}[\mathcal{M}]$. The outputs generated by the Draw block are drawn according to (1), (4) and the distribution of $v(t)$. The model block uses respectively equations (2) and (3) to compute the $y^k(t)$ and $z^k(t)$ that correspond to the $x^k(t)$ and the $v^k(t)$ received in inputs. The PF block represents a particle filter that computes an estimate $\hat{z}_{\mathcal{M}}^k(t)$ of $z^k(t)$ from previous intermittent measurements $\{y^k(\tau)\}_{\tau \leq t, \tau \in \mathcal{M}}$. The MSE block computes the mean squared of the difference between the inputs, it is $\hat{\mathbb{E}}_{\text{MSE}}[\mathcal{M}]$ defined in equation (6).

2.3. Monte Carlo algorithm

A first challenge to solve problem (5) is to compute the expectation. To tackle this problem, we estimate this expectation using a Monte Carlo approach. For a given set of measurement times \mathcal{M} , one can simulate K realisations $\{x^k(t)\}_{t=0, \dots, T}^{k=1, \dots, K}$, $\{y^k(t)\}_{t \in \mathcal{M}}^{k=1, \dots, K}$ and $\{z^k(t)\}_{t=0, \dots, T}^{k=1, \dots, K}$ drawn according to (1)-(4). From each simulated sequence of measurements $\{y^k(t)\}_{t \in \mathcal{M}}$, the particle filter computes estimates $\{\hat{z}_{\mathcal{M}}^k(t)\}_{t=0, \dots, T} := \{PF[\{y^k(\tau)\}_{\tau \leq t, \tau \in \mathcal{M}}]\}_{t=0, \dots, T}$ of $\{z^k(t)\}_{t=0, \dots, T}$. These quantities are used to estimate the expectation in problem (5). The Monte Carlo estimator of problem (5), denoted by $\hat{\mathbb{E}}_{\text{MSE}}[\cdot]$, is given by

$$\mathbb{E}_{\text{MSE}}[\mathcal{M}] = \mathbb{E} \left[\frac{1}{T+1} \sum_{t=0}^T \|z(t) - \hat{z}_{\mathcal{M}}(t)\|^2 \right] \approx \frac{1}{K(T+1)} \sum_{k=1}^K \sum_{t=0}^T \|z^k(t) - \hat{z}_{\mathcal{M}}^k(t)\|^2 =: \hat{\mathbb{E}}_{\text{MSE}}[\mathcal{M}]. \quad (6)$$

With this notation, problem (5) is approximately equiva-




Fig. 2. Evolution of the average and minimum cost $\hat{\mathbb{E}}_{\text{MSE}}$ with respect to the number of cost function evaluations for both the genetic algorithm (GA) and the random trials optimizer (RT). It illustrates the quality of minimization for given computational resources. One generation of the genetic algorithm corresponds to 50 evaluations.

lent to

$$\min_{\mathcal{M} \subseteq \{0, \dots, T\}} \hat{\mathbb{E}}_{\text{MSE}}[\mathcal{M}] \text{ subject to } |\mathcal{M}| = N. \quad (7)$$

The fact that we only have access to an approximation of the objective function requires a particular attention during the optimization. The nesting of the Monte Carlo algorithm and the particle filter is represented in figure 1.

2.4. Genetic algorithm

Solving the combinatorial optimization problem (7) corresponds to finding the set $\mathcal{M} \subseteq \{0, \dots, T\}$ of cardinality N that minimizes $\hat{\mathbb{E}}_{\text{MSE}}[\mathcal{M}]$. An exhaustive search would require to test all admissible \mathcal{M} which represents $\frac{(T+1)!}{(T+1-N)!N!}$ possibilities. It is computationally intractable for large N and T .

A genetic algorithm [13] is used to find an approximate solution of problem (7). Genetic algorithms are generally used for unconstrained optimization. In our case, to deal with constraint $|\mathcal{M}| = N$, a count preserving crossover is implemented [14].

If measurement times are widely spaced for an individual of the genetic algorithm, it can happen that all particles of the particle filter have zero weights (due to degeneracy problem [2]). In such a case, the corresponding individual is killed and will not be used for next generations.

As mentioned in the previous section, the objective function $\mathbb{E}_{\text{MSE}}[\mathcal{M}]$ can only be evaluated approximately as $\hat{\mathbb{E}}_{\text{MSE}}[\mathcal{M}]$. It makes the minimization difficult if too sensitive to bad cost function estimations. To face this issue, our

genetic algorithm returns the best individual of the last generation instead of the best individual among all generations.

In Section 3, the performance of the genetic algorithm is compared to a random trial optimizer. It will show that the ability of evolutionary algorithms to broadly sample a population makes the genetic algorithm better suited for the addressed problem.

Our implementation of the genetic algorithm uses stochastic universal sampling and sigma scaling with unitary sigma coefficient [13]. Crossover probability is 1 and mutation probability per gene is 0.003.

3. RESULTS AND DISCUSSION

In order to illustrate the performances obtained by our approach, the following commonly studied model is used [2, 15, 16, 17],

$$x(t+1) = \frac{x(t)}{2} + \frac{25x(t)}{1+x(t)^2} + 8 \cos(1.2t) + w(t), \quad (8)$$

$$y(t) = \frac{x(t)^2}{20} + v(t) \text{ for } t \in \mathcal{M}, \quad (9)$$

$$z(t) = x(t), \quad (10)$$

$$x(0) \sim \mathcal{N}(0, 5^2), \quad (11)$$

where $w(t) \sim \mathcal{N}(0, 1)$ and $v(t) \sim \mathcal{N}(0, (\sin(0.25t) + 2)^2)$ are zero mean independent Gaussian noise processes. Equations (8) and (10) hold for $t = 0, \dots, T-1$ and $t = 0, \dots, T$, respectively.

The problem is to find the best set \mathcal{M} of $N = 21$ measurement times in the range of 0 to $T = 60$ such that $\hat{\mathbb{E}}_{\text{MSE}}[\mathcal{M}]$ is minimized. Even for a problem this size, an exhaustive search would require to test over $1.2 \cdot 10^{15}$ admissible \mathcal{M} .

In the following, the particle filter uses 500 particles, the Monte Carlo algorithm uses $K = 1,000$ draws, and the population size and the number of generations of the genetic algorithm are respectively 50 and 25.

Figure 2 illustrates the evolution of the average and minimum $\hat{\mathbb{E}}_{\text{MSE}}$ with respect to the number of cost function evaluations (one generation of the genetic algorithm corresponds to 50 evaluations, i.e. the population size). The average and minimum $\hat{\mathbb{E}}_{\text{MSE}}$ decrease over generations until reaching quasi-convergence. As the genetic algorithm has reached a quasi-convergence state, most of the individuals are identical, which shows good convergence behaviour. The set of measurement times returned by the genetic algorithm is denoted \mathcal{M}_{GA} .

In addition, the genetic algorithm is compared to a random trial optimization method. It samples measurement times randomly and evaluates their corresponding costs. Figure 2 indicates the average and minimum costs of the random trials with respect to the number of cost function evaluations, i.e. the number of trials. One can observe that for a same computational cost, i.e. a same number of cost function evaluations,




Fig. 3. Comparison of true value $z(t)$ with the values $\hat{z}_{\mathcal{M}}(t)$ filtered by the regular particle filter (RPF) and the optimal intermittent particle filter (IPF) over a single draw. The relative gain on the complete sequence is $g = 28.3\%$. Optimal measurement times \mathcal{M}_{GA} are indicated with red stars (*) and regularly spaced measurement times \mathcal{M}_{REG} are indicated with blue plus (+). Results simulated from model (8)-(11).

our genetic algorithm significantly outperforms the random trial optimizer.

Now that the genetic algorithm performance has been demonstrated, our *optimal intermittent particle filter* (IPF) method is compared with a *regular particle filter* (RPF). It is a particle filter with regularly spaced measurement times, defined by

$$\mathcal{M}_{REG} := \left\{ \text{Round} \left[\frac{kT}{N-1} \right] \mid k = 0, \dots, N-1 \right\}, \quad (12)$$

where $\text{Round}[\cdot]$ is the rounding operator.

For a given measurement times set \mathcal{M} , one defines the random variable $\text{MSE}[\mathcal{M}] := \frac{1}{T+1} \sum_{t=0}^T \|z(t) - \hat{z}_{\mathcal{M}}(t)\|^2$. With this notation, one can define the relative gain, $g := \frac{\text{MSE}[\mathcal{M}_{REG}] - \text{MSE}[\mathcal{M}_{GA}]}{\text{MSE}[\mathcal{M}_{REG}]}$. This relative gain is positive when our IPF method outperforms the RPF.

Figure 3 compares the estimates produced by RPF and IPF to the exact $z(t)$ for one particular realization. Optimal measurement times \mathcal{M}_{GA} and regularly spaced measurement times \mathcal{M}_{REG} are indicated with red stars (*) and blue plus (+), respectively. One can observe better filtering performances using IPF. Quantitatively, the relative gain obtained using IPF instead of RPF is $g = 28.3\%$. Due to the nonlinearities of the problem, the obtained measurement times can not be easily interpreted.

The histogram presented in figure 4 approximates the probability density function of the relative gain g . It is obtained by running on 100,000 draws both the IPF and the RPF and computing the corresponding gain g . The mean relative




Fig. 4. Histogram of the relative gain $g = \frac{\text{MSE}[\mathcal{M}_{REG}] - \text{MSE}[\mathcal{M}_{GA}]}{\text{MSE}[\mathcal{M}_{REG}]}$. It has been obtained by running both the regular particle filter and the intermittent particle filter on the same 100,000 draws. The vertical red line corresponds to a null gain. The average is 27.7% and g is positive in 87.5% of the cases.

gain is 27.7% and in 87.5% of the cases, our IPF method outperformed the RPF.

4. CONCLUSION

The problem of selecting optimal measurement times for particle filtering over a finite time horizon was presented. Then, an algorithm nesting a genetic algorithm, a Monte Carlo method and a particle filter was proposed to find these optimal measurement times.

Firstly, a numerical example demonstrated that our genetic algorithm significantly outperforms a random trial optimizer. Then, we demonstrated that in comparison to regularly spaced measurements, the optimal choice of intermittent measurement times led to better filtering performance in 87.5% of the cases. On average, the relative gain was 27.7%.

Further work will consider selecting the measurement times *online* instead of fixing them *a priori*: after measurements have been acquired, one can recompute the next optimal measurement times, thereby incorporating all the then-available information. Extensions to continuous-time systems with discrete measurements will also be investigated. Finally, robustness analyses will be performed to determine how model uncertainties affect the performance of the proposed method, for situations in which the dynamics of the system are only approximately known.

Overall, our results demonstrate the added value of *a priori* selecting measurement times based on system dynamics for the optimal estimation of a state from limited measurements.

5. REFERENCES

- [1] Ludwig Arnold and René Lefever, *Stochastic Nonlinear Systems in Physics, Chemistry, and Biology: Proceedings of the Workshop Bielefeld, Fed. Rep. of Germany, October 5–11, 1980*, vol. 8, Springer Science & Business Media, 2012.
- [2] M Sanjeev Arulampalam, Simon Maskell, Neil Gordon, and Tim Clapp, “A tutorial on particle filters for on-line nonlinear/non-gaussian bayesian tracking,” *IEEE Transactions on signal processing*, vol. 50, no. 2, pp. 174–188, 2002.
- [3] Carine Hue, J-P Le Cadre, and Patrick Pérez, “A particle filter to track multiple objects,” in *Proceedings 2001 IEEE Workshop on Multi-Object Tracking*. IEEE, 2001, pp. 61–68.
- [4] Jung Uk Cho, Seung Hun Jin, Xuan Dai Pham, Jae Wook Jeon, Jong Eun Byun, and Hoon Kang, “A real-time object tracking system using a particle filter,” in *2006 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2006, pp. 2822–2827.
- [5] Henrik Andreasson, André Treptow, and Tom Duckett, “Localization for mobile robots using panoramic vision, local features and particle filter,” in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*. IEEE, 2005, pp. 3348–3353.
- [6] Gregory C Sharp, Steve B Jiang, Shinichi Shimizu, and Hiroki Shirato, “Prediction of respiratory tumour motion for real-time image-guided radiotherapy,” *Physics in Medicine & Biology*, vol. 49, no. 3, pp. 425, 2004.
- [7] Antoine Aspeel, Damien Dasnoy, Raphaël M Jungers, and Benoît Macq, “Optimal intermittent measurements for tumor tracking in x-ray guided radiotherapy,” *Medical Imaging 2019: Image-Guided Procedures, Robotic Interventions, and Modeling*, vol. 10951, pp. 109510C, 2019.
- [8] A. Sano and M. Terao, “Measurement optimization in optimal process control,” *Automatica*, vol. 6, no. 5, pp. 705–714, 1970.
- [9] Alexandre Aksenov, Pierre-Olivier Amblard, Olivier Michel, and Christian Jutten, “Optimal measurement times for a small number of measures of a brownian motion over a finite period,” *arXiv preprint arXiv:1902.06126*, 2019.
- [10] Sebastian Thrun, “Particle filters in robotics,” in *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 2002, pp. 511–518.
- [11] Fredrik Gustafsson, Fredrik Gunnarsson, Niclas Bergman, Urban Forssell, Jonas Jansson, Rickard Karlsson, and P-J Nordlund, “Particle filters for positioning, navigation, and tracking,” *IEEE Transactions on signal processing*, vol. 50, no. 2, pp. 425–437, 2002.
- [12] AA Abd Rahni, E Lewis, MJ Guy, B Goswami, and K Wells, “A particle filter approach to respiratory motion estimation in nuclear medicine imaging,” *IEEE Transactions on Nuclear Science*, vol. 58, no. 5, pp. 2276–2285, 2011.
- [13] Melanie Mitchell, *An introduction to genetic algorithms*, MIT press, 1998.
- [14] AJ Umbarkar and PD Sheth, “Crossover operators in genetic algorithms: A review.,” *ICTACT journal on soft computing*, vol. 6, no. 1, 2015.
- [15] Bradley P. Carlin, Nicholas G Polson, and David S Stoffer, “A monte carlo approach to nonnormal and non-linear state-space modeling,” *Journal of the American Statistical Association*, vol. 87, no. 418, pp. 493–500, 1992.
- [16] Genshiro Kitagawa, “Monte carlo filter and smoother for non-gaussian nonlinear state space models,” *Journal of computational and graphical statistics*, vol. 5, no. 1, pp. 1–25, 1996.
- [17] Visakan Kadirkamanathan, Ping Li, Mohamed H Jaward, and Simon G Fabri, “Particle filtering-based fault detection in non-linear stochastic systems,” *International Journal of Systems Science*, vol. 33, no. 4, pp. 259–265, 2002.